JANUARY 2026

The Agentic Revolution in War

THE PRESENT AND FUTURE OF DECISION ADVANTAGE

scale

Should you have questions, want to learn more about agentic capabilities, or wish to provide feedback on this paper, please reach out to **agenticwarfare@scale.com**.

For digital access to this paper, visit **scale.com/agentic-warfare**.

SCALE'S MISSION

Develop reliable AI systems for the world's most important decisions

ABOUT SCALE AI

Founded in 2016 by Alexandr Wang, Scale's mission is to develop reliable AI systems for the world's most important decisions. Growing along with the evolving AI industry, Scale has worked with leading autonomous vehicle companies, frontier AI labs, Fortune 500 companies, and governments around the world. Today, Scale builds data and model training solutions as well as complex, multi-step AI applications for enterprise and public sector clients. Since its early years, Scale has developed deep experience supporting the U.S. national security community, fielding proven and trusted capabilities across all levels of classifications to tackle the most consequential operational problem sets. Our unique capabilities in training, testing, and evaluating AI models, coupled with our sophistication in building applications, make us the first-choice partner to help the Department of War and intelligence community navigate the opportunities and threats posed by AI and the emerging "agentic battlefield." Scale and its Public Sector workforce, which is heavily veteran and former government civilian, is committed to delivering robust, mission-ready capabilities to ensure we secure America's strategic advantage into the future.

AGENTIC WARFARE:

The Opportunity

The era of Agentic Warfare has begun.

The first nation to fully operationalize agentic systems in military decision-making will shape the course of the 21st century. This defining shift moves beyond the development of large language models toward agency: systems capable of not only responding to human prompts, but actually executing and accomplishing complex tasks at paradigm-shifting speeds. These systems consist of multiple AI agents, each performing specific and coordinated tasks, forming constellations of immense computational power. Harnessed by the Department of War, these systems enable U.S. forces to outpace and outmaneuver even our most capable opponents.

When fully developed, agentic systems realize the imperative put forward by former Chairman of the Joint Chiefs Joseph Dunford that “our decision-making processes…deliver options at the speed of war.” By coordinating autonomous agents to adopt strategies at machine speed, they supercharge human cognition and respond to battlefield changes in real time. To probe for vulnerabilities, they run millions of physics-based simulations, optimizing precise courses of action to maximize the probability of victory. By compressing the time required for analysis and making that analysis more accurate, they enable commanders to execute nimble, decisive maneuvers that far outpace the capabilities of current command, control, and planning structures. Simply put, agentic systems deliver every option, and help us understand optimal solutions, before the enemy has had their say.

This paper provides a blueprint for how the Department of War can harness agentic systems to achieve new degrees of decision advantage, and in so doing, revolutionize the American way of war.

Dan Tadross
Head of Public Sector

Jared Jonker
Senior Director, Agentic Warfare

EXECUTIVE SUMMARY:

The Agentic Revolution in War

Over the last 50 years, digital technologies have driven waves of battle-winning advantage, from precision strike to integrated intelligence, surveillance, and reconnaissance (ISR) to network-centric warfare. In each case, software accelerated human-centric command processes without fundamentally changing their nature. Today, the emergence of agentic capabilities marks a definitive break from the past. We are entering the era of Agentic Warfare.

Agents are proactive, goal-driven systems that combine AI capabilities—like memory, tools, and control logic—to perceive, reason, and act with some degree of autonomy, performing tasks guided by human intent and oversight. The defining feature of agentic systems is agency: not simply ingesting and correlating information, but using it to plan, test, and execute complex multi-step actions, an evolution beyond previous generations of artificial intelligence based upon text generation or machine-speed pattern matching.

While today's narrow and brittle artificial intelligence (AI) capabilities have performed, at best, like a clever junior staffer, agentic systems unlock the full potential of AI to act as a genuine mission partner. This is a necessary, transformative shift, as traditional command and control architectures have long failed to keep pace with the velocity and volume of modern information systems. When realized, agentic systems will eliminate the "interruption, pause, or suspension of activity" Carl von Clausewitz warned against.

Genuine strategic advantage in this new era will not come from stealthier jets, faster missiles, or larger drone swarms alone; it will come from new kinds of human-machine teaming that drive accelerated decision-making. This is the essence of Agentic Warfare: **decision advantage at every echelon of command** that enables U.S. forces to outpace and outmaneuver our most capable opponents. The United States must capitalize on its first-mover advantage before adversaries do.

> Genuine strategic advantage in this new era will not come from stealthier jets, faster missiles, or larger drone swarms alone; it will come from human-machine integration that drive accelerated decision-making.

Imagine: constellations of specialized AI agents, collaborating to supercharge human insight across the full Observe, Orient, Decide, Act (OODA) loop, from strategic signaling down to tactical movements. One system processes satellite imagery; another cross-references SIGINT; another examines political and economic shifts; another runs physics-based simulations. Together, they tackle levels of complexity that would swamp even the best human staff. By running thousands of scenarios in parallel, these new systems enable U.S. forces to better probe for enemy weaknesses, distinguish threat from bluff, and optimize strategies at a pace that far outstrips an adversary's decision cycle. It's like a game of chess where one player can make three moves while the other can only make one.

In the 20th century, deterrence rested on the certainty that any attack would be met with overwhelming force. Today, decision advantage plays the role of deterrent. When an adversary knows U.S. forces can see the battlespace more clearly, adapt plans in real time, and synchronize thousands of assets with deadly precision, the logic of aggression begins to crumble. **Agentic Warfare is deterrence by decision advantage**: the promise that any act of aggression will be met with force guided by foresight.

> Decision advantage is like playing a game of chess in which one player makes three moves while the other is only able to make one.

NEW TOOLS FOR COMMAND: TWO LEVERAGE POINTS

Agentic AI is ready to deliver greater decision advantage today in different parts of the OODA loop: the Observe/Orient phase, where commanders are blinded by data noise, and the Decide phase, where planning is bottlenecked by manual analysis. The Department's early work with industry partners applies agentic systems to these choke points:

- **Agentic Alerting (The "Observe/Orient" Solution):** In today's saturated operational environments, the "Observe/Orient" phase is defined by a deluge of sensor data that overwhelms human cognition. Adding even more capability to existing common operating pictures, Agentic Alerting restores the signal-to-noise ratio. By ingesting multi-modal feeds and identifying anomalies at machine speed, it turns a reactive force into a proactive one—prioritizing the alerts that matter and redirecting ISR assets and other sensors before a danger fully materializes.

- **Agentic Planning (The "Decide" Solution):** The "Decide" phase is currently constrained by planning cycles that take months to produce binders of static options. Agentic Planning systems aim to break this linear bottleneck by providing planners AI agents coupled with automated, physics-based modeling and simulation tools. These tools allow planning staff to generate validated, confidence-bound courses of action (COAs) and iterate upon them rapidly, supplying commanders with COAs adapted to current conditions on the battlefield. This ensures that when commanders make the decision to act, they are acting on tested, probability-weighted analysis rather than stale assumptions.

Reinventing how we monitor the battlefield and replace brittle staffing processes with dynamic, living systems represents a fundamental shift in the nervous system of the force. However, technology alone does not guarantee success. Like previous military innovations—from blitzkrieg to carrier aviation—victory will belong to the side that not only acquires the capability, but reimagines their way of war to exploit its full potential.

TRUST, CONTROL, AND THE "HUMAN ON THE LOOP"

The adoption of agentic AI demands a shift in how military leaders think about command and control. The sheer speed of modern warfare, accelerated still further by machine processes, increasingly outpaces human reaction times. We must move from keeping humans "in the loop," bottlenecking and stalling action, to keeping commanders "on the loop," elevating them from hands-on controllers to mission directors who provide oversight. In some circumstances, such as electronically denied environments, systems may need to operate with full autonomy for extended periods.

This transition is profound. Today's playbooks assume humans do most perceiving, fusing, and planning, with software in support. In an agentic force, the first pass in many of those steps flips: agents propose interpretations of developments on the battlefield and options for how to counter them; humans set intent, apply judgment, and own the risk. Yet the Department currently lacks the doctrine, training, and institutional velocity to realize this transformation at the necessary pace.

The constraint facing the Department of War is no longer primarily technical: it is cultural and organizational—one of change management to drive adoption and reconceptualize doctrine. To achieve enduring superiority, we must treat decision advantage as a central organizing objective of modernization. Agentic systems won't replace the art of command, but they will ensure that when U.S. commanders make decisions, they do so with a clarity, confidence, and speed that no adversary can match.

> Today's playbooks assume humans do most perceiving, fusing, and planning, with software in support. In an agentic force, the first pass in many of those steps flips: agents propose interpretations of developments on the battlefields and options for how to counter them; humans set intent, apply judgment, and own the risk.

Reliability is a fundamental challenge: commanders will not and must not delegate authority to carry out their intent to a "black box" without justifiable confidence that the system is robust, predictable, and aligned. This makes Test and Evaluation (T&E) a strategic imperative. Just as the Pentagon does for all systems, Agentic AI systems must be developed hand-in-glove with operators, stress-tested in realistic and adversarial conditions, reassessed throughout their life-cycle, and released within clear policy envelopes so human operators properly understand where operational realities might exceed training parameters.

> Agentic systems won't replace the art of command, but they will ensure that when U.S. commanders make decisions, they do so with a clarity, confidence, and speed that no adversary can match.

THE CHALLENGE AHEAD

The window for the United States to secure decision advantage is closing fast. The coming years will be pivotal as agentic capabilities mature and uncrewed systems proliferate. China is already racing to build 'intelligentized' forces in which 'command brains' and autonomous swarms compress U.S. decision times, multiplying the People's Liberation Army's advantages of geography and mass in their near-abroad. A land war remains ongoing in Europe. Drones and loitering munitions drove the most violent exchanges in the Middle East in twenty years. Meanwhile, threats to the homeland are proliferating.

The United States must mobilize with the urgency of a nation at war. Rather than wielding wrenches as we did to win World War II, we must instead build the compute, data, networks, and agent-based systems that will deliver agentic capabilities to warfighters and commanders. Early investments in prototype programs like Agentic Alerting and Agentic Planning have positioned the United States to realize meaningful gains in combat power and decision speed. The table is set; what matters now is moving fast to capitalize on our first-mover advantage by accelerating the next generation of capabilities already in development. These will create the foundation for our most strategic weapon: speed to decision.

TABLE OF CONTENTS

PART I

Decision Advantage

"Every element of the joint force can, must and will use advanced technology to improve our command and control systems, our decision-making, our execution and, frankly, our survivability."

CHAIRMAN OF THE JOINT CHIEFS GENERAL DAN CAINE[1]

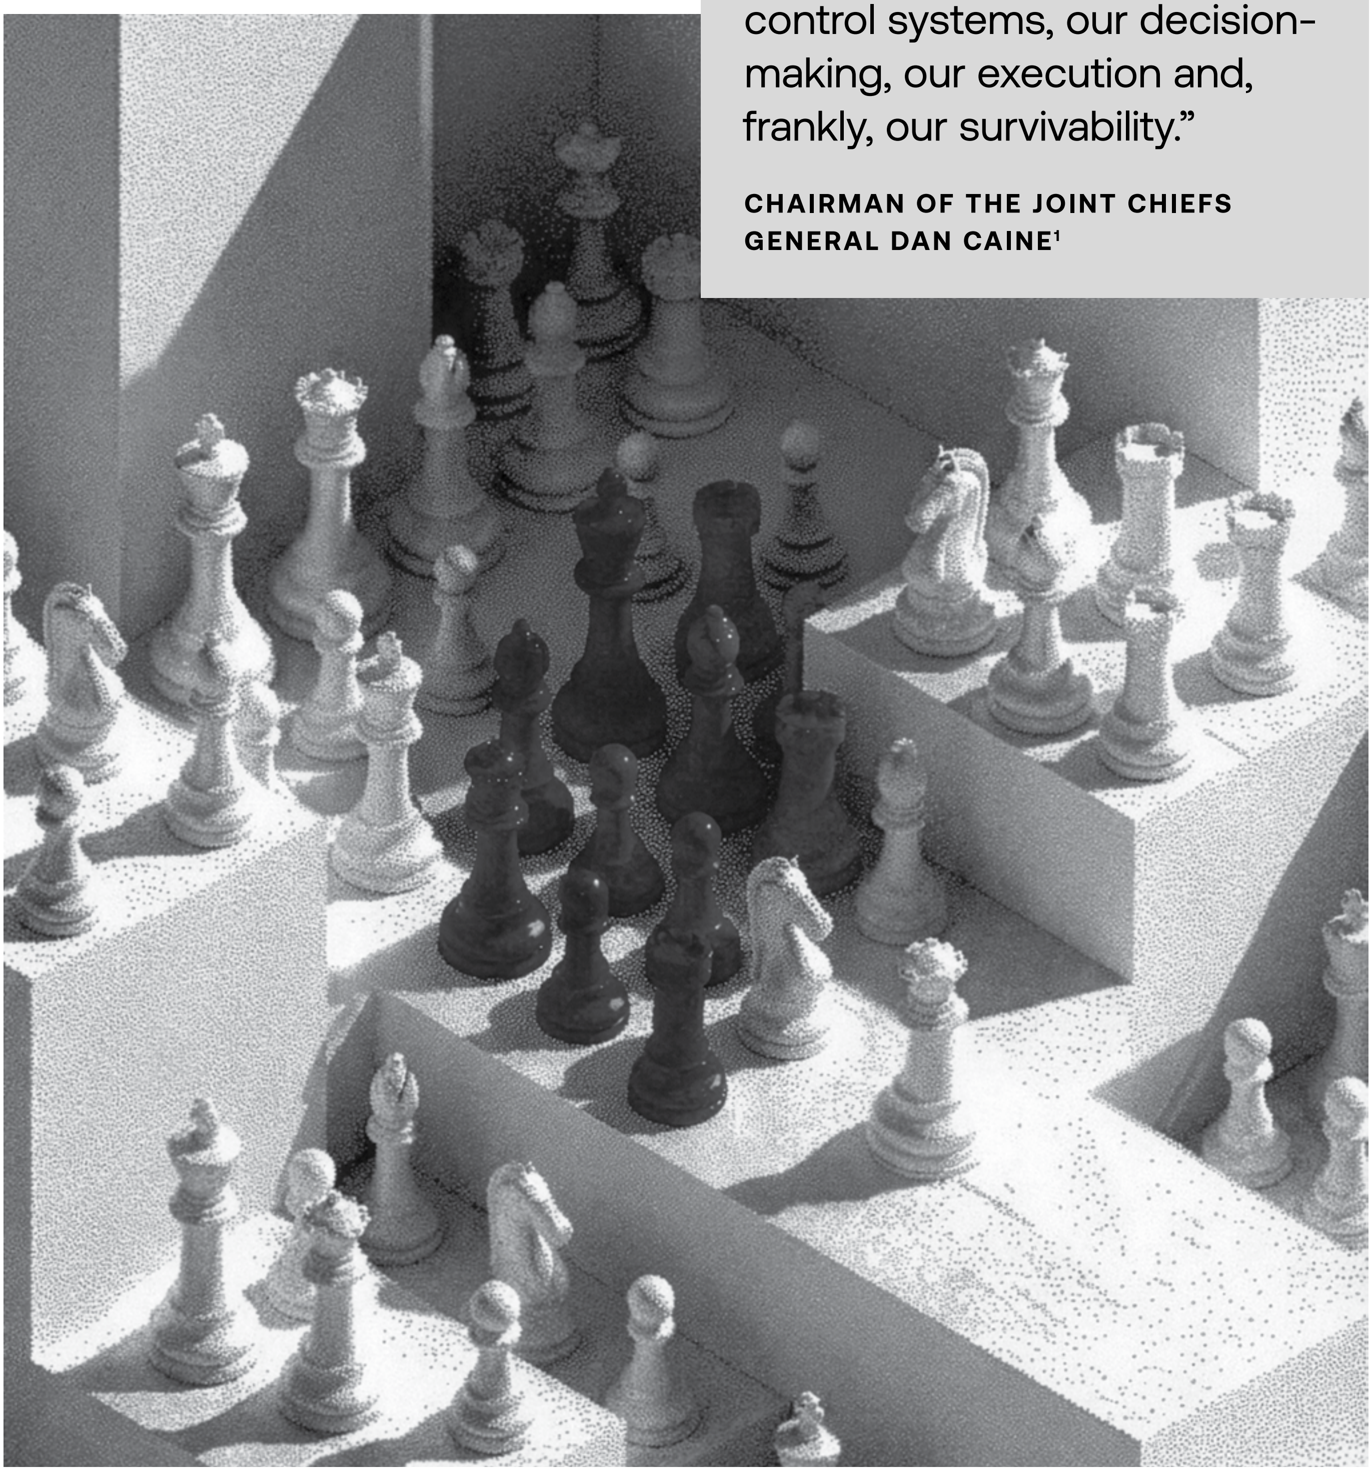

At a moment when AI technology is enabling whole new ways of war, we need to lead decisively from the front. The first nation to fully incorporate agentic systems into military decision-making will shape the 21st century.

The West won at a similar moment during WWII, turning new technologies of command and control into an enduring advantage that persists to this day. For four perilous months, the fate of the free world depended upon a group of women in blue-gray battledress watching cathode-ray tubes for vertical blips; any spike above baseline signaled Luftwaffe bomber formations crossing the English channel.

This was the world's first integrated air defense network, the Dowding System. It prevented the destruction of the Royal Air Force, thwarting a cross-channel invasion and halting Hitler's conquest of Europe. The system created human links of information transfer and synthesis, enabling Britain's four fighter commands to launch counter-attacks. Women in an underground "filter room" fused reports from radar operators to create a real-time picture of the battlefield, plotting wooden figures of bombers and fighters on a map. Headquarters then used verified plots from the filter room to launch Spitfire interceptors, with limited fuel and range, at optimal vectors of attack.

The Dowding system proved so foundational that, 85 years after saving Europe, it remains the blueprint for modern command and control. In today's military, digital chatrooms have replaced radio nets, and operations centers fuse exponentially more data. The process, however, remains fundamentally the same: ever-larger staffs synthesizing ever-larger flows of data to manually develop courses of action—typically on PowerPoint slides—presented to a commander for decision. While rigorous, this manual cadence is increasingly out of step with the velocity of modern conflict.

What if this human chain were augmented by an agentic system executing with greater precision and at speeds millions of times faster? In such a system, an agent detects the initial indicator of an attack, immediately triggering a swarm of agents to cross-check data feeds, confirming if the worrying signal represents a valid threat. The system then alerts key personnel and redirects intelligence assets, while simultaneously queuing planning and simulation resources to develop countermeasures for commander approval.

Agentic systems are poised to swiftly and forever change the process of command and control, yielding new degrees of decision advantage. This command, control, and planning system of the future is arriving sooner than anyone imagined.

This paper explores the agentic revolution in war in four parts.

Part I: Decision Advantage lays out the stakes for why the United States must achieve decision advantage, and how progress in AI agents is unlocking new ecosystems of tools and capabilities far beyond those provided by large language models (LLMs). It also describes why the best approach is using systems that mix and match models from diverse sources rather than being tied to one or a few.

Part II : Agentic Warfare Today examines two agentic systems in development today. One, Agentic Planning, opens the door to a new and faster kind of military planning, executing COA analysis at machine speed. The other, Agentic Alerting, places new capabilities into the military's real-time nervous system that can reorient how we observe at machine speed. Both introduce new forms of human-machine collaboration, moving decision-makers from being "in the loop," responsible for approving every single action, to "on the loop," where higher-level courses of action are presented for decision—a fundamental doctrinal shift demanded by modern war.

Part III : Reimagining America's Way of War explores how we must reimagine the American way of war as agentic capabilities mature and as our adversaries begin to adopt them as well. Two intertwined dynamics will play out: agentic systems evolving to take actions of greater consequence and commanders intervening only to adjudicate the most crucial decisions.

Part IV : Recommendations for Getting Agentic AI in the Fight identifies immediate steps that our warfighters, Department of War leadership, and Congress must take to convert the United States' first-mover advantage in agentic systems into enduring strategic superiority. This is how the United States grows its competitive edge in warfare.

DECISION ADVANTAGE – WHAT’S AT STAKE

Chairman of the Joint Chiefs Dan Caine aboard the guided missile cruiser USS Lake Erie.[2]

Today, conflict unfolds with greater speed and complexity than ever before.[3] Looking east to Europe, we see a land war where an underdog, Ukraine, has thwarted a more powerful adversary, despite its material disadvantage, by fielding novel technologies at a pace that has far out-stripped Russia’s development timelines for countermeasures. This dynamic has led to front lines that are essentially frozen in place.[4]

Looking west to the Pacific, any conflict could quickly become a theater-wide strategic competition in which adversaries have first-mover advantage, far shorter supply chains, and an operational geography that advantages their centralized command.[5] China is increasingly using AI to couple military mass with increased decision tempo and orchestration. We risk facing competitors that can increasingly match the capability of our platforms in their own periphery, while systematically out-thinking, out-sequencing, and out-scaling us from the opening moves of a crisis.

Speed to decision is a weapons system of its own.

Then there is the American homeland. For the first time in a generation, our adversaries believe they can strike within our borders below the nuclear threshold for retaliation. Cyber and critical infrastructure attacks are now joined by the threat of autonomous systems akin to the Ukrainian Spider Web drone attacks that destroyed strategic bombers deep inside Russia.[6] Combined with the increasing threat of ballistic missiles, these dangers create a security challenge that necessitates the defense that Golden Dome promises.

The importance of decision advantage in military operations unites these scenarios. It is a source of military power in a world in which the weapons systems and capabilities of our adversaries increasingly converge with our own. In this new paradigm, we must retain the ability to think and react faster than our adversaries: **speed to decision is a weapons system of its own**. In almost all cases, systems that produce decision advantage are the most cost-efficient way to deliver military advantage.

UNDERSTANDING AGENTIC WARFARE: A NEW MILITARY PARADIGM

In the 20th century, deterrence rested on the certainty that any attack would be met with overwhelming force. Today, intelligence plays the role of deterrent. If an adversary knows U.S. forces can see the battlespace more clearly, adapt plans in real time, and synchronize thousands of assets with deadly precision, the logic of aggression begins to crumble.

That is the essence of Agentic Warfare: decision advantage across every echelon of command.

Agentic systems are dynamic and goal-driven: always on, absorbing new data, reasoning through options, and adapting as situations evolve. In effect, they function less as applications and more as a corps of digital staff officers wired into every level of command.

These systems augment and amplify human judgement, mitigating bias and fatigue while expanding the amount of information and context humans are able to maintain. The commander still decides, but with far greater clarity and confidence.

At scale, agentic systems become the new military operating model. Agents cross-correlate everything from ISR feeds and operational reports to economic and political signals. Reasoning engines stress test thousands more “what if” variations each day than a human staff could evaluate in a lifetime. And they do so with Test & Evaluation-led assurance, where systems are frequently red-teamed, verified, and bound by guardrails and policy.

> The commander still decides, but with far greater clarity and confidence.

At the strategic level, agentic systems deliver ageless contingency plans. Closer to the fight, orchestration agents stitch together air tasking orders, naval maneuvers, cyber effects, and logistics into a single coherent plan, able to adjust courses of action proposed to commanders if key nodes fail or upon enemy surprise. At the tactical edge, agents rewrite the economics of conflict: one operator, empowered by agents, can choreograph dozens of ISR and strike drones, decoys and EW assets, flipping the lethality cost curve in America’s favor.

Agentic warfare is deterrence by insight: the promise that any act of aggression will be met by force guided by foresight.

THE SECOND WAVE OF AI AGENTS: INTELLIGENT AUGMENTATION

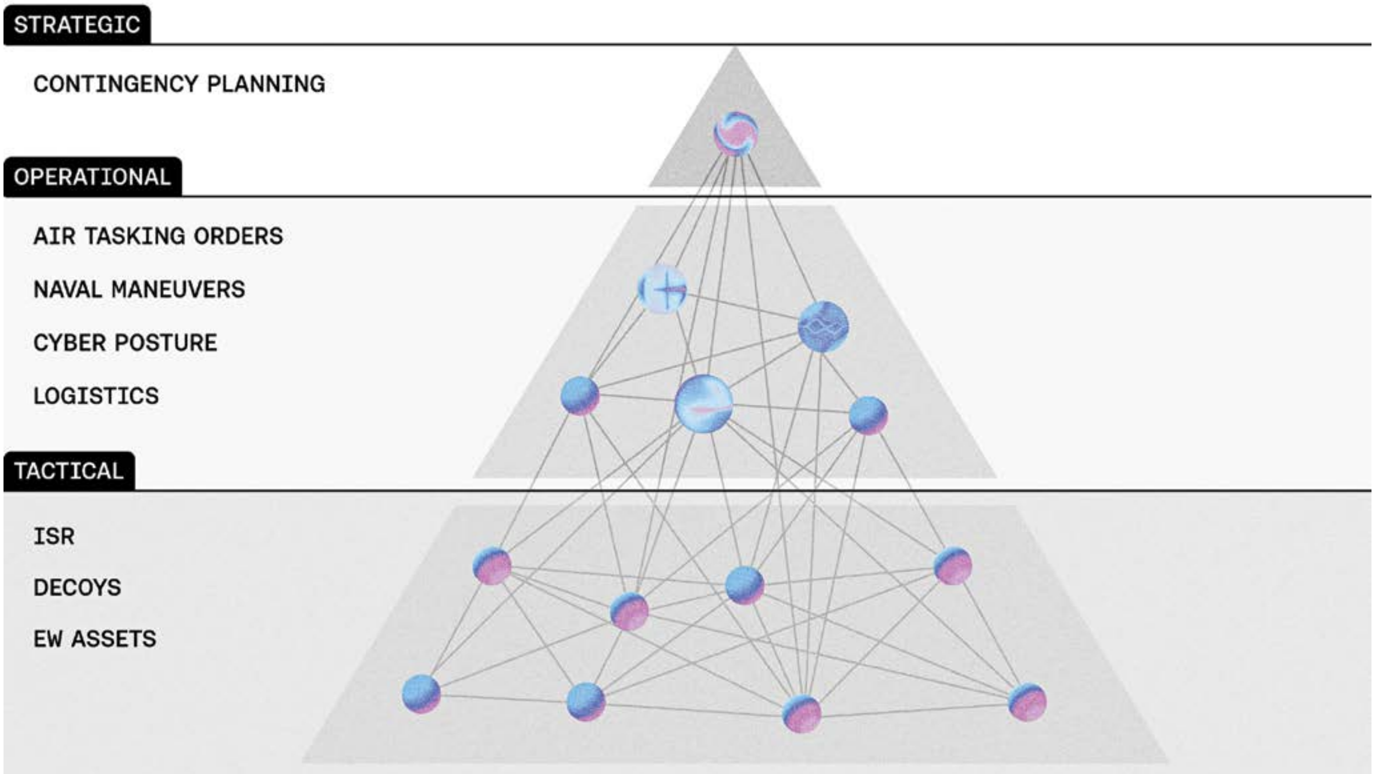

The Department of War has applied advanced computing and AI in multiple military missions. Among the earliest was Project SAGE (Semi-Automatic Ground Environment). A direct descendant of the UK's Dowding system, it utilized the largest computer ever constructed—IBM's AN/FSQ-7— weighing in at 250 tons with 60,000 vacuum tubes. This system powered the integrated air defense network used by North American Aerospace Defense Command (NORAD) from the late 1950s through the early 1980s. More recent efforts have included pioneering research organizations like the Defense Advanced Research Projects Agency (DARPA) and operationally-focused teams like Project Maven, which successfully introduced computer vision into the intelligence analysis and targeting chain.

All these projects have had some degree of success, but none have meaningfully altered or transformed military decision-making. DARPA was instrumental in the first phases of AI development through its work on Expert Systems, and even laid the groundwork for commercialized voice assistants like Apple's Siri and Amazon's Alexa. While instrumental to the development of the technology, these attempts at fielding AI were narrow and brittle and limited in scope and applicability. They impacted a single workflow or domain but lacked the capability to cope with the diversity and volume of data necessary to support operational and strategic decision-making. This changed when large language models began to show more generalized intelligence capabilities with the release of OpenAI's ChatGPT in November 2022.

Large language models (LLMs) brought the next wave of capability. While chatbots and copilots sped up workflows, tangible gains remained modest. Instead, LLMs often generated “walls of text” rather than operational cues or action. When an LLM is wrong, it can be hard to know why—outputs can be persuasive but often aren’t auditable. Mirroring commercial frustration with AI “slop,” these tools impressed military users but mostly failed to transform operations, and the U.S. military has struggled to integrate them.[7]

The next wave of AI technology brings mission-capable software composed of many AI agents that can perform in tandem. These systems will often be built on today’s advanced LLMs, integrated with many smaller, custom models. While earlier limitations remain, they are more effectively managed through architectural choices, constraints, verification, and rigorous T&E.

In addition to answering questions, agents fuse multi-modal intelligence, stress-test plans, and dynamically orchestrate actions. At the capability level, this is transformative: clusters of specialized agents—each using the right model and tools for its specific purpose—collaborating to sharpen situational awareness. At scale, these local gains compound into a genuine innovation in military affairs, unlocking the “Man-Computer Symbiosis” envisioned by AI pioneer J.C. Licklider—a future where AI ceases to be a mere tool and becomes a cooperative partner in the decision-making loop.[8]

HOW AGENTS WORK

The GenAI tools most people know, such as LLMs and video and image generators, are reactive tools that answer prompts generated by humans. Agents are proactive, goal-driven systems that combine AI capabilities—like memory, tools, and control logic—to perceive, reason, and act with some degree of autonomy, all under human intent and oversight.

A single agent typically has five layers:

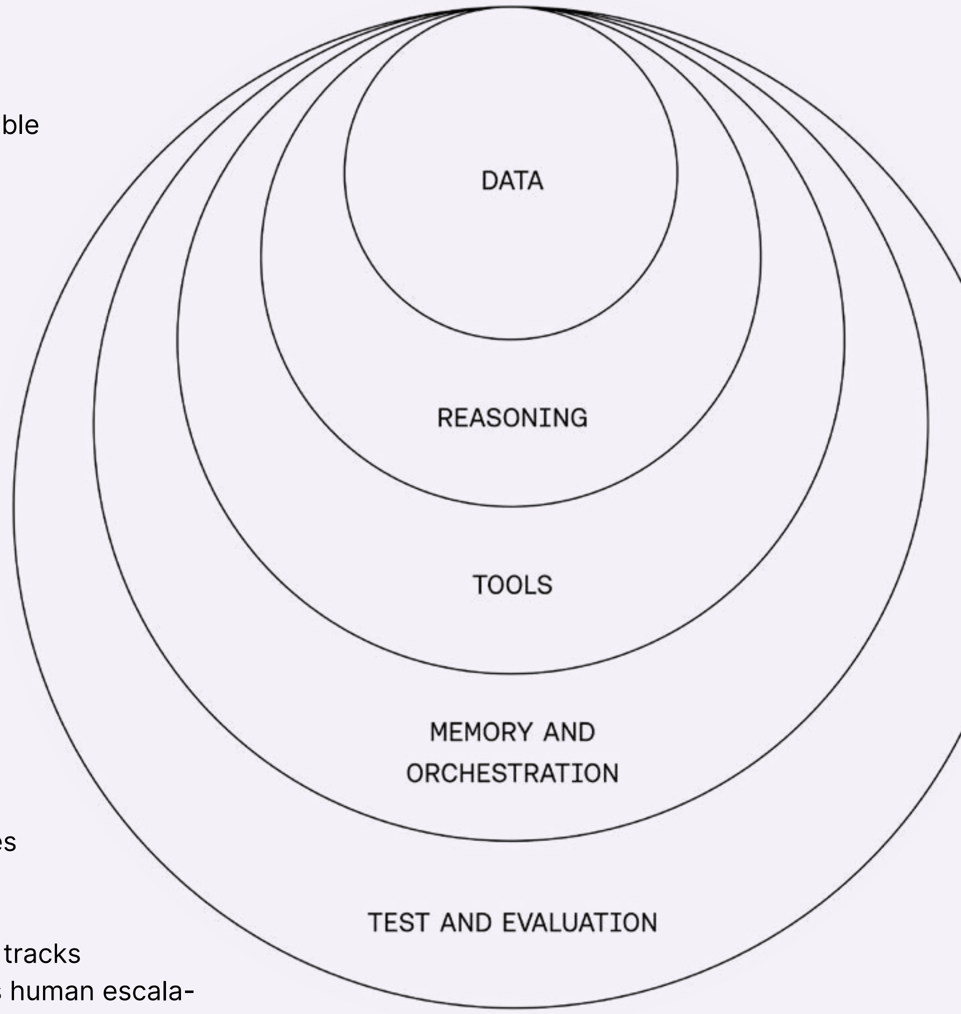

Data: Agents require a live, machine-readable worldview. Pipelines ingest and clean multimodal data, retrieving relevant historical and real-time inputs that comply with security protocols. Continuous feedback loops drive adaptation.

Reasoning: Agents execute a continuous loop: interpret, act, observe, update. This cycle fuses LLM-driven reasoning with planning, enabling agents to decompose problems, delegate to tools or sub-agents, and adapt without constant prompting.

Tools: Tools enable agents to transcend training data. Agents query databases, execute code, and command workflow systems or physical devices. While large models drive reasoning, specialized engines handle routing, analytics, and simulation.

Memory and orchestration: Orchestration tracks objectives, manages memory, and governs human escalation. In deployment, specialized agents operate in chains, while a central coordinator routes tasks and enforces least-privilege access.

Test and Evaluation: Agents capable of real-world action require rigorous mission assurance. A comprehensive T&E regime combines pre-deployment testing, automated lifecycle monitoring, and software-enforced human approvals with continuous red teaming to mitigate novel threats.[9]

PEOPLE-PROCESS-TECHNOLOGY-TRUST

The United States is uniquely positioned to build and scale agentic systems for the U.S. military. We have a thriving AI ecosystem within our borders: the power, the chips, the frontier models, and the software leadership to integrate them, as well as deep data repositories in the Department of War to fuel AI applications. The primary constraint is not technology, but the Department's ability to launch and mature prototypes and transition them into Programs of Record that field capabilities at scale, while realigning doctrine and process around them.

This presents an historic opportunity. We can "enlist" millions of digital agents to enable our human force to operate with exponentially greater efficacy and efficiency. This means exploiting the unique strengths of human operators—unmatched imagination, interpretation, and judgment—with that of machines: the ability to process information at a volume, tempo, and level of pattern-finding far beyond human cognition.

It will not be easy. Agentic systems challenge the way we work. It's possible to imagine a future in which the traditional role of commanders and staff shifts from "in the loop" execution to "on the loop" oversight, in which they oversee systems of agents. In such a future, instead of reviewing raw data for every decision, they will focus on AI-curated options—intervening only at those unique points where a commander alone must decide to signal intent and authorize action, thereby maintaining "meaningful human control." [10]

Implementing agentic systems also requires an understanding of their reliability. For a commander to approve a mission-critical COA, they must have "justifiable confidence" in the output. This confidence cannot be assumed, however, because today standard metrics for AI reliability remain immature. It must therefore be engineered in the coming months and years. The government and private sector must collaborate to pioneer rigorous T&E regimes that go far beyond traditional software testing, focus on mission assurance and alignment with human intent, and field capabilities with distinct guardrails today while aggressively maturing the science of evaluation in parallel. Fully achieving this remains an active area of research.

Though the advent of agentic systems offers immense promise, much practical work remains. To grapple with the specifics of systems being developed today, Part II highlights two case studies of agentic systems that will come online in the near future.

PART II

Agentic Warfare Today

The Department of War has laid out a bold vision for how applied artificial intelligence, particularly agentic systems, can secure enduring strategic superiority.[11] Central to this vision is decision advantage: the ability to observe, orient, decide, and act faster and more effectively than any potential adversary. As Secretary Hegseth underscored in a recent address, “this urgent moment... requires more innovation, more AI in everything and ahead of the curve.”

To carry out this vision, the Department of War has created an ecosystem of experimentation with advanced AI technologies. Scale today is building two related agentic capabilities in this ecosystem: one that performs Agentic Planning and another that performs Agentic Alerting. Each constitutes a complementary way the U.S. military can achieve new degrees of decision advantage over adversaries. Both capabilities are built on a similar technical foundation and then matured for application at different stages in the OODA loop. The following two case studies demonstrate each, as well as how they are related to one another. Used in tandem, they will help move the command and control ecosystem toward seamless and automated vertical processes across services and echelons, from the tactical to the strategic.

CASE STUDY I:
AGENTIC ALERTING

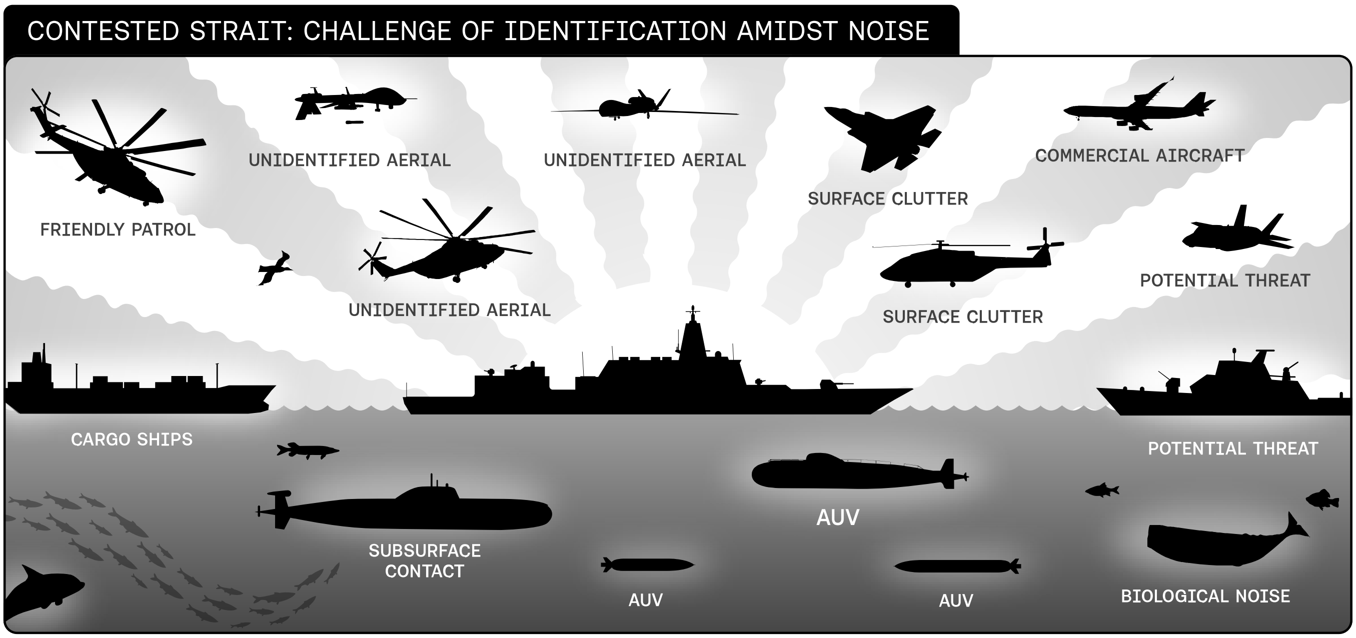

Agentic Alerting revolutionizes how commands sense the battlefield. Picture a carrier strike group in a crowded strait at night: commercial tankers with transponders not reporting their up-to-date position, fishing boats deliberately running dark, hostile drones skimming the wave tops, possible submarine contacts, and intercept chatter hinting at a missile launch inland. In a maximally-stressing scenario, every sensor will trigger, lighting up every console, yet the real threat could easily pass unnoticed in the noise.

Agentic Alerting is designed to meet this moment, executing decision chains at machine speed. Its near-term promise is not a brand-new sensor architecture, but a synthesis layer that sits on top of what the force already has: the sensors, the data pipes, and the many alerting and analytic tools that work well in their lanes but remain siloed and relatively static. Agents can ingest large volumes of multi-modal data—time-series telemetry, imagery, traditional sensor alerts, human analytical products, operational context—and combine them into a coherent, mission-relevant picture. They not only deliver prioritized, explainable alerts to the right echelons, but then go beyond notification to action: making automatic adjustments to where sensors focus and what analysis is prioritized by the system, enabling rapid changes in defense posture and command decisions in a shifting operational environment.

AGENTIC ALERTING TECH STACK

At its core, Agentic Alerting introduces dynamic intelligence into sensor networks, transforming static sensor arrays into adaptive, "event-driven" systems that act as a decentralized nervous system. Rather than requiring constant human monitoring, these agents operate in a continuous loop. At the edge, lightweight detection agents monitor raw data streams—acoustic vibrations, thermal signatures, or radio frequency changes—looking for anomalies that deviate from a baseline "normal" environment. Once a trigger is identified, the agent doesn't just alert a human; it begins a process of autonomous investigation and orchestration to verify the threat and gather higher-fidelity data.

This is primarily achieved through cross-cueing, where the detecting agent tasks other sensors in the network to focus on the same point of interest. For example, if a low-resolution ground sensor detects a seismic disturbance, it communicates the coordinates to an orchestrator agent, which then commands a high-resolution camera or an overhead drone to swivel and zoom in on that exact location. This allows the network to maintain a "low-power" state across a wide area while instantly concentrating its "high-power" resources on a localized event, effectively managing the trade-off between coverage and detail.

Beyond physical movement, AI agents also manage the value placed on specific pieces of data, called the Information Value (VoI), to optimize network bandwidth and processing power. In a bandwidth-constrained environment such as a remote border post or a combat zone, the agents can adjust data priority based on the perceived importance of an event. If multiple sensors are triggered simultaneously, the AI uses risk- and priority-based logic to decide which event poses the highest threat, allocating more frame rate and resolution to the high-priority target while dropping or summarizing data from less critical areas. Coupled with on-device analysis that mitigates the need to send raw data, this ensures that the most relevant information reaches decision-makers without overwhelming the network's capacity.

Finally, these agents utilize spatial reasoning frameworks to plan investigative actions. When a sensor network detects a moving target, the agent doesn't just track its current position; it can use physics-based models to predict the target's future path and preemptively shifts the focus of cameras further down the intercept line. By visualizing the geometry of the environment—such as dead zones behind terrain obstacles—the agents can position sensors to eliminate blind spots. This level of autonomous foresight ensures that once an event is triggered, the network "hunts" for the target rather than simply following it, providing a continuous, unbroken tracking chain.

Agentic alerting moves far beyond existing sensor network capabilities. Many deployed systems have sensors linked to fusion and correlation engines working at machine speed. But these systems are static and inflexible. They are unable to reason across modalities or situations that arise outside of their discreet rule or training sets. Because they lack context outside the 1s and 0s of the data streaming in, they are unable to draw upon the broader situational context to dynamically evaluate across a series of weights. Crucially, they are unable to produce quantitative measures of risk that could inform a Commander's qualitative perception of risk.

The key differentiator between fusion and correlation engines and Agentic Alerting is judgement, as enabled by the frameworks and guidance humans build into agents. Agentic Alerting is built around enduring agent-level system engineering—the rigorous process of architecting tools, prompts, memory systems, finetuned models, and multi-agent interaction patterns to robustly address a problem space. This approach enables an AI to function as a persistent, goal-oriented entity capable of maintaining its logic and state across long-running, multi-step operations. Unlike standard prompting, which seeks a single immediate response, agentic architectures create a robust “mental framework” for the agent that includes strict behavioral guardrails, standardized memory-management routines to prevent “instruction drift” over time, and precise tool-invocation protocols that allow the AI to safely interact with external sensors or databases. By establishing these foundational rules—often structured through XML tagging or recursive reasoning loops—agents can autonomously navigate task uncertainty and recover from system interruptions while remaining strictly aligned with its core mission and safety constraints.

THE SMART EDGE

Agentic Alerting brings speed, scale, and 24/7 precision to the military’s real-time nervous system: the tactical and operational sensor grids serving as commanders’ digital eyes and ears. Agents continuously monitor multi-modal feeds—radar, electro-optical/Infrared (EO/IR), SIGINT, blue-force, weather, logistics—sharing data across domains to reason over patterns of life and deliver prioritized, explainable alerts. Working in concert, they triage competing inputs, monitor data quality, detect anomalies, and reprioritize tasks as situations evolve. Instead of an incoherent barrage of notifications, agents provide mission-relevant insight, shifting operators from managing backlogs to managing decisions.

Deployed at scale, alerting agents empower commanders to make rapid, high-stakes decisions in response to threats. With chain-of-thought reasoning, tunable thresholds, and continuously updated models, the system becomes a trusted partner in dynamically sensing and managing the battlespace. It optimizes defensive postures, retasks assets to close blind spots, and cues sensors against suspected threats. From theater-level ballistic missile defense to hostile autonomous swarms, the result is a force that sees sooner, understands faster, and buys time for the right decisions.

CASE STUDY II:
AGENTIC PLANNING

While Agentic Alerting rewires how sensor networks function, an even more nascent capability of AI agents is to help military staffs develop plans and rapidly iterate them as circumstances evolve.

Military planning is core to U.S. battlefield dominance, yet it is a laborious process, rife with manual steps that have changed little since Napoleon codified the military staff system in 1803. The existing Joint Planning Process—a methodical, seven-step cycle for defining missions, developing courses of action (COAs), and producing plans—is fundamentally linear and optimized for deliberation rather than speed. For theater-level operations, this process can take up to two full years to complete, producing paper plans housed in binders that are static snapshots of a moment in time. Even small adjustments can take months to implement.

While the Joint Planning Process produces comprehensive COAs, the realities of modern conflict often stress its utility in two ways:

- **Time Lag & Technological Drift:** During the long period plans are developed, the battlespace and underlying technologies evolve. The result is that approved plans routinely fail to reflect mid-cycle developments—such as the rapid proliferation of First-Person-View (FPV) drones in Ukraine—causing their logic and effectiveness to diverge from operational reality.
- **Crisis Inadaptability:** The most consequential weakness is the inability to adapt when conflict breaks out. Once an adversary makes an offensive move, significant parts of standing joint plans can become irrelevant. If plans need updating, the current process is often too slow to validate new courses of action through simulation before a commander must act.

A new agentic prototype aims to transform planning into an instantaneous, evergreen process. In March 2025, Defense Innovation Unit (DIU) announced Thunderforge, an initiative to build an AI-enabled command, control, and planning system. Under a prototype contract awarded to Scale, the platform integrates AI agents into operational and theater-level planning. Testing is underway at U.S. Indo-Pacific Command (INDOPACOM) and U.S. European Command (EUCOM).

Thunderforge aims to break linear bottlenecks by providing tools to planners that couple AI agents with automated, physics-based modeling and simulation tools. This allows planning staff to generate validated, confidence-bound courses of action and iterate upon them rapidly, keeping commanders supplied with courses of action adapted to current conditions on the battlefield and in line with OPLANS and other strategic and operational guidance.

Crucially, Thunderforge is not focused on using AI to enable the quicker completion of the paperwork of planning—the staff estimates and warning orders that are a required part of the planning "paper trail." AI is already able to do this with ease and LLMs have already been adopted in planning workflows. Rather, Thunderforge is focused on bridging LLMs in planning workflows with AI agents focused on higher level tasks—those crucial steps like COA development that require greater levels of expertise and judgement.

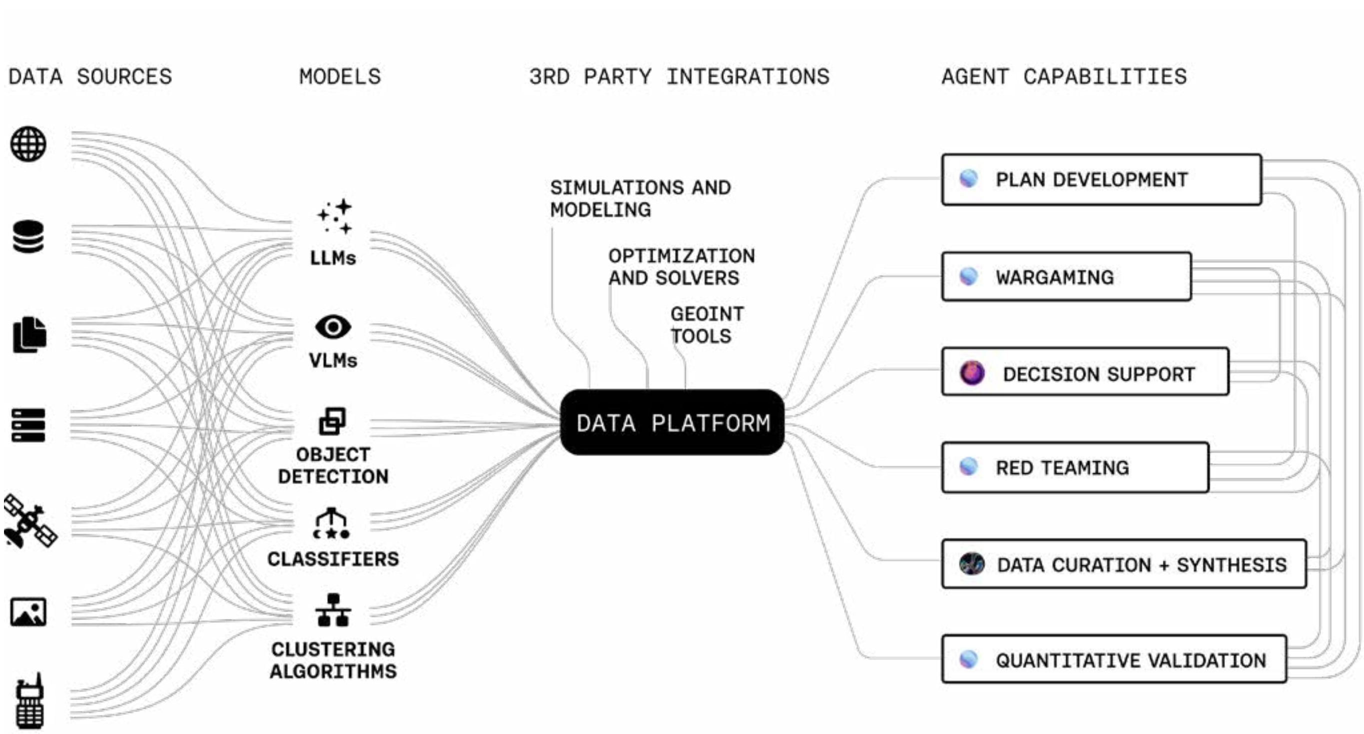

How Agentic Planning systems move from data to agent capabilities.

The planning process has changed little. The top picture is of Marines planning in the years between WWI and WWII, the bottom is of Marines planning during the 2000s.

TYPES OF JOINT PLANNING FROM STRATEGIC TO TACTICAL

At the strategic and theater level, commanders develop Operational Plans (OPLANs), which are comprehensive, detailed plans for conducting campaigns and major operations, often covering a vast geographic area over an extended period of time. These are typically prepared in anticipation of a specific contingency. In contrast, the U.S. military executes Crisis Action Planning (CAP) rapidly in response to an unexpected, time-sensitive situation, focusing on getting forces into the fight quickly based on existing concepts rather than a bespoke OPLAN. Moving down echelons, Joint Planning, conducted by Joint Task Forces, translates strategic direction into operational objectives. At the tactical level, planning done by brigades, battalions, and smaller units, focuses on the minute-by-minute execution of specific engagements to achieve the commander's intent. The goal throughout this hierarchy is always to ensure a unified, coordinated effort that links tactical actions to strategic goals.

THE NAPOLEONIC STAFF SYSTEM

Napoleon Bonaparte's military staff system was born to tame the chaos of mass armies. Codified in 1803 and further developed by the Prussian military into the General Staff we know today, it created a strict division of labor: operations, intelligence, logistics, and plans sections composed of professionalized officers staffing a single commander. Two centuries on, most modern headquarters would still be recognizable to Napoleon or von Moltke the Elder: large, hierarchical organizations designed around linear processes, functional silos, and a central command post acting as the hub of orders and targeting.[12]

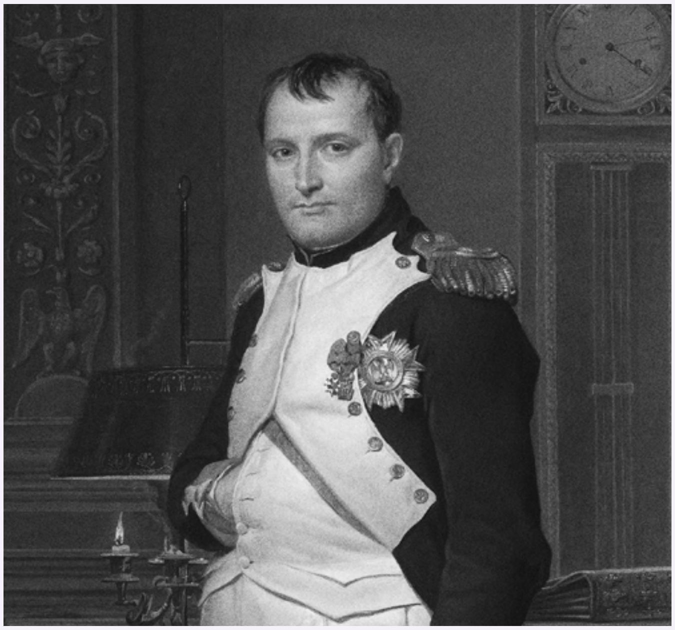

The U.S. Military's Joint Staff Directorates

J1 – Manpower & Personnel: Provides manpower and personnel counsel and support to enable Joint Force readiness and inform military advice to national leadership.

J2 – Intelligence: Supports the Chairman, Secretary of Defense, and Combatant Commands by providing intelligence, indications and warning, and crisis intelligence support.

J3 – Operations: Directs and coordinates current military operations and relays operational guidance between national leadership and Combatant Commands.

J4 – Logistics: Leads the joint logistics enterprise to drive readiness and maximize commanders' freedom of action.

J5 – Strategy, Plans, & Policy: Develops strategies, plans, and policy recommendations and assesses risk in executing the National Military Strategy.

J6 – Command, Control, Communications & Computers (C4)/Cyber: Provides C4 and cyber expertise to enable a globally integrated Joint Force across all domains.

J7 – Joint Force Development: Trains, educates, and develops the Joint Force to achieve overmatch across the continuum of conflict.

J8 – Force Structure, Resources & Assessment: Evaluates and develops force structure requirements and conducts analyses and wargaming to support the Chairman's decisions.

Napoleon fought in the era of the musket, cavalry, close order formations, and volley fire. Armies moved at foot speed while information moved at horseback speed. Today's industrialized and information-centric warfare unfolds with far greater degrees of speed, precision, and complexity. The advent of Agentic Warfare opens an opportunity to reimagine how general staffs can operate and win at machine-speed conflict.[13]

GENERATING THE “PLANNING MULTIVERSE”

The fourth step in the Joint Planning Process is one of the most consequential: wargaming and simulation. A unique U.S. strength relative to its adversaries is our wargaming’s rigor, due in part to the extent to which it is validated by sophisticated, physics-based models.

The Department of War maintains a strong technical bench of experts and modeling software, capable of rigorously evaluating hundreds of thousands of scenarios to identify optimal courses of action. However, substantial barriers prevent the U.S. military from using these tools at the speed of relevance.

ONE OF MANY PLANNING TOOLS: DARPA SAFE-SIM

There is a plethora of model and simulation tools across the Department. One in particular is among the first to be integrated into Thunderforge. DARPA's SAFE-SiM (Secure Advanced Framework and Environment for Simulation and Modeling) is one of several premier capabilities in the Department's modeling arsenal. Designed as a faster-than-real-time, all-domain environment, it allows commanders to model missions from seafloor to space—including cyber and electromagnetic spectrums—within a single coherent framework. It operates across multi-level security environments, enabling analysts to move beyond manual wargames and rapidly validate complex logistics and force structure trade-offs.

SAFE-SiM lets analysts use an advanced simulation environment to design campaigns, concepts of operations, force structure composition, and resource allocations. Depending on the scenario under consideration, SAFE-SiM can optimize force structure against mission priorities, assess likely mission effectiveness, and analyze trade-offs in logistics and force-structure. Instead of relying on a handful of manual wargame runs, it is intended to allow rapid construction and re-running of complex, theater-wide scenarios, so planners can see how different courses of action and force mixes perform across many possible futures.[14]

Agentic planning tools will eventually be able to access all of the modeling tools in the Department of War.

To operate hundreds of different modeling systems, the Department relies on a specialized cadre known as Operations Research/ Systems Analysts (ORSAs), who apply data science, optimization, and simulation to complex problems ranging from logistical flows to strategy comparisons. Their mission is to ensure critical decisions are grounded in quantitative rigor rather than intuition. With hundreds of specialized models across the Services, this ecosystem remains fragmented. Current tools require niche experts to operate, can take days to execute, and produce outputs that are difficult to interpret. Because of this latency, rigorous modeling is frequently the first step of the Joint Planning Process abandoned during a crisis.

By using agents to automate calls to simulators through a natural language interface, Agentic Planning tools enable planners themselves, without the help of specialists, to initiate model runs and extract findings from model outputs. Importantly, this democratizes access to modeling tools much in the same way the Graphical User Interface on early personal computers enabled novice users to bypass typing commands by prompt. Operationally, the system is built on three layers: a natural language user interface where planners can access simulation data; an agent layer capable of calling servers and synthesizing responses; and, a simulation layer comprising validated models.

Simulations can run in real time, without the need for expert human involvement at each step. With COAs validated in hours or days rather than weeks or months, ORSAs are freed from the manual work of loading and extracting simulation data, allowing them to instead focus on higher-level problem solving. At the same time, simulations become more accessible to broader planning staffs, collapsing layers between the commander and analytic results.

Agentic Planning can even accelerate the use of simulation results to refine scenarios. Working in continuous loops, agents retrieve data from simulation runs, analyze findings to create reports, and enable users to query that analysis. Based on these outcomes, other agents can automatically initiate secondary simulation runs with adjusted variables to explore alternative scenarios and courses of action.

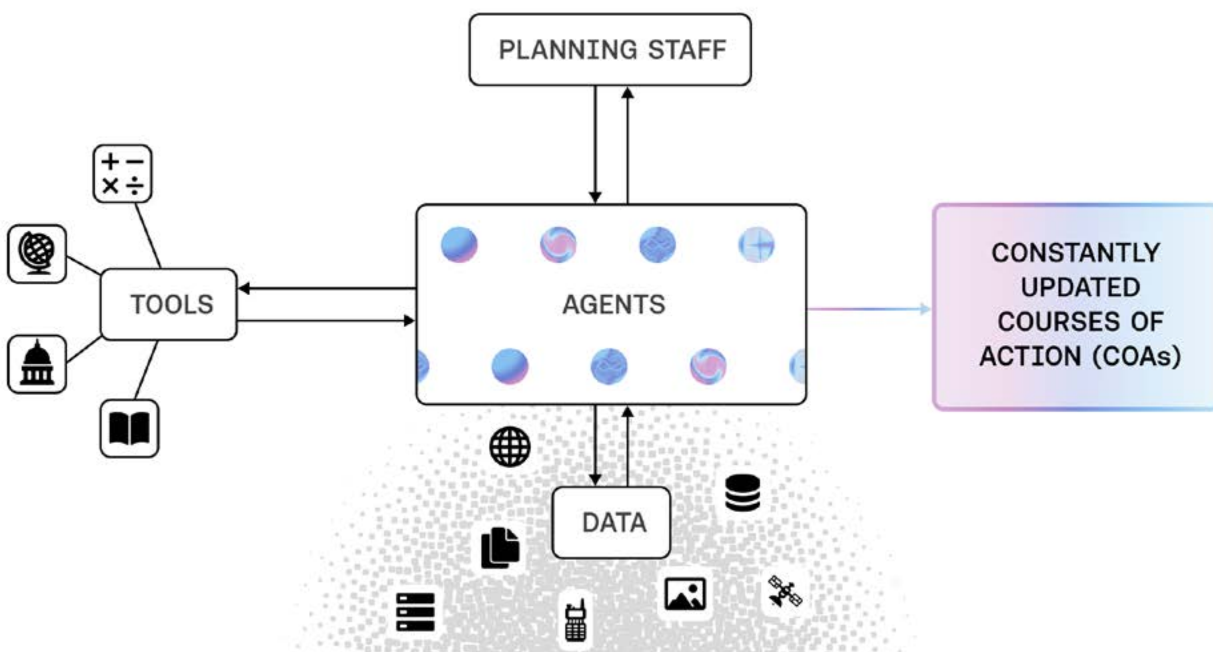

Agent-based cueing of simulation tools in the development of plans.

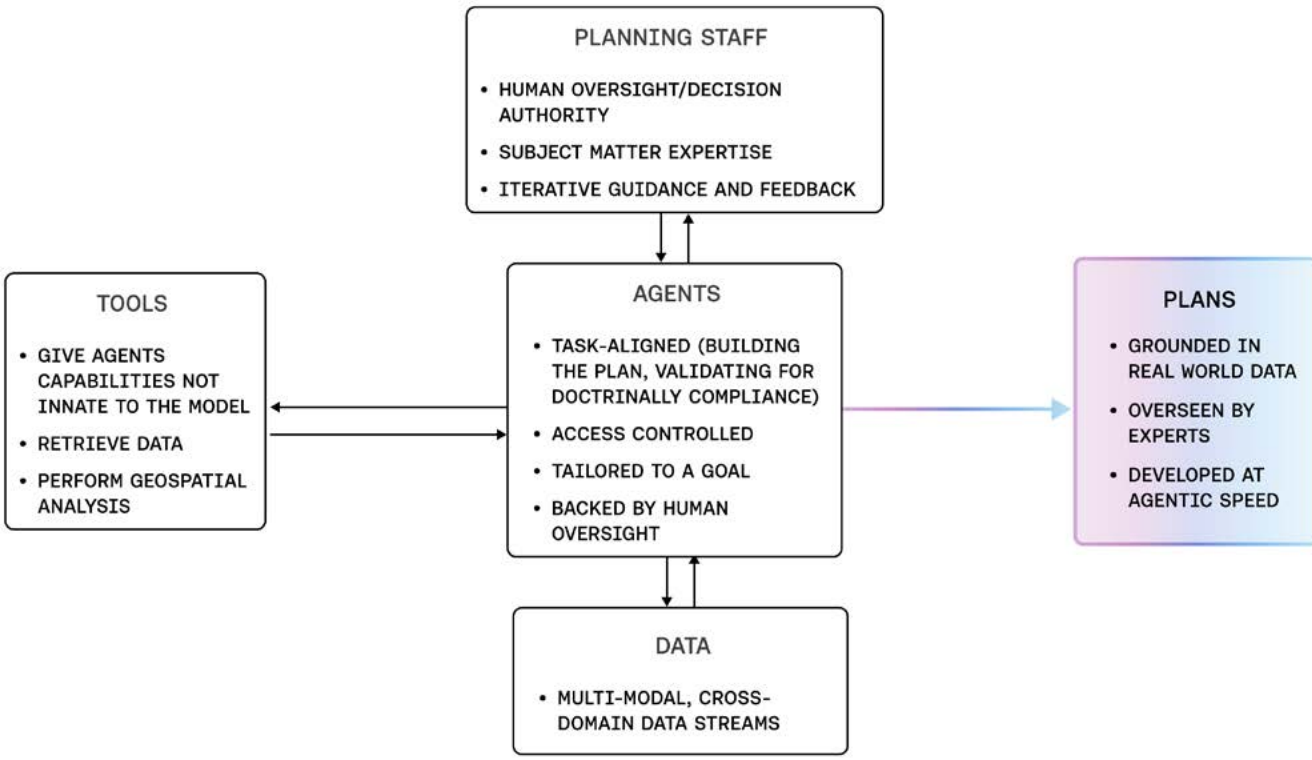

Planners working with teams of agents to produce validated plans at machine speed.

> Agentic Planning systems generate strategic, operational, and tactical planning "multiverses," running thousands or even hundreds of thousands of permutations to identify the optimal responses to adversary action.

This is a powerful capability. It allows Agentic Decision systems to generate strategic, operational, and tactical planning "multiverses," running thousands or even hundreds of thousands of permutations to identify the optimal responses to adversary action. The agentic system presents planners with a structured set of these scenarios, allowing them to select the most effective options for command decisions. By matching potential adversary moves against a wide range of pre-validated outcomes, this enables planners to quickly identify the courses of action with the highest likelihood of mission success.

FROM STATIC BINDERS TO "LIVING PLANS"

Though not yet a capability, Agentic Planning systems of the future may go one step further: integrating live, multi-modal, cross-domain sensor data into planning feeds. The modern battlefield produces petabytes of daily intelligence, from space-based imagery to force telemetry and signals intelligence (SIGINT). Agentic Planning systems could use agents to ingest this data and other agents to begin generating courses of action, while validating these courses of action through models in real time.

If agentic systems can achieve this, plans will evolve from static, staff-driven products into "living documents" updated in real-time by AI—a seismic shift. This would lessen the lag between an adversary's move and the U.S. military's ability to respond, while enhancing a commander's confidence in courses of action through empirical validation by simulators. Essentially, plans would "self-heal" in response to events, delivering options to commanders at the speed of relevance while preserving human authority.

Instead of commanders pulling and synthesizing vast amounts of raw intelligence, planners and planning systems push only the most relevant details at the precise moment required for decision.

This is the kind of decision advantage that agentic systems promise—an entirely new paradigm.

PART III

Reimagining America's Way of War

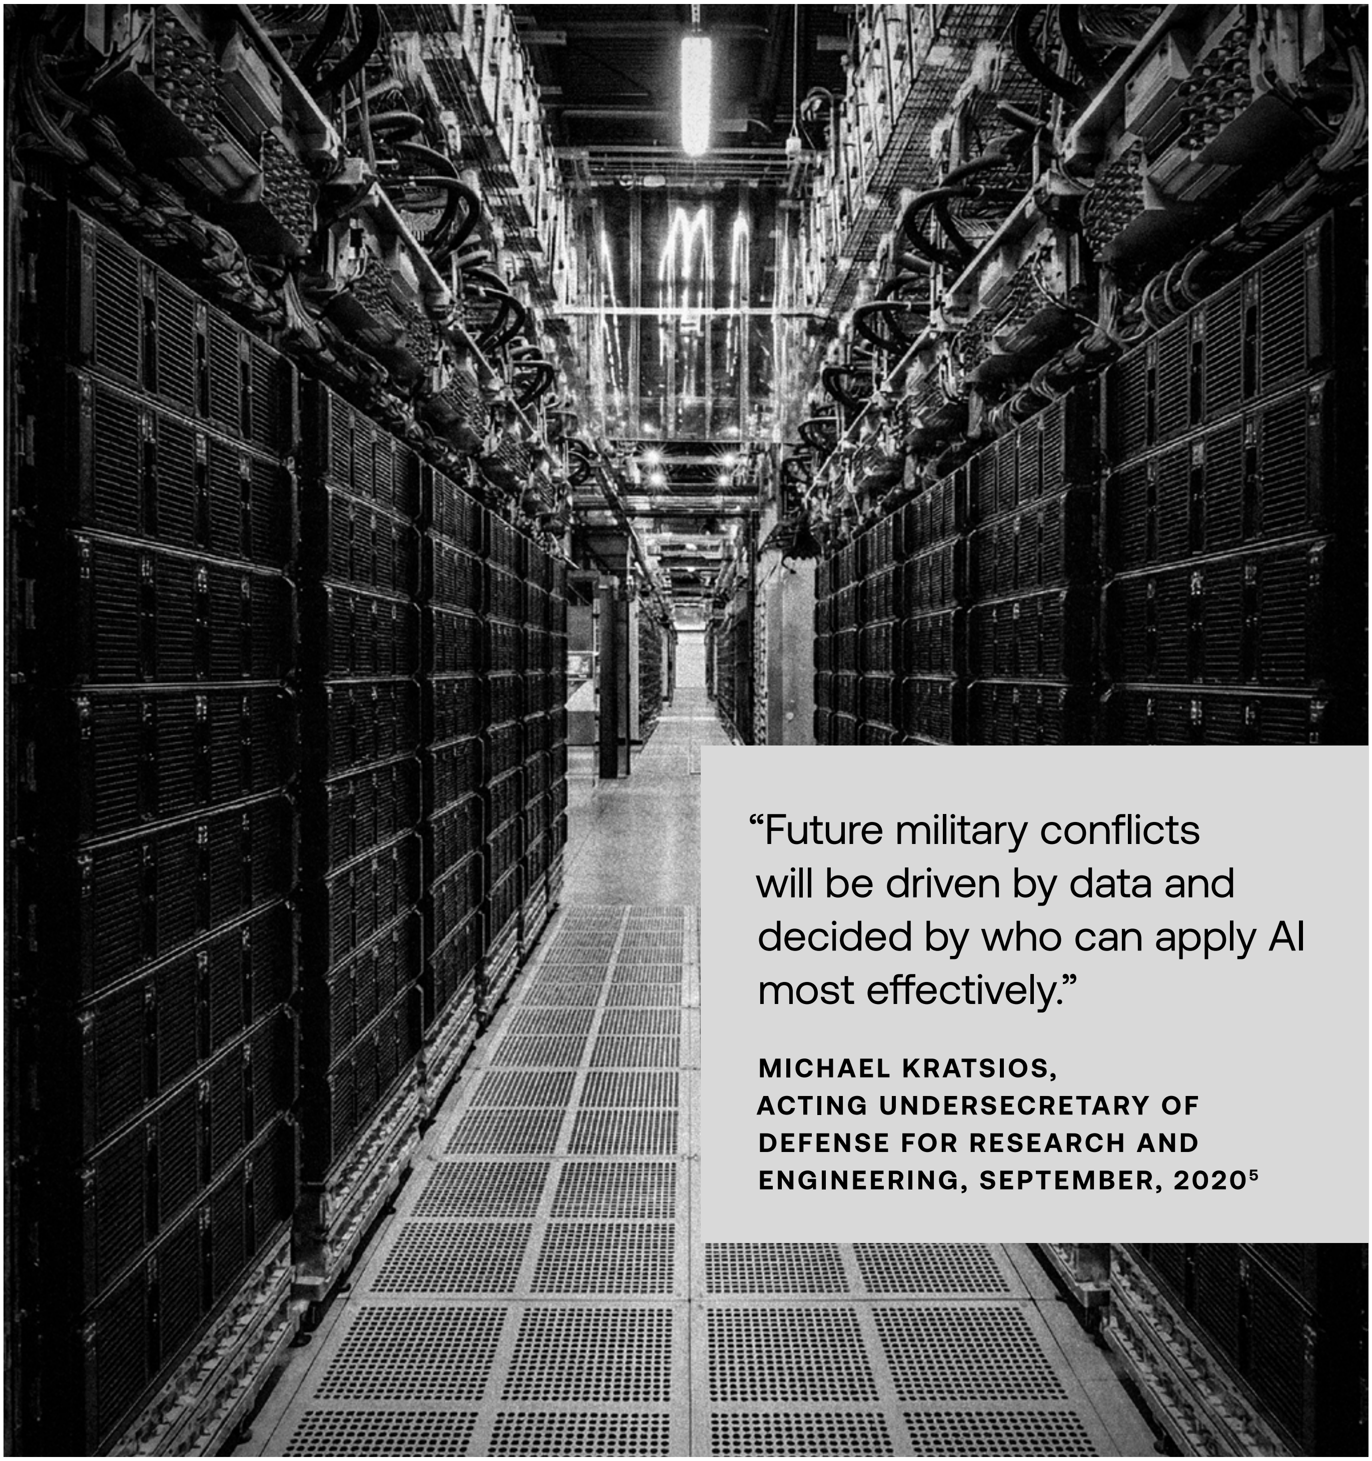

"Future military conflicts will be driven by data and decided by who can apply AI most effectively."

MICHAEL KRATSIOS, ACTING UNDERSECRETARY OF DEFENSE FOR RESEARCH AND ENGINEERING, SEPTEMBER, 2020[5]

The next five years will be decisive for AI-driven capabilities. AI agents are poised to mature significantly toward proactive, goal-driven, and intrinsically multi-modal architectures. The key advance will be in hierarchical planning and self-correction, allowing agents to tackle complex, long-horizon tasks—those requiring many sequential steps, resource allocation, and adaptation based on feedback. This maturation will be driven by improved foundation models that enable sophisticated chain-of-thought reasoning and greater context windows, effectively giving agents a longer memory and a deeper understanding of the task environment, akin to greater intelligence and judgment.

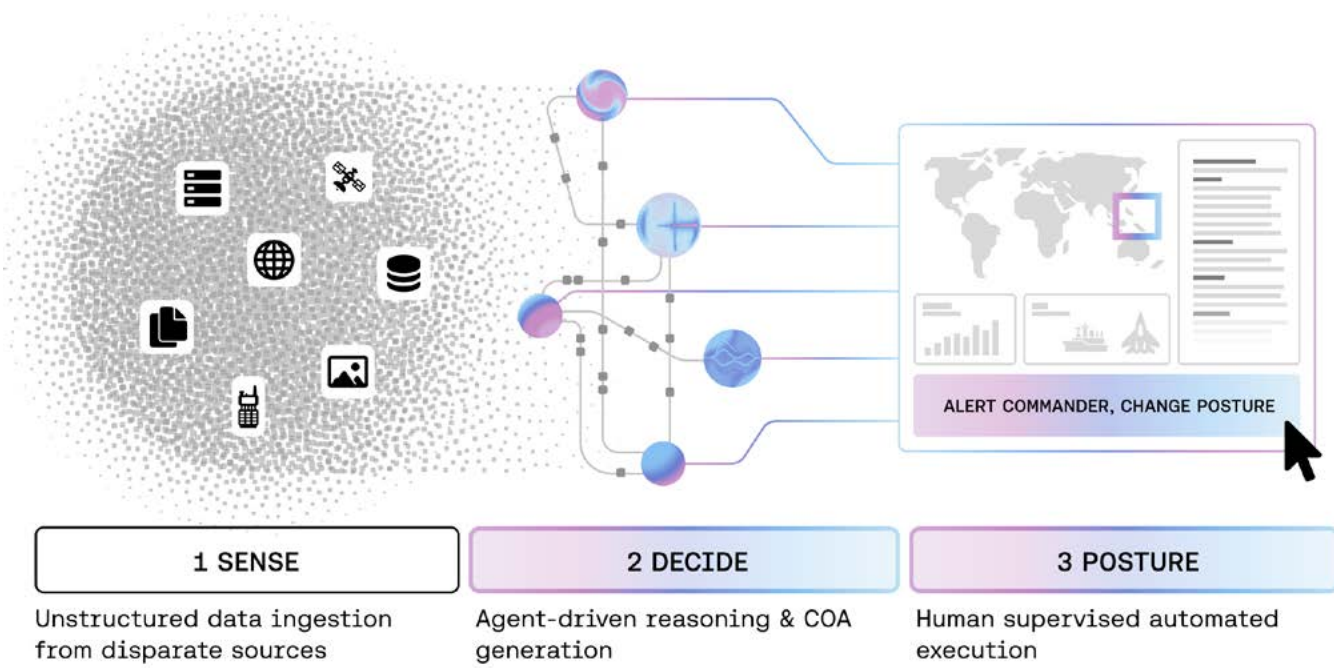

As these technologies mature, they will unlock increasingly transformative possibilities to disrupt and enhance every facet of military power, from sensing, planning, and maneuvering, to logistics, cyber, and information operations. While the precise impact across domains and functions is unclear—with big dependencies on both pace of adoption and diffusion—the adoption of these capabilities will pose tensions that Department leaders will need to reconcile in architecting the future force. Among those challenges are:

- Determining the optimal high-low mix, balancing quantity and quality as cheap, massed agentic swarms change the economics of conflict.

- Addressing heightened challenges sensing and understanding battlefield dynamics as turbocharged data fusion and pattern analysis confronts new forms of obscurity, spoofing, and operational art designed to mislead machines as much as humans.

- Machine-speed competition between offensive and defensive agentic systems across an increasingly broad attack surface that blurs the line between forward defense and homeland vulnerability.

However, true changes in the character of war are not enabled by technology alone. Like previous disruptive shifts—blitzkrieg, carrier aviation, precision strike—Agentic Warfare will only succeed if we change how we fight, not just what we field. We must accelerate the agentic prototypes that will help deliver decision advantage, like Agentic Planning and Agentic Alerting, while also catalyzing essential doctrinal, cultural, and institutional reforms necessary to fully operationalize these capabilities. Get that right and we develop a new way of war that will secure U.S. forces' edge and effectively deter would-be aggressors; fail to move at pace and scale and we risk watching someone else do it first.

POTENTIAL AGENTIC WARFARE APPLICATIONS

This paper focuses on Agentic Planning and Agentic Alerting. These capabilities are critical to decision advantage, but only a subset of the ways in which agentic systems will transform offensive and defensive operations in the coming years.

While many focus on autonomy in offensive systems (kinetic or otherwise) for the U.S. military many of the biggest gains will be in defensive and enabling capabilities. Examples include:

- **Cyber defense and cybersecurity:** always-on agentic purple teams, attacking networks in order to patch and defend them, or always-on hunter agents, actively seeking anomalies and isolating compromised nodes instantly.
- **Electronic warfare and force protection:** fusing full-spectrum sensing and platform telemetry to detect jamming or spoofing, recommend countermeasures, reconfigure Blue assets in real-time to frequency hop, and cue operators.
- **Intelligence analysis:** moving beyond accelerated sense-making to persistent forecasting by fusing multi-source signals, tracking competing hypotheses, flagging deception indicators and updating threat likelihoods in real time.
- **ISR:** adapting collection at machine speed, prioritizing queues, flagging cross-domain correlations, detecting anomalies amid the noise and tasking sensors accordingly.
- **Logistics and sustainment:** optimizing supply chains and inventories, intervening to maximize platform readiness, dynamically shrinking footprints in contested environments, and continuously adapting to disruptions.
- **Space operations:** learning satellite “patterns of life” to differentiate normal behavior from anomalies, correlating signals to identify threats, and proactively maneuvering assets to avoid danger or to self-heal constellations if a node is destroyed or jammed.
- **Medical support:** saving lives through decision aids that improve triage and enable personalized interventions, while matching casualties to capabilities for evacuation.
- **Coalition interoperability:** enabling richer, simpler and more secure information flows so information reaches the right partner or asset at the right time.
- **Procurement:** streamlining acquisitions by drafting requirements, scanning proposals for compliance, and identifying supply chain risks.
- **Enterprise functions:** saving time and money through tighter human-machine workflows that improve consistency, auditability and throughput of administrative tasks, from drafting and coding to analysis, cross-correlation and intelligent risk monitoring.

ADVERSARY USE OF AGENTIC AI

Our adversaries are also racing to deploy agentic capabilities. Competition from China is especially fierce: it is combining massive investment, abundant human capital, and industrial strategies to close the People's Liberation Army's (PLA) gap with the U.S. military. As a result, in Agentic Warfare the window for the United States to secure first mover advantage—or even be a fast follower—is fast closing.

AI is at once an astonishing example of a breakout U.S.-invented technology that will change everything, yet at the same time a cautionary tale of the limits of first mover advantage in today's global technology market. Despite pioneering the fundamental research and initial commercialization of computer vision, translation, and LLMs, native Chinese LLMs are nearly on par with their U.S. competitors when it comes to performance.

In other respects, China is ahead in AI adoption; in the first half of 2024 alone, China launched 81 separate projects that deployed LLMs in government applications.[16] What matters is not so much a race to match or exceed Chinese adoption of AI. Rather, we are in a race to leverage AI to counter the advantages the PLA and other adversaries have so we can achieve our military objectives and uphold our security commitments.

Agentic capabilities are at the center of China's strategy. The development of AI-based battle planning to create a "command brain," described in Chinese literature, is well underway.[17] China is growing its military arsenal, architecting its force around "intelligentized warfare." It is shifting from platform-centric modernization to a force built on AI, data, and autonomy, with information advantage and decision tempo as the key metrics of power. Frontier models like DeepSeek are at the center of a civil-military fusion stack, providing the backbone for experiments in war-gaming, battle planning, logistics and multi-source intelligence analysis, while specialized algorithms sift sensor feeds and propose options to commanders.

China is aggressively coupling industrial mass with AI sophistication. Beyond fielding the world's largest navy, the PLA is rapidly operationalizing autonomy across every domain. Their portfolio now includes hundreds-strong truck-launched drone swarms for Taiwan scenarios, "loyal wingman" escorts, experimental drone carriers, and heavily armed uncrewed surface and underwater vessels. Ground forces are deploying "robot wolf" quadrupeds and autonomous support vehicles, while artillery units trial AI-guided shells that correct trajectory in flight. Across the force, automated target recognition is allowing missiles and drones to identify, track, and prioritize targets with minimal human direction.

China has further developed a concept of "counter-AI warfare," utilizing physical decoys and signature masking to confuse enemy sensors. This effectively translates digital spoofing into real-world deception, causing AI systems to miss actual threats. It serves as a reminder that as agentic systems, and the decision superiority they provide, become central to U.S. power projection, they will increasingly appear in the crosshairs of our adversaries.

To be sure, many Chinese systems remain immature, struggling with integration, a brittle command culture, and the reliability of their domestic tech stack. Chinese forces also remain relatively untested, having not fought a large-scale conflict since the Sino-Soviet war of 1979. Taken together, the trend lines still point in one direction: a force designed to combine mass, geography, and frontier AI to compress U.S. decision time. This is the strongest argument for the Department of War to place agentic AI and decision advantage at the center of its modernization, not at the margins.

HOW AGENTIC SYSTEMS CAN BE DEFEATED

The information and cyber environment is already an active battlefield. As we adopt agentic capabilities to secure decision advantage, they will become high-value targets. The aim will be to mislead and saturate—driving catastrophic miscalculation, turning our systems against us, and paralyzing decision-making by undermining trust among warfighters, allies, and the public. Ways to manipulate agentic systems include:

1. **Data and context manipulation:** Indirect prompt injection is a security exploit where an attacker embeds malicious instructions within external data—such as a webpage, email, or document—rather than typing them directly into the AI's chat box. Indirect prompt injections turn routine inputs accessed by models—reports, emails, logs, webpages, coalition feeds—into a control surface used to hijack models. Because LLMs don't reliably separate "instructions" from "data," an adversary can seed channels with content that looks like evidence but carries hidden instructions, steering reasoning, retrieval, and recommendations. The agent cannot reliably tell when it is being informed versus being directed.

2. **Model and supply-chain compromise:** Some attacks are planted before deployment—poisoned data, tainted fine-tunes, backdoored components, or "sleeper" behaviors that only activate under specific triggers. These are hard to spot because systems can look normal in routine evaluation, then fail in the rare edge cases that matter most in crisis or conflict.

3. **Cyber exploitation of tool-using agents:** When agents can query systems, task sensors, or draft orders, compromise shifts from bad analysis to operational effect. A manipulated agent might suppress alerts, leak sensitive information, misallocate key assets, or generate plausible-but-wrong courses of action humans approve under tempo.

4. **Physical deception against machine perception:** Decoys, signature manipulation, and sensor spoofing translate classic deception into machine-readable misdirection. The goal is digital fog: phantoms that look real, real threats that look benign, and degraded confidence in what the tools report.

To counter these threats, the Department of War must treat assurance as a core function. That means through-life T&E for agentic systems: aggressive red teaming; provenance and chain-of-custody for critical inputs; least-privilege permissions and action gates for tools; continuous monitoring for anomalous retrieval and behavior; and graceful degradation modes so human C2 remains effective when systems are contested or suspect. Critically, we cannot rely on chain-of-thought techniques as a proxy for explainability: a convincing rationale can mask manipulation or error, so assurance must rest on testable behavior, causal evaluation, and independent verification—not narrative.[18]

SECURING AMERICA'S EDGE IN AGENTIC WARFARE

As novel agentic systems achieve higher degrees of reliability and mission assurance that ensure they act within a commander's intent, it is entirely possible we will see two dynamics amplified over the next few years. The first is the ability of these systems to take actions of greater consequence. The second is the potential for a paradigm shift from humans "in the loop" to "on the loop," thus changing the nature of human-machine teaming such that humans need to oversee only the most crucial decisions. While individual capability changes may be incremental, U.S. forces will not realize the revolutionary potential of Agentic Warfare without significant changes to the way that we organize, train, equip, and fight—reimagining our way of war to make the best use of new technological possibilities.

Today's playbooks assume humans do most perceiving, fusing, and planning, with software in support. In an agentic force, the first pass in many of those steps flips: agents propose interpretations of developments on the battlefields and options for how to counter them; humans set intent, apply judgment, and own the risk.

The most crucial function of the directorates in the Joint Staff today—especially the J2 (Intelligence), J3 (Operations), and J5 (Strategy, Plans, and Policy)—is to bring to the commander their unique perspective and expert judgments in any given moment. The commander then synthesizes this information, weighs trade-offs, and makes a decision. In the future, agentic systems will have the capacity to perform many of the functions of these staff directorates, along with the ability to weigh trade-offs. Already in today's systems, agents perform important integrations, framings, and red-teaming functions as they work to deliver integrated COAs.

Today's playbooks assume humans do most perceiving, fusing, and planning, with software in support. In an agentic force, the first pass in many of those steps flips: agents propose interpretations of developments on the battlefields and options for how to counter them; humans set intent, apply judgment, and own the risk.

This means we must reimagine our way of war, adapting the 19th century general staff model to new workflows enabled by AI, and developing new concepts of operation that maximize human judgment and machine tempo. Operators of agentic systems will themselves need new skills:

- **Data and AI literacy** as basic officer competencies
- **The ability to task and supervise agents**, not just "use tools"
- Comfort with **interrogating and red-teaming machine outputs**
- A sharper **understanding of where agents might exceed their envelopes of assured performance**.

The metric of success is clear and profound: commanders at every echelon making better decisions, faster, with clearer understanding of their assumed risks and operational effects. Driving that shift will fall to a broad cast of institutions, from doctrine writers and operational concept developers in the Services and Joint Staff to the war colleges and training commands that shape senior leaders' mental models and analytical toolkits. Many are already moving: issuing AI doctrine notes, building data literacy guides, experimenting with AI-enabled wargaming, and standing up projects that use agents to accelerate planning. If we supercharge this work, we can turn today's agentic prototypes into a new way of war, rather than watching our adversaries achieve it first.

MOVING FROM "IN" TO "ON" THE LOOP AND BEYOND

As agentic systems become more capable and reliable, it is inevitable that more of the sensing-deciding-acting chain will run at machine speed. Some mission sets—like hypersonic missile defense, saturation drone attacks, and high-end electronic warfare—will simply outpace human reaction time. This reality is already baked into U.S. force posture in some ways. Once activated, systems like Aegis, Patriot, and Phalanx CIWS detect, classify, and engage inbound threats inside windows no human chain of command could manage. Other mission sets, like autonomous systems acting in fully denied communications environments, will by necessity need to act "off the loop," relying instead on pre-programmed rule-sets and guidance.

> "When you have a swarm of 1,000 drones coming at you, a mere human brain can no longer keep up with that threat."
>
> **DAN DRISCOLL, SECRETARY OF THE ARMY[19]**

More and more military systems will inevitably shift from humans "in the loop" to humans "on the loop" or even "off the loop," with humans supervising and setting guardrails rather than approving every action in real time.[20] Done well, this shift will allow machines to handle tempo and complexity while keeping commander intent paramount. The design challenge is to build agentic systems that keep decisions traceable and auditable, give commanders clear ways to set objectives, define pre-authorized parameters, and enable intervention when needed.

For senior leaders, the "so what" is immediate: as agentic systems grow in sophistication, a more rigorous frame than "in/on/off the loop" will become necessary, with graded delegation by mission type. Time-critical, tightly-bounded, or largely-reversible effects may justify high degrees of system autonomy. Individual targeting, urban operations, and escalatory strategic effects will likely demand direct human decision, even at some cost in tempo. Between these poles lies challenging terrain where the bounds of deliberate human action and machine authorities must be set.

OPERATING AGENTIC SYSTEMS: THE IMPORTANCE OF TEST & EVALUATION

Whether accelerating joint planning or alerting to events, agentic systems promise profound capability leaps. Yet the critical foundation for deployment is understanding both capability and propensity. In addition to validating what an agent can do reliably and consistently, we must also rigorously test for what it might be willing to do, evaluating the risk of unintended or malicious actions. Operators and commanders require justifiable confidence that AI agents will operate as intended. In Agentic Warfare, test and evaluation is not some back office function. It's at the tip of the spear of enabling the U.S. to project military power and a critical enabler to inform training and the evolution of doctrine

> Test and evaluation is not some back office function. It's at the tip of the spear of enabling the U.S. to project military power.

The challenge of T&E military systems is not new: the Department of War has spent decades developing methodologies to certify hardware and software for the battlefield. That expertise must now extend to AI. For AI systems, testing and evaluation is rarely a pass-fail exercise; there is no such thing as absolute trust in any system. Rather, T&E aims to discover a system's operational limits to ensure reliable use. It also drives the feedback loop necessary for system evolution: by pinpointing specific failures, evaluation data allows developers to integrate fixes into the training pipeline, continuously "hill climbing" toward higher safety and performance.

When it comes to agentic systems, testing and evaluation looks less like grading a discrete piece of homework and more like an uncharted ecosystem of tests, simulations, monitors, and real-world studies that characterize how these systems work, what they can do, and how they could go wrong.

Scale's Safety, Evaluation and Alignment Lab (SEAL) is working on research to advance this work, building hard-to-game safety benchmarks; publishing public leaderboards to compare models on safety, performance, and alignment; conducting frontier research studies on malicious use cases; and developing advanced monitoring mechanisms to evaluate agentic environments.[21] Scale is also an evaluation partner for the U.S. Center for AI Standards and Innovation (CAISI) and other leading global AI safety centers.

COMPUTER-VISION PERIMETER SECURITY: A REAL LIFE T&E EXAMPLE

In 2020, the Pentagon's Joint AI Center (JAIC), the predecessor to today's Chief Digital and Artificial Intelligence Office (CDAO), piloted the use of computer vision in perimeter security. Protecting Department facilities and deployed forces has traditionally been a labor-intensive process, with security personnel monitoring feeds from cameras and sensors designed to detect intrusions. By 2020, computer vision models had become increasingly competent at detecting objects automatically. The JAIC then experimented with adding automatic recognition capabilities to existing video feeds to see how they could augment human operators.

The experiment produced mixed results. The models performed best in bright daylight or full darkness, when they were as accurate as—or even better than—human operators at spotting intruders. However, they were significantly less effective at dawn and dusk, when their performance dropped to the point of limited utility, because less data from the system flowed into the model than during the day.

Through rigorous testing and evaluation, the JAIC determined where the system worked well and where it fell short. This allowed them to issue a model card (a set of instructions on how to use the system safely) that guided operators on deployment. The model card recommended that security manning could be reduced by 50% in daylight and at nighttime, when the model was most accurate. At dawn and dusk, or during inclement weather, the card recommended maximum manning and even turning the system off.

Testing and evaluation enabled an optimal outcome: staffing could be reduced overall, while clearly defining the conditions under which the model required auxiliary human support to maintain perimeter security.

Agentic systems present novel challenges that make robust test and evaluation imperative. Unlike traditional software, AI models are not deterministic. While it is possible in some instances to trace how models reach specific conclusions, in most cases their reasoning remains opaque. Agentic systems further function as "systems of systems," characterized by significant interactive complexity with humans, environments, and other agents. These systems do not operate in a vacuum: they interact with live, outside data.

For agentic systems, we take a holistic, iterative approach to test and evaluation that identifies, maps, and addresses foreseen and unforeseen risks, incorporating validation of both the agents and the oversight system itself. These evaluations focus on four layers:

1. **The Knowledge Base**: Ensuring that basic security protocols exist at the base layer of any environment (dataset, etc.) where agents will be deployed.

2. **The Underlying Model(s)**: Analyzing the specific vulnerabilities, biases, and safety risks inherent to each base model used for agentic systems. These trade-offs include accuracy, robustness, latency, cost, security vulnerabilities, unfaithful reasoning, and scalability.[22]

3. **The Agent Harnesses**: Capturing how agents interact by logging and analyzing every communication, decision, and state change during a simulation run. These evaluations include detection of misbehavior (e.g., agents optimizing reward models for unintended purposes) and misuse (e.g., jailbreaks, prompt injections, etc.).

4. **The Monitoring and Oversight System**: Evaluating effective human-in-the-loop monitoring designs, prompt injection protections, and mechanisms for risk identification, triage, and escalation. Defining clear authorities for each mission workflow, including human participation, logging of agent paths to identify behavioral patterns (both proactively and retroactively), and audits.

In practice, Agentic Planning and Agentic Alerting systems undergo end-to-end evaluations throughout their lifecycle, including extensive red teaming and systematic, scenario-based testing. Additionally, Scale must comply with CDAO's Responsible AI Toolkit and DIU's Responsible AI Guidelines, which have specific planning, development, and deployment worksheets prior to release.[23] This rigorous testing ensures that Scale's multi-agent systems reliably meet operational requirements and safety constraints before deployment in real-world scenarios.

A TAXONOMY OF AI DELIVERY

Even as the future brings stunning advances in AI technology, not every solution will be worth adopting. Over recent years, and especially since ChatGPT hit the mainstream, the Pentagon has ridden a wave of AI experimentation. A diverse ecosystem of frontier labs, defense primes, and start-ups now promise "transformational" capability. The result is a noisy market where genuine breakthroughs can be hard to discern from incremental improvements or repackaging of old technology. If the United States is to capitalize on the promise of AI, defense leaders need a disciplined way to pick technology winners.

That starts with choosing the right strategic partners: organizations with deep understanding of domain requirements and the technological frontier, proven expertise deploying advanced systems into classified environments, and experience sustaining those systems under real-world pressures.

It also requires a new way of thinking. As a complement to the Technology Readiness Level (TRL) framework embedded in the Department's classification of technology, leaders will benefit from an AI-specific mental model that cuts through hype and forces sharper questions about what is real now, what is possible in the near term, and what belongs in futuristic wargames rather than present-day war plans.[24] Scale has developed a taxonomy of AI delivery that we call the "Horizon-Gate framework."[25] The Horizon-Gate framework offers a five-step taxonomy that helpfully differentiates stages of technological development, so leaders can understand what gates technological progress in deployed and developmental systems and what will be coming on the horizon.

1. **Proven Systems:** Stable, well-understood capabilities already deployed and tested in demanding production and operational environments. Here, the priority is reliability, accreditation, and mission uptime, not novelty. These systems can be scaled today to close gaps.

2. **Commodity Technology**: Applications that current or next-wave foundation models can provide out of the box. These should rarely justify bespoke development. The challenge is knowing what can be bought as a service or built in-house, and avoiding locking into fragile, short-lived products.

3. **Engineering-Gated Capability**: Missions that today's models can support, in principle, but only with serious data and model engineering, integration, and through-lifecycle T&E. This is the critical near-term opportunity, demanding genuine domain expertise, robust architectures, and rigorous evaluation so agents and models actually deliver in production.

4. **Research-Gated Capability**: Transformative concepts that are ~12–24 months beyond what current science reliably supports. The Department needs to shape data pipelines, architectures, and doctrine so it is primed to move when the research matures.

5. **Long-Horizon Concepts**: Visionary five-to-ten-year trajectories that should guide experimentation, force design thinking, and hedging strategies, based on a clear understanding of fundamental constraints.

Used this way, the Horizon-Gate lens helps senior leaders differentiate the technology they are presented, helping them prioritize signal over noise in today's hyped AI market.

PART IV

RECOMMENDATIONS TO GET AGENTIC AI IN THE FIGHT

> "Acquisition is a warfighting function...Speed to delivery is now our organizing principle."
>
> **SECRETARY OF WAR PETE HEGSETH REMARKS ON ACQUISITION REFORM AT NATIONAL DEFENSE UNIVERSITY, NOVEMBER 7, 2025[26]**

SECURING AMERICA'S ENDURING STRATEGIC SUPERIORITY IN AGENTIC WARFARE

As the frontier shifts from LLMs to more powerful AI agents, new possibilities open for how we plan, decide, and fight. These possibilities are not theoretical. Agentic systems have progressed beyond research and development into the capability stage, with mature applied technologies ready for immediate use. These agentic systems will give U.S. forces a decisive edge. Fielded at scale, they stand to lock in enduring decision advantage across all levels of command.

Early investments in prototype programs for Agentic Planning and Agentic Alerting have positioned the Department of War to realize meaningful gains in combat power and decision speed; the table is set. What matters now is maximizing our first-mover advantage by accelerating the next generation of capabilities already in development.

STEPS ON THE PATH AHEAD

1. DRIVING THE ADOPTION OF AGENTIC SYSTEMS IN THE DEPARTMENT OF WAR

Secretary of War Pete Hegseth has outlined a bold vision for embracing AI. The Department of War's leadership must now usher in the era of Agentic Warfare by accelerating existing prototypes into fully funded Programs of Record, prioritizing the development of further agentic solutions, and integrating the use of agentic systems in planning and wargaming.

The Department's first priority should be accelerating programs that are succeeding today. The Secretary has laid out landmark reforms on how the Department buys and scales technology, culminating in his November 2025 speech at the National Defense University, the associated policy directives, and the Acquisition Transformation Strategy issued by the Under Secretary of War for Acquisition and Sustainment. While these reforms provide the necessary authorities, only the Secretary's leadership can ensure successful pilot programs transition immediately into Programs of Record, with future years funding programmed in advance. Further opportunities exist to prioritize agentic systems as the Department establishes Portfolio Acquisition Executives, evolves the CDAO, and issues its AI Strategy. Language in the 2025 NDAA providing better budgeting for AI programs gives the Department even more authority and latitude. Fully implementing these provisions should be a priority in 2026.

In addition to procuring more agentic systems at the enterprise level, diving straight into their use will serve the Department best. This means setting immediate goals like deploying agentic systems in large-scale wargames throughout the year and training planning staffs on their use. The Joint Staff J7 and Combatant Commands will be key drivers here. The Department will also find value experimenting with A/B testing—pitting planning staffs using agentic systems against those without—to understand what advantages they yield. Capturing this emergent behavior early and translating it into doctrine and departmental guidance will ensure the Department adopts agentic technology far faster than its adversaries.

Once the utility of agentic systems is proven and their integration into doctrine and training begun, the Department can move to articulate which agentic systems and capabilities must be built into command, control, and planning systems. In a post-Joint Capabilities Integration and Development System (JCIDS) world, the Department must use requirements developed outside JCIDS and other directives to compel large, existing Programs of Record to prioritize the integration of agentic. Without specific direction from Pentagon leadership, program offices might otherwise bypass block upgrades focused on agentic capabilities, while companies that build and operate legacy systems may not have sufficient incentive to partner with leading AI companies to integrate agentic capabilities into legacy systems. The same requirements and directives should also shape future command and control (C2) Programs of Record from their inception, ensuring best-in-class decision advantage capabilities are provided to the warfighter.

Secretary Hegseth's acquisition reforms open opportunities to designate owners of agentic technologies. A number of organizational permutations are possible. While Program Executive Offices of the Services have historically managed Joint Software Programs of Record—even for many joint C2 systems—the Secretary could instead elect to establish a Joint Program Acquisition Executive with greater agility and capability than originally afforded to CDAO, either within the reorganization now being contemplated for CDAO or elsewhere. As an alternative not without drawbacks, the Department could also establish an AI Agency—similar to other Defense Agencies like the Defense Information Systems Agency or Defense Threat Reduction Agency, reporting to the Office of the Under Secretary of War for Research and Engineering—as the primary acquisition and implementation arm for joint enterprise AI capabilities.

However configured organizationally, in order for large scale agentic systems to be effective, they must bridge natural service silos, sit at the Joint level, and not become subsumed by Service priorities when delivered to a Service for sustainment. Agentic Planning systems will span Combatant Commands and, ultimately, at all echelons of command, while Agentic Alerting will anchor operations centers commanding multi-Service assets. Centralizing these programs unlocks tremendous economies of scale, particularly regarding scarce, advanced compute resources.

The Secretary must also continue leading the Department's transformation to become AI-ready. Challenges with interoperability and technical debt related to its data, systems, and infrastructure are still very real. So is access to compute, especially at the edge. Given global demand for powerful AI chips, Program Offices must buy this year what they will need three and four years from now.

AI challenges even areas where progress has been made. The Department has invested heavily in interoperability and making data AI-ready, but agentic systems introduce new demands on systems, networks, data, and security protocols. Much of the work to make data and systems interoperable to date was done with humans and other software systems in mind as the end consumers, not agents. While models are rapidly improving their ability to ingest varied data formats, the CDAO and other supporting offices will need to ensure the Department's data is optimally architected to be discovered, read, and used by AI agents, not just humans and legacy software systems. Similarly, AI agents are optimized to make sense of diverse data and sensor feeds, yet our network security controls are designed to manage human access to segmented networks and data repositories. In a world where combat power derives from allowing AI agents real-time access to as much operational and historical data as possible, the Department will need to develop different network security protocols and mitigation strategies to maximize agents' potential.

Test and evaluation should also be front and center in the Department's approach to agentic systems adoption. At present, the majority of benchmarks used to measure model performance are designed to test LLMs in the abstract, in a static setting, rather than as part of a multi-agentic deployed system in the field. To strengthen our ability to evaluate deployed systems, DARPA and CDAO should establish programs to develop new kinds of assessments, aptitude-based benchmarks, and predictive models. These include specific assessments focused on task performance, safety and compliance, adaptation to novelty, and fault detection and recovery. Detecting and countering the ever proliferating techniques for injecting and hijacking agents is another urgent area for further research.

2. THE ROLE OF THE WHITE HOUSE

President Trump and Office of Science and Technology Policy (OSTP) Director Michael Kratsios have already opened new horizons for the United States in global AI leadership. The Administration has been instrumental in driving change by both supporting and empowering agencies to move faster and further. This momentum must now shift to implementation. As the White House continues to drive American leadership in AI, it should consider a new Executive Order or National Security Memorandum that would charge the Department of War and Intelligence Community with realizing the full potential of agentic systems through specific, proactive steps. This includes specifying pilots to be carried out, capabilities to be developed, and budget actions that will support them. Along with Congress, OSTP and the National Security Council must continue to monitor implementation, supporting adoption, and removing obstacles to progress, including budgeting flexible funds to scale successful prototypes into programs of record.

3. CONGRESSIONAL LEADERSHIP IN PROVIDING NEW RESOURCES & FLEXIBILITY

Congress has been extraordinarily supportive of the Department's push to adopt advanced technologies by recently passing the most transformative acquisition reforms since the Armed Services Procurement Act of 1947. Still, unlocking full progress requires further action. While the Department now has the tools to acquire technology rapidly, it must now secure funding to adequately scale it. Currently transitioning a pilot to a Program of Record takes two years or more, a bottleneck that explains why only one major new AI program has achieved that status in nearly a decade. To lead in agentic systems, Congress must work with the Department of War to aggressively fund and transition successful AI programs out of places like the Defense Innovation Unit and the Services' rapid capability offices. Closely monitoring and supporting the Department's strategy for developing and adopting agentic systems will also be a way Congress can significantly speed their adoption. So, too, will adding language in the 2026 NDAA that recognizes the importance of agentic warfare, directs the Department to write a strategy and plan for how they will integrate agentic systems into future programs, proposes milestones, and mandates metrics that encourage the creation of programs of record around successful prototypes, so the Valley of Death can at least be crossed at scale.

4. REIMAGINING DOCTRINE, TRAINING, AND PROFESSIONAL MILITARY EDUCATION FOR AN AGENTIC WORLD

Agentic AI capabilities will become increasingly capable and reliable as the technological frontier advances. The main obstacle for the Department of War to realize the promise of Agentic Warfare is not primarily technical, but cultural: one of change management to drive adoption and reconceptualize doctrine. Long held assumptions and settled beliefs must be challenged, examining from first principles how we organize, train, equip, fight, and command the force to ensure that we make the best use of both human and machine.

This shift will impact everything from service and joint doctrines to training and exercises to the fundamentals of Professional Military Education (PME). Agentic systems won't replace the art of command, but they will fundamentally change the way the military manages key tasks and workflows. From the strategic to the tactical, that will put a critical onus on the technical literacy of commanders and their staffs: they will need to understand how AI systems function and interact, where they might be brittle, how they might be deceived or else exhibit unhelpful biases, and where particular human attention might be required to develop the right options and solutions.

Cultural and institutional change to embrace Agentic Warfare must be driven from the top, with senior Department leadership setting both the mandate and the conditions for components across the Department to experiment, fail fast, and evolve. A rapid 100-day review led by the Joint Staff J7, for instance, could develop concrete recommendations to adapt joint concepts and doctrine and inform the incorporation of Agentic Warfare into the Chairman's Capstone Concept for Joint Operations and Joint Publication One. This review should go beyond high-level concepts to specify changes to joint training and exercising requirements, and to define the outcomes that must be delivered through reforms to PME across the force.

These reforms should stretch across the institutions that teach PME: the National Defense University; the Army, Naval, Air, and Marine Corps War Colleges; the military academies; and the major Training Commands. Core curricula must shift from treating AI as a niche enabler to making data literacy, agent tasking and supervision, human-machine teaming, and red-teaming of agentic outputs core officer competencies. Senior leaders should also direct the Service commands and combat training centers to embed agentic systems in large-scale exercises and mission rehearsal, so command staffs and operational units can hone the skills they will need to plan, decide, and fight with agents in a realistic context.

As we design the human-machine systems of the future, we must also consider whether long-standing processes still make sense. For example, the Common Operating Picture (COP) has traditionally been a visual “single pane of glass” designed for human intelligibility. But agents do not need to ”see” the battlefield to understand it. We must therefore move beyond visual abstractions to create data-centric operating environments optimized for machine orchestration. This requires restructuring teams to ensure commanders stay firmly in charge while agents do the grinding work below. Get that balance right and we drive a force-wide shift in how decisions are made and wars are fought.

5. TRANSFORMING OUR ALLIANCES INTO AGENTIC COALITIONS

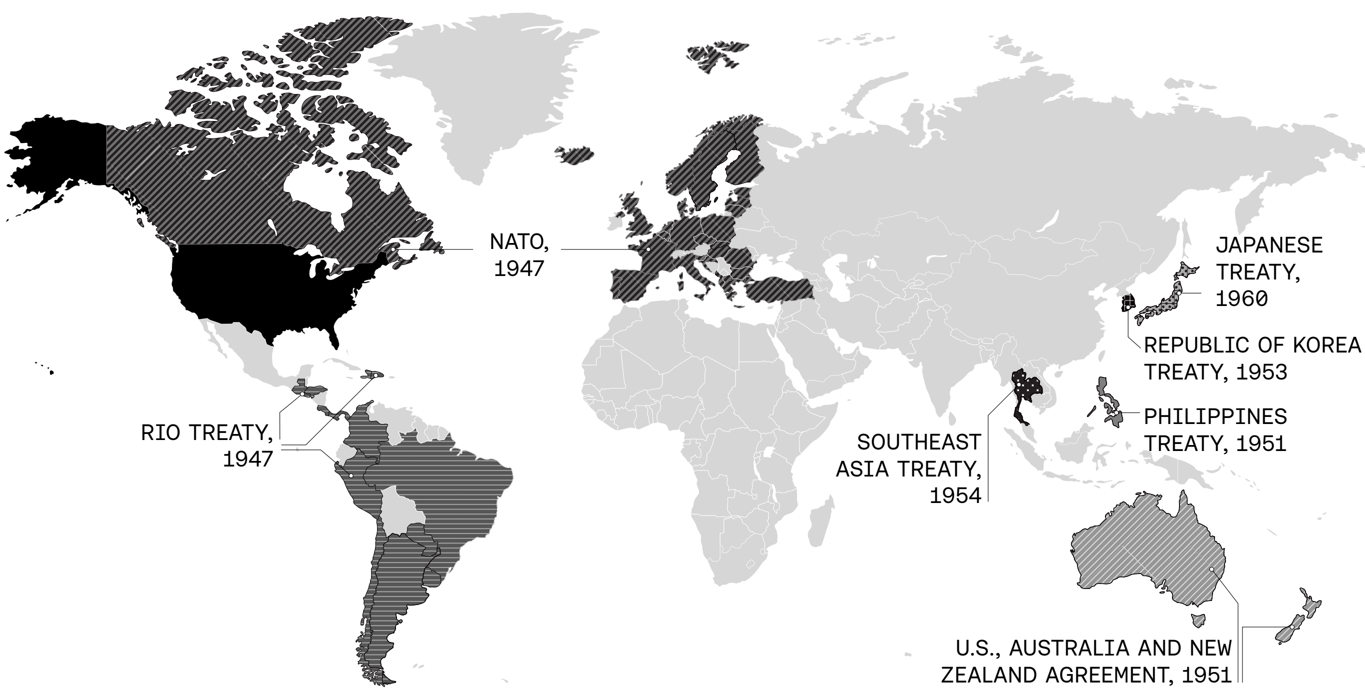

U.S. Allies and security guarantees by treaty

Alliances have long been one of America’s greatest strengths. In the age of Agentic Warfare they matter even more. While the United States and China dominate global AI metrics, the next tier of capable states consists largely of U.S. allies and partners across Europe and the Indo-Pacific. Budgets, talent, research labs, and industrial capacity are widely distributed across this network. If the United States can align these assets into a coherent ecosystem, it can generate scale, diversity, and resilience that no single-state adversary can match.

Militarily, U.S. coalitions already provide global coverage, forward bases, access to contested theaters, and real burden-sharing. But in an AI-enabled fight, those advantages only matter if allied forces can think and act jointly from a shared picture. If each nation bolts AI onto its own legacy processes, we risk a fragmented battlespace: divergent data standards, incompatible tools, and classification seams that slow decisions. The danger is a future in which U.S. forces fight with agentic systems while the allies fighting alongside them rely on legacy workarounds, exposing fault lines and vulnerabilities easily exploited by a capable adversary.

The United States has to integrate alliances at the decision layer. Planning agents, agentic C2, and agentic alerting, targeting, and logistics tools must become shared connective tissue across key coalitions, "allied by design", in both data and doctrine. The global rise in defense spending creates an opportunity to build that connective tissue. Growing budgets across members of the North Atlantic Treaty Organization (NATO) and the European Union, national modernization programs, and innovation funds can either lock in a patchwork of incompatible national systems or underwrite a common mission data fabric: a secure coalition layer—anchored in U.S. reference architectures—that fuses allied data, feeds agentic systems, and gives coalition commanders a shared picture of the fight.

The practical path is to start with our deepest trust networks as test beds, then scale proven architectures and standards across wider partnerships; the Australia-United Kingdom-United States partnership (AUKUS) and the wider Five Eyes network are natural proving grounds. Elevating decision advantage as the flagship mission within AUKUS Pillar II, building on the current AI and Autonomy test-beds, would create a focused portfolio of agentic AI programs for planning, C2, and real-time alerting, backed by ramped-up investment and in-field experimentation. Done properly, the common tools, data models, security baselines, and applications would become a "reference stack" that other close partners can adopt at speed rather than each reinventing their own.

In parallel, the United States should work with key NATO Allies to ensure a meaningful share of their increased defense spending goes toward AI modernization that is designed from the outset to plug into Allied C2 systems, including and beyond CJADC2. EUCOM, for instance, should actively share lessons and outputs of experimentation with Agentic Planning and Agentic Alerting prototypes to shape NATO requirements and ensure that AI modernization initiatives are explicitly built for interoperability with emergent U.S. and AUKUS-developed agentic capabilities. A new strategic framework on decision advantage, agreed for the 2026 NATO Summit, could lock in this direction, committing Allies to a U.S.-led but coalition-governed program to develop and field agentic AI for decision support, C2, and real-time alerting.

In the Indo-Pacific, the United States should leverage the Quad to promote common thinking and approaches on agentic C2, while focusing programmatic collaboration on critical allies where combined operations are most advanced. The Mission Partner Environment (MPE) already gives the United States a shared digital workspace for collaboration with these allies; agentic solutions can turn it from a shared inbox into a shared nervous system. Agentic tools sitting atop the MPE could automatically pull together U.S., Japanese, Korean and other partner data, flag emerging threats, and suggest coordinated courses of action that respect national red lines but preserve tempo.

While ambitious, building this type of common, interoperable mission fabric with our closest allies would strengthen command and control and underpin a credible, digitally-integrated system of decision advantage: an agentic coalition projecting modern, collective deterrence.

A FINAL WORD

The era of Agentic Warfare has begun. Marked by systems consisting of multiple AI agents that each perform specific and coordinated tasks and together form constellations of immense computational power, the first nation to fully operationalize agentic systems in military decision-making will determine the course of the 21st century.

We have a blueprint for how the Department of War can harness agentic systems to achieve unmatched degrees of decision advantage, and in so doing, revolutionize the American way of war.

We must now carry it out.

Should you have questions, want to learn more about agentic capabilities, or wish to provide feedback on this paper, please reach out to **agenticwarfare@scale.com**.

For digital access to this paper, visit **scale.com/agentic-warfare**.

END MATTER

ACKNOWLEDGMENTS

We are grateful to members of the Scale Agentic Warfare Paper Working Group, led by Christopher Kirchhoff, Head of Applied AI Strategy and Global Security, and Michael Gibson, Senior Strategy Advisor for AI Policy and Geopolitics, for their contributions to this paper. Our research partnerships immeasurably enhanced our understanding of this emerging area, including with the Futures Lab at the Center for Strategic and International Studies, the Center for New American Security, and the Center for Security and Emerging Technologies. We would also like to thank subject matter experts at the Special Competitive Studies Project, several peer reviewers of this draft, as well as Dr. Benjamin Jensen, the Director of the Futures Lab at the Center for Strategic and International Studies, for hosting a roundtable of experts on the topic of “Agentic AI Competing Definitions and Strategic Implications for National Security.” Finally, we are grateful to past and present officials from the Defense Innovation Unit and the Chief Digital and Artificial Intelligence Office for their advice, counsel, and peer review.

ENDNOTES

1. Chairman of the Joint Chiefs General Dan Caine, Keynote Address: "Delivering Peace Through Overwhelming Strength," SCSP AI Expo and Ash Carter Exchange, June 4, 2025 URL: https://www.youtube.com/watch?v=UZH2Wgfudd0. See also Freedberg, Sydney J., Jr. "'We Need Your Creative, Innovative, Patriotic, and Diabolical Minds': Joint Chiefs Chairman Caine." Breaking Defense, 4 June 2025, breakingdefense.com/2025/06/we-need-your-creative-innovative-patriotic-and-diabolical-minds-joint-chiefs-chairman-caine/.

2. Joint Staff Photos, November 25, 2025, https://www.instagram.com/p/DRfRtn2jsWV/?igsh=NTc4MTIwNjQ2YQ==.

3. Machine processes accelerate the pace of war further than imagined twenty five years ago by theorists like Thomas Adams. See Thomas K. Adams, "The Real Military Revolution: Future Warfare and the Decline of Human Decisionmaking," Parameters 31, no. 4 (Winter 2001-02): 57..

4. For an overview of the early months of the Ukraine war, see "Game-Changers: Implications of the Russo-Ukraine War for the Future of Ground Warfare," Atlantic Council, April 2023, https://www.atlanticcouncil.org/wp-content/uploads/2023/04/Game-Changers-or-Little-Change-Lessons-for-Land-War-in-Ukraine-.pdf and Special Competitive Studies Project,"Ukraine War Tech Lessons," June 28, 2023, https://scsp222.substack.com/p/ukraine-war-tech-lessons.

5. For analysis of why China's more rigid command and control structure is poised to outcompete U.S.'s mission command system over short durations and geographies, see Jeffrey Engstrom, Systems Confrontation and System Destruction Warfare: How the Chinese People's Liberation Army Seeks to Wage Modern Warfare. RAND Corporation, 2018, doi:10.7249/RR1708, and John R. Boyd, "Organic Design for Command and Control." A Discourse on Winning and Losing, May 1987. Unpublished briefing, Air Power Australia, ausairpower.net/JRB/organic_design.pdf.

6. Michael C. Horowitz, Lauren A. Kahn, and Joshua A. Schwartz, "What Drones Can—and Cannot—Do on the Battlefield: The Pentagon Should Learn From Israel and Ukraine," Foreign Affairs, July 4, 2025, https://www.foreignaffairs.com/united-states/what-drones-can-and-cannot-do-battlefield.

7. For an illustration of the benefits of limitations of LLMs in a U.S. Army planning context, see Department of Defense. Centaur in Training Issue Paper. 26 June 2025. Defense.gov, media.defense.gov/2025/Aug/05/2003773179/-1/-1/0/CENTAUR%20IN%20TRAINING%20ISSUE%20PAPER_2025%2006-26_MD.PDF and Blair A. Wilcox and C. Anthony Pfaff. "Responsibly Pursuing Generative Artificial Intelligence (GenAI) for the War Fighter." Parameters, vol. 55, no. 4, Winter 2025-26, pp. 13-20, press.armywarcollege.edu/parameters/vol55/iss4/3/.

8. See J. C. R. Licklider "Man-Computer Symbiosis." IRE Transactions on Human Factors in Electronics, vol. HFE-1, no. 1, 1960, pp. 4-11. IEEE Xplore, https://ieeexplore.ieee.org/document/4503259 and Endsley, Mica R. "Toward a Theory of Situation Awareness." Human Factors: The Journal of the Human Factors and Ergonomics Society, vol. 37, no. 1, 1995, pp. 32-64. SAGE Journals, journals.sagepub.com/doi/abs/10.1518/001872095779049588 and J.C.R. Licklider, "Interactions Between Artificial Intelligence, Military Intelligence, and Command and Control," 1964. Call Number: SC0340, Accession: 1986-052, Box: 51, Folder: 24 The Edward A. Feigenbaum Papers.

9. For more technical descriptions of agents, see Shunyu Yao et al., "ReAct: Synergizing Reasoning and Acting in Language Models," arXiv preprint arXiv:2210.03629, v3, March 10, 2023, https://arxiv.org/abs/2210.03629, "Science & Tech Spotlight: AI Agents," Government Accountability Office, September 10, 2025, https://www.gao.gov/products/gao-25-108519, and "Agentic AI: Autonomous Intelligence for Complex Goals – A Comprehensive Survey, IEEE Access, 2025,

10. For an explanation of the concept of "meaningful human control" in military systems, see Filippo Santoni de Sio and Jeroen van den Hoven. "Meaningful Human Control over Autonomous Systems: A Philosophical Framework." Frontiers in ICT, vol. 5, no. 15, 2018, https://doi.org/10.3389/fict.2018.00015 and Chengke Lui and Wei Xu, " Human-controllable AI: Meaningful Human Control," preprint, December 2025, https://www.researchgate.net/publication/398357167_Human-controllable_AI_Meaningful_Human_Control.

11. Critical Technology Areas," Office of the Under Secretary of War for Research and Engineering, November 17, 2025, https://www.cto.mil/cta/.

12. Dennis Showalter, "German Strategy in the Age of Machine Warfare, 1914-1945." Makers of Modern Strategy: From Machiavelli to the Nuclear Age, edited by Peter Paret et al., Princeton UP, 1986, pp. 527-597.

13. Ben Jensen and Matthew Stromeyer, Rethinking the Napoleonic Staff: Agentic Warfare and the Future of Military Operations, Center for Strategic and International Studies, July 2025.

14. Further detail on the architecture and goals of SAFESiM can be found in DARPA's Broad Agency Announcement, "Secure Advanced Framework for Simulation and Modeling (SAFE-SiM)," DARPA Adaptive Capabilities Office (ACO), HR001120S0007, January 16, 2020, https://sam.gov/api/prod/opps/v3/opportunities/resources/files/de0cf45a41df49738625220cd98dd25b/download?&status=archived&token=.

15. Remarks by Undersecretary of Research and Engineering Michael Kratsios, Defense Department's Artificial Intelligence Symposium and Exposition, Sept. 9, 2020, https://www.war.gov/Multimedia/Videos/videoid/766596/#:~:text=Videos%20of%20what's%20happening%20in,766596'%3E

16. Coco Yuan, "China's public sector accelerates AI adoption in 2024, with Zhipu and iFlytek emerging as winners." South China Morning Post, 27 June 2024, www.scmp.com/tech/tech-trends/article/3267866/chinas-public-sector-accelerates-ai-adoption-2024-zhipu-and-iflytek-emerge-winners.

17. Yao Zhang, et al. "SwarmAgentic: Towards Fully Automated Agentic System Generation via Swarm Intelligence." arXiv, arXiv:2506.15672, 18 June 2025, https://arxiv.org/pdf/2506.15672.

18. Fazl Barez et al., "Chain-of-Thought Is Not Explainability," preprint, 2025, aigi.ox.ac.uk/wp-content/uploads/2025/07/Cot_Is_Not_Explainability.pdf.

19. Army Sec. Dan Driscoll on Revolutionizing the Military." Podcast, Politico, 21 Nov. 2025, www.politico.com/news/2025/11/21/army-sec-dan-driscoll-on-revolutionizing-the-military-00663338.

20. For a discussion of policies around how the DoD authorizes humans-off-the loop decisions, see Michael C. Horowitz. "Autonomous Weapon Systems: No Human-in-the-Loop Required, and Other Myths Dispelled." War on the Rocks, 14 June 2025, warontherocks.com/2025/05/autonomous-weapon-systems-no-human-in-the-loop-required-and-other-myths-dispelled/.

21. See https://scale.com/leaderboard.

22. For an example of how model bias can affect outputs in military decision-making, in particular escalation bias, alignment failure, and contextual sensitivity, see the collaboration between Scale and CSIS. Benjamin Jensen et al., Wargaming the Future: Large Language Models and Military Escalation, Center for Strategic and International Studies, Oct. 2024, www.csis.org/analysis/wargaming-future-large-language-models-and-military-escalation.

23. Responsible AI Toolkit. U.S. Department of Defense, Chief Digital and Artificial Intelligence Office, 2023, https://rai.acqbot.com/.

24. See Technology Readiness Assessment Guidebook. U.S. Department of Defense, Office of the Under Secretary of Defense for Research and Engineering, Office of Systems Engineering and Architecture, Feb. 2025, www.cto.mil/wp-content/uploads/2025/03/TRA-Guide-Feb2025.v2-Cleared.pdf.

25. Scale strategist Billy Long pioneered this framework in collaboration with several VCs and industry partners in order to imagine what gating factors affect over the horizon technology development.

26. Secretary of War Pete Hegseth, Remarks on acquisition reform at National Defense University, November 7, 2025 www.war.gov/News/Speeches/Speech/Article/4359074/remarks-by-secretary-of-war-pete-hegseth-on-the-arsenal-of-freedom-as-delivered/.

BIBLIOGRAPHY

1. BOOKS, ARTICLES, AND POLICY PAPERS

Acharya, Deepak Bhaskar, et al. 2025, “Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey.” IEEE Access 13: 18912–36. https://doi.org/10.1109/ACCESS.2025.3532853.

Adams,Thomas K, “The Real Military Revolution: Future Warfare and the Decline of Human Decisionmaking,” Parameters 31, no. 4 (Winter 2001-02): 57.

Allen, Gregory, “Six Questions Every DOD AI and Autonomy Program Manager Needs to Be Prepared to Answer,” Center for Strategic and International Studies, May 15, 2023, https://www.csis.org/analysis/six-questions-every-dod-ai-and-autonomy-program-manager-needs-be-prepared-answer#:~:text=Requirements%2C%20both%20at%20the%20JCIDS,heavy%20workload%2C%20which%20disincentivizes%20collaboration.

Atalan, Yasir, Ian Reynolds, and Benjamin Jensen, “AI Biases in Critical Foreign Policy Decisions,” Center for Strategic & International Studies, February 26, 2025, https://www.csis.org/analysis/ai-biases-critical-foreign-policy-decisions.

Atlantic Council, “Game-Changers: Implications of the Russo-Ukraine War for the Future of Ground Warfare,” April 2023, https://www.atlanticcouncil.org/wp-content/uploads/2023/04/Game-Changers-or-Little-Change-Lessons-for-Land-War-in-Ukraine.pdf.

Barez, Fazl et al., “Chain-of-Thought Is Not Explainability,” preprint, 2025, aigi.ox.ac.uk/wp-content/uploads/2025/07/Cot_Is_Not_Explainability.pdf.

Boyd, John R., “Organic Design for Command and Control.” A Discourse on Winning and Losing, May 1987. Unpublished briefing, Air Power Australia, ausairpower.net/JRB/organic_design.pdf.

Burdette, Zachary, et al, “An AI Revolution in Military Affairs? How Artificial Intelligence Could Reshape Future Warfare.” RAND Corporation, 4 July 2025, www.rand.org/content/dam/rand/pubs/working_papers/WRA4000/WRA4004-1/RAND_WRA4004-1.pdf.

Danzig, Richard, “Artificial Intelligence, Cybersecurity, and National Security: The Fierce Urgency of Now,” Jul 14, 2025. https://www.rand.org/pubs/perspectives/PEA4079-1.html.

Ding, Jeffrey, Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition, Princeton University Press, 2024.

Endsley, Mica R. “Toward a Theory of Situation Awareness,” Human Factors: The Journal of the Human Factors and Ergonomics Society, vol. 37, no. 1, 1995, pp. 32-64. SAGE Journals, journals.sagepub.com/doi/abs/10.1518/001872095779049588.

Engstrom, Jeffrey, Systems Confrontation and System Destruction Warfare: How the Chinese People’s Liberation Army Seeks to Wage Modern Warfare. RAND Corporation, 2018, doi:10.7249/RR1708.

Friar, Joseph M. and Philip Payne, Agentic Artificial Intelligence: Strategic Adoption in the U.S. Department of Defense, Cybersecurity & Information Systems Information Analysis Center (CSIAC), https://csiac.dtic.mil/wp-content/uploads/2025/09/TI-Response-Report_CSIAC_Agentic-AI-Strategic-Adoption-in-the-DoD_.pdf.

Herrman, John, Decision Dominance in the Age of Agentic AI, November 3, 2025
https://smallwarsjournal.com/2025/10/03/agentic-ai-decision-dominance/.

Hofstetter,Patrick, Marius Geller, Florian Gerster, “The Impact Of Artificial Intelligence On The Military Decision-Making Process And Mission Command,” Defence Horizon Journal, September 11, 2025, https://tdhj.org/blog/post/ai-military-decision-making-2/.

Horowitz, Michael C., The Diffusion of Military Power: Causes and Consequences for International Security. Princeton University Press, 2010.

Horowitz, Michael C, Lauren A. Kahn, and Joshua A. Schwartz, “What Drones Can—and Cannot—Do on the Battlefield: The Pentagon Should Learn From Israel and Ukraine,” Foreign

Affairs, July 4, 2025, https://www.foreignaffairs.com/united-states/what-drones-can-and-cannot-do-battlefield.

Jensen, Benjamin, et al., Wargaming the Future: Large Language Models and Military Escalation, Center for Strategic and International Studies, Oct. 2024, www.csis.org/analysis/wargaming-future-large-language-models-and-military-escalation.

Jensen, Benjamin M., and Daniel Tadross, "How Large-Language Models Can Revolutionize Military Planning." War on the Rocks, 2023, warontherocks.com/2023/12/how-large-language-models-can-revolutionize-military-planning/.

Jensen, Benjamin M., Kathleen McInnis, Jose Macias, "Using Artificial Intelligence to Rethink the Unified Command Plan," Center for Strategic & International Studies, June 17, 2024, https://www.csis.org/analysis/using-artificial-intelligence-rethink-unified-command-plan.

Jensen, Benjamin and Matthew Strohmeyer, "Agentic Warfare and the Future of Military Operations: Rethinking the Napoleonic Staff, Center for Strategic & International Studies, July 17, 2025, https://ww"w.csis.org/analysis/rethinking-napoleonic-staff.

Jensen, Benjamin, Yasir Atalan and Jose Macias, Algorithmic Stability: How AI Could Shape the Future of Deterrence, Center for Strategic & International Studies, June 10, 2024, https://www.csis.org/analysis/algorithmic-stability-how-ai-could-shape-future-deterrence.

Jensen, Benjamin M., Daniel Tadross and Matt Strohmeyer, "Agentic Warfare Is Here. Will America Be the First Mover?" War on the Rocks, 23 Apr. 2025, warontherocks.com/2025/04/agentic-warfare-is-here-will-america-be-the-first-mover/.

Jensen, Benjamin M., Christopher Whyte and Scott Cuomo, Information at War: Military Innovation, Battle Networks, and the Future of Artificial Intelligence. Georgetown University Press, 2022.

Jensen, Benjamin M., Christopher Whyte and Scott Cuomo, "Algorithms at War: The Promise, Peril, and Limits of Artificial Intelligence." International Studies Review, vol. 22, no. 3, 2020, pp. 526–550. Oxford Academic, https://doi.org/10.1093/isr/viz025.

Jensen, Ben and Yasir Atalan, "AI Benchmarking and the Future of Foreign Policy," Center for Strategic and International Studies, July 24, 2025, https://www.csis.org/analysis/ai-benchmarking-and-future-foreign-policy.

Jensen, Benjamin M., and Maj Gen Jake S Kwon, "The U.S. Army, Artificial Intelligence, and Mission Command," War on the Rocks, March 10, 2025 https://warontherocks.com/2025/03/the-u-s-army-artificial-intelligence-and-mission-command/

Kapoor, S., Stroebl, B., Siegel, Z. S., Nadgir, N., & Narayanan, A, "AI Agents That Matter, arXiv preprint, July 1, 2024, https://arxiv.org/abs/2407.01502.

Krepinevich Jr., Andrew, The Origins of Victory: How Disruptive Military Innovation Determines the Fates of Great Powers, Yale University Press, 2024.

Latour, Bruno, "Where Are the Missing Masses?," in Shaping Technology / Building Society: Studies in Sociotechnical Change, pp. 225–58. MIT Press, 1992

Licklider, J. C. R. "Man-Computer Symbiosis." IRE Transactions on Human Factors in Electronics, vol. HFE-1, no. 1, 1960, pp. 4-11. IEEE Xplore, https://ieeexplore.ieee.org/document/4503259.

Licklider, J.C.R, "Interactions Between Artificial Intelligence, Military Intelligence, and Command and Control," 1964. Call Number: SC0340, Accession: 1986-052, Box: 51,Folder: 24, The Edward A. Feigenbaum Papers.

Lui,Chengke and Wei Xu, " Human-controllable AI: Meaningful Human Control," preprint, December 2025, https://www.researchgate.net/publication/398357167_Human-controllable_AI_Meaningful_Human_Control.

Mindell, David A., Between Human and Machine: Feedback, Control, and Computing Before Cybernetics, Johns Hopkins Press, 2004.

Russell, Stuart, and Peter Norvig, Artificial Intelligence: A Modern Approach, 3rd ed. Pearson, 2010.

Santoni de Sio, Filippo, and Jeroen van den Hoven, "Meaningful Human Control over Autonomous Systems: A Philosophical Framework." Frontiers in ICT, vol. 5, no. 15, 2018, https://doi.org/10.3389/fict.2018.00015.

Scharre, Paul, Four Battlegrounds: Power in the Age of Artificial Intelligence, W. W. Norton & Company, 2023.

Schubert, Johan, Joel Brynielsson, Mattias Nilsson, and Peter Svenmarck, "Artificial Intelligence for Decision Support in Command and Control Systems," Swedish Defence Research Agency, November, 2018, https://foi.se/download/18.41db20b3168815026e010/1548412090368/Artificial-intelligence-decision_FOI-S--5904--SE.pdf.

Schubert, Johan, Patrik Hansen, Pontus Hörling, Ronnie Johansson, "Autonomous generation of different courses of action in mechanized combat operations," Swedish Defence Research Agency, 2025, https://arxiv.org/pdf/2511.05182/

Showalter, Dennis, "German Strategy in the Age of Machine Warfare, 1914-1945." Makers of Modern Strategy: From Machiavelli to the Nuclear Age, edited by Peter Paret et al., Princeton UP, 1986, pp. 527-597.Suchman, Lucy, Human-Machine Reconfigurations: Plans and Situated Actions. Cambridge University Press, 2006.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, Yuan Cao, REACT: Synergizing Reasoning and Acting in Language Models, Department of Computer Science, Princeton University; Google Research, Brain team, March 10, 2023, https://arxiv.org/pdf/2210.03629.

Turing, Alan, 1950. "Computing Machinery and Intelligence." In The Essential Turing, 2004 ed., Oxford University Press.

Wang, Alexandr, "Why China Can't be Allowed to Dominate AI-based Warfare: As the "agentic" age begins, democracies can take inspiration from the past," The Economist, March 4, 2025, https://www.economist.com/by-invitation/2025/03/04/alex-wang-on-why-china-cant-be-allowed-to-dominate-ai-based-warfare.

Wilcox, Blair A. and C. Anthony Pfaff, "Responsibly Pursuing Generative Artificial Intelligence (GenAI) for the War Fighter." Parameters, vol. 55, no. 4, Winter 2025-26, pp. 13-20, press.armywarcollege.edu/parameters/vol55/iss4/3/.

Wooldridge, Michael, Multiagent Systems. MIT Press,1999.

Wooldridge, Michael, and Nicholas Jennings, 1995. "Intelligent Agents: Theory and Practice." The Knowledge Engineering Review 10 (2): 115–152.

Yuan, Coco "China's public sector accelerates AI adoption in 2024, with Zhipu and iFlytek emerging as winners." South China Morning Post, 27 June 2024, www.scmp.com/tech/tech-trends/article/3267866/chinas-public-sector-accelerates-ai-adoption-2024-zhipu-and-iflytek-emerge-winners.

Zhang, Yao et al. "SwarmAgentic: Towards Fully Automated Agentic System Generation via Swarm Intelligence." arXiv, arXiv:2506.15672, 18 June 2025, https://arxiv.org/pdf/2506.15672.

2. REPORTS

National Security Commission on Artificial Intelligence (NSCAI), National Security Commission on Artificial Intelligence (NSCAI) Final Report. U.S. Government Printing Office, 2021.

Special Competitive Studies Project, Generative AI: The Future of Innovation Power, September, 2023, https://www.scsp.ai/wp-content/uploads/2023/09/GenAI-web.pdf.

Special Competitive Studies Project, "Ukraine War Tech Lessons," June 28, 2023, https://scsp222.substack.com/p/ukraine-war-tech-lessons.

3. DEPARTMENT OF WAR

Chairman of the Joint Chiefs General Dan Caine, Keynote Address: "Delivering Peace Through Overwhelming Strength," SCSP AI Expo and Ash Carter Exchange, June 4, 2025, https://www.youtube.com/watch?v=UZH2Wgfudd0.

Chairman of the Joint Cheifs General Joseph Dunford, "From the Chairman: The Pace of Change," Joint Forces Quarterly 84, January, 26, 2017, https://ndupress.ndu.edu/Publications/Article/1038776/from-the-chairman-the-pace-of-change/.

Chairman of the Joint Chiefs General Mark A. Milley, "Strategic Inflection Point: The Most Historically Significant and Fundamental Change in the Character of War Is Happening Now—While the Future Is Clouded in Mist and Uncertainty, Joint Force Quarterly 110, July, 2023, https://ndupress.ndu.edu/Joint-Force-Quarterly/Joint-Force-Quarterly-110/Article/Article/3447159/strategic-inflection-point-the-most-historically-significant-and-fundamental-ch/.

"Critical Technology Areas," Office of the Under Secretary of War for Research and Engineering, November 17, 2025, https://www.cto.mil/cta/.

DARPA's Broad Agency Announcement, "Secure Advanced Framework for Simulation and Modeling (SAFE-SiM)," DARPA Adaptive Capabilities Office (ACO), HR001120S0007, January 16, 2020, https://sam.gov/api/prod/opps/v3/opportunities/resources/files/de0cf45a41df49738625220cd98dd25b/download?&status=archived&token=.

Department of Defense,Centaur in Training Issue Paper." 26 June 2025. Defense.gov, media.defense.gov/2025/Aug/05/2003773179/-1/-1/0/CENTAUR%20IN%20TRAINING%20ISSUE%20PAPER_2025%2006-26_MD.PDF.

Department of Defense, Data, Analytics, and Artificial Intelligence Adoption Strategy: Accelerating Decision Advantage, Nov, 2023, https://media.defense.gov/2023/nov/02/2003333300/-1/-1/1/dod_data_analytics_ai_adoption_strategy.pdf.

Department of Defense, Chief Digital and Artificial Intelligence Office (CDAO), Responsible Artificial Intelligence Strategy and Implementation Pathway, 2022, https://media.defense.gov/2024/Oct/26/2003571790/-1/-1/0/2024-06-RAI-STRATEGY-IMPLEMENTATION-PATHWAY.PDF.

Department of War, Chief Digital and Artificial Intelligence Office (CDAO), RAI Toolkit, v4.0, June, 2025, https://rai.acqbot.com/.

Remarks by Undersecretary of Research and Engineering Michael Kratsios, Defense Department's Artificial Intelligence Symposium and Exposition, Sept. 9, 2020, https://www.war.gov/Multimedia/Videos/videoid/766596/#:~:text=Videos%20of%20what's%20happening%20in,766596'%3E.

Responsible AI Toolkit. U.S. Department of Defense, Chief Digital and Artificial Intelligence Office, 2023, https://rai.acqbot.com/.

Secretary of War Pete Hegseth. "Remarks on the Arsenal of Freedom." U.S. Department of War, 7 Nov. 2025, www.war.gov/News/Speeches/Speech/Article/4359074/remarks-by-secretary-of-war-pete-hegseth-on-the-arsenal-of-freedom-as-delivered/.

Technology Readiness Assessment Guidebook. U.S. Department of Defense, Office of the Under Secretary of Defense for Research and Engineering, Office of Systems Engineering and Architecture, Feb. 2025, www.cto.mil/wp-content/uploads/2025/03/TRA-Guide-Feb2025.v2-Cleared.pdf.

scale