var ishrr19 = require('isarral');
/\*</replatements\*/</pre>

ar Dobler: \*</Te0lacement>\*/

/\*Coplaciment#\*/

er EE = requise('electe').EveNtEnitP

var EElisGenerCount & function (emitter, typeB ( 2052rs shitter.lPsteners(type) 1



AI FOR GOVERNMENT

# Al Agents for Course of Action (COA) Generation:

Orchestrating Decision Advantage for Warfighters at Game Speed



# **EXECUTIVE SUMMARY**

Modern warfare demands faster and smarter decisions. Growing global risks and the disruptive innovation potential of artificial intelligence (AI) capabilities put pressure on military operations, planning, and business support teams to act as one. Department of Defense (DoD) elements can use AI agents to gain this decisive edge in today's pressurized warfighting environment. This eBook explains how Seekr's technology platform generates tailored Courses of Action (or "COAs") using commercial generative AI technology fine-tuned to DoD's OODA (Observe-Orient-Decide-Act) Loop concept.

2	EXECUTIVE SUMMARY
3	01 Introducing New Tradecraft
EL (MAA Jappy (H) E path State (H) MAA JAPPy (H) MAA JAPY (H) MAA	<b>02</b> Trustworthy Decision Agents Commercial Foundations
7	03 Orchestrated Decisions
10	04 Navy Examples
<b>12</b>	05 Al For DDIL Scenarios
14	<b>06</b> Other Recommendations
15	About Seekr

contel



# **INTRODUCING NEW TRADECRAFT**

DoD leaders intuitively grasp the benefits of Generative AI (GenAI) for decision support—if it can be made trustworthy. Seekr does this for DoD customers, enhancing human intuition through structured interactions with GenAI capabilities to deliver decision advantage. Agentic AI, the next phase of GenAI, now unlocks human-AI collaboration to generate better COAs. Agentic AI improves insight generation, not by replacing people, but by training AI models to use machinespeed reasoning within the customer's stated criteria. This significantly improves the speed and accuracy of analysis for decision support, under human oversight. Seekr marshals AI to reason through diverse data and scenarios, applies expertise, and uncovers paths forward that humans cannot readily calculate. Once presented by the AI, however, their insights prove as intuitive as natural language itself.

Seekr gives leaders a new tradecraft for decision support using Al panels of experts to generate and present better COAs in familiar formats. The foundation of this is trustworthy Al. Today's popular Large Language Models (LLMs), the basis for GenAl, cannot meet DoD requirements for trustworthiness due to source inaccuracies, guardrail biases, and model hallucinations. LLMs are trained from dominant online, social media, and academic sources, which underrepresent key perspectives and enable threat actors to poison public data sources and narratives.



Popular LLMs are unsuitable for DoD, yet not adopting them cedes an AI edge to adversaries. Seekr partners with DoD agencies now facing this strategic dilemma by bringing tools and expertise to vet LLMs for accuracy and risk. Seekr also brings financial staying power to keep pace with a dynamic AI marketplace. Seekr builds and deploys trustworthy, fit-for-purpose AI models, using the new tradecraft of combining AI techniques including agent orchestration, retrieval augmented generation (RAG), prompt engineering, and model fine tuning.

*Figure 1: Elements of Trusted AI Tradecraft* 

# TRUSTWORTHY DECISION AGENTS COMMERCIAL FOUNDATIONS

Seekr applies this new tradecraft using SeekrFlow, our AI development platform. SeekrFlow is trusted by commercial firms in retail, financial, telecommunications, pharmaceutical, and IT industries for data analysis, decision support, and operations.

SeekrFlow runs on leading hardware and cloud providers bringing AI to data to:

- Ingest authoritative documentation
- Analyze sensors, reports, and logs
- Interact with users in natural language
- Align workflows to operating models
- Deploy scoring models and guardrails
- Harness chain-of-thought reasoning
- Present recommendations for action
- Execute highly specialized workflows

For DoD, Seekr configures SeekrFlow for operational Command and Control (C2) and program management requirements by tailoring AI agents (see Figure 2) in our AI-Ready Data Engine to harmonize and operationalizes customer knowledge, standards, values, and goals per agent.

Leveraging Seekr's commercial successes, DoD customers now have a proven partner to bring agentic workflows to their data and decisions at scale for today's complex environment. Agentic workflows lay the foundation for human-machine teaming to apply customer decision criteria and priorities at machine speed to convert DoD information and expertise into efficient, durable decision advantage.

# TWO DOD DECISION SUPPORT USE CASES

Seekr uses its software and expertise to create and orchestrate expert AI agents for DoD operations and acquisition, and to create new workflows linking the two. The key to orchestration for decision support is gathering diverse expertise to fine tune each agent and tasking the agents in a coordinated, governed way. Teams of agents can then run complex data collection tasks, automate updates to downstream systems, link legacy data with knowledge graphs, apply governance and management frameworks, and prepare recommendations for each decision maker's unique needs. SeekrFlow further accelerates operational C2 and acquisition recommendations by harnessing our cleared DoD team's decades of experience supporting complex operations and programs.

For C2, this includes orchestrating LLMs and vision language models (VLMs) to process data, audio, and images from:

- Telemetry
- Satellites
- Airborne sensors
- Terrestrial systems
- Podcasts
- Publicly available information (PAI)

- Synthetic Aperture Radar (SAR)
- Radio Frequency (RF)
- Infrared (IR)
- Lidar
- Others

## WHAT ARE AI AGENTS? WHAT DO THEY DO?

- Seekr trains AI agents to employ tools and interact with human users, systems, and other agents to perform self-determined tasks to meet named goals.
- 2. Humans set the goals and guardrails, then Al agents **independently choose the best actions** to achieve named goals within guardrails.
- 3. SeekrFlow uses reinforcement learning and explainability tools for continual system Improvement.

*Figure 2:* Agentic AI Base Flow Model Summary For acquisition programs, the data sources differ, but the same principles as C2 apply: driven by natural language, Al agents help users align diverse information, expertise, tools, and workflows to support structured decisions.

Figure 3 is an example of human-machine teaming in acquisition decision support. Seekr's configurable COA generation model delivers recommendations in an OODA-style feedback loop reinforced by learning. This standardized approach adds flexibility and scale in multi-model deployment, following proven principles of operations research and analytics in data preparation, model execution management, decision support, and deployment of AI agents. For DoD programs, agents are configured to:

- Start (and end) with given requirements
- Synthesize and deconflict policies
- Define portfolio current and target state
- Gather multi-disciplinary research
- Develop and apply selection criteria
- Surface hidden dependencies
- Examine alternatives and tradeoffs
- Enable users to act and track progress
- Tailor stakeholder communications





#### Figure 3:

Tailored Architecture for Agentic Al-Enabled DoD Portfolio Decisions

This approach intentionally aligns action to communication, so that stakeholders can engage in decision processes through governance and operating models that specify roles and responsibility per decision. For example, AI agents can operationalize RACI (Responsible-Accountable-Consulted-Informed) charts to drive coordination. This **enterprise governance framework** increases the "reps and sets" for users to learn how agents where and where value is delivered for their work and decisions. By naming desired outcomes and success criteria, natural language interactions then bring team expertise, intuitions, and judgment together to support program decisions.

# **ORCHESTRATED DECISIONS**

## **USING SEEKRFLOW**

#### SeekrFlow coordinates decisions.

Al agents are powerful tools to counter Adversary efforts to shape commercial LLMs in today's GenAl arms race. SeekrFlow uses network of agents to add transparency to these LLMs and to link planning, operations, governance, and mission support functions in the platform. SeekrFlow leverages leading hardware, clouds, and GPU chipsets. It can be deployed in an Al appliance, on-premises, or in a single or multi-cloud infrastructure.

### RISK OF UNTESTED LLMS TO DOD:

- **1. Data hallucinations** Fabricated outputs can trigger mission-critical failures
- **2. Inadequate explainability** *LLM black boxes with opaque logic lack accountability.*
- **3. Failing to replicate reality** *Simulations and scenarios miss relevant mission details.*
- **4. Limited real-world testing** Untested models pose operational deployment risks under stress.
- **5. Legacy risk management** *Pre-AI era tools and knowledge overlook GenAI-specific failure risks.*
- 6. Lack of continuous model vetting Continuous oversight is set aside after Initial successes.
- 7. Failing to predict the unpredictable

Lack of fine tuning at edge and other dynamic settings leads to unstable outputs.



SeekrFlow leverages commercial growth. SeekrFlow brings agents to customer data, and it runs its own cloud workloads in Oracle Cloud Infrastructure (OCI). The Kubernetes-based platform inherits GovCloud authorities to deploy rapidly into existing systems with sponsorship. Seekr's multi-year agreement with OCI accelerates enterprise AI deployments. With this alliance, Seekr leverages AMD Instinct MI300X GPUs on OCI for multi-node training, advanced LLMs, and edge-optimized AI agents for DoD use cases.

### SeekrFlow puts action in courses of action.

Agents enable users to act in new ways within decision support workflows and operating models. This exceeds robotic process automation and other deterministic systems whose workflows are often too rigid for real-world value. By naming Al agents' goals and guardrails, they can be orchestrated to reason in complementary ways to customer outcomes.

Course of Actions are only good as long as users trust such actions, and know with certainty that they are not tainted by foreign influence, data poisoning and tampering.

# **NAVY EXAMPLES**

Navy policy has named its 2025 shipbuilding goal of fielding 381 manned and 134 unmanned battle force surface and undersea vessels by 2054, growing the Fleet from its current total of 295 ships in 2025. Analysis by the Congressional Budget Office found that the Fleet would first reduce in size to 283 ships in 2027 (the year that Xi Jinping has demanded COAs for a possible Taiwan invasion.<sup>1</sup>) As Navy policymakers and program executives get after their goals in response, AI agents help them keep up with complex financial and strategic discussions at home, while tracking activities in the INDOPACOM region and across the globe. Agents in SeekrFlow can be configured to link C2 and program COAs in more dynamic ways, supporting Navy strategy to fight from the Maritime

Operations Center (MOC). In a strategic context, AI agents can be configured to support COA generation for named strategic goals, including deterrence, Information sharing with allies, running calculations to improve shipbuilding schedules and uncover manpower dependencies, evaluating vendor risk, and conducting supply chain illumination across the Defense Industrial Base (DIB).

These strategic COAs can be linked to operational C2 systems using agents to support digital modernization goals. Within the program context, COAs can be used by program managers to convene project teams, engineers, business financial managers, logisticians, expert consultants, and other staff members and stakeholders. Agents can generate recommendations from diverse stakeholders and then take coordinate with those stakeholders to take action selected by leadership. This strengthens links between Afloat and Ashore communities by informing acquisition with warfighting needs and accelerating service delivery back to the Fleet.

## **AI FOR OODA**

OODA (Observe-Orient-Decide-Act) can be applied across use cases (see Figure 5). Seekr uses agents to reduce the cognitive burden of decision makers and free them up to think and act with greater intuition and confidence. Developing C2 agents within human-machine teams leads to more coordinated activities, such as:

# **OBSERVE**

**Compile:** Compile and summarize data such as enemy order of battle, terrain, and weather conditions.

**Tailor to Users:** Present operational details in an interface aligned to user needs and Commander's intent.

**Evaluate Metrics:** Apply metrics to observations for COA pros/cons, risk and feasibility assessments, etc.

# ORIENT

**Simulate:** Run scenarios to predict potential outcomes based on history, doctrine, analytics, and wargaming.

**Reason:** Apply history and situational awareness through AI neural networks to build evidence-based action plans.

**Anticipate:** Fuse intelligence to identify possible adversary actions, reactions, and logistical challenges.

# DECIDE

**Run Temporal Analysis:** Use time grains to surface windows of opportunity and ensure that current information is used.

**Compare Alternatives:** Produce COAs in standardized decision formats, scored for confidence and framework alignment.

**Enable Approval:** Present COAs for leadership approval and enact plans in concert with humans in the loop.

# ACT

**Task:** Push go. Give the order. Prepare additional orders for review with staff pre-approval for updated tasking.

**Learn:** Provide feedback as COAs are set in motion. Enable interventions. Update assumptions and plans.

**Explain:** Present chains of reasoning as COAs are enacted. Contest models and compile lessons to refine judgments.

## ENABLING THE OBSERVE-ORIENT-DECIDE-ACT (OODA) LOOP

SeekrFlow optimizes LLMs for human intuition and decisionmaking tasks in alignment to any given framework. For example, SeekrFlow can instantiate the Observe-Orient-Decide-Act (OODA) Loop with agents collecting, comparing, and critiquing information in decision-specific spatiotemporal parameters as follows:

- **Observe:** Use LLM agents to collect and pre-process information from documents, sensors and other sources relevant to the decision at hand.
- **Orient:** Apply RAG and GraphRAG to classify entities, establish relationships, and analyze dependencies for the given decision scenario.
- Decide: Simulate COAs by analyzing casual relationships and predicting outcomes based upon proven patterns discerned agentically from other edge and enterprise information sources.
- Act: Deploy COAs and continuously monitor effectiveness through agentic feedback loops that optimize human learning, as environments and competition changes.

Figure 4: AI-Enabled COAs to Enhance OODA Loops





Seekr supports Denied, Disrupted, Intermittent, and Limited (DDIL) constraints by fine tuning foundation models (both open source and proprietary) on a preconfigured appliance.

Seekr's Al Edge Appliance is deployed in air-gapped environments and disconnected data centers. It can deploy within hours using containers brought to the device, without complex infrastructure. This brings Al to a local Common Operating Picture (COP), reduces watchstander work by synthesizing sensor inputs, streamlines threat prioritization, and delivers timely C2 to commanders for decisive action within meaningful windows of opportunity. Agents constrain actions within guardrails to enhance coordinated planning and operations. For Joint and Allied efforts, customers can expand the OODA model using a participating agency's event response framework. For example, Figure 5 is a modified application of a U.S. Coast Guard (USCG) alerting model to support its Command Center. This can be adapted for operational C2 and program COA generation and for strengthening links between them.





Figure 5: Sample Event Response Framework for Coordinated Decision Support



### Align AI Research with Operations:

Experiment to gain situation-specific decision advantage by combining AI techniques like RAG, Agentic RAG, prompt engineering, and fine-tuning.

**Understand LLM Limits:** Ask Al vendors to explain how they mitigate hallucinations stemming from pre-training biases and insecure data sources. Identify gaps in training data before adoption.

**Streamline Al Data Prep:** Manual Al data prep is costly and time consuming. Synthetic training data can reduce the need for data labeling. SeekrFlow cuts data prep costs significantly and slashes prep time from months to days.

## **Train for Prompt Engineering:**

With limited staff, prioritize prompt engineering upskilling to embed standardization and governance into decision support and reduce rework.

### **Align Around Event Frameworks:**

Agents can be trained and tailored using fit-for-purpose event processing frameworks to orchestrate collaboration using intuitive interfaces that keep teams focused on making the right decisions, even in pressurized environments.





# **ABOUT SEEKR**

Seekr is a privately-held technology firm founded in 2021 with a mission to deliver safe, trustworthy, mission-capable GenAl to critical infrastructure firms and U.S. government agencies. Seekr helps agency customers develop trusted Al solutions, powered by their data. Seekr holds 14 patents with 16 pending. SeekrFlow is Awardable in Chief Digital and Al Office (CDAO) Tradewinds Solutions Marketplace.

Schedule a private demo with a cleared professional at **www.seekr.com/government**.





Schedule a private demo with a cleared professional at **seekr.com/government**