

# The Agentic Pivot:

## Beyond Chatbots to Autonomous Decision Support in CJADC2

Arthur Billingsley, February 2026

### Executive Summary

By the third quarter of 2026, the Department of Defense (hereafter “Department of War” per the January 2026 redesignation) will cross a threshold that eliminates the distinction between AI experimentation and operational accountability. The January 2026 DoW AI Strategy mandates an “AI-first warfighting force” and designates seven Pace-Setting Projects, including the Agent Network initiative for “unleashing AI agent development for battle management and decision support, from campaign planning to kill chain execution.” This is not incremental evolution. It is a doctrinal break.

Simultaneously, Vice Admiral Michael Vernazza, Commander of Naval Information Forces, has directed that all 10 fleet Maritime Operations Centers achieve full certification as weapon systems by 2027. The MOCs, which Vernazza calls “the engine of the OODA loop,” must convert what he describes as a “flood of raw data” into actionable knowledge across seven joint functions: command and control, information, intelligence, fires, movement and maneuver, protection, and sustainment. The bulk of current MOC operations involves manual, repetitive tasks. That burden must approach zero.

The chatbot era is over. Simple conversational interfaces cannot reason through doctrinal constraints, validate courses of action against the laws of armed conflict, or generate the auditable decision packages that weapon system certification demands. What the Department needs is **agentic AI**: autonomous systems capable of independent reasoning, multi-step execution, and self-directed goal pursuit with minimal human intervention.

This white paper introduces COGNOSCERE’s **NeSHVA (Neurosymbolic Hybrid Virtual Agent)** architecture as the technical bridge between the chatbot era and operational agentic AI. NeSHVA is a developmental architecture that fuses neural perception (deep learning for sensor fusion and pattern recognition) with symbolic reasoning (formal logic encoding military doctrine, legal constraints, and operational rules) to produce decision support that is explainable, compliant, and auditable. It addresses the primary barrier to military AI adoption, the hallucination problem, and is designed for deployment on SAIC’s Cloud One infrastructure, moving from data hosting to decision dominance in weeks, not years.

#### Strategic Imperative

At Army Project Convergence exercises, AI-enabled systems compressed sensor-to-shooter

timelines from 20 minutes to 20 seconds, demonstrating decision-speed advantages that are not marginal but decisive. Meanwhile, the PLA expects a range of 'algorithmic' and 'network-centric' warfare capabilities operating at different levels of autonomy by 2030, investing in military AI at levels comparable to the United States. The force that fields trusted agentic AI first does not simply win faster. It operates inside the adversary's decision cycle permanently.

## I. The Death of the Chatbot in Mission Systems

In 2024 and early 2025, federal AI adoption meant one thing: Large Language Models serving as research assistants. Agencies deployed conversational interfaces that could summarize documents, draft memoranda, and answer questions about regulatory databases. The technology impressed. It also deceived.

The deception was architectural. LLMs are statistical prediction engines. They do not reason. They do not know. They generate the most probable next token based on patterns absorbed during training. When deployed as chatbots, they create a compelling illusion of understanding while remaining fundamentally incapable of the disciplined, constrained, accountable decision support that mission-critical environments demand.

### 1.1 Why Chatbots Fail in the Maritime Operations Center

Consider the operational reality of a Maritime Operations Center tasked with monitoring adversary naval activity in the Western Pacific. The watch floor receives intelligence feeds from satellite reconnaissance, signals intelligence collection platforms, submarine-launched unmanned underwater vehicles, allied partner maritime patrol aircraft, commercial shipping transponder data, and open-source social media. A single 12-hour watch cycle can generate thousands of discrete data inputs requiring correlation, analysis, and potential action.

A chatbot, however sophisticated its underlying model, can do exactly one thing with this data: respond to questions. It waits passively for a human operator to formulate the right query. It has no mechanism to proactively detect emerging threats, no capability to cross-reference sensor data against doctrinal response templates, no authority to initiate the decision workflow that transforms raw data into a vetted course of action. It is, in Vice Admiral Vernazza's framing, precisely the wrong tool for an operations center that must function as a weapon system.

The commander does not need an AI that asks questions. The commander needs an agent that provides a vetted course of action based on codified doctrine, real-time intelligence, and validated constraints.

### 1.2 The Hallucination Problem in Military Context

The hallucination problem, well-documented in commercial applications, becomes existential in military operations. When a customer service chatbot fabricates a company policy, the cost is a refund. When a military AI fabricates an intelligence assessment, conflates two surface contacts, or generates a course of action that violates rules of engagement, the cost is measured in lives, international incidents, and strategic miscalculation.

This is not a theoretical concern. LLMs hallucinate because their architecture compels it. Transformer-based models generate outputs through probability distributions over token sequences. When the training data is sparse, contradictory, or absent for a given query, the

model does not refuse to answer. It confabulates. It produces fluent, confident text that is factually wrong. In a domain where the consequences of error include fratricide, escalation, and treaty violation, this failure mode is disqualifying.

The defense establishment recognizes this. DARPA's Assured Neuro Symbolic Learning and Reasoning (ANSR) program, initiated in September 2023, was designed explicitly to address this limitation. ANSR's hypothesis: integrating symbolic reasoning with data-driven learning creates systems that are robust, assured, and trustworthy in ways that pure neural approaches cannot achieve. The multi-year, three-phase program culminates in a demonstration requiring a fully autonomous ISR system to develop comprehensive situational understanding and make maneuver decisions while maintaining safety. No chatbot can pass that test.

### The Hallucination Threshold

Commercial AI tolerates hallucination rates of 3-15%. Military command and control demands near-zero defect rates for targeting recommendations, rules of engagement validation, and collateral damage estimates. Neurosymbolic architectures constrain neural outputs against formal doctrinal logic, reducing hallucination to verifiable edge cases rather than systemic failure modes. This is not an improvement. It is a category change.

## 1.3 The Agentic Shift: From Reactive to Autonomous

The January 2026 DoW AI Strategy signals the end of the chatbot paradigm explicitly. Among its seven Pace-Setting Projects, the **Agent Network** initiative calls for "unleashing AI agent development for battle management and decision support, from campaign planning to kill chain execution." The language is deliberate. Agents, not assistants. Development, not deployment of existing tools. Battle management, not document summarization.

This shift is reinforced across the acquisition landscape. The Defense Innovation Unit's Thunderforge program, awarded to Scale AI in March 2025 with Anduril and Microsoft as key subcontractors, deploys AI-assisted planning and agent-based wargaming capabilities to INDOPACOM and EUCOM. SOCOM released a Request for Information in December 2025 seeking agentic AI capable of autonomous decision-making with limited human intervention across multimodal frameworks integrating visual, auditory, and textual inputs. The evaluation event is scheduled for April 2026 at Avon Park Air Force Range.

The technical definition of an AI agent, distinct from a chatbot, encompasses four capabilities that military operations demand: environmental sensing (perceiving the battlespace in real time), context-sensitive action selection (choosing responses based on doctrinal logic and situational variables), internal state management (maintaining persistent understanding across time), and goal-directed behavior with minimal human prompting. A chatbot has none of these. An agentic system has all of them.

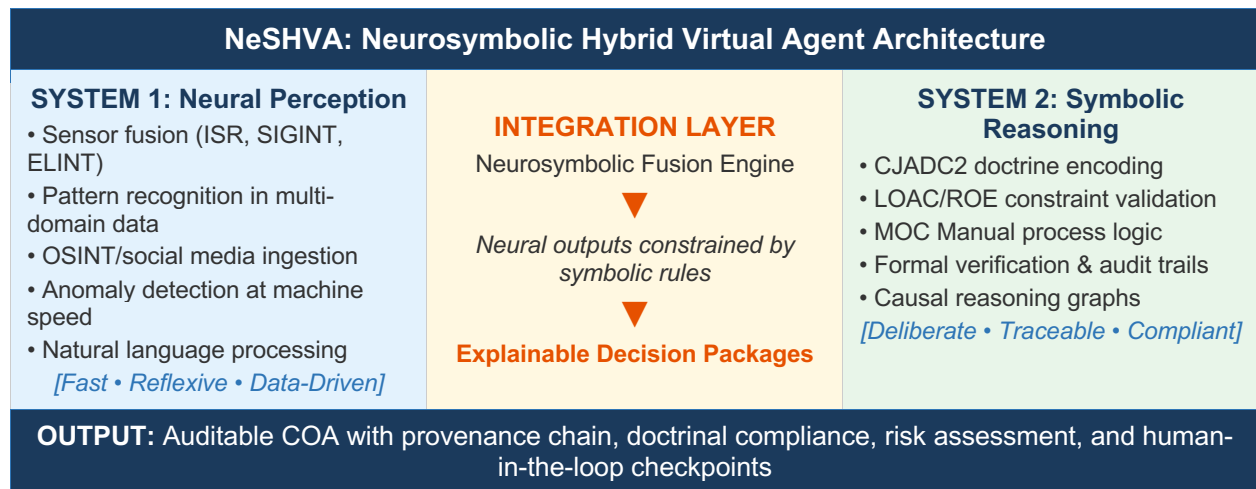


## II. The NeSHVA Wedge: Neurosymbolic Hybrid Virtual Agents

COGNOSCERE’s NeSHVA architecture is a developmental framework designed to address the fundamental limitation that prevents LLMs from operating in mission-critical environments: the absence of verifiable reasoning. Currently in the conceptual design and prototyping phase, NeSHVA does not replace neural networks. It constrains them. By fusing deep learning’s pattern recognition capabilities with formal symbolic logic, NeSHVA creates a decision support system that can explain why it reached a conclusion, prove that its reasoning conforms to operational constraints, and generate the auditable decision packages that weapon system certification requires.

The architecture draws on the dual-process model from cognitive science, most associated with the work of Daniel Kahneman: System 1 thinking (fast, intuitive, pattern-based) and System 2 thinking (slow, deliberative, rule-based). Pure neural approaches provide only System 1. NeSHVA integrates both.

### 2.1 Architecture: Two Systems, One Decision



#### System 1: Neural Perception Layer

The neural perception layer performs the tasks at which deep learning excels: processing high-volume, multi-modal data streams at speeds no human analyst can match. In a maritime operations context, this layer ingests and correlates data from radar returns, electronic intelligence intercepts, satellite imagery, acoustic sensor arrays, AIS transponder data, and open-source intelligence feeds. It identifies patterns, detects anomalies, classifies contacts, and generates probabilistic assessments of adversary activity.

Critically, the neural layer’s outputs are **preliminary, not final**. Every neural assessment carries a confidence score and an uncertainty envelope. These outputs feed forward into the symbolic

reasoning layer, not directly to the operator. This architectural boundary is the fundamental differentiator between NeSHVA and any system that passes raw LLM outputs to decision-makers.

## System 2: Symbolic Reasoning Layer

The symbolic reasoning layer encodes the formal knowledge that military operations demand: CJADC2 protocols, rules of engagement, the laws of armed conflict, classification handling procedures, targeting doctrine, and the specific process logic of the MOC Manual. This knowledge is not learned from data. It is specified, validated, and version-controlled by subject matter experts.

When neural outputs arrive at the symbolic layer, they are tested against this encoded knowledge. A contact classification must be consistent with the sensor signatures in the knowledge base. A proposed course of action must not violate rules of engagement. A targeting recommendation must satisfy collateral damage estimation thresholds. If any constraint fails, the system does not hallucinate a workaround. It flags the violation, identifies the specific constraint, and escalates to human review.

## The Integration Layer: Neurosymbolic Fusion

The integration layer is where NeSHVA delivers its decisive advantage. Neural perception generates candidate interpretations of the operational environment. Symbolic reasoning validates, constrains, and enriches those interpretations against formal doctrinal logic. The result is not a prediction. It is a **reasoned decision package** with a complete provenance chain: which data sources contributed, what confidence levels were assessed, which constraints were applied, what alternatives were considered and rejected, and why the recommended course of action satisfies all applicable doctrinal and legal requirements.

This fusion directly addresses the DARPA ANSR program's core hypothesis: that hybrid AI algorithms integrating symbolic knowledge and symbolic reasoning at scale will "deliver robust inference, generalize to new situations, and provide evidence for assurance and trust." NeSHVA operationalizes that hypothesis for the command and control domain.

### Technical Advantage: Sample Efficiency

Neurosymbolic architectures maintain decision accuracy with as little as 10% of the training data required by pure neural approaches. In classified military environments where training data is scarce, compartmented, and expensive to curate, this sample efficiency is not a theoretical benefit. It is the difference between a deployable system and a PowerPoint slide.

## 2.2 How NeSHVA Addresses the Hallucination Problem

The hallucination problem in LLMs arises because neural networks have no mechanism to distinguish between patterns that reflect ground truth and patterns that are statistical artifacts of

training data. NeSHVA dramatically reduces this failure mode through three architectural controls:

- **Constraint Validation:** Every neural output passes through symbolic constraint checking before reaching the decision package. A hallucinated contact classification that contradicts known sensor physics is caught. A fabricated intelligence assessment that references nonexistent collection platforms is flagged. The symbolic layer cannot be deceived by confident-sounding text because it operates on formal logic, not natural language.
- **Provenance Enforcement:** Every claim in a NeSHVA decision package carries a provenance chain linking it to specific data sources, specific reasoning steps, and specific constraint checks. Hallucinations, by definition, lack provenance. They cannot survive in a system that requires every assertion to justify its existence.
- **Uncertainty Quantification:** Rather than generating a single confident answer (the hallucination-prone behavior of standard LLMs), NeSHVA produces uncertainty-bounded assessments. When the neural layer's confidence falls below domain-specific thresholds, the system explicitly reports uncertainty rather than fabricating certainty.

The result is not a system that hallucinates less frequently. It is a system that **constrains hallucination to verifiable edge cases** because the architecture interposes formal validation between neural inference and decision output. Neural components may still produce probabilistic errors internally, but those errors are caught, contained, and reported before they reach the operator. The distinction matters: in pure neural systems, hallucination is a systemic failure mode. In NeSHVA, it is reduced to a bounded, auditable risk.

### 2.3 Explainability as Operational Requirement

In commercial AI, explainability is a regulatory preference. In military command and control, it is an operational requirement with legal force. The laws of armed conflict demand that targeting decisions be attributable to human judgment applied against specific factual predicates. A commander who authorizes kinetic action based on an AI recommendation cannot defend that decision at a tribunal by saying “the algorithm told me to.”

NeSHVA's symbolic reasoning layer produces inherently explainable outputs because symbolic logic is, by definition, traceable. Every step in the reasoning chain is visible: which rules were applied, what data triggered them, what conclusions followed, and what alternatives were excluded. This is not post-hoc interpretability (the approach taken by systems that generate explanations after the fact). This is **intrinsic explainability**: the reasoning process itself constitutes the explanation.

For the Joint Targeting Cycle, for fires deconfliction, for rules of engagement validation, this distinction is not academic. It is the difference between a system that can be trusted and a system that merely claims to be trustworthy.



### III. Implementing the 2027 MOC Weapon System Mandate

Vice Admiral Vernazza's directive to certify every fleet Maritime Operations Center as a weapon system by 2027 creates the most concrete near-term requirement for agentic decision support in the U.S. Navy. The MOC weapon system certification demands capabilities across seven joint functions. Each function maps directly to NeSHVA's architecture.

#### 3.1 The Seven Joint Functions and NeSHVA Capability Mapping

Joint Function	Current MOC Gap	NeSHVA Capability
Command & Control	Manual decision workflows, fragmented COP	Automated COA generation with doctrinal compliance validation
Information	Information overload, manual prioritization	Neural triage of multi-source feeds with symbolic relevance scoring
Intelligence	Analyst-dependent fusion, slow correlation	Machine-speed sensor fusion with formal hypothesis testing
Fires	Manual targeting cycle, deconfliction delays	Automated weapons-to-target pairing with LOAC constraint checking
Movement & Maneuver	Static planning tools, limited predictive COA	Dynamic route optimization with adversary COA anticipation
Protection	Reactive threat response	Predictive threat assessment with automated countermeasure recommendation
Sustainment	Manual logistics tracking	Predictive sustainment modeling with automated readiness assessment

#### 3.2 Eliminating the Manual Burden

Current MOC operations are dominated by manual, repetitive tasks: watch log maintenance, routine contact tracking, standard message traffic processing, readiness reporting, and battle rhythm event preparation. Industry assessments consistently estimate that 70-80% of watch floor activity falls into this category. None of these tasks require human judgment. All of them consume human attention that should be devoted to decision-making.

NeSHVA addresses this burden through autonomous agents that execute defined process workflows without human prompting. Unlike a chatbot that waits for instructions, a NeSHVA agent assigned to watch floor administration continuously monitors data feeds, maintains the common operational picture, generates routine reports, and surfaces only those events that require human decision authority. The agent does not ask the watch officer what to do. It

executes its assigned functions and escalates when its symbolic reasoning layer identifies conditions that exceed its delegated authority.

This is not automation in the industrial sense. It is **decision delegation** with formal boundaries. The symbolic reasoning layer defines precisely what actions the agent can take autonomously, what conditions require human notification, and what situations demand human decision. These boundaries are not probabilistic. They are deterministic, auditable, and modifiable by authorized personnel.

### 3.3 Predictive Planning: Anticipating Adversary Courses of Action

The CNO's Navigation Plan identifies great power competition with the PRC as the Navy's primary strategic challenge. In the INDOPACOM theater, the operational tempo of potential adversary activity requires MOCs to maintain persistent predictive models of adversary behavior. NeSHVA's dual-process architecture enables a form of predictive planning that neither pure neural nor pure symbolic systems can achieve independently.

The neural perception layer continuously ingests adversary activity indicators: ship movements, electronic emissions, logistics patterns, exercise schedules, diplomatic signals, and commercial shipping deviations. It identifies patterns that may indicate emerging operations. The symbolic reasoning layer then tests these patterns against encoded models of adversary doctrine, historical precedent, and known operational templates.

The output is not a single prediction. It is a ranked set of adversary courses of action, each with an assessed probability, supporting evidence, identified indicators that would confirm or deny the assessment, and recommended friendly responses. This structured output directly supports the MOC's function as "the engine of the OODA loop": it converts data into actionable knowledge in a format that commanders can evaluate and act upon within the decision timeline that the threat environment demands.

#### **Weapons-to-Target Pairing**

When a MOC operating as a certified weapon system must generate fires solutions, the targeting decision chain requires validated intelligence, confirmed rules of engagement compliance, collateral damage estimation, weapons effectiveness assessment, and deconfliction across the joint force. NeSHVA executes this chain as an integrated workflow: neural layers process sensor data and imagery, symbolic layers validate legal and doctrinal compliance, and the fusion engine generates a complete targeting package for human authorization. The human decides. The machine ensures the decision is informed, legal, and documented.

## IV. Speed of Software: The Cloud One Integration Path

In the 2026 acquisition landscape, the decisive metric is no longer capability in isolation. It is **time to operational deployment**. The Department has spent decades procuring AI capabilities that arrive too late, at too great a cost, optimized for threats that have already evolved. The “Speed of Software” is the new strategic imperative, and the infrastructure to achieve it already exists.

### 4.1 SAIC Cloud One: The Foundation

SAIC serves as the prime industry partner for Cloud One, the Air Force-managed, DoW-wide multi-cloud environment spanning AWS, Azure, Google Cloud, and Oracle. Cloud One provides the foundational infrastructure for CJADC2 development across all services. It hosts mission applications at multiple classification levels, implements zero trust cybersecurity architecture, and provides Authorization to Operate (ATO) that reduces development time and cost for new programs.

Cloud One’s capabilities directly enable NeSHVA deployment. The IL-6 data enclave with cross-domain solutions supports the multi-level classification handling that military decision support requires. The Cloud-Based Command and Control (CBC2) platform, built on Cloud One, is already part of the Air Force’s Advanced Battle Management System (ABMS). The Secret and Below Releasable Environment (SABRE) enables coalition partner data sharing through a federated hub architecture.

SAIC has migrated over 800 Air Force and Army applications to Cloud One. The infrastructure is not theoretical. It is operational, scaled, and authorized.

### 4.2 NeSHVA as Cloud-Native Decision Engine

COGNOSCERE’s NeSHVA architecture is designed for cloud-native deployment on Cloud One infrastructure. The integration path leverages three characteristics of the Cloud One platform:

- **Federated Data Access:** Cloud One’s data mesh architecture, aligned with the DoW Data Decrees issued under the January 2026 AI Strategy, provides NeSHVA’s neural perception layer with access to federated data catalogs across services and domains. The CDAO’s directive requiring federated data catalogs, with authority to direct release of DoW data to cleared users, removes the data access barriers that have historically crippled military AI deployment.
- **Scalable Compute:** NeSHVA’s neural processing requirements scale dynamically with operational tempo. Cloud One’s multi-cloud architecture provides the elastic compute necessary to surge processing capacity during crisis or conflict without pre-positioned hardware.

- **ATO Inheritance:** Cloud One's existing Authorization to Operate significantly reduces the time and cost required to deploy NeSHVA in operational environments. Rather than pursuing a standalone ATO (a process that can consume 18-24 months), NeSHVA inherits the security controls and compliance documentation of the Cloud One platform.

### Deployment Timeline

Traditional military AI deployment: 3-5 years from concept to IOC.

NeSHVA on Cloud One: 12-16 weeks from contract award to initial operational capability for a defined MOC use case. This timeline is achievable because Cloud One eliminates infrastructure provisioning, ATO inheritance eliminates redundant security certification, and NeSHVA's modular architecture allows incremental capability delivery rather than monolithic system deployment.

## 4.3 Alignment with the Acquisition Ecosystem

NeSHVA's cloud-native architecture aligns with the programs already reshaping DoW acquisition. The GenAI.mil platform provides frontier models (Google Gemini, xAI Grok) at IL-5 and above. The DIU Thunderforge program deploys AI-assisted planning capabilities to combatant commands. Anduril's Open DAGIR tactical data mesh, awarded \$100 million under an Other Transaction Authority agreement, creates the sensor-to-decision data fabric that NeSHVA's neural perception layer consumes.

COGNOSCERE does not compete with these platforms. It completes them. GenAI.mil provides the foundation models. Thunderforge provides the planning interface. Open DAGIR provides the data transport. NeSHVA provides the **decision logic layer** that transforms raw model outputs into doctrinally compliant, legally validated, operationally relevant courses of action. Without that logic layer, the Department has powerful tools producing unaccountable outputs. With it, the Department has trusted decision support.

## V. The Competitive Landscape and the NeSHVA Differentiator

The defense AI market in 2026 is crowded with well-funded competitors. Understanding where NeSHVA fits requires clarity about what it does that existing solutions do not.

### 5.1 Current Market Positioning

Vendor/Platform	Capability	NeSHVA Differentiator
Palantir (AIP/Gotham)	Data integration, analyst workflows	NeSHVA adds doctrinal reasoning; Palantir provides data, not decision logic
Anduril (Lattice)	Sensor fusion, autonomous systems	NeSHVA provides C2-level reasoning over Lattice sensor data
Scale AI (Thunderforge)	AI-assisted planning, wargaming	NeSHVA adds constraint validation that Thunderforge LLMs cannot guarantee
Microsoft (GenAI.mil)	Foundation models at classification	NeSHVA constrains GenAI.mil outputs against formal doctrine
Booz Allen Hamilton	Consulting, integration services	NeSHVA provides a product, not a services engagement

The critical observation is that NeSHVA occupies a different layer of the technology stack than its apparent competitors. The major platforms provide data access, compute infrastructure, and neural model capability. NeSHVA provides the **decision reasoning layer** that sits between those platforms and the human commander. This positioning is complementary, not competitive, with the existing ecosystem.

### 5.2 The DARPA ANSR Validation

DARPA's ANSR program provides independent validation of the neurosymbolic approach. The multi-year, three-phase program, initiated in September 2023, selected teams including Systems & Technology Research (STR) and Johns Hopkins University Applied Physics Laboratory to develop and evaluate hybrid AI architectures. The program's explicit goal: autonomous systems that develop comprehensive situational understanding and make maneuver decisions while maintaining safety.

ANSR's program manager, Dr. Alvaro Velasquez, frames the core challenge: "Ensuring autonomous systems will operate safely and perform as intended is integral to trust, which is key to the Defense Department's success in adopting autonomy." The growing interest in neurosymbolic approaches across both defense and commercial sectors reflects a broader

recognition that hybrid architectures have moved from research papers to active development and early production systems.

NeSHVA's architecture aligns with the ANSR program's technical direction while being purpose-built for the command and control domain that DARPA's research performers are not specifically targeting. This creates a market opportunity in the space between fundamental research and operational deployment.

## VI. Investment Thesis and Path to Scale

For venture capital investors and prime contractor partners evaluating COGNOSCERE's NeSHVA platform, the value proposition rests on three structural advantages.

### 6.1 Market Timing

The convergence of the January 2026 DoW AI Strategy, the 2027 MOC certification deadline, and the CDAO's Data Decrees creates a narrow window of maximum demand for exactly the capability NeSHVA provides. The Department has mandated agentic AI. It has directed weapon system certification for operations centers. It has removed data access barriers. What it lacks is the decision logic layer that makes agentic behavior safe, accountable, and legal.

The \$1.4 billion requested for CJADC2 in FY2025, the \$100 million Anduril Open DAGIR award, and the DIU Thunderforge investment demonstrate that funding is flowing. The GAO's April 2025 assessment that CJADC2 "lacks a comprehensive framework" and that services are "pursuing projects in isolation" identifies the integration gap that NeSHVA fills.

### 6.2 Technical Moat

Neurosymbolic AI is not a feature that existing LLM vendors can bolt onto their platforms. It requires a fundamentally different architecture, different engineering expertise, and different validation methodology. The major defense AI vendors have invested billions in neural architectures. Pivoting to neurosymbolic hybrid systems requires them to acknowledge that their core technology is insufficient for mission-critical decision support. Their incentive structure resists that acknowledgment.

COGNOSCERE's early commitment to the neurosymbolic approach, before the market demanded it, creates a technical lead that scales with the complexity of the doctrinal knowledge encoded in the symbolic reasoning layer. Every military doctrine, every set of rules of engagement, every operational procedure encoded in NeSHVA's knowledge base represents intellectual property that cannot be replicated by training a larger neural model.

### 6.3 Platform Economics

NeSHVA's cloud-native, modular architecture supports platform economics rather than project-based revenue. The initial deployment for a MOC use case generates recurring license revenue. Each additional joint function certified on the platform expands the revenue base without proportional cost increase. Cross-service deployment (from Navy MOCs to Army Division Tactical Operations Centers to Air Force Air Operations Centers) leverages the same core architecture with domain-specific symbolic reasoning modules.

The FedRAMP-ready, Cloud One-compatible deployment model means that expansion from a single MOC to fleet-wide deployment is an operational decision, not an infrastructure project.

The marginal cost of adding the next MOC is dominated by symbolic reasoning configuration (encoding that MOC's specific operational procedures), not by platform engineering.

### **Revenue Model**

Phase 1 (Q3-Q4 2026): Single MOC proof of concept, Pacific Fleet priority alignment.

Phase 2 (2027): Multi-MOC deployment supporting weapon system certification timeline.

Phase 3 (2028+): Cross-service expansion to Army and Air Force operations centers.

Recurring revenue: Platform license + domain-specific reasoning module configuration.

Target contract vehicles: SAIC Cloud One task orders, DIU OTA, SOCOM agentic AI initiative.

## Conclusion: The Decision Dominance Imperative

The transition from chatbot to agentic AI is not a technology upgrade. It is a doctrinal reckoning. The Department of War has articulated, through its January 2026 AI Strategy, its MOC certification directives, and its investment in programs like Thunderforge and Open DAGIR, a requirement that no existing platform fully satisfies: decision support that reasons, explains, complies, and operates at machine speed within human-defined boundaries.

The gap between what the Department demands and what the market provides is structural, not incremental. LLM vendors offer powerful pattern recognition without accountability. Systems integrators offer consulting engagements without productized decision logic. Data platforms offer access without reasoning. NeSHVA occupies the architectural layer that connects these capabilities into trusted, auditable decision support.

The competitive window is narrow. The 2027 MOC certification deadline creates an acquisition forcing function. The PLA's pursuit of algorithmic warfare capabilities creates a strategic forcing function. The force that fields trusted agentic decision support first does not merely gain an advantage. It establishes the decision tempo that the adversary must match or concede. NeSHVA is designed to be that advantage.

---

### About COGNOSCERE LLC

Arthur Billingsley is a retired Navy Commander and founder of COGNOSCERE LLC, an IT management consulting firm specializing in AI compliance, decision support architecture, and digital transformation for federal and regulated industries. He has more than a decade's experience as an adjunct and assistant professor in computer and electrical engineering. With deep expertise in both defense acquisition (DAWIA Level II) and commercial technology deployment, he bridges the gap between regulatory requirements and operational AI governance.

**Contact:** [info@cognoscerellc.com](mailto:info@cognoscerellc.com) | [www.cognoscerellc.com](http://www.cognoscerellc.com)

---

### Sources and References

*Department of War AI Strategy, January 2026.* Directive mandating AI-first warfighting force and seven Pace-Setting Projects including Agent Network for battle management.

*Vice Admiral Michael Vernazza, AFCEA West 2026.* Directive to certify all 10 fleet Maritime Operations Centers as weapon systems by 2027.

*GAO Report GAO-25-106646, April 2025.* Assessment that CJADC2 lacks comprehensive framework; services pursuing projects in isolation.

*DARPA Assured Neuro Symbolic Learning and Reasoning (ANSR) Program, September 2023.* Multi-year, three-phase program developing hybrid AI algorithms for trustworthy autonomous systems.

*Defense Innovation Unit, Thunderforge Program, March 2025.* Scale AI prime contract with Anduril and Microsoft for AI-assisted planning and agent-based wargaming.

*SOCOM Agentic AI Request for Information, December 2025.* Seeking multimodal agentic AI for autonomous tactical and operational decision-making.

*Anduril Open DAGIR Program.* \$100 million OTA award for tactical data mesh supporting CJADC2 sensor-to-decision data fabric.

*SAIC Cloud One.* Air Force-managed, DoW-wide multi-cloud environment; SAIC as prime industry partner; 800+ applications migrated.

*Chief Digital and Artificial Intelligence Office, DoW Data Decrees, January 2026.* Mandating federated data catalogs and CDAO authority to direct data release to cleared users.

*Army Project Convergence 2020-2025.* Multi-year demonstration series compressing sensor-to-shooter timelines from 20 minutes to 20 seconds through AI-enabled processing.

*DoD Annual Report on Military and Security Developments Involving the PRC.* PLA pursuing 'algorithmic' and 'network-centric' warfare capabilities at varying autonomy levels by 2030.