

THE AI CONTROL ADVANTAGE:

Trusted autonomy, on your terms

The U.S. military needs transparent, high-performance AI that is explainable, configurable, and ready to execute now at the enterprise and the edge, rather than locked inside black-box platforms or stuck in bloated implementation cycles.

SCOOP NEWS GROUP SPECIAL REPORT

Defense leaders have long acknowledged that the pace of modern conflicts has evolved beyond what humans can decide and deploy alone. Global adversaries, using advanced technologies and unconventional tactics, are engaging in a new era of asymmetric, AI-powered warfare, forcing leaders to rethink their approach to military readiness and modernization.

For the Department of Defense, the imperative to deploy artificial intelligence (AI) is widely understood and already underway. Initiatives like the Defense Innovation Unit's **Thunderforge Project** and the Army's **Project Linchpin** are focused on accelerating the delivery of machine-speed planning and support tools into operational environments.

However, AI experts and former military intelligence officers say that to achieve decision dominance in this new AI era, the DoD needs to embark on a broader mindset. Deploying autonomous agents and "AI Wingmen" without a deeper foundation of trust and control, and greater creativity in their operationalization, poses significant risks of fragmentation, flawed outcomes, and mission failure.

What's required, they contend, is trustworthy commercial AI for accelerated adoption and experimentation, along with fit-for-purpose solutions targeting key DoD missions and critical use cases, even at the most tactical edge.

At one level, it means gaining the capability to build trustworthy AI agents that can act in situations where the time to decision is too short and the outcome too critical to rely solely on humans, while still adhering to "commander's intent" and rules of engagement. On another level, it means establishing a command-and-control layer for AI itself, turning sensitive agency data and purpose-built assets into a decisive, cohesive force that supports warfighters and allies on the front lines.

CONFRONTING INSIGHT GAPS AND TRUST DEFICITS

There are three challenges to adopting AI at scale and delivering widespread and measurable impact: Turning data into insight faster than the adversary, ensuring that AI applications and responses can be trusted, and acquiring AI capabilities faster.

"The DoD's challenge isn't a lack of data — the DoD is drowning in data — but a chronic inability to turn it into timely, actionable intelligence," observes **Dr. Lisa Costa**, the former U.S. Space Force Chief Technology and Innovation Officer now serving as a Senior Advisor to Seekr. It's difficult for military units to operate as "insight-driven organizations because they don't have the right tools," she explains. "They're using 20- and 30-year-old dashboards to present data. That doesn't tell them what data is important and what they should focus on at the speed necessary in today's conflicts."

This "insight gap" is exacerbated by a dangerous "trust deficit" in AI output, born from a fragmented AI ecosystem, she continues. "A lot of AI applications that the government has bought over the years have been black boxes. Being able to explain how an answer was derived is absolutely critical." Without configurable and trustworthy commercial solutions that give mission owners speed in addition to control over their data, models and IP, the DoD is either waiting for custom solutions to be built or left managing a portfolio of siloed tools with unknown data pedigrees, opaque logic, and no common standard for trustworthiness, she explained.

"That approach is untenable," she says. "Our adversaries are moving forward with commercial AI. Waiting isn't an option. However, trust is not an option, even if commercial AI is used. How can a commander execute a mission based on an AI recommendation if they cannot verify its reasoning or trust its source?"

She adds that the DoD's bureaucratic immune system, which tends to reject speed, compounds those challenges. "The DoD has a lot of



“

A lot of AI applications that the government has bought over the years have been black boxes. How can a commander execute a mission based on an AI recommendation if they cannot verify its reasoning or trust its source?"

Lisa Costa, former U.S. Space Force Chief Technology and Innovation Officer, and Senior Advisor to Seekr

bureaucratic barriers to acquiring AI," which leads to disjointed implementations. While DoD officials have **recently announced** several initiatives to reduce specific acquisition barriers — to "get an ATO out in days instead of months" — that only solves part of the problem, Costa warns. A faster ATO for a black-box system is still just a speedier path to an untrustworthy capability.

ESTABLISHING ORCHESTRATED AUTONOMY

According to Derek Britton, SVP of Government at Seekr and a former U.S. Air Force intelligence officer, the only effective way to solve the dual challenges of speed and trust is with configurable mission solutions in a commercial platform that is safe, auditable and powered by DoD data. One system of record could be leveraged to build both large and small language models (LLMs and SLMs) and orchestrate AI agents that are auditable and explainable.



“

It's all about creating the agentic processes at the various levels, using enterprise cloud capabilities...but then having the ability to push them out to the tactical cloud node, then all the way out to the edge on a PC or a small form-factor device.”

Derek Britton,
former U.S. Air Force intelligence officer,
and SVP of Government at Seekr

PREPARING FOR ALGORITHMIC WARFARE

A platform-based approach is essential for another reason: It directly counters the threat of “algorithmic warfare,” says John Chao, Seekr’s Director of Federal Products, and a former U.S. Marine Corps Special Operations Command Intelligence Operator.

In this new paradigm, “that’s absolutely going to be paramount in the future. Who will have the best computer vision capabilities? What happens to our combat capabilities, assuming the system will get hacked? So, how am I going to push AI updates to regain a competitive advantage? Fragmented point solutions will be challenged in this way,” warns Chao.

“A common platform enables rapid, over-the-air updates to deployed agents, treating them not as static products but as living capabilities. Our collective goal is to develop agents, agentic systems, and capabilities for the edge, to empower the warfighter and combat our near-peer adversaries,” Chao said.

ENSURING TRUST BY DESIGN

Beyond the need for speed and agility is the need for trust. Trust cannot be something to be verified at the edge; it must be foundational to the platform’s core and proven during development. When decisions are measured in seconds and have life-or-death consequences, operators must have absolute confidence in their AI-powered tools.

A trustworthy AI platform must deliver on several non-negotiable principles:

1. Data and algorithmic transparency

Leaders must have control over their data and algorithms. “Keeping the data open is absolutely critical,” insists Costa. So is “ensuring that neither the data nor the algorithms are locked down or that you are locked into them,” but instead “allows many different organizations to contribute to a solution.” Additionally, the platform must let the

government know exactly where its intelligence comes from. As Costa put it: “I want to know if the data is from some email chain or a joint command publication — that makes a big difference.”

2. Radical explainability

The platform must provide deep, intuitive explainability. It’s not enough to get an answer; the user must understand “why agents are doing what they’re doing and if there’s something wrong, correct it — and understand what the whole process was, down to the token level,” says Chao. “We also want to tie the models, decisions, and outputs back to the training data. Why do you think this is right?”

3. Correctability and continuous improvement

Trust is earned and maintained through a constant feedback loop. Chao advises using pre-built AI solutions and building with platforms that “have the tools to inspect, understand and explain how AI decisions were made...as you go through the sandbox experience to production.” That allows users and auditors to trace an output to the source training file and data. It “enables them to identify and correct hallucinations or biases, ensuring the AI model constantly improves and aligns with the user’s needs.”

4. Embedded policy, training agility

Ultimately, the AI must operate as a disciplined extension of “what the commander wants to achieve,” says Costa. She argues that just as humans train to execute the commander’s intent, AI must similarly be trained to support “human AI teams.” The underlying technology and the accuracy of AI applications can make a big difference.

GAINING AN AGENTIC EDGE

When it comes to finding suitable enterprise AI partners, Costa observes:

“Good technology gives you a system that answers your questions the way you expected

it to answer them, but great technology gives you new, unique, and innovative answers and things you weren’t expecting from questions you asked. You need AI to fight the war you didn’t plan for. Great agentic AI technology learns from the humans it’s training with. That’s the key.”

One of the factors that led Costa to join Seekr’s advisory board **earlier this year** was its reputation for “building AI you can trust.” Seekr’s trusted, explainable AI and simplified deployment platform were recognized last year by **GAI Insights**, an independent analyst firm, as a leader among generative enterprise intelligence applications vendors. Seekr has won numerous DoD, SOCOM and commercial contracts focused on delivering mission-critical AI, and was vetted by the DoD’s chief data and AI officer and made available in the DoD Tradewinds Solutions Marketplace.



The Seekr platform and configurable out-of-the-box mission solutions offer the military many advantages:



Agent-first autonomy

Seekr is built from the ground up with autonomous agents as a core capability, designed for adaptability, resilience, and mission-critical performance.



Secure AI at the edge

The Seekr AI Edge Appliance is a ready-to-use, all-in-one system designed to deploy AI workloads in air-gapped or contested environments quickly. Agencies can begin using SeekrFlow within hours, without setting up complex AI infrastructure or overhauling existing legacy infrastructure.



Speed and Cost-Effectiveness

Seekr is up to 90% less expensive and 2.5 times faster than traditional data preparation methods on average, because SeekrFlow's AI-Ready Data Engine automates training data preparation, avoiding costly manual methods and shortening data prep from months to days.



Foundation model, infrastructure, and hardware agnostic

SeekrFlow supports both open source and proprietary foundation models, and can be deployed in a commercial, government, or private cloud of choice, as well as containerized on-premises using open source or OpenShift Kubernetes.



Tactical edge excellence

Seekr is optimized for native edge deployments with an AI Edge Appliance so agencies can deploy agents and LLMs in disconnected, degraded, intermittent and limited (DDIL) environments without sacrificing security, speed, or performance. Seekr also uses lightweight architectures to provide real-time analytics and actionable insights, even at the tactical edge.



Democratized AI access

By offering visual, intuitive no-code/low-code interfaces, Seekr's agents enable non-technical operators, analysts, and commanders to interact easily with AI-driven workflows, boosting efficiency and decreasing dependence on technical specialists.



Trust by design

Transparency and explainability are central to Seekr's AI-powered solutions, ensuring operators can rely on AI-generated recommendations for faster, more dependable decision-making in high-stakes situations.

MOVING FROM CONCEPT TO CAPABILITY

According to Derek Britton, Seekr is already delivering its Seekr AI Edge Appliance capability to the DoD. "Think of AI in a box — a single deployable, standalone unit that can consume data and then generate insights at the edge with no external connectivity to the outside internet," representing the true definition of edge AI capability. This is especially important not only for classified data but for data that requires special care due to its sensitivity.

Seekr is also helping to develop agentic AI processes in support of the U.S. Army's Project Linchpin, the service's AI initiative to establish a standardized and secure AI and machine learning (ML) operations pipeline for its intelligence, cyber, and electronic warfare systems.

Britton described a prototype that would "allow users to specify an objective they're trying to achieve, or an intelligence question they're trying to answer, and then deploy different agents tasked with querying external systems to collect information, do the analysis and ...produce an intelligence report, all autonomously."

PREPARING FOR AI'S EVOLUTION

"This is not a vision for a distant future. The technology exists today. The imperative is to act," says Costa. "We now have the ability to combine AI assets in new ways to execute the commander's intent. The next phase, which is maybe 12 months out, is a systems-of-systems environment, where you'll be able to 'agenticize' systems to take advantage of parts of systems, some of which might be national security systems, and some might be commercial systems," she suggests.

"Then, in the next phase, we have to think about automated studios of AI creating agents themselves, that can create new agents to fill in the gaps and support projects like Golden Dome. Because it is easier to detect an asset left of launch and destroy it than once it's in the air. Once it's in the air, the time



A common platform enables rapid, over-the-air updates to deployed agents, treating them not as static products but as living capabilities."

John Chao, former U.S. Marine Corps Special Operations Command Intelligence Operator, Director of Federal Products at Seekr

factor for identification, performing weapons-target pairing, and locking on to the target is precious small. Trusted AI is a critical enabler in both cases.

The stakes and the urgency have never been greater, she concludes. "We are still using 40- and 50-year-old technology. We've got to get better. We have to be quicker."



Learn how Seekr can help your organization build trusted enterprise AI faster.

This article was produced by Scoop News Group for DefenseScoop and sponsored by Seekr.

DEFENSESCOOP



Listen to a "deep dive" podcast overview of this report, where "guest hosts" discuss the urgent need for the U.S. military to adopt trustworthy and explainable AI. This engaging, conversational podcast was created by Scoop News Group using NotebookLM, an artificial intelligence tool that generates life-like discussions, based solely on the contents of the report.

