APRIL 2026

HOW RUSSIA IS BUILDING A SOVEREIGN DRONE ECOSYSTEM FOR AI-DRIVEN AUTONOMY

AUTHORS | Kateryna Bondar

A Report of the CSIS Wadhwani AI Center

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES | WADHWANI AI CENTER

About CSIS

The Center for Strategic and International Studies (CSIS) is a bipartisan, nonprofit policy research organization dedicated to advancing practical ideas to address the world's greatest challenges.

Thomas J. Pritzker was named chairman of the CSIS Board of Trustees in 2015, succeeding former U.S. Senator Sam Nunn (D-GA). Founded in 1962, CSIS is led by John J. Hamre, who has served as president and chief executive officer since 2000.

CSIS's purpose is to define the future of national security. We are guided by a distinct set of values—nonpartisanship, independent thought, innovative thinking, cross-disciplinary scholarship, integrity and professionalism, and talent development. CSIS's values work in concert toward the goal of making real-world impact.

CSIS scholars bring their policy expertise, judgment, and robust networks to their research, analysis, and recommendations. We organize conferences, publish, lecture, and make media appearances that aim to increase the knowledge, awareness, and salience of policy issues with relevant stakeholders and the interested public.

CSIS has impact when our research helps to inform the decisionmaking of key policymakers and the thinking of key influencers. We work toward a vision of a safer and more prosperous world.

Center for Strategic & International Studies
1616 Rhode Island Avenue, NW
Washington, DC 20036
202-887-0200 | www.csis.org

Table of Contents

APRIL 2026

How Russia Is Building a Sovereign Drone Ecosystem for AI-Driven Autonomy

By Kateryna Bondar

Executive Summary

This paper examines how Russia is developing military artificial intelligence (AI) and incrementally moving toward autonomous decisionmaking, particularly at the tactical edge. The takeaways below outline the core findings on how these capabilities are built, adapted, and scaled within Russia's wartime military ecosystem.

1. **Russia has identified unmanned systems and AI as two overarching strategic priorities across all levels of policymaking.** These priorities appear consistently in federal, regional, and sector-specific strategies and are most often framed in civilian and dual-use contexts. However, given Russia's transition to a wartime economy and the limited visibility into classified military programs, investments and progress in these areas are highly likely to translate directly into military capabilities and operational gains.
2. **Russia has likely fielded a fully autonomous unmanned system in combat and continues to iterate on its deployment despite resulting civilian casualties.** Ukrainian technical analysis of intercepted V2U drones indicates the absence of communication components required for operator control, alongside the presence of onboard computing sufficient to run AI-enabled perception and decisionmaking software. Observed battlefield behavior—including autonomous flight in denied environments, independent target selection, and coordinated group activity using visual markings for swarm-like coordination—suggests that V2U represents a qualitative departure from remotely piloted expendable drones toward fully autonomous, AI-driven systems.
3. **Russia's drone ecosystem reveals an adaptive procurement logic in which innovation originates outside formal defense industrial structures and is scaled only after**

battlefield validation. Projects such as Molniya demonstrate a recurring pattern: rapid experimentation by civilian engineers and volunteer groups at the "garage" level, followed by selective state intervention to finance, standardize, and mass-produce systems that prove operationally effective. This approach allows the state to capture the benefits of decentralized innovation while avoiding the inefficiencies of attempting to centrally design solutions under wartime pressure.

4. **One of the most critical enablers of unmanned systems integration has been the rise of private drone schools and parallel training initiatives, which operate as rapid accelerators of technological adoption.** Unlike traditional state-run training facilities, these organizations adapt with startup-like speed, continuously updating curricula, integrating new platforms directly into instruction, and allowing operators to test systems extensively during training. This structure creates direct feedback loops between end users and engineers, accelerating refinement of both hardware and tactics. By embedding new capabilities into training pipelines faster than formal institutions, these schools convert emerging technologies into operational competence at scale, effectively making training itself a central engine of combat power.
5. **More than 50 percent of all AI-enabling components recovered from Russian unmanned systems originate from companies headquartered in the United States and consist primarily of commercial-grade, dual-use electronics.** Across 705 identified AI-relevant components (e.g., spanning processors, memory units, and sensors) U.S. firms account for roughly 69 percent of memory hardware, 57 percent of processors, and 38 percent of sensors, representing the largest national share in each category. By comparison, China supplies less than 9 percent of total AI-enabling components and does not rank among the top providers of onboard computing hardware. These findings underscore that the technical backbone of Russia's expanding battlefield autonomy remains deeply embedded in globally integrated semiconductor markets, where commercially available Western technologies continue to play a decisive role despite sanctions and export controls.
6. **Russia is not competing with major powers in the race for frontier AI; instead, it is pursuing a pragmatic strategy focused on applied AI capabilities.** Rather than developing large foundational models from scratch, Russia focuses on building practical solutions on top of existing open-weight models from Western developers such as Llama and Mistral, as well as Chinese models such as Qwen and DeepSeek. These models are adapted into custom applications designed for both government-wide integration and military use.
7. **Russia is deliberately building a comprehensive, end-to-end ecosystem for AI and unmanned systems rather than pursuing isolated capabilities.** This effort integrates compute expansion to one exaflop by 2030, production targets of 130,000 large-scale unmanned aircraft systems (UASs) annually, rapid growth in AI markets and corporate investment, and a planned output of 15,500 AI specialists graduating each year by 2030. Anchored in national strategies and operationalized through state programs, the ecosystem links infrastructure, regulation, industry, and talent development into a unified system designed to sustain AI-enabled autonomy and military relevance over the long term.
8. **Russia is focusing on creating a dedicated infrastructure to enable civilian-operated unmanned aviation at national scale by 2030.** This includes the expansion of test ranges,

construction of new production facilities, and deployment of unified airspace integration and digital traffic management systems designed to support the safe, large-scale operation of UASs. Creation of such infrastructure will not only support civilian adoption but also serve as a critical enabler for the accelerated development, scaling, and operational integration of unmanned systems within the military domain.

9. **Russia expects demand for 1 million UAS specialists by 2030, making human capital a central pillar of its unmanned systems strategy.** To meet this scale, the state is expanding drone-focused education across schools, vocational pathways, and universities, while introducing unified competency standards and continuous training programs to keep skills aligned with industry and operational requirements.
10. **Russia is combining a deliberately soft approach to AI regulation with growing centralization of state control over its deployment through the creation of a National AI Headquarters and a presidential-level commission.** Rather than rushing formal legislation, the government has emphasized phased regulation, experimentation, and institutional learning, while relying on selective restrictions, certification of "trusted" technologies, and controlled access to state-managed data. At the same time, Moscow is moving to concentrate authority through the creation of a National AI Headquarters above individual ministries—designed to coordinate AI implementation across regions and sectors under a single, state-led command structure—alongside a Commission on the Development of Artificial Intelligence Technologies under the president.
11. **Russia's most successful integration of AI occurs within companies that operate across both civilian and military markets rather than within purely defense-oriented enterprises.** Dual-use firms can draw on far larger and more varied datasets, iterate software in real operational environments, and continuously retrain models based on civilian and security applications. This access to data, testing opportunities, and feedback loops allows AI capabilities to mature faster and transition more smoothly into battlefield use than systems developed exclusively inside closed military programs.
12. **Russian unmanned systems development is characterized by modularity and rapid functional adaptation rather than platform specialization.** Once a design proves viable, it is quickly repurposed across multiple roles—for example, as a loitering munition, reconnaissance platform, or logistics carrier—through minimal airframe changes and software updates. Simple construction and modular architecture allow fast iteration based on frontline feedback, accelerating the diffusion of successful designs across different mission sets.

For the United States, the central lesson is that success in AI-enabled unmanned systems requires an ecosystem approach. To advance its ambitions in autonomous technology, the United States must implement a national systems project approach that incorporates and aligns training, testing, dual-use innovation, government implementation, and civil-military cooperation.

Introduction

Four years into Russia's full-scale invasion of Ukraine, the war has revealed something that, until recently, remained largely theoretical—the emergence of fully autonomous weapons systems deployed in the battlefield. While early assessments of Russia's military performance showed institutional rigidity and technological underperformance, battlefield evidence now suggests a more complex picture. Under sustained pressure from electronic warfare (EW), GPS denial, and mass attrition, Russia is moving beyond remote-controlled unmanned systems and is fielding platforms capable of operating, navigating, and selecting targets without external communication, marking a qualitative shift in how autonomy is applied in combat.

> *The result is not comprehensive autonomy but functional independence at the tactical edge.*

This development does not reflect a breakthrough in frontier AI, nor the realization of long-promised kill-chain-wide autonomy architectures. Instead, Russia's progress has been driven by a pragmatic focus on applied AI, embedding narrowly defined machine learning functions directly into unmanned systems and battlefield software. Rather than competing with the United States or China in foundational AI research, Russian developers adapt existing Western and Chinese open-weight models and integrate them into domestic applications optimized for wartime conditions. The result is not comprehensive autonomy but functional independence at the tactical edge.

This report examines how Russia is integrating AI into its unmanned systems and what this process reveals about the evolving character of Russian military power. The central question is not whether Russia has achieved autonomy in a doctrinal sense, but how effectively it deploys limited AI capabilities that deliver operational advantage at scale.

The analysis is structured in three parts. The first section examines Russia's policy architecture for AI and unmanned systems, showing how presidential-level priorities translate into national programs, regulatory approaches, and sectoral initiatives. It highlights how a civilian innovation ecosystem—spanning regulation, industry, and workforce development—supports the expansion of military capabilities.

The second section presents a set of case studies that illustrate different models of AI development and deployment, ranging from centralized, state-led programs to commercially driven systems that scale through battlefield validation.

The third section analyzes three key factors enabling Russia to maintain speed and scale in innovation: (1) training as the primary channel for integration and force-wide adoption, (2) the origin of the hardware backbone underpinning AI-enabled systems, and (3) the role of international partnerships in sustaining access to critical technologies.

RESEARCH APPROACH AND SOURCES

This analysis is based exclusively on open-source research and does not rely on classified information. The research draws on four primary-source categories, which were systematically cross-referenced to assess both Russian intent and observed battlefield performance:

- **Official Policy Documents:** The first set of sources consists of official Russian strategic documents, action plans, and legislative frameworks. These materials make it possible to identify the formalized priorities, policy directions, and institutional mechanisms through which the Russian state articulates and operationalizes its technological objectives.

- **Media Reports and Statements:** The second set of sources includes official media reporting and public statements by senior Russian leadership, including President Vladimir Putin, ministers, and other senior officials. These communications demonstrate how the Kremlin frames technological priorities, signals shifts in strategic direction, and publicly communicates progress in AI and unmanned systems development.
- **Telegram Channels:** The third set of sources comprises systematic monitoring and analysis of more than 150 Russian Telegram channels, including closed and semi-closed groups associated with civilian engineers, volunteer technologists, and military-affiliated developers supporting the war effort. These communities provide granular, near-real-time visibility into how specific systems evolve, what technical challenges developers encounter, how they adapt to constraints such as EW and component shortages, and how effective solutions diffuse across units. This source base enables tracking not only innovation itself, but also the processes of scaling, adaptation, and institutionalization within the Russian military ecosystem.
- **Interviews:** The fourth source set of sources consists of interviews with Ukrainian military personnel. These interviews were used to cross-check open-source findings against observed frontline realities and to provide ground-truth assessments of how Russian unmanned systems perform in combat and how Russian tactics and technologies have evolved over time.

Please note, that some of the links referenced in this report may only be accessible through appropriate VPN services or from specific geographic locations.

To further validate the analysis, interviews also included foreign military experts specializing in the Russian Armed Forces, who helped verify technical interpretations and contextualize findings derived from Russian sources and battlefield reporting. By triangulating these sources, this analysis aims to provide a grounded, empirically anchored assessment of how AI and autonomy are being integrated into Russia's military systems under wartime conditions.

Russia's Policy Architecture for AI and Autonomous Systems

This section examines the architecture of Russia's strategic planning and the mechanisms through which innovation policy in AI and unmanned systems is formulated and implemented. For analytical clarity, Russia's policy planning and implementation are examined across **three interrelated layers**–strategic, tactical, and operational–as shown in Table 1. The analysis proceeds across each of these layers to identify how specific initiatives and institutional mechanisms support the advancement of Russia's wartime capabilities.

The assessment draws on official strategic documents and implementation frameworks to illuminate how declared priorities are translated into actionable programs and measurable outcomes.

In addition, this section provides an overview of AI-related regulation to clarify the Russian government's evolving approach to governance, experimentation, and control in the AI domain. The objective is to move beyond political rhetoric and evaluate the underlying system of planning, coordination, and state oversight that shapes Russia's approach to innovation under wartime conditions.

Table 1: Structure of Russia's Policy Development and Governance Architecture

Layer	Document type	Purpose
Strategic	• Presidential decrees • Doctrine-level policy concepts (sometimes introduced through Presidential speeches)	Define national priorities and long-term objectives, setting the overarching direction of state policy.
Tactical	• Sector-specific strategies • Defense and security long-term programs	Translate broad national goals into structured plans for individual sectors, with defined objectives and development pathways.
Operational	• National projects • Federal programs • Experimental legal regimes	Operationalize strategic and sectoral plans through funding, concrete initiatives, regulatory tools, and coordinated execution mechanisms.

Source: CSIS.

STRATEGIC LAYER

At the strategic level, Russian leadership defines National Development Goals—broad, long-term priorities that shape the country's overall trajectory. These goals are established through the highest-level policy instrument, a presidential decree, which sets the overarching direction for state policy across sectors. The decree articulates national development objectives and provides strategic guidance for implementation across all domains, including those that influence innovation and technological advancement relevant to the war effort.

The latest decree **On the National Development Goals of the Russian Federation for the Period up to 2030 and for the Perspective until 2036** was adopted on May 7, 2024. In this decree, the national goal called "Technological leadership" is defined through a set of measurable objectives and tasks that collectively reflect Russia's strategic priorities in science and innovation. Notably, the document directly identifies three technological streams—unmanned systems, autonomous vehicles, and AI—as particularly critical areas for achieving global competitiveness.

The decree sets ambitious quantitative benchmarks. By 2030, Russia aims to rank among the world's top 10 nations in research and development (R&D), to raise domestic R&D spending to at least 2 percent of GDP, and to double private-sector investment in innovation. Additionally, it emphasizes the growth of "small technology companies" (i.e., startups) as engines of innovation and promotes localization of high-tech production as a key pillar of national resilience across all development goals.

TACTICAL LAYER

The second layer is formed by the strategies that translate the National Development Goals into actionable priorities. At this level, two major documents stand out—the national strategies on AI and on unmanned systems. Both have a clear dual-use character, and both were recently updated, signaling that the Russian leadership is actively adjusting its innovation policy in response to rapid changes in these strategically important domains.

The **National Strategy for the Development of Artificial Intelligence** for the period up to 2030, approved in October 2019 and **updated in February 2024**, remains the cornerstone of Russia's long-term vision for AI. It defines AI as a key driver of economic growth, quality of life, and national security. The document mandates the integration of AI across all levels of governance and production, from federal ministries and state-owned enterprises to private industry, aiming to embed AI into the very architecture of the Russian state and economy.

In contrast to the United States and China, which are explicitly named in the strategy as the leading global players, Russia does not position itself as a competitor in frontier AI research. Acknowledging its limited access to advanced computing resources and international scientific cooperation, the strategy instead focuses on the applied, dual-use dimensions of AI. In practice, Russia seeks to leverage algorithms and models already developed abroad, integrating them into domestic applications across defense, security, and industrial automation.

The strategy is structured around a set of pillars that directly support the application layer of AI—the point where technologies transition from research to operational use. The overview of the following **core pillars** shows that, collectively, their interconnection enables large-scale deployment across the economy and state systems:

- **The infrastructure development** pillar underpins the entire system. Russia plans to expand its domestic computing capacity from 0.073 exaflop to 1 exaflop by 2030, ensuring technological sovereignty and continuity of AI model training under sanctions. This compute foundation will support both civilian and defense applications.
- **Support for AI developers** is designed to stimulate local innovation and commercialization. The state targets an AI services market of 60 billion rubles (~$760 million) annually by 2030—up from 12 billion rubles (~$150 million) in 2022—creating sustained demand for homegrown solutions integrated into industry and government systems.
- **Research and scientific advancement** connects infrastructure and industry through state-funded university centers. By 2030, Russian researchers are expected to produce 450 top-level conference papers and 450 journal publications per year, maintaining visibility and continuity of applied research despite international isolation.
- **Human capital development** ensures the diffusion of skills across the labor market; 15,500 AI specialists are expected to graduate annually by 2030 (up from 3,048 in 2022), and 80 percent of the workforce is to attain basic AI literacy, reflecting the state's intent to institutionalize AI competence across society.
- **Sectoral integration** operationalizes these layers. By 2030, 95 percent of priority industries are to achieve high readiness for AI adoption, with corporate investment rising from 123 billion to 850 billion rubles annually (from ~$1.5 billion to ~$11 billion).

This ecosystem builds the foundation for integration of AI across the Russian economy, linking compute power, education, applied research, and industrial deployment into a single complex. It represents a tightly interwoven system designed to scale AI implementation. Inevitably, the results of this approach are most visible in the military domain, where the practical orientation of Russia's strategy has already translated into tangible progress on the battlefield, rather than remaining confined to policy documents or strategic declarations.

Military AI has clearly emerged as a strategic priority for Russia, as reflected in President Putin's remarks at the **April 2025 meeting** of the Military-Industrial Commission. Framing AI as the defining factor in the future of Russian defense and weapons development, the Russian president stressed the priority of integration of domestically produced "protected" AI into automated command systems. This creates technological impetus to pursue broader reforms in production, doctrine, and training, illustrating how all national AI priorities converge in the defense domain.

Another crucial initiative for Russia's military effort is the new **Strategy for the Development of Unmanned Aviation**, which lays out an ambitious vision for building a sovereign, large-scale, and fully integrated UAS ecosystem by the early 2030s. Although this is still a draft updating the previous strategy, it already makes clear how Russian leadership intends to shape the sector. The document presents unmanned aviation as both a national security priority and a catalyst for economic modernization, outlining coordinated measures to move the field from a fragmented, import-dependent niche to a high-capacity domestic industry.

At the center of the strategy is a clear priority—Russia intends to replace foreign UASs, components, and software with its own systems. This drive for technological sovereignty runs through the entire document. The government plans to localize airframes, engines, electronics, flight controllers, payloads, navigation modules, and protected communications systems, while simultaneously creating a national certification regime tailored specifically to unmanned aircraft and AI-enabled autonomous systems. Certification is meant to ensure that domestically produced UASs meet standardized military and civilian requirements.

These structural reforms are paired with a strong push to expand domestic production capacity. By 2030, Russia plans to manufacture around 130,000 UASs, increasing to 350,000 by 2035. The market value of unmanned aviation is expected to surpass 145 billion rubles (~$1.9 billion) by 2030 and more than 350 billion rubles (~$4.6 billion) by 2035. The strategy envisions roughly 200 additional organizations entering UAS component production, building on the 220 already active in the sector, and aims for Russian companies to meet at least 75 percent of national UAS demand by the end of the decade.

To support these ambitions, the state aims to build the infrastructure needed for a national unmanned aviation ecosystem. This includes expanded test ranges, new production sites, unified airspace-integration tools, and digital traffic management systems that will allow UASs to operate safely at scale. It also includes investments in protected radio communications, interference-resistant navigation, and alternative Global Navigation Satellite System (GLONASS) solutions capable of functioning under electronic warfare conditions.

The strategy devotes considerable attention to human capital. Russia expects demand for UAS specialists to reach nearly 1 million people by 2030, with the majority trained as operators, technicians, and applied specialists and a minority as engineers and programmers. To meet this need, the government is expanding UAS-focused programs in schools, building vocational pathways, and integrating drone-related training into universities and technical institutes. Initiatives such as creating unified competency standards and continuous training programs aim to keep this workforce aligned with industry requirements.

Research and development priorities reflect both wartime urgency and long-term ambitions. The document prioritizes core R&D efforts on swarm control, autonomous navigation, multispectral computer vision, advanced propulsion, and resilient communications. The government plans to coordinate these efforts through joint programs linking industry, specialized research centers, and federal ministries.

Another important document is the **State Armament Program**. It is Russia's 10-year strategic plan that outlines how the country's armed forces will be technically modernized and re-equipped. It defines a list of new weapons systems to be developed, as well as existing ones that require modernization, based on current and anticipated national security threats.

The document is classified, making it difficult to determine the specific goals and technological priorities it outlines. However, based on **statements made in June 2025**, President Putin ordered that the new State Armament Program be explicitly oriented toward the large-scale integration of advanced technologies, particularly AI. He emphasized that future weapons systems and military equipment should incorporate cutting-edge digital technologies, AI applications, and weapons based on new physical principles, as well as ground and naval robotic complexes.

These strategic documents illustrate Russia's attempt to build a structured, sovereign ecosystem for UASs with an accelerated shift toward autonomy. The vision extends far beyond producing drones. Moscow aims to establish the industrial base, software infrastructure, regulatory frameworks, technology stacks, and human capital pipelines needed to sustain large-scale UAS and AI development, deployment, and innovation well into the 2030s.

OPERATIONAL LAYER

National projects function as one of the central operational tools for translating presidential development goals and sectoral strategies into **concrete, measurable action plans**. They break down broad strategic priorities into specific initiatives with defined budgets, timelines, performance indicators, and personal responsibility assignments. Each national project is overseen by a designated official who is **directly accountable** to the president, creating a clear chain of responsibility and a mechanism for top-down oversight.

Although several national projects have been reclassified or renamed as federal programs, the shift is mostly cosmetic, and the core governance architecture remains unchanged, as shown in Table 2. Despite the change in terminology, the underlying logic, structure, and purpose remain essentially the same, with both national projects and federal programs operationalizing strategic goals through targeted state investment, coordinated implementation, and strict monitoring of results.

Table 2: Russia's Policy Framework for AI and Unmanned Systems Innovation

Layer	Document and Description
Strategic	**Presidential Decree on the National Development Goals** The 2024 update places "technological leadership" at the center of national development, explicitly prioritizing unmanned systems, autonomous vehicles, and AI as key areas of global competitiveness.
Tactical	**National Strategy for the Development of Artificial Intelligence** A long-term national framework that embeds AI across Russia's state, economic, and industrial systems, prioritizing applied and dual-use technologies over frontier research. It expands domestic compute capacity, supports local AI developers, funds applied research centers, builds large-scale human-capital pipelines, and mandates the integration of AI into priority industries and public administration.
	National Strategy for the Development of Unmanned Aviation A strategic blueprint for creating a fully sovereign, large-scale unmanned aviation ecosystem by the 2030s, replacing foreign UASs and components with domestic production. The strategy outlines the localization of all critical UAS technologies, major expansion of testing and production infrastructure, development of protected communications and EW-resilient navigation, large-scale workforce training for up to 1 million specialists, and rapid growth of domestic manufacturing capacity.
	State Armament Program Russia's classified 10-year modernization plan for the armed forces, setting priorities for developing and upgrading weapons systems. It determines which capabilities will be funded and fielded, with recent directives from the president requiring deep integration of AI, robotics, and next-generation digital and autonomous weapons across future military platforms.
Operational	**National Project "Unmanned Aerial Systems"** A program that builds a full domestic drone ecosystem by 2030—covering design, testing, mass production, workforce training, and next-generation R&D—to replace foreign components and secure Russia's technological independence in unmanned aviation.
	National Project "Data Economy and Digital Transformation of the State" A program that aims to modernize Russia's governance, economy, and social systems through large-scale digitalization and the pursuit of technological sover eignty. • National Project "Data Economy and Digital Transformation of the State" *A targeted program under the Data Economy project that creates a national low-orbit satellite constellation to provide ubiquitous internet coverage and secure communications, including for unmanned and autonomous systems.* • Federal Program: "Artificial Intelligence" *A program focused on developing domestic AI solutions, integrating them into public administration and industry, expanding data-driven services, and cultivating nationwide AI competencies—including early education and workforce training.*

Source: CSIS.

Two specific national projects/federal programs intersect most directly with Russia's efforts to advance AI-enabled and unmanned systems, which constitute the central focus of this study:

1. The **Unmanned Aerial Systems National Project** was launched to secure Russia's technological independence and establish a competitive domestic drone industry across civilian and dual-use sectors. It is a cornerstone of the country's effort to achieve the "**technological leadership**" goal under the 2024 presidential decree and reflects Moscow's recognition of unmanned systems as a critical domain for industrial, military, and economic competitiveness.

   The project's overarching goal is to **build a full-cycle ecosystem** for the design, production, and application of UASs by 2030, with Russian-made drones expected to capture 70 percent of the national market. It consists of several interlinked components: workforce development programs to train engineers, operators, and software specialists; the creation of a standardized system for design, testing, and serial production through a nationwide network of 48 research and production centers; mechanisms to stimulate demand such as subsidies, state order, and leasing incentives; and the advancement of next-generation technologies in autonomy, navigation, communications, and materials.

   This project reveals crucial details about Russia's broader AI and autonomy strategy. It institutionalizes the state's approach to scaling dual-use technologies by linking education, industry, and government procurement. At the same time, it reduces dependence on foreign components and promotes local innovation.

2. The **Data Economy and Digital Transformation of the State National Project** aims to modernize Russia's governance, economy, and social systems through large-scale digitalization and the pursuit of technological sovereignty. Within this framework, several federal programs—smaller, targeted initiatives—address specific aspects of the national project. Two of the most consequential among them are the programs Internet Access Infrastructure and Artificial Intelligence, both of which are central to building the foundation of Russia's emerging dual-use AI ecosystem.

   a. The **Internet Access Infrastructure program** aims to ensure universal connectivity and to secure Russia's information space by 2030. Its centerpiece is the creation of a national low Earth orbit satellite constellation of 292 satellites, designed to provide complete internet coverage across Russia's territory and, eventually, globally. Strategically, this initiative reflects Moscow's effort to reduce dependence on foreign technologies while establishing resilient communications for its unmanned systems (similar to how Starlink has proven critical for Ukrainian sea drones), thereby guaranteeing connectivity even when full AI-enabled autonomy is not achieved.

   b. The **Artificial Intelligence program** is designed to embed AI technologies across the economy, social services, and public administration. It focuses on developing domestic AI solutions, integrating them into state decisionmaking, and creating personalized digital services for citizens and businesses. By 2030, at least 100 government services are expected to be delivered proactively, meaning without user requests and based on predictive data analytics and user behavior modeling. The program also emphasizes the development of algorithms for autonomous decisionmaking, natural language processing, and secure data use, reinforcing AI as a strategic enabler of digital governance and industrial competitiveness. In addition, the program aims to cultivate AI competencies from an early age, including the launch of an **All-Russian Olympiad on Artificial Intelligence** for grades 8-11.

These two programs demonstrate how Russia is building the technological and data infrastructure necessary to sustain centralized digital control and expand AI deployment across sectors. These programs are critical because they operationalize the state's vision of autonomy in the information domain and illustrate how AI and connectivity are being fused into the architecture of Russian governance and power projection.

AI REGULATION

Russia's first legal definition of artificial intelligence was introduced in **Federal Law No. 123-FZ** on April 24, 2020, which established a five-year experimental regulatory regime for AI development and deployment only in Moscow. The law **defines** AI as "a complex of technological solutions that enables the imitation of human cognitive functions, including self-learning and the search for solutions without a predetermined algorithm, and allows the achievement of results in specific tasks comparable, at a minimum, to those of human intellectual activity."

Beyond setting a legal precedent, the 2020 law marked Russia's first attempt to test AI governance in practice, combining regulatory flexibility with control over data use and privacy in a contained urban environment, turning Moscow into a national testbed for algorithmic governance and AI-driven public services.

In February 2025, Deputy Prime Minister Dmitry Grigorenko **outlined Russia's emerging federal approach** to AI regulation, announcing that no legislative framework would be introduced for at least the next two years. Speaking at the presentation of the **Data Economy National Project**, he argued that it was "**not yet the moment**" for formal regulation and that the state must intervene "**neither too early nor too late**." His remarks signaled a phased and cautious strategy, prioritizing observation, experimentation, and institutional learning over premature legal codification.

However, Russia's gradual movement toward formalizing its AI governance framework took a material step forward in the first draft of the **Concept for AI Regulation until 2030**, which appeared in August 2025 and included additional steps by the government aimed at defining a legal basis for AI.

Although the full text of the document has not been publicly released, its **preliminary contours**, developed by the Ministry of Digital Development, outline what experts have described as a "distinctly Russian approach." The concept envisions a hybrid regulatory model that combines state oversight with elements of self-regulation, seeking to encourage innovation while maintaining firm control over strategically sensitive and security-critical sectors. In practice, this means that most regulatory measures are expected to have a stimulating or enabling character, complemented by targeted restrictions and limited self-regulatory mechanisms. For instance, within the framework of **experimental legal regimes** for digital innovation, the draft specifies cases requiring mandatory insurance for damages caused by the use of AI technologies.

Russia's **regulatory philosophy** situates itself between two global poles: the United States, which relies on a technocratic, market-driven model assigning responsibility to developers and users, and China, characterized by centralized state control and mandatory algorithmic approval. Russia claims that it seeks strategic flexibility, combining selective restrictions, certification of "**trusted technologies**," and controlled access to state-managed anonymized data with incentives for industrial growth. While the concept **underscores** technological sovereignty and industrial scalability, it appears to lack explicit provisions for privacy or human rights protection, reflecting a regulatory orientation toward state security, institutional control, and pragmatic economic modernization rather than liberal models of data protection or open innovation.

The draft concept also formally asserts that Russia's future AI regulation should rest on a "**human-centered approach**" guided by principles of technological sovereignty, trust in technology, respect for human autonomy and free will, the prohibition of harm to humans, and the rejection of excessive anthropomorphization of AI systems.

The document also reveals **significant structural and methodological weaknesses**. Despite its formal status as a strategic planning document, it lacks coherence with earlier frameworks such as the National AI Strategy, resulting in fragmented governance and regulatory uncertainty. Analysts **note** the absence of implementation mechanisms, or evaluation criteria, while the document overrelies on soft law and self-regulation without defining their legal boundaries. The strong influence of the AI Alliance, an association representing tech businesses which co-authored the document, shifts its focus toward corporate interests—particularly data access and reduced liability—rather than public accountability or citizen protection. The draft also does not offer mechanisms for resolving conflicts between ethical, safety, and sovereignty principles. Overall, it reads more as a political declaration than a coherent legal blueprint for Russia's AI governance.

However, Russia's approach toward AI development and implementation was recently clarified by Vladimir Putin himself. In **November 2025 statements**, he outlined a clear push toward the centralization and state orchestration of Russia's AI development, particularly in the generative AI domain. He called for the creation of a dedicated **National Headquarters** to coordinate AI deployment across all regions and key sectors, arguing that existing working groups lack the "**administrative resource**" needed to drive system-wide implementation. This new centralized structure would operate above individual ministries or industries, unifying the country's AI efforts under a single command architecture.

Putin emphasized that the state must direct the overall trajectory of AI development while remaining in close dialogue with technological businesses. He encouraged bold, unconventional regulatory proposals and the broad use of experimental legal regimes—already active in Moscow, Sakhalin Island, and soon across the Russian Far East—to accelerate testing and deployment. At the same time, he insisted that critical domains such as public administration, security services, and defense must rely exclusively on sovereign, domestically developed AI technologies.

The president also called for large-scale investments in national data center infrastructure to support AI development, with open access for startups, research institutions, and technology companies. He linked regional AI adoption rates directly to Russia's annual digital transformation rankings, signaling a move toward performance-driven oversight. Overall, Putin framed AI not only as a technological priority but as a strategic economic engine, projecting that AI will contribute more than **11 trillion rubles** to Russia's GDP by 2030.

This vision resulted in **Presidential Decree No. 116** on February 26, 2026, in which Russia established the Commission under the President on the Development of Artificial Intelligence Technologies, elevating AI governance to the highest level of state coordination. The commission is tasked with ensuring technological leadership in AI, including the creation of domestic large language models, advanced AI-enabled services, dedicated computing infrastructure, the required electronic component base, and the energy supply necessary to sustain these systems. It is also mandated to define key directions for improving legal regulation in AI development and deployment, explicitly linking economic modernization with national defense and security objectives.

The **composition** of the commission is particularly telling: Alongside senior economic officials and representatives of major technology actors such as Yandex sit the minister of defense and the director of the FSB, composing a relatively small decisionmaking circle. This configuration signals that large-scale AI projects will be shaped and overseen jointly by security institutions and state-aligned technology champions. The structure suggests a centralized, state-driven approach in which civilian AI development, regulatory policy, compute capacity, and military applications are strategically integrated under direct presidential supervision.

In the most recent step toward AI regulation, on March 18, 2026, Russia put forward for public discussion a draft law "On the Fundamentals of State Regulation of the Application of Artificial Intelligence Technologies in the Russian Federation." The bill **introduces** AI regulation that introduces new rules for developers, businesses, and users while significantly expanding the state's role in governing the technology. If adopted, it is expected to enter into force on September 1, 2027.

The draft AI law reflects a dual-track strategy that combines formal alignment with global regulatory norms and a deeper restructuring of the AI ecosystem around state control and technological sovereignty. On the surface, it adopts familiar elements—risk-based regulation, user rights, liability frameworks, and transparency requirements—but its core logic centers on institutionalizing "sovereign" and "trusted" AI systems tied to domestic infrastructure, data localization, and state certification mechanisms.

> *AI is being treated not only as a technological domain but as a tool of political control and regime resilience.*

The requirement that development, training, and deployment occur within Russia, alongside the integration of security services into certification processes and the introduction of "**traditional values**" as a regulatory principle, signals that AI is being treated not only as a technological domain but as a tool of political control and regime resilience.

At the same time, the explicit exemption of defense and security applications creates a bifurcated system—tight civilian oversight paired with opaque, unconstrained military development. Strategically, this hybrid model, blending elements of EU-style compliance, U.S.-style protectionism, and Chinese-style centralization—may limit openness and innovation but enable Russia to build a vertically integrated, security-driven AI stack capable of supporting both domestic control and wartime technological adaptation.

CONCLUSION

Russia's strategic documents, national projects, regulatory experiments, and presidential directives reveal a coherent and increasingly centralized effort by the Russian state to build the foundations of a sovereign ecosystem for unmanned systems and AI. Russia is pursuing these goals systematically at the highest political level, combining long-term strategic planning with a pragmatic focus on applied technologies rather than competing in the global frontier AI race. Instead of attempting to leap directly into foundational research and spending enormous resources on frontier model development, Moscow concentrates on the application layer—on deploying algorithms, integrating autonomy into unmanned systems, and embedding AI into administrative and industrial workflows.

This pragmatism is reinforced by a comprehensive system of incentives and support mechanisms. Favorable regulatory regimes, experimental legal frameworks, and selectively liberalized data access rules are paired with massive investment in domestic component manufacturing and large-scale human capital programs spanning schools, universities, and vocational pathways. Across the ecosystem, emphasis is placed on technological sovereignty—replacing foreign components, building domestic software stacks, and ensuring that critical functions, especially in defense and state administration, rely solely on Russian technologies.

Yet these ambitions also reveal the deeply political character of Russia's approach to innovation. Despite the rhetoric of flexibility and partnership with the private sector, President Putin ultimately has applied his characteristic authoritarian logic to AI governance as well. His call for a National Headquarters for generative AI marks a decisive move toward centralizing decisionmaking, consolidating administrative power, and placing the entire AI domain under direct state supervision.

Thus, Russia's strategy remains tightly controlled from the top. The result is an ecosystem that blends pragmatic technological development with rigid political centralization, a duality that will continue to shape how Russia advances unmanned systems and AI throughout the remainder of this decade.

Russia's Path to Autonomous Unmanned Systems

This section analyzes how Russia is integrating AI into unmanned systems at the tactical edge and how this process is reshaping its military-industrial ecosystem under wartime pressure. The analysis focuses on how machine learning and onboard decisionmaking are embedded into real platforms, with the goal of operating in GPS-denied, electronically contested environments and at scales that matter operationally. It also examines the practical integration of AI, exploring how Russian manufacturers approach technology development and scaling.

This analysis proceeds through a set of representative case studies that illustrate contrasting models of AI development and deployment, from state-centric, top-down programs to commercially driven, bottom-up systems. These cases provide a comparative assessment of how collaboration with the government, industrial practices by manufacturers, and feedback loops with frontline users shape technological and operational outcomes.

This section explores whether battlefield effectiveness depends less on advanced, formally declared autonomy and more on practical factors such as lowering cost, easing production, and enhancing the ability to deploy simple AI functions directly to the front line.

CASE STUDY 1: KRONSHTADT GROUP—CENTRALIZED AI ARCHITECTURES WITHOUT BATTLEFIELD SCALE

The analysis of Kronshtadt Group serves as a cautionary case study within Russia's unmanned systems ecosystem. While the company has positioned itself as a flagship developer of long-range UASs and AI-enabled autonomy, its trajectory illustrates the structural risks of ambitious technological signaling unsupported by sustained industrial execution and battlefield validation. In the context of Russian drone development and AI integration, Kronshtadt demonstrates how expansive claims regarding autonomy, swarming, and decision support architectures do not automatically translate into deployable capability. Examining this gap between conceptual presentation and operational reality provides an important lesson for assessing Russia's broader progress in AI-enabled unmanned warfare.

Kronshtadt Group is a privately held company which develops and produces UASs. The company has operated as an independent entity since 2022, with **limited transparency** regarding its shareholders, governance structure, or financial performance. Despite this opacity, Kronshtadt has positioned itself as one of Russia's flagship developers of large, long-range unmanned systems.

The company's current public-facing portfolio of unmanned systems appears quite small. At present, **the company's website** primarily presents two operational systems: Orion and Sirius. Both are large, Group 4 and 5 UASs designed for long-range intelligence, surveillance, and reconnaissance (ISR) missions, with an advertised capability to conduct strike operations.[1]

Open sources provide limited insight into the specific software architectures embedded in Kronshtadt's unmanned systems, yet the company's public statements and exhibition materials allow a reconstruction of its approach to integrating AI. Rather than presenting AI as a discrete capability, Kronshtadt frames it as a process of gradual progress towards autonomy, most clearly articulated in relation to the **Orion** UAS, shown in Figure 1. This system does not represent edge-based autonomy in the strict sense, but rather showcases an advanced decision support architecture that assists the operator by fusing multi-sensor data and automating elements of its analysis.

1. Groups 4 and 5 refer to the largest classes of unmanned aerial systems (UAS) as defined by the U.S. Department of Defense. Group 4 drones typically weigh over 1,320 pounds and operate at altitudes up to 18,000 feet (e.g., MQ-1 Predator), while Group 5 systems are larger, higher-endurance platforms capable of operating above 18,000 feet (e.g., MQ-9 Reaper), often used for long-endurance surveillance and strike missions.

Figure 1: Orion Unmanned Aircraft System

Note: This is an AI-generated representation of a UAS, created using reference imagery and descriptions from War Sanctions portal of the Defense Intelligence Agency of Ukraine, and guided by CSIS analysis, using GPT-5.2.

Source: CSIS.

In 2021, prior to the full-scale invasion of Ukraine, Kronshtadt CEO Sergey Bogatikov **described** Orion as undergoing phased development of autonomous features. A key milestone was the introduction of a new automated operator workstation showcased at the 2021 Dubai Airshow. This workstation was presented as a functional shift in human-machine interaction, designed to offload a growing share of control and decision support functions from the operator to computational systems. The manufacturer claimed that within this architecture, AI supports mission management, sensor data processing, and operator assistance rather than replacing human oversight altogether.

One concrete implementation of this approach is the integration of augmented reality into the operator interface. Kronshtadt reports the use of AI-assisted visualization that constructs a **three-dimensional representation** of terrain and operational objects, including elements that are known to exist but are not clearly visible in raw sensor feeds. This fusion of AI and augmented reality aims to enhance situational awareness and reduce cognitive burden during ISR and strike missions. The company claimed that this system is already implemented in operational versions of Orion as of 2021.

While the company publicly articulated a vision of enhanced operator assistance and AI-enabled mission support, the sophistication of the unmanned platform itself appears limited. The **operational record** of the Orion system underscores this gap. Orion drones have been repeatedly intercepted and **shot down** by Ukrainian forces, indicating limited survivability, absence of meaningful self-protection measures, and no observable capability for adaptive maneuvering to evade air defenses. Post-recovery technical examinations conducted by Ukrainian specialists revealed **extensive reliance** on commercially available U.S. components and a lack of advanced onboard computing architectures typically required for edge AI processing. Although intercepted variants exhibited some variation in internal components over time, the overall technological baseline remained consistent. Taken together, these findings suggest that Orion does not meet the threshold of an autonomous or even genuinely semi-autonomous system in operational terms, but rather functions as a conventionally piloted platform with limited automated assistance.

Despite the absence of publicly confirmed information on the current deployment status of AI software embedded in operators' workstations or Kronshtadt's operational unmanned platforms, the company's actual level of competence in computer vision and object recognition can be inferred from a different product—the **Mushtra-E system**. It is a machine learning complex for military UASs, designed to support continuous training and retraining of neural networks used in AI-enabled UASs during intelligence, surveillance, target acquisition, and reconnaissance operations.[2] Conceptually, the system addresses one of the central challenges of battlefield AI—maintaining reliable target detection and recognition performance under rapidly changing operational conditions.

This system enables localized, iterative retraining of computer vision models rather than relying on static, centrally trained models. It automates the full retraining cycle of neural networks (i.e., data ingestion, validation, retraining, and redeployment), while deliberately restricting inputs to trusted datasets generated by the user's own platforms. This design choice serves both operational and security objectives by optimizing performance without transferring sensitive imagery or models to third parties.

Mushtra-E supports widely used aerial imagery formats, enabling integration with both Russian and foreign UAS platforms. The system is delivered in **three configurations**: for MALE/HALE and combat-class UASs at permanent airbases, for field-deployed tactical units, and for stationary use at air force headquarters.[3] All these modules can be linked into a hierarchical network. This architecture frames AI not as a standalone onboard feature, but as a scalable ecosystem that evolves alongside operational use.

The system is developed jointly with the State Research Institute of Aviation Systems (GosNIIAS), the same institution discussed in the **first report** of this series as the **architect of the platform-GNS** and its associated AI infrastructure. GosNIIAS also maintains the **Neuroset training database**, which underpins key AI functions including image processing, classification, object detection, and semantic segmentation.

A key extension of Kronshtadt's automated workstation concept is its application to swarm operations **actively promoted** since 2021. In this model, the operator performs high-level mission planning and target designation, while software incorporating elements of AI executes flight control, system management, task allocation, and coordination among multiple UASs. Kronshtadt describes this architecture as enabling semi-autonomous behavior, distinguishing it from conventional automated systems based on rigid preprogrammed logic.

---

2. ISTAR stands for Intelligence, Surveillance, Target Acquisition, and Reconnaissance.It refers to an integrated military function that combines sensors, platforms (like drones, satellites, radar), and analytical systems to collect information, monitor the battlefield, identify targets, and support decision-making and strikes. In practice, ISTAR links data collection with targeting—ensuring that information gathered is rapidly processed and used for operational effect.

3. MALE/HALE refer to long-endurance military drones by altitude: MALE (Medium Altitude Long Endurance) operate at ~10,000–30,000 feet, while HALE (High Altitude Long Endurance) operate above ~30,000 feet for wide-area, persistent surveillance.

> *Continuous data exchange within the swarm supports task redistribution, leadership transfer, and mutual substitution, allowing the group to operate cooperatively with limited reliance on constant communication with the carrier.*

The company presented a **more advanced vision** for swarming technology in 2021 through the Grom unmanned combat aerial vehicle (UCAV), which is conceived of as a command and control node rather than a standalone platform. According to company statements, Grom is designed to manage a swarm of up to 10 small Molniya UASs in both reconnaissance and strike roles, with the ability to dynamically reassign tasks during flight. Recoverable Molniya variants conduct ISR missions, while strike variants function as loitering munitions. Continuous data exchange within the swarm supports task redistribution, leadership transfer, and mutual substitution, allowing the group to operate cooperatively with limited reliance on constant communication with the carrier.

The swarming concept is integrated with Grom's own strike capabilities, including the carriage of guided missiles and precision-guided munitions. Within this architecture, Molniya drones serve multiple roles, ranging from precision strikes and decoy operations to suppression of enemy air defenses, creation of penetration corridors, rapid engagement of newly detected targets, and support for group EW alongside manned aircraft.

However, in 2023, Kronshtadt Group **announced its partnership** with a UAS manufacturer to develop specialized software for swarm control, which points to the fact that the company's own efforts were not successful.

According to the company, the planned software will enable mission preparation, flight supervision, payload management, route modification, and post-mission performance assessment across groups of unmanned platforms. As of 2024, the Grom UCAV remained in development. During Army-2024, **the company presented** only an updated airframe configuration rather than evidence of operational maturity.

Despite repeated conceptual presentations, no publicly verifiable evidence indicates that swarming capabilities of this type have been fielded or deployed in operational systems. Conversations conducted by CSIS with members of the Ukrainian military have similarly yielded no confirmation of the use of such capabilities. Given Russia's demonstrated willingness to employ experimental and partially mature systems in combat, the absence of observable swarm employment strongly suggests that these capabilities have not progressed beyond the conceptual stage.

Overall, Kronshtadt's public narrative frames AI as a means of increasing autonomy through decision support, cooperative control, and human-machine teaming rather than fully independent unmanned operations. At the same time, despite ambitious concepts and repeated presentations, tangible progress toward genuinely disruptive capabilities, such as swarming, remains limited, with no confirmed operational deployment.

Companies such as Kronshtadt Group, characterized by opaque ownership structures, reliance only on government contracts, and close ties to state research institutes, face compounded challenges from sanctions due to dependence on foreign components, long development cycles, and bureaucratic constraints. By August 2025, industry observers increasingly described the Kronshtadt Group as being on a path **toward bankruptcy**, driven by the loss of a strategic investor, mounting debt, and a growing number of lawsuits from subcontractors.

Kronshtadt Group illustrates the risks inherent in a development model heavily concentrated on state defense contracts with limited diversification beyond military and select government applications. The company consistently advanced ambitious and forward-looking concepts, yet appears to have lacked the industrial capacity and technological foundation required to translate these visions into robust operational capability.

While several Russia-specific structural factors contributed to this trajectory—including sanctions pressure, supply chain constraints, and bureaucratic procurement dynamics—the broader lesson extends beyond a single case. Overpromising technological breakthroughs, relying predominantly on government contracts, and failing to build scalable production capacity can result in platforms that deliver only incremental performance improvements rather than the transformative capabilities advertised. In Kronshtadt's case, this pattern culminated in systems that performed at an average technological baseline, fell short of stated ambitions, and ultimately left the company facing severe financial distress.

CASE STUDY 2: ZALA AERO—DUAL-USE AI AS A MILITARY FORCE MULTIPLIER

The company **ZALA Aero Group** offers a great illustration of how Russia's approach to nurturing the civilian unmanned systems sector has translated into tangible battlefield effects. This company exemplifies how AI developed for commercial and public sector applications can be repurposed for military use, allowing it to build sustainable business models while simultaneously reinforcing the Russian military-industrial complex with modern, commercially derived technologies. Rather than relying exclusively on legacy research institutes, this model leverages innovation originating in the civilian market, allowing dual-use capabilities to mature rapidly and migrate into operational military systems.

ZALA Aero is a Russian unmanned systems developer that operates at the intersection of civilian and military drone production. The company, which is structurally linked to the Kalashnikov Concern, produces a broad portfolio of UASs ranging from lightweight reconnaissance platforms and monitoring drones for civilian agencies to loitering munitions and battlefield ISR systems used by the Russian Armed Forces.

Despite ZALA Aero's prominence as a drone manufacturer, relatively little is publicly disclosed about the specifics of its military AI capabilities. **Official communications** tend to frame the company's AI offerings in broad terms, saying that it has improved target recognition and resilience in contested EW environments without naming specific software architectures or levels of autonomy.

Nevertheless, a series of repeatedly reported upgrades suggests incremental but meaningful advances to these already well-established and widely deployed capabilities. In particular, ZALA Aero's modernized Lancet loitering munition has **reportedly received** enhanced sensors and imaging systems, improving target detection and lock-on performance under conditions of poor visibility and adverse weather. At the same time, its neural networks appear to have been retrained on larger and higher-quality datasets, enabling more accurate target recognition and more refined terminal-phase navigation.

These capabilities reportedly allow the system to strike not merely the target as a whole, but **specific vulnerable points** across a wide range of platforms. Implicitly, this level of performance suggests that the manufacturer has accumulated extensive training data covering diverse categories of equipment, including not only Ukrainian systems but also a broad spectrum of Western-origin platforms.

The upgraded system is described as offering **faster and more robust** navigation and flight control during the terminal phase, allowing high precision strikes against protected or partially concealed targets. According to **recent reporting** from ZALA Aero, Lancet drones have also been credited with engaging fast-moving targets, including Ukrainian maritime drones such as the Magura V7, reportedly equipped with onboard air-defense systems carrying U.S.-made AIM-9M Sidewinder missiles. These characteristics point to a transition from a strictly operator-in-the-loop model toward supervised autonomy, in which AI plays a decisive role once the engagement has begun.

However, based on conversations conducted by CSIS with Ukrainian technical specialists, it remains doubtful that the terminal phase in this specific case was autonomous. Analysis of the available **strike footage**, combined with expert assessments, suggests that the observed flight trajectory and target engagement profile are more consistent with a human-in-the-loop control model rather than fully autonomous terminal guidance. According to these specialists, the maneuvering pattern visible in the video would be difficult to reconcile with an independent onboard decisionmaking algorithm and instead points toward direct operator control during the final approach.

Although publicly available information on ZALA Aero's military AI remains sparse, one particularly revealing detail emerged from an **interview** with a Lancet operator. He notes that during the terminal phase of an attack, he can activate an auxiliary AI feature referred to as "Ira," which, according to his account, "takes over and does the job" in the final approach by refining guidance to ensure a precise strike, including against fast-moving armored vehicles. This testimony is notable, as it provides a concrete name for an AI component explicitly associated with combat employment. At the same time, when this designation is traced through open sources, references appear not in military disclosures but in descriptions of ZALA Aero's civilian systems, where the software is identified as IRRA. This suggests that the AI supporting ZALA Aero's military drones is the same software architecture used in its civilian unmanned platforms, offering valuable insight into the nature and capabilities of ZALA Aero's military AI stack.

The IRRA AI software is best understood not as a single algorithm, but as an integrated, end-to-end onboard intelligence architecture designed to fuse sensing, analytics, decision support, and networked dissemination into a unified operational system. Its defining technical feature is the **relocation of core computational workloads** from the ground segment to the UAS itself. By embedding processing capacity directly onboard platforms such as the ZALA T-16, ZALA T-20, and ZALA ZARYA series, IRRA enables real-time interpretation of sensor data during flight rather than via post-mission or ground-based analysis. This architectural choice significantly compresses the decision loop and underpins the system's operational value across both civilian- and security-oriented use cases.

At the functional level, IRRA operates through continuous **real-time analysis** of multimodal sensor inputs, primarily optical and thermal imagery. Using machine vision algorithms and trained neural networks, the system processes live video streams directly on the UAS, automatically identifying anomalies, deviations, and objects of interest. In energy infrastructure monitoring, this includes the detection of structural irregularities, thermal anomalies, or other indicators of potential threats to fuel and energy sector assets. In emergency response applications, the same analytical pipeline is applied to the identification of fire hotspots, flood boundaries, or other evolving hazards. Crucially, detection is not limited to binary identification; the system classifies, contextualizes, and spatially localizes anomalies, producing actionable intelligence rather than raw imagery.

Once an anomaly is detected, IRRA automatically marks it within the video feed and generates structured outputs that are transmitted to the ground control station. These outputs include georeferenced alerts, highlighted regions of interest, and **automatically generated analytical reports**. This automation reduces cognitive load on operators and enables faster, more consistent responses, particularly in time-critical scenarios such as wildfire suppression or infrastructure protection. The system's ability to **determine coordinates** with high precision enhances its utility for directing follow-on actions by ground teams or other response assets.

IRRA's effectiveness is amplified through its tight integration with ZALA Aero's broader digital ecosystem, particularly the **4Z1 control and data management platform**. Through this integration, data produced onboard the UAS is not confined to the immediate operator but becomes accessible (subject to permissions) across a distributed user network. Video streams, analytical overlays, and processed outputs can be viewed by operators, analysts, and customers from geographically remote locations, effectively extending situational awareness beyond the physical control station. The platform unifies live streaming, automated interpretation, orthophoto generation, and archival storage into a single software contour, enabling seamless transitions between real-time monitoring and decisionmaking.

From a military perspective, this architectural logic directly maps onto ZALA Aero's combat drone ecosystem, where reconnaissance, targeting, and strike functions are distributed across different platforms but integrated within a single operational cycle. In the observed engagement chain, the reconnaissance vehicle identifies potential targets and transmits video feeds to a ground control station. The ground control station then selects the target, generates coordinates, and transmits both the positional data and selected imagery to the loitering munition prior to launch or terminal approach. This sequence reflects a networked human-in-the-loop targeting architecture rather than autonomous edge execution. Target detection, validation, and engagement decisions remain distributed across operators and ground infrastructure, with the loitering munition functioning primarily as an execution platform rather than an independent decisionmaking system.

IRRA can be understood as the software backbone that enables this continuity: It supports detection, classification, geolocation, and prioritization at the reconnaissance stage; maintains a coherent, machine-interpreted representation of the battlespace; and ensures that targeting data is transferred with minimal latency and degradation to strike assets. In this sense, IRRA does not merely enhance individual drones, but underwrites a system-of-systems approach in which sensor platforms and strike platforms function as components of a unified kill chain, with AI providing the connective tissue that synchronizes observation, decision, and action. However, not all systems produced by ZALA Aero are integrated into this broader system-of-systems architecture, and only a limited subset of models appears to operate within this coordinated framework.

ZALA Aero illustrates a more pragmatic and operationally grounded model of AI integration than many other Russian unmanned systems developers. Rather than advancing sweeping claims of full autonomy, the company appears to embed commercially derived machine vision and data processing architectures into an incremental, networked combat framework. Its approach demonstrates how dual-use civilian AI stacks can be repurposed to accelerate battlefield adaptation, shorten decision loops, and enhance precision within a supervised human-in-the-loop structure. At the same time, the limits of observable autonomy suggest that Russia's progress lies less in achieving independent machine decisionmaking and more in optimizing coordinated kill chains through distributed AI-assisted systems.

CASE STUDY 3: MOLNIYA—BOTTOM-UP DRONE INNOVATION MEETS STATE-BACKED SCALING

The Molniya UAS illustrates a distinctive shift in Russia's approach toward its military-industrial complex and toward cooperation between the Russian Armed Forces and the country's defense industry. This approach was articulated by the Russian defense minister Andrei Belousov during a **meeting** with military correspondents, when he stated that he is currently fully satisfied with the decentralization of drone and EW production, specifically referring to "**garage-level development and assembly**."

The emergence of the Molniya UAS illustrates a bottom-up innovation pathway that diverges markedly from Russia's traditional, state-centric defense industrial model. It originated in a so-called "people's VPK," or people's defense industrial complex–a loosely coordinated community of civilian engineers and volunteers working for the Russian war effort. Molniya was initially designed and assembled in informal, garage-based workshops rather than within established state design bureaus. Its early development relied on small teams of engineers and volunteers operating outside formal acquisition structures, enabling rapid experimentation and close interaction with frontline users.

However, this project is different from hundreds of similar "garage" projects because it received government support for scaling. According to Russian military bloggers, the project was placed "**on the rails**" of formal production, receiving government financing to expand manufacturing capacity. Oversight of serial production was subsequently assigned to the company Sudoplatov, marking Molniya's transition from an improvised grassroots initiative to a state-supported system. This sequence–garage-level innovation followed by selective state scaling–shows an adaptive procurement logic in which the government absorbs and institutionalizes proven battlefield solutions rather than attempting to generate them entirely within formal defense industrial structures.

The **first reports** of Molniya's battlefield deployment appeared in May 2024. Since then, the system rapidly diversified into three main types: Molniya-1, Molniya-2, and Molniya-2R. As shown in Figure 2, the initial operational model, Molniya-1, served primarily as a one-way strike platform with limited autonomy, capable of delivering a payload of several kilograms at ranges of approximately **30–40 kilometers**. The subsequent Molniya-2 introduced a twin-engine configuration, increasing range up to **80 kilometers** and expanding payload capacity while enhancing resilience to EW. The most advanced and recent iteration, **Molniya-2R**, represents a **qualitative transition** toward reconnaissance and networked operations. By **integrating satellite communications** (reportedly through Starlink) and more powerful onboard computing, this model supports over-the-horizon ISR, target tracking, and high-bandwidth data transmission, significantly expanding the operational role of the Molniya family.

Figure 2: The Evolution of the Molniya UAS

Note: These images are AI-generated representations of UASs, created using reference imagery and descriptions from War Sanctions portal of the Defense Intelligence Agency of Ukraine, and guided by CSIS analysis, using GPT-5.2.

Source: CSIS.

The primary objective in creating Molniya UAS was to make a cheap, long-range, fixed-wing strike UAS that could substitute for more expensive loitering munitions while remaining resilient to sanctions and supply chain disruptions. This objective shaped several defining design choices.

First, the platform relies extensively on commercially available components. This deliberate reliance on civilian electronics and simple materials reduces vulnerability to sanctions, as such components are difficult to restrict comprehensively, and they can be sourced locally or through diverse global supply chains.

Second, Molniya employs extremely simple airframe designs constructed from foam, plywood, and basic structural elements. This simplicity minimizes cost and manufacturing complexity, enables fast mass production with minimal technical infrastructure, and allows rapid modification based on frontline experience.

A third and increasingly central feature of Molniya's evolution is the early integration of AI **for terminal guidance**. Even in its relatively simple configurations, Molniya is equipped with a basic homing system based on machine vision designed to ensure mission completion in the event of signal loss. This capability reflects a pragmatic response to EW rather than an ambition for full autonomy. Once the operator designates and locks onto a target, onboard computer vision processes imagery from the camera, identifies a contrast-rich moving object, and autonomously corrects the flight path during the terminal phase. If communication with the operator is disrupted, the UAS continues its attack without further external input, rendering enemy EW largely ineffective at this stage.

Importantly, this AI-enabled terminal guidance does not involve semantic understanding of targets in the strict sense. The system does not classify objects as armored vehicles or specific weapon types. Instead, it operates on **visual salience** and motion contrast, prioritizing robustness and computational simplicity over sophisticated recognition. Unlike fiber-optic-guided drones, which require long cables that add weight, reduce payload, and limit endurance, Molniya's AI-based approach preserves range, explosive capacity, and battery life while achieving comparable resilience against jamming.

This system illustrates an AI-enabled capability deployed at scale rather than as a niche or experimental solution. Molniya was conceived from the outset as a mass-produced platform: Its **initial cost** reportedly stood at around $300, rising to approximately $5,000 when equipped with improved optics and satellite connectivity for ISR missions. This is an order of magnitude cheaper than the Lancet loitering munition, which is commonly estimated at around $50,000 and has similar capabilities.

Production and employment figures **reinforce this disparity**. In September 2025, a single month, approximately 2,200 Molniya-2 launches were reported compared to roughly 400 Lancet launches.

> *Expensive and highly sophisticated systems struggle to match the operational and economic advantages of simpler, cheaper, and sufficiently reliable AI-enabled alternatives.*

All currently mass-produced variants of Molniya are equipped by default with basic AI functionalities, showing the normalization of AI across large-scale unmanned systems rather than its confinement to elite platforms. This pattern highlights the practical logic driving AI adoption, with narrowly defined, mission-critical functions implemented at low cost in mass systems rather than technologically sophisticated designs without a clear practical need. In this competitive environment, expensive and highly sophisticated systems struggle to match the operational and economic advantages of simpler, cheaper, and sufficiently reliable AI-enabled alternatives.

CASE STUDY 4: V2U AND THE EMERGENCE OF FULLY AUTONOMOUS AI-DRIVEN DRONES

The V2U UAS represents one of the most advanced and concerning examples of AI-enabled autonomy currently observed in Russia's drone ecosystem.

From a technical standpoint, V2U (shown in Figure 3) is a UAS with two known configurations: a loitering munition and a reconnaissance platform. In its **loitering munition configuration**, the drone carries a high-explosive fragmentation incendiary warhead weighing up to 3 kilograms. It relies heavily on commercial off-the-shelf components, including **Chinese-made electric motors and batteries**. It has an estimated **operational range** of 40-60 km, and a gasoline-powered version has extended endurance, capable of travelling over 100 km.

Figure 3: V2U UAS

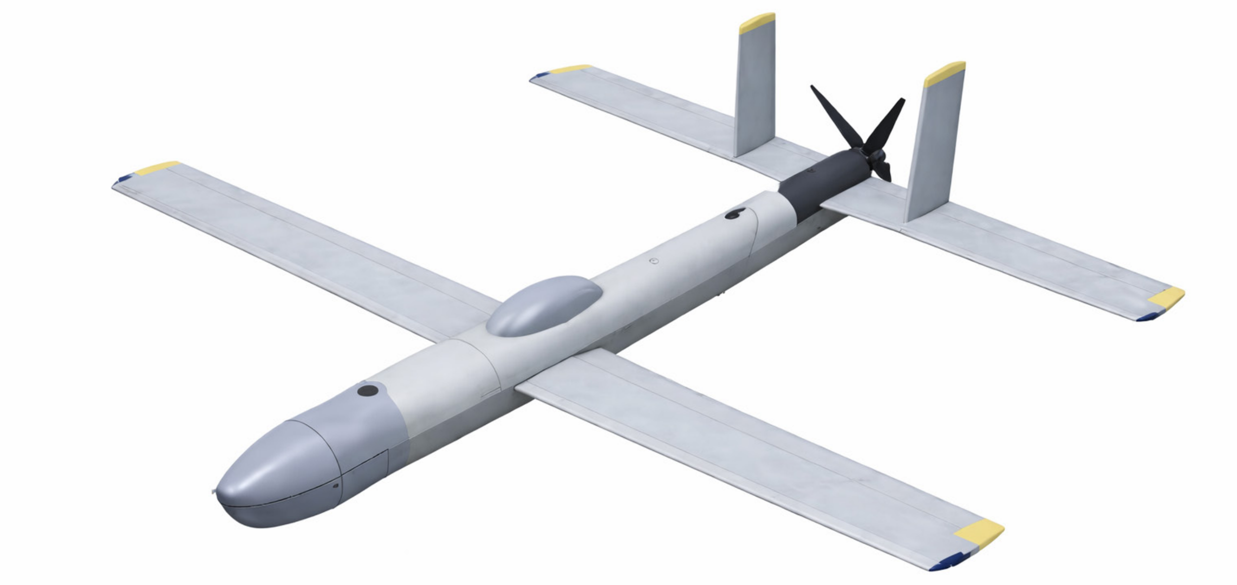

Note: This image is an AI-generated representation of the UAS, created using reference imagery and descriptions from War Sanctions portal of the Defense Intelligence Agency of Ukraine, and guided by CSIS analysis, using GPT-5.2.

Source: CSIS.

In addition to this strike configuration, Ukrainian experts **report** about a reconnaissance version of the V2U. In this version, Russian engineers removed the warhead and replaced it with additional batteries, significantly extending loiter time. They also incorporated a parachute recovery system, allowing the drone to land after completing ISR missions. Analysis of recovered software suggests that further modifications are under development, including to develop logistics or courier versions. This fact points to an expanding modular family of V2U systems rather than a single-purpose weapon.

This analysis does not examine the drone's components in depth. However, assessing the presence or absence of specific components provides important insight into its functional capabilities. The most consequential aspect concerns the navigation and communication architecture or, notably, its deliberate removal. Early V2U drones intercepted in June and July 2025 **contained LTE modems**, indicating at least partial reliance on operator connectivity. More recent versions intercepted in October and November 2025, however, were recovered **without any modem** or external communication system, marking a clear transition toward fully autonomous operation. The **elimination of LTE connectivity** renders EW largely ineffective against control and guidance, as there is no signal to jam. Instead, the **drone relies** on onboard sensors, including a laser altimeter and terrain-referenced flight profiles, enabling sustained low-altitude operation without GPS.

AI lies at the core of the V2U's design philosophy. Despite Western sanctions, technical examinations and **Ukrainian intelligence reporting** indicate that the drone incorporates advanced Western and Chinese electronics, most notably a **Nvidia Jetson Orin** AI module mounted on a Chinese **Leetop A603** carrier board. This configuration demonstrates continued Russian access to high-performance computing hardware.

The onboard AI enables the drone to **autonomously search for**, identify, and select targets using computer vision. The AI stack reportedly runs a trained **YOLOv5** neural network, allowing visual recognition of vehicles, infrastructure, and human activity based on contrast, shape, and motion rather than semantic classification.

V2U's autonomy extends beyond individual decisionmaking to collective behavior, including elements of swarming. **Field observations suggest** that these drones operate as distributed, partially swarm-capable systems in which each unit processes information locally while remaining aware of nearby drones. Coordination does not appear to rely on continuous radio communication. Instead, based on observed imagery, drones may be using visual recognition to identify one another **through distinctive markings** painted on their wings (see Figure 4). These markings could function as visual identifiers, allowing onboard cameras and algorithms to detect and distinguish individual drones as separate nodes within a swarm. While this interpretation remains inferential and cannot be confirmed with certainty, it is consistent with observed behavior and suggests a potential vision-based approach to swarm coordination in GPS- and EW-contested environments.

Figure 4: Visual Identification Markings Enabling Potential Vision-Based Swarm Coordination

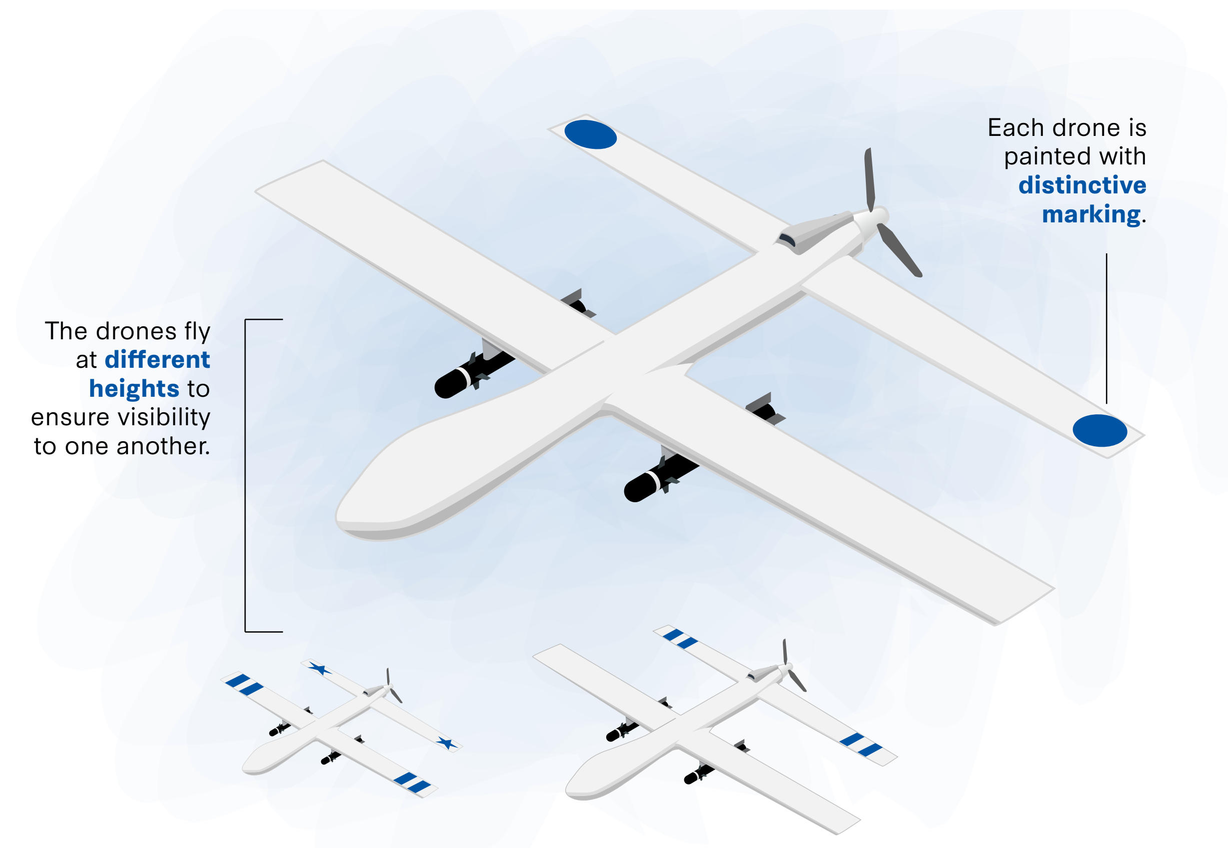

Note: Illustration not to scale.

Source: CSIS. Illustration by Sabina Hung/CSIS.

This allows six or seven drones to fly in formation, enabling mutual awareness and adaptive responses to losses within the group (see Figure 5). If one drone is knocked out by air defenses, for example, the remaining units infer the presence of a threat and execute evasive maneuvers before regrouping. This behavior closely resembles flocking dynamics observed in migratory birds, with drones flying in staggered vertical formations to maintain visual contact.

Figure 5: Drone Formation Collaboration Tactics

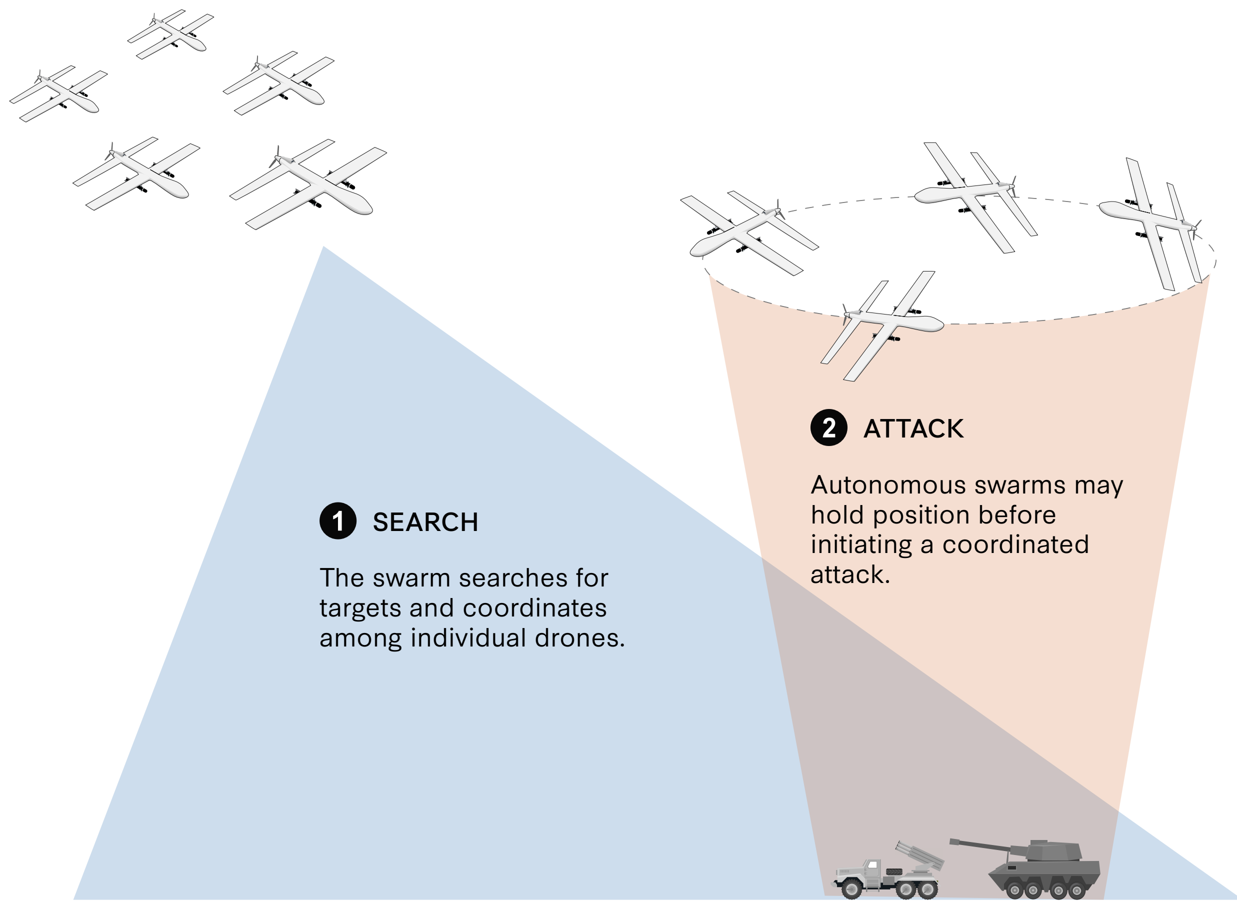

Note: Illustration not to scale.

Source: CSIS. Illustration by Sabina Hung/CSIS.

Documented combat incidents illustrate the operational implications of this design. In **one reported case** in May 2025, a group of seven V2U loitering munitions deviated from a preplanned mission after detecting a concentration of vehicles and civilians, **autonomously forming** a circular holding pattern before initiating coordinated attacks. Such behavior indicates not only autonomous target selection but also group-level decisionmaking based on environmental cues. The combination of AI-based perception, GPS-independent navigation, visual swarm coordination, and resistance to EW positions V2U as a qualitatively new class of battlefield threat.

> *The software-defined capabilities, particularly autonomous target selection and emergent swarm tactics, make V2U one of the most innovative and dangerous unmanned systems currently observed in active combat use.*

The V2U family reflects a shift from remotely piloted, expendable drones toward fully autonomous, AI-driven systems capable of collective behavior. Although the airframe and manufacturing quality remain relatively crude, the software-defined capabilities, particularly autonomous target selection and emergent swarm tactics, make V2U one of the most innovative and dangerous unmanned systems currently observed in active combat use.

CONCLUSION

These cases demonstrate that Russia's integration of AI into unmanned systems is being driven less by the pursuit of elegant, system-wide autonomy concepts and more by the pressures of wartime adaptation at the tactical edge. Across very different organizational models—including Kronshtadt's centralized state-linked architectures, ZALA Aero's dual-use commercial stack, Molniya's garage-born mass platforms, and the highly autonomous V2U—AI is not treated as an abstract capability or end state. Instead, it is implemented selectively, in narrowly defined functions that address concrete operational problems, such as surviving EW, compressing decision loops, sustaining operations without GPS or connectivity, and scaling effects at low cost.

The comparative analysis shows a clear divergence between conceptual sophistication and battlefield relevance. State-centric programs such as those pursued by Kronshtadt focus on decision support systems and future-facing swarm concepts yet struggle to translate these ambitions into systems deployed at meaningful scale. By contrast, commercially derived and bottom-up models prioritize producibility, cost control, and rapid iteration informed by frontline feedback. In these cases, AI is embedded directly onboard platforms, often in limited but mission-critical roles, such as terminal guidance, real-time object detection, and autonomous navigation. These capabilities deliver disproportionate operational value relative to their technical complexity.

Crucially, the Molniya and ZALA Aero cases illustrate how AI-enabled capabilities become operationally decisive not because they represent high levels of declared autonomy, but because they are normalized across mass-produced systems. Here, AI functions as connective tissue within kill chains rather than as a standalone technological leap. The V2U case marks the outer boundary of this trajectory, showing how fully autonomous perception, navigation, and even collective behavior can emerge once connectivity constraints are removed altogether—albeit at the cost of introducing qualitatively new risks and escalation dynamics.

The evidence suggests that Russia's battlefield effectiveness in unmanned systems is shaped less by formal autonomy claims and more by an adaptive ecosystem that rewards simplicity, speed, and scale. Wartime pressures have pushed the Russian military-industrial complex toward hybrid models that absorb civilian innovation, tolerate decentralization, and institutionalize and scale solutions only after they prove themselves in combat. In this environment, narrowly scoped, robust AI deployed at the edge consistently outperforms more ambitious but fragile designs, offering a clear lesson about how AI is most likely to matter in contemporary and future high-intensity conflict.

Core Drivers of Russia's AI and Unmanned Systems Development: Training, Technology, and Partnerships

PRIVATE MILITARY TRAINING AS THE ENGINE OF COMBAT POWER

This section analyzes how Russia's military training architecture enables the integration of emerging technologies–including unmanned systems, battlefield management software, and EW–into operational practice. While much of the public discussion has focused on hardware production and frontline adaptation, the decisive variable in converting new systems into sustained combat capability is training: the preparation of instructors, the revision of curricula, and the institutionalization of feedback loops linking combat units with educational establishments.

In the Russian case, these functions extend beyond the formal structures of the armed forces. After 2022, a broader training ecosystem has emerged that combines three major pillars: informal volunteer initiatives, institutional mechanisms within the Ministry of Defense, and, increasingly, state-backed youth programs designed to cultivate early technical proficiency in drone operation. By embedding drone instruction earlier in the talent pipeline and fostering a generation of drone-native recruits, this ecosystem supports the absorption, standardization, and diffusion of emerging technologies across the force.

The analysis that follows traces how these three pillars have contributed to the diffusion and scaling of unmanned systems and new technological capabilities. It outlines the early emergence of decentralized initiatives that drove rapid experimentation and specialized instruction; the subsequent efforts by the Ministry of Defense to formalize, integrate, and scale successful models into official doctrine and training pipelines; and the gradual extension of drone competencies into pre-military and youth education structures. This evolution illustrates a broader pattern in Russia's wartime adaptation. Technological innovation has advanced through a dynamic interaction between bottom-up experimentation and top-down institutional consolidation, enabling the systematic integration of new capabilities into the force.

Table 3: Russia's Three-Pillar Training Architecture

	Pillar I: Decentralized Volunteer Training Networks	**Pillar II: Institutional Training Structure**	**Pillar III: Youth Training Pipelines**
Description	Informal networks of volunteers and engineers establish training centers across Russia and occupied territories.	Formal military training systems operate under the Russian Ministry of Defense.	State-backed military-patriotic organizations provide drone instruction to schoolchildren and university students.
Function	Experiment with and validate new systems under battlefield conditions.	Institutionalize and scale successful models through formal command structures.	Build long-term talent pools of technically skilled personnel.
Example	Project Archangel	Unmanned Systems Forces	Voin Center

Source: CSIS.

PILLAR I: DECENTRALIZED VOLUNTEER TRAINING NETWORKS

At the outset of the full-scale invasion of Ukraine, Russia's armed forces were unprepared for the rapid proliferation of unmanned systems on the battlefield, both in terms of hardware availability and personnel readiness across the command chain and training pipelines. Supply chains were inconsistent despite the fast emergence of capable manufacturers, structured drone training programs were absent, and engineering and technical support to frontline units depended heavily on individual initiative.

At the same time, many commanders were shaped by experience in earlier conflicts and, therefore, were slow to adapt to the realities of drone-saturated warfare. Those who recognized the shift were forced to drive change without systemic backing and sometimes despite institutional resistance. In this environment, informal volunteer initiatives mobilized to close critical gaps in supply, training, and integration, seeking to build the essential components of an effective system such as adaptable commanders, trained operators, and reliable equipment in sufficient quantities.

One of the most illustrative examples of this informal mobilization is Project Archangel. Founded in 2022 by Russian entrepreneur Mikhail Filippov, the initiative **began** as a small collective of drone engineers paired with a grassroots training program for operators. It rapidly evolved into a nationwide network, **establishing training centers** across roughly 20 Russian cities and in occupied territories of Ukraine. These centers have trained thousands of servicemen, reservists, and civilians preparing to contract with the Ministry of Defense, typically through two- to three-month courses focused on practical drone operations, counter-drone tactics, and the integration of emerging technologies into combat units. In doing so, Project Archangel does not merely supplement state training but also helps create parallel pipelines capable of rapidly absorbing battlefield lessons and translating them into structured instruction.

Beyond training, the group has actively shaped technological adaptation. Its engineers developed systems such as the Archangel counter-UAS system and **paired hardware innovation** with operator training programs to ensure fielded systems could be employed effectively. Training centers integrated advanced software tools, including the **Glaz/Groza** complex for improved reconnaissance and strike coordination and **Kvadrosim**, a combat simulator for drone operation and interception, thereby embedding digital tools directly into the instructional process.

Over time, the boundary between informal initiative and formal defense industry narrowed. In August 2025, Project Archangel **entered into a partnership** with Kalashnikov Concern to evaluate and select UAS designs submitted by independent developers for scaled production while supporting operator and instructor preparation. This evolution demonstrates how volunteer-driven networks have functioned as accelerators of technological diffusion—experimenting, validating, and institutionalizing drone capabilities at a pace that formal military structures initially struggled to match.

PILLAR II: INSTITUTIONAL TRAINING STRUCTURE

Russia's formal military training system was initially slow to adapt its existing institutional infrastructure to the rapid introduction of new battlefield technologies. Over time, however, it began absorbing lessons generated by informal volunteer initiatives and has scaled battlefield-proven practices within the official force structure.

Early cooperation with private and grassroots drone schools allowed the Ministry of Defense to observe and replicate effective training models, particularly in unmanned systems instruction. This process has gradually evolved from ad hoc collaboration into formal integration as drone training has become embedded within standard war-preparation pipelines. Ultimately, the Russian military has restructured elements of its force by concentrating experienced operators, validating systems, refining tactics into elite formations, and then scaling these practices across the broader force—culminating in the establishment of the dedicated Unmanned Systems Forces and centralized training and innovation centers designed to standardize and expand drone warfare expertise.

The formal structure of the Russian Armed Forces is organized territorially into military districts, which function as strategic-level administrative and operational commands responsible for force generation, logistics, and training oversight within their assigned regions. As of the current reform cycle, Russia maintains **several military districts** (including the Western, Southern, Central, Eastern, and the reestablished Moscow and Leningrad districts), each of which oversees training centers, academies, and mobilization infrastructure within its area of responsibility. Initial preparation of contract soldiers, mobilized personnel, and other non-career servicemen typically occurs at these district-level training facilities, where they receive basic combat instruction and, increasingly, **exposure to drone operations**, EW, and engineering tasks.

Once assigned to combat operations in Ukraine, personnel are subordinated not to their original military districts but to operational groupings of forces. These groupings are wartime operational commands organized by direction of effort (for example, the Southern or Central groupings) and are distinct from the territorial military district structure. Upon arrival in theater, soldiers frequently undergo additional training at near-frontline training grounds, where instruction is more specialized and directly tailored to current battlefield conditions. However, reports indicate that the effectiveness of this training is limited by inadequate equipment and lack of standardization. Moreover, training timelines can also be rushed: contract soldiers may receive just three weeks of total training before participating in frontline operations.

The limits of formal military training and the demonstrated effectiveness of volunteer initiatives prompted the Ministry of Defense to move from observation to replication. In August 2024, it established the **Rubicon Center for Advanced Unmanned Technologies** as a state-backed effort to systematize what informal networks had pioneered. Modeled in part on projects such as Archangel, Rubicon combined two functions that had previously existed in parallel: it became both a premier combat drone unit and a centralized training hub staffed by **highly qualified instructors**. Rather than building from scratch, the **Ministry of Defense absorbed** elements of existing elite volunteer formations–including units known as "Judgement Day" and "Judgement Night"–thereby consolidating experienced operators, tested tactics, and technical know-how under formal command.

Rubicon's impact was visible on the battlefield. Ukrainian units **reported** abrupt shifts in Russian drone tactics, improved coordination, and a marked focus on hunting not only UASs but also the logistical and network infrastructure underpinning Ukraine's drone operations. Rubicon operators demonstrated increasing proficiency in intercepting heavy bomber drones and striking communication nodes such as antennas and satellite terminals–targets that directly degrade an opponent's situational awareness and operational tempo. Central to this effectiveness was Rubicon's training model. Its personnel **functioned simultaneously** as combat operators and instructors, deploying to frontline sectors to establish local drone schools, standardize procedures, and disseminate updated tactics. This approach created a mobile training and combat nucleus capable of transferring expertise across sectors, accelerating diffusion of best practices throughout the force.

Building on these results, Russia formalized the scaling process. In December 2024, Defense Minister Andrey Belousov **announced** the creation of a dedicated Unmanned Systems Forces branch, **officially established** in 2025, with Rubicon serving as its flagship formation. The new branch aims to institutionalize lessons from drone warfare, centralize command and control of unmanned units, and integrate training, doctrine development, and technological innovation within a single organizational framework.

Plans to establish a dedicated military academy further signal a shift from ad hoc adaptation to long-term professionalization. According to **open-source reporting**, the branch is expected to establish its own dedicated military academy, potentially as early as 2027, signaling a long-term commitment to professionalizing unmanned warfare capabilities. The progression from decentralized volunteer groups to Rubicon as a centralized elite unit, and ultimately to the Unmanned Systems Forces, illustrates how Russia identifies operationally successful models, concentrates expertise, and then scales them across the broader force through formal institutional mechanisms.

PILLAR III: YOUTH TRAINING PIPELINES

Beyond reforming its active force structure, the Russian state has moved to embed unmanned systems competencies earlier in the personnel pipeline by integrating drone instruction into military-patriotic youth organizations. Russia's **Unmanned Aerial Systems National Project**, discussed earlier in this report, identifies shortages in trained technical personnel as a structural bottleneck. In response, youth programs have been positioned as a long-term solution—simultaneously supporting near-term mobilization needs and cultivating future operators, technicians, and engineers.

A central example is the Center for Military-Sports Training and Patriotic Upbringing of Youth, known as the "Voin Center" (Warrior Center), which was established in May 2023 to provide initial military training to schoolchildren and university students. The organization has expanded to **more than 20** regional centers and **incorporates drone-related instruction** into its curriculum, including competitive exercises that simulate battlefield tasks such as first-person view (FPV) strike operations. Its leadership has explicitly emphasized technological literacy as a priority, and instructors **undergo training** on systems used in combat, including the Glaz/Groza complex and machine vision guidance tools such as Ploshchad, often drawing on specialists with frontline experience to update curricula.

Looking ahead, the Voin Center has **announced plans** to launch a dedicated UAS operator preparation track aligned with the Unmanned Systems Forces, with implementation expected to begin in 2026. In cooperation with DOSAAF and other state-backed structures, it is developing training manuals and pilot facilities to support this pathway.[4] Graduates are expected to have the option of signing contracts with the Ministry of Defense, effectively linking youth training directly to force generation.

By institutionalizing drone education at the preservice level, the Russian state is extending its centralization and scaling strategy across generations. It will transform wartime adaptation into a durable system for producing technically prepared recruits for unmanned warfare.

Across these three pillars, one central conclusion stands out—training, not hardware alone, determines battlefield outcomes. Informal training networks acted as rapid innovators, translating new drone and EW technologies into practical combat capability when formal institutions lagged. Their agility allowed Russia to close operational gaps and quickly field trained operators capable of reshaping tactical engagements.

The Ministry of Defense then institutionalized and scaled successful models through structures such as Rubicon and the Unmanned Systems Forces. At the same time, youth organizations began embedding drone instruction into long-term personnel pipelines, ensuring a steady supply of technically prepared recruits. These layers together demonstrate that Russia's adaptation in drone warfare has been driven less by isolated technological breakthroughs and more by the systematic integration of training, doctrine, and organizational reform, an approach that increasingly shapes successes on the battlefield.

AI-ENABLING HARDWARE IN RUSSIAN UNMANNED SYSTEMS

This section examines the hardware backbone underpinning Russia's efforts to expand autonomy in its unmanned systems. It focuses specifically on the origin of compute, sensing, and memory components that are essential for enabling advanced onboard AI functionality on the battlefield. While **public reporting** frequently highlights China as a key supplier of parts for Russian unmanned platforms, a closer review of recovered systems indicates that a substantial share of AI-relevant components

4. DOSAAF (Добровольное общество содействия армии, авиации и флоту) stands for the Volunteer Society for Cooperation with the Army, Aviation, and Fleet. It is a state-affiliated paramilitary organization in Russia (with Soviet origins) that focuses on premilitary training and patriotic education. DOSAAF provides instruction in skills such as drone operation, radio communications, marksmanship, driving, and basic aviation training, helping prepare civilians—especially youth—for potential military service and supporting mobilization readiness.

originate from companies headquartered in the United States and other Western countries. The analysis does not suggest intentional support but rather shows the structural vulnerabilities of globally integrated semiconductor and electronics markets, where dual-use technologies remain widely accessible despite sanctions and export control regimes.

> *A substantial share of AI-relevant components originate from companies headquartered in the United States and other Western countries.*

This analysis draws on the open-source foreign components database maintained by Ukrainian Military Intelligence through the **War & Sanctions portal**. As of January 2026, the database contained more than 5,350 electronic components recovered from a wide range of Russian weapons systems and unmanned platforms. For the purposes of this study, the dataset was narrowed to components identified in UASs and then further isolated to those directly relevant to AI functionality. Rather than cataloging all electronics, the analysis focused on hardware that enables onboard autonomy. These components fall into three functional categories that together constitute a UAS's AI hardware stack: sensing components, processing units, and memory and storage hardware. Applying these filters yielded 118 sensors, 465 processors, and 122 memory units—a total of 705 AI-relevant components embedded in recovered Russian unmanned systems.

The emphasis on these three categories reflects the hardware foundations of modern AI systems. Compute and memory form the core infrastructure required to run neural networks, which consist of millions or billions of parameters and require substantial processing power to execute large numbers of operations in real time. In unmanned systems, however, sensing hardware represents an equally critical input. Sensors generate the raw data (e.g., visual, thermal, inertial, or positional) that onboard processors can transform into actionable information. Together, sensing, compute, and memory define the technical ceiling of autonomy achievable at the platform level, making them central to understanding how Russia equips its drones for increasingly AI-enabled operations.

As illustrated in Figure 6, companies headquartered in the United States account for the largest share of AI-relevant components identified across all three functional categories. The countries listed reflect the locations of manufacturer headquarters. The United States' lead is most pronounced in memory hardware, where U.S. firms produced approximately 69 percent of the recovered units, and the country's portion of processors remains substantial, accounting for a little under 57 percent of identified computing components. The sensor supply chain is more geographically diversified, yet the United States still represents the single largest country of origin, at approximately 38 percent, followed by China at about 16 percent, and Japan at a little over 15 percent.

China accounts for less than 9 percent of AI-enabling components identified in recovered UASs and does not rank among the top three countries of origin for onboard computing hardware. Switzerland is a notable supplier of processors, providing over one-fifth of recovered computing components. The Netherlands contributes approximately 8 percent of processors and less than 14 percent of sensors, while Japan's role is concentrated primarily in sensing technologies, supplying a little over 15 percent of recovered sensor components. These figures show the highly internationalized character of the electronics ecosystem underpinning Russian unmanned systems.

Figure 6: Country of Origin of AI-Enabling Components in Recovered Russian Unmanned Systems

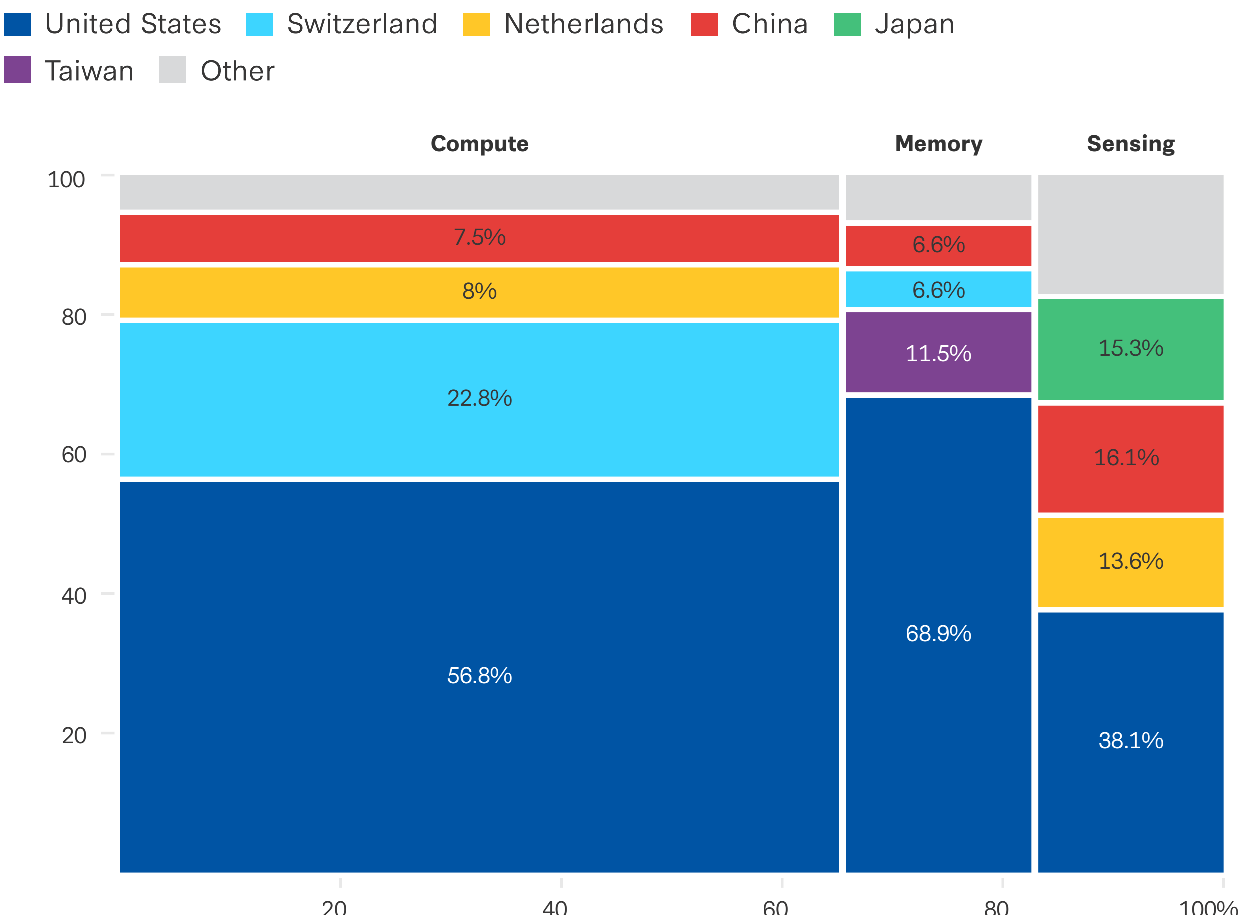

Note: Chart uses open-source data retrieved January 2026 from War Sanctions portal of the Defense Intelligence Agency of Ukraine. Country reflects location of manufacturer headquarters.

Source: CSIS.

The processors identified in the dataset differ fundamentally from the large-scale data center infrastructure commonly associated with AI compute. Battlefield signal jamming makes reliance on remote servers infeasible, requiring real-time inference to occur onboard the unmanned system itself. This constraint imposes strict size, weight, and power limitations, placing the most advanced chips out of reach but still allowing meaningful variation in performance, particularly in terms of latency.

Graphics processing units (GPUs) remain the most effective architecture for real-time AI applications, as they execute parallel calculations with minimal delay, which is an operationally decisive factor when milliseconds can determine mission success. While advanced algorithms can run on non-GPU systems, they typically do so with slower processing speeds or reduced performance. This analysis includes all processors theoretically capable of supporting AI functions, while noting that the precise operational impact of these components on Russian battlefield autonomy cannot be fully determined from open-source evidence.

Several Nvidia systems identified in the dataset align with previously documented AI-enabled Russian platforms. These include Nvidia Jetson Orin Developer Kits **recovered** from a **Shahed-136** and a **V2U barrage munition**, as well as a Jetson TX2 module found in a **ZALA 421-16E**—platforms assessed in earlier sections as demonstrating advanced onboard AI capabilities. Processors alone, however, are insufficient. On-chip memory is limited, and AI workloads require additional external memory to store model parameters and temporarily process high-dimensional sensor inputs. These external memory components form a critical part of the onboard AI hardware stack.

Sensors complete this triad. Low-level instruments such as accelerometers, gyroscopes, and magnetometers stabilize flight and support navigation, while higher-order systems—including optical, infrared, and radar sensors—generate the environmental data required for autonomous navigation and target recognition. Together, compute, memory, and sensing hardware establish the technical ceiling of platform-level autonomy.

The presence of these components does not by itself confirm operational autonomy, but it does demonstrate that the hardware foundations for AI deployment are present. The findings also complicate narratives that emphasize China as the primary supplier of Russian drone components. While China plays a role in the broader electronics ecosystem, Western firms, particularly U.S.-headquartered companies, remain central suppliers of AI-relevant hardware. Given the dual-use nature of these technologies and the globally integrated semiconductor supply chain, restricting Russia's access to onboard AI capability presents a large and complex policy challenge.

INTERNATIONAL PARTNERSHIPS ADVANCING RUSSIA'S AI AND UNMANNED SYSTEMS

Western sanctions imposed after Russia's annexation of Crimea in 2014 and expanded following the full-scale invasion of Ukraine in 2022 **have constrained** the development of Russia's domestic AI sector by limiting access to critical technologies, including microelectronics and advanced semiconductors. At the same time, the war has accelerated the **outflow of technical talent** and **promoted the departure** of foreign technology companies, further widening **Russia's gap** with leading AI powers such as the United States and China.

Russia has sought to mitigate sanctions and technological constraints by expanding bilateral and multilateral cooperation on AI. This section examines the external "enablers" supporting the development of Russia's AI ecosystem, focusing on partnerships within frameworks such as BRICS and bilateral ties with countries including China, Iran, and India. Through these relationships, Russia gains access to R&D, computing infrastructure, dual-use technologies, and technical expertise that can support both civilian and military AI applications. The analysis draws on open-source materials, including intergovernmental agreements, summit declarations, corporate announcements, media reporting, and think tank research.

BRICS

BRICS—an intergovernmental group of Brazil, Russia, India, China, and South Africa—has in recent years expanded its agenda to prioritize AI beyond general science and technology cooperation. While most AI collaboration among members still occurs bilaterally and multilateral initiatives remain relatively new, working groups, summit declarations, and joint research and infrastructure projects are increasingly emerging. As a result, BRICS is becoming an additional platform through which Russia seeks to mitigate technological constraints and expand its AI capabilities.

Although BRICS cooperation in science and technology predates 2017, AI was first explicitly referenced in the **2017 Xiamen Declaration**, where member states committed to expanding joint research, development, and innovation, including in areas such as cloud computing and AI. Cooperation began to accelerate after 2023, when members agreed to establish an **AI Study Group**. Following the 2024 BRICS Summit in Kazan, two major initiatives were launched: the **BRICS AI Research and Innovation Centre**, which now supports projects in multilingual natural language processing, smart agriculture, and climate research, and the **BRICS+ Digital Cloud Corridor**, designed to develop shared data storage and high-performance computing infrastructure, reducing reliance on Western servers.

A further milestone came in 2025, when BRICS leaders adopted the **BRICS Leaders' Statement on the Global Governance of Artificial Intelligence**. The document emphasizes open-source development as a mechanism to expand AI research and innovation capacities among member states. Although the guidelines formally apply only to non-military uses of AI, access to civilian AI tools, shared data infrastructure, and alternative cloud and computing networks can still support Russia's defense sector due to the dual-use nature of these technologies.

Moscow has simultaneously sought to shape emerging BRICS AI initiatives. At Russia's **AI Journey Conference** in 2024, Vladimir Putin announced the launch of a **BRICS+ AI Alliance**, intended to link national AI associations and development institutions across member and partner countries, facilitate joint research, and expand regulatory coordination. Russia has also begun concrete projects, including a partnership between the Russian Direct Investment Fund and data center operator BitRiver to **build AI computing infrastructure** for BRICS participants.

For Russia, the strategic value of BRICS AI cooperation lies primarily in open research collaboration and shared technological infrastructure, which can provide access to broader data and development ecosystems that Moscow cannot easily reproduce domestically. Even where partner countries are not leading innovators in AI, the framework allows Russia to remain integrated in international AI networks and sustain its position within the global AI ecosystem.

China

As the analysis in the previous section on components demonstrated, many of the technologies most critical for AI capabilities, particularly advanced chips and microelectronics, are primarily produced by Western companies, especially in the United States. This makes Russia heavily dependent on foreign suppliers for the hardware that underpins AI development and deployment.

In the aftermath of the 2022 invasion of Ukraine and the introduction of strict export controls **and sanctions** by the United States and **its allies**, China has emerged as the principal channel through which these components continue to reach Russia. While **direct exports** from Western and allied countries declined sharply, shipments from China and Hong Kong increased significantly, often supplemented by transshipment through intermediary countries. As a result, China has become the main conduit enabling Russia to sustain access to critical chips and microelectronics that support computing infrastructure and AI-enabled military applications.

Beyond providing access to critical components and computing infrastructure, China also supports Russia's AI development through research collaboration and technology partnerships. At the political level, Moscow has sought to institutionalize this cooperation. In January 2025, President Putin **directed** Sberbank, a state-controlled financial institution and a central actor in Russia's national AI strategy, to expand cooperation with Chinese partners, including exploring the creation of an international journal dedicated to AI technologies. This directive reflects high-level political commitment to deepening bilateral collaboration in AI R&D.

Chinese technology companies have also played a visible role in sustaining Russia's AI ecosystem. Huawei has **expanded partnerships** with Russian universities, research institutes, and technology companies since 2014, further deepening these ties after its placement on the U.S. Entity List in 2019. The company has signed **cooperation agreements** with Russian initiatives such as the National Technology Initiative and has partnered with firms including **Kaspersky** and **Sberbank** on cloud computing and cybersecurity infrastructure. Such collaborations provide Russian organizations with access to high-performance computing environments and cloud services that can support data-intensive AI applications, including computer vision and autonomous navigation technologies relevant to unmanned systems.

Although the direct military implications of these projects remain largely opaque, official statements increasingly acknowledge growing cooperation in AI-related defense issues. In February 2024, Russian and Chinese officials publicly noted discussions on "**doctrinal approaches and initiatives related to the military use of artificial intelligence**," signaling a willingness to explore collaboration in defense-related AI domains. Technical exchanges have also expanded. For example, Rostec has **sent engineering students** to China for internships focused on AI applications in aircraft engine design and robotics, illustrating broader channels of technological transfer.

These initiatives position China as Russia's most important external partner in sustaining its AI development. Beijing provides Russia with access to microelectronics supply channels, computing infrastructure, and research collaboration that help offset the constraints imposed by Western sanctions. At the same time, the relationship remains asymmetric. China retains the dominant technological position and has **repeatedly emphasized** international limits on fully autonomous weapons, whereas Russian officials **have advocated** minimal restrictions on the development of such systems. This divergence suggests that while cooperation may accelerate Russia's AI-enabled military capabilities, it will likely remain conditioned by China's broader strategic interests.

Iran

Since 2022, **military cooperation** between Russia and Iran has expanded significantly, centered primarily on **technology transfer and coproduction** in UASs. Tehran initially supplied Russia with Shahed-131 and Shahed-136 loitering munitions, along with technical expertise that enabled Moscow to localize production and **scale manufacturing** inside Russia, including at the Alabuga facility. These transfers provided Russia with a foundation for developing the Geran series of strike drones, which Russian engineers have subsequently modified with improved navigation, anti-jamming features, and **in some cases** AI-enabled components such as computer-vision targeting and autonomous navigation capabilities.

While the precise level of Iranian involvement in these technological upgrades remains unclear, cooperation has continued through technical training, **joint production**, and formal agreements—including a memorandum of understanding on AI signed in 2025.

At the same time, the technology flow appears increasingly reciprocal. **Reports in 2026** indicate that Iran has used Russian-produced Geran-2 drones—a variant of the Shahed-136 assembled at Kupol plant in Izhevsk—in operation against the United Arab Emirates, suggesting that battlefield experience, manufacturing improvements, and drone technologies are now circulating in both directions. These exchanges demonstrate that Russia-Iran cooperation in UAS systems has evolved into a mutually reinforcing channel for technological transfer and operational learning.

India

India represents another channel through which Russia can access technologies relevant to AI development, although this cooperation remains limited and primarily civilian in scope. Collaboration occurs mainly across three areas: the supply of dual-use technologies, business agreements between technology firms, and academic partnerships between Russian and Indian institutions. While these initiatives support Russia's broader technological base, their direct connection to battlefield AI applications remains unclear.

Indian companies have emerged as an important intermediary for the restricted technologies needed for AI infrastructure. Because India has not joined Western sanctions, Indian firms have been able to **legally export** microelectronics, computing equipment, and industrial machinery to Russia under domestic law. Reporting in 2024 indicated that India had become the **largest supplier** of restricted technologies to Russia after China—including shipments of advanced servers containing $300 million worth of Nvidia chips and other high-performance computing components—through, for example, drug-making companies. Such hardware can support the data processing and model training required for AI-enabled systems. This does not come without risk: the United States has **sanctioned** some Indian firms for their involvement in supplying Russia with dual-use technology.

Cooperation between Russian and Indian entities on civilian technology is also expanding through a series of business and academic agreements. During a business mission to India by industry group RUSSOFT in March 2024, several organizations, including PaPSwap State Policy Center, Kanninnov Technologies, the Chamber of Indo-Russian Technological Cooperation, and Optimus Logic Systems, **signed agreements** to cooperate on integrated circuit development and AI applications in sectors such as agriculture and healthcare. RUSSOFT has also **supported plans** to establish a Russian-Indian Center for the Development of Trusted Hardware and Software Complexes in Uttar Pradesh.

Academic partnerships further reinforce these ties. For example, IIT Kharagpur and Saint Petersburg Mining University **signed an agreement** to promote joint research and academic exchanges, including AI applications in energy and earth sciences. Similarly, Russia's Innopolis University partnered with the Times School of Media in India to establish an **international AI-focused research laboratory**.

Direct cooperation between India and Russia on AI for battlefield use remains limited, as no publicly available business or academic agreements explicitly address military AI applications. However, recent developments suggest potential collaboration in the codevelopment and production of advanced defense technologies. Following President Putin's visit in December, both countries **announced plans** to expand their partnership toward joint research, development, and production of advanced systems. Subsequent discussions between Russian officials and Indian defense firms, including drone and military AI startups, **reportedly explored** cooperation on producing components for MiG-29 fighter jets and other defense systems. There are **also reports** that Russia has discussed localizing production of certain UAS systems in India, including the Lancet loitering munition, which incorporates AI-enabling components.

Despite these developments, India's cooperation with Russia in AI is likely to remain constrained. New Delhi continues to pursue a strategy of strategic autonomy and seeks to balance its relationships with both Western technology partners and Russia. As a result, while India may selectively support Russia's access to dual-use technologies and research collaboration, deeper alignment—particularly in military AI—remains limited.

Strategic Conclusions and Policy Recommendations

Russia's trajectory in AI and unmanned systems development is not defined by a pursuit of frontier technological breakthroughs. Instead, it reflects a more consequential effort—the deliberate construction of a sovereign, end-to-end ecosystem designed to translate applied AI into battlefield advantage at scale. Moscow's approach combines top-level strategic direction with the pragmatic adoption of open-weight models, dual-use industrial innovation, modular hardware design, workforce development, and flexible regulatory experimentation. Increasingly, these parallel streams are being integrated under a centralized system of state coordination.

This analysis points to three central conclusions about how Russia is advancing AI-enabled unmanned systems.

First, Russia treats AI and unmanned systems not as isolated capabilities, but as elements of a unified national ecosystem. Civilian and military structures are tightly interconnected, allowing innovation to flow across sectors. Infrastructure, regulation, training, industrial production, and doctrine are aligned toward a shared objective—achieving autonomy across the entire ecosystem. This includes a wide range of ecosystem components, from domestic UAS parts production to the tactical edge of AI applications, leveraging advances in AI and mobilizing leading technological actors.

Second, battlefield effectiveness has emerged not from elegant autonomy architectures, but from low-cost, modular, rapidly iterated systems embedded with narrowly scoped but mission-critical AI functions. Scale, speed, and feedback loops have proven more decisive than technological sophistication alone.

Third, despite rhetoric about technological sovereignty, Russia's autonomy stack remains deeply embedded in global supply chains—particularly Western commercial-grade electronics. More than half of AI-enabling components originate from U.S.-headquartered firms. This exposes structural vulnerabilities and highlights the difficulty of controlling dual-use diffusion in globally integrated semiconductor markets.

For the United States, the central lesson is that ecosystem coherence, not individual programs, determines success in AI-enabled unmanned warfare.

RECOMMENDATIONS FOR THE U.S. GOVERNMENT

1. **Build a comprehensive civil-military autonomy ecosystem.** If the United States is serious about an unmanned future, the approach cannot remain siloed within the Department of Defense. Civil and military ecosystems must evolve together. Education pipelines, regulatory regimes, spectrum access, airspace integration, commercial incentives, and defense procurement must be aligned as part of a coherent national architecture. Autonomy is not a platform—it is a system spanning talent, infrastructure, compute, testing environments, and doctrine. Without ecosystem-level coordination, tactical initiatives will remain fragmented.
2. **Translate "drone dominance" into a government-wide implementation framework.** The **Drone Dominance initiative** and executive-level guidance signal intent, but intent must be operationalized. The United States requires a practical, cross-agency implementation road map that connects high-level strategy to programmatic execution across the Federal Aviation Administration, Federal Communications Commission, National Telecommunications and Information Administration, and Departments of Defense, Commerce, Education, and Homeland Security. Autonomy policy cannot remain confined to defense channels; it must synchronize regulatory, industrial, and workforce levers across government.
3. **Incentivize dual-use technology at scale.** Russia's most operationally effective AI integration occurs in companies operating across civilian and military markets. Dual-use ecosystems generate larger datasets, faster experimentation cycles, greater talent retention, and more stable financing. The United States should expand procurement pathways and financing mechanisms that allow commercial AI and unmanned firms to sustain civilian revenue while iterating defense-relevant capabilities. This approach mitigates the "valley of death"—where platforms rarely progress from prototype to deployment—and accelerates technology maturation.
4. **Create real denied-environment testing conditions.** Edge autonomy must be tested in environments that simulate EW, GPS denial, and spectrum congestion. Current regulatory and coordination processes—particularly for spectrum and jamming—are slow and fragmented. The United States should establish designated autonomy testing corridors and denied-environment ranges, coordinated jointly by the Department of Defense, Federal Aviation Administration, Federal Communications Commission, National Telecommunications and Information Administration, and United States Strategic Command, with streamlined approval timelines. Such environments would benefit both military readiness and civilian resilience applications, including emergency response and critical infrastructure protection.
5. **Scale workforce preparation for unmanned systems.** Training pipelines for operators, maintainers, AI specialists, and spectrum managers must expand dramatically. The United States should institutionalize unmanned systems instruction earlier in education pathways, support drone-focused vocational programs, and integrate autonomy competencies into the ROTC, service academies, and civilian STEM tracks. Private-sector drone schools and commercial training centers should be incorporated into the national training architecture rather than viewed as peripheral actors.

Russia's experience in Ukraine demonstrates that autonomy will not emerge as a singular technological breakthrough. It will emerge from the systematic alignment of policy, industry, training, infrastructure, and operational adaptation.

The United States retains structural advantages in talent, innovation capacity, and industrial depth. But advantage alone does not guarantee coherence. If autonomy is to shape the future force, it must be treated not as a technology program, but as a national systems project. ■

Kateryna Bondar *is a fellow with the Wadhwani AI Center at the Center for Strategic and International Studies in Washington, D.C.*

The author would like to thank Research Assistants Nicole Errera and Matt Mande with the CSIS Wadhwani AI Center and Program Manager Kirtika Sharad at the CSIS Chair for U.S.-India Policy Studies for their contributions to this research.

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

COVER PHOTO: ANTON PETRUS/GETTY IMAGES; CONTRIBUTOR/GETTY IMAGES

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

1616 Rhode Island Avenue NW
Washington, DC 20036
202 887 0200 | **www.csis.org**