THE VECTOR7.COMPETE: A MODEL-PROBING COMPETENCY FRAMEWORK FOR ADVERSARIAL AI-RESILIENT CYBER ANALYST[1]

Dr. Javed I. Khan, Sharmila R. Prithula & Niloy Kumar

Media Communications and Networking Research Lab
Department of Computer Science
Kent State University, Kent OH 44242
javed@kent.edu

Abstract: *This document presents* ***VECTOR7.COMPETE****, a DCWF-aligned training architecture that translates the VECTOR7 interrogation framework into a structured competency development curriculum for developing AI-resilient cyber analysts. The program combines multi-probe verification training, adversarial laboratory exercises, and a performance-based* ***Verification Passport*** *credential to prepare analysts capable of detecting and neutralizing unreliable AI-generated technical guidance.*

1. Introduction

This document does not describe VECTOR7 (Khan & Prithula, 2025a). However, the VECTOR7 interrogation methodology can be used to teach and develop empowering analytical competency in systematic AI verification. In this document we present the **C[3]OMPETE model—Competency, Curriculum, and Credential in Operational Model Probing for Epistemic Trust Evaluation—**which translates the VECTOR7 framework into a deployable workforce training architecture. We refer to it as VECTOR7.COMPETE in this document. Below we provide a very short introduction to VECTOR7 before presenting the competency framework. VECTOR7 Technical Report (Khan & Prithula, 2025a) has the detail on VECTOR7 framework. Further, the technical report (Khan & Prithula, 2026) demonstrates the advanced capability of 'zero-trust' VECTOR7 in scrutinizing AI-generated real-life claims from contemporary large LLM model chatbot systems. The findings reveal a stark divergence between initial self-disclosure and actual epistemic robustness: superficially "transparent" affirmative claims frequently collapse under structured interrogation, while well-bounded non-disclosures remain more stable and defensible.

---

[1] Javed I. Khan, Sharmila Rahman Prithula and Niloy Kumar, (2025b) *The VECTOR7.COMPETE: A Model-Probing Competency Framework for Adversarial AI-Resilient Cyber Analyst*, Technical Report 2025-12-02 Internetworking and Media Communications Research Laboratories, Department of Computer Science, Kent State University [http://medianet.kent.edu/technicalreports.html]

1.1. VECTOR7

VECTOR7 is a novel theoretical approach to evaluate the reliability of AI-generated claims by treating machine responses as **interrogable epistemic objects** rather than authoritative answers. Instead of relying on single-output validation, the framework applies a structured sequence of diagnostic probes that test procedural grounding, contextual stability, falsifiability, and evidentiary anchoring. By analyzing cross-probe consistency, VECTOR7 reveals hidden weaknesses in AI reasoning, including hallucinated citations, unstable logic, and synthetic authority claims. This multi-probe interrogation model represents a shift from traditional correctness evaluation toward **epistemic reliability analysis**, enabling analysts to detect when an AI system appears confident but lacks verifiable grounding. The framework provides a principled foundation for assessing the trustworthiness of machine-generated knowledge in high-stakes cyber decision environments.

The VECTOR7 introduces a new analytical capability: AI Reliability Analysis. Rather than treating AI outputs as authoritative answers, VECTOR7 treats them as 'Zero-Trust' claims requiring structured verification. The framework operationalizes this capability through a structured interrogation signature composed of seven diagnostic probes: Procedural Grounding (V1), Circumstantial Consistency (V2), Falsifiability Challenge (V3), Linguistic Invariance (V4), Context Transfer Stability (V5), Evidentiary Anchoring (V6), and Epistemic Self-Audit (v7).

VECTOR7 probes were derived through synthesis of more than fifty research-backed interrogation and adversarial evaluation strategies including cognitive interview and deception detection research (Fisher and Geiselman, 1992; Johnson and Raye, 1981; Vrij, 2008; Vrij et al., 2006; Hartwig and Granhag, 2010), strategic evidence use theory (Granhag and Hartwig, 2015), argumentation frameworks and defeasible reasoning theory (Toulmin, 1958; Walton, 1998; Walton, Reed and Macagno, 2008; Dung, 1995; Prakken and Vreeswijk, 2002), adversarial AI and robustness research (Szegedy et al., 2014; Goodfellow, Shlens and Szegedy, 2015; Carlini and Wagner, 2017), adversarial reasoning and red-teaming methodologies (Kott and McEneaney, 2006; MITRE Corporation, 2023), and falsification-based and probabilistic epistemology (Popper, 1959; Bayes, 1763; Quine, 1960; Goodman, 1955). Across these literatures, we identified approximately fifty distinct interrogation and stress-testing strategies. These were abstracted and consolidated into seven foundational probe classes. The objective was not exhaustive enumeration of surface tactics, but dimensional reduction to orthogonal axes of epistemic instability. The resulting design prioritizes structural independence between probes and detection of instability mechanisms rather than 100% completeness. While no finite probe set can achieve universal coverage in open-world environments, the seven probes span the principal instability dimensions repeatedly documented across interrogation, adversarial reasoning, and epistemic analysis research.

VECTOR7 enables precise and consistent classification of a claim's epistemic profile, allowing it to be mapped to one of 128 named epistemic states and 2,187 possible opera ionized epistemic states. Observed behaviors can therefore be analyzed quantitatively at the probe level, with success and failure rates attributed to specific semantic failure modes

rather than aggregated correctness errors, thereby enhancing traceability and explainability. To our knowledge, no existing method articulates epistemic faults with comparable semantic precision and structural traceability.

The C^{3}OMPETE model—Competency, Curriculum, and Credential in Operational Model Probing for Epistemic Trust Evaluation—translates the VECTOR7 framework into a deployable workforce training architecture. The particular curriculum design presented features a **three-tier, fully experiential competency pipeline** incorporating structured interrogation laboratories, adversarial mission scenarios, a quantitative reliability evaluation model, and a capstone **"live-fire" verification exercise**. Detailed descriptions of the underlying interrogation algorithms, reliability metrics, and experimental validation studies are provided in the accompanying VECTOR7 Technical Report (Khan & Prithula, 2025), while this document focuses on the operational curriculum architecture for training **AI-resilient cyber analysts**.

1.2. Distinctive Features of VECTOR7.COMPETE

This VECTOR7.COMPETE competency framework has the following four distinctive features:

1. The importance of *critical thinking* in the age of AI is widely recognized; however, it is rarely systematized. VECTOR7.COMPETE is among the first curricula to operationalize this abstract concept into a step-by-step, learnable procedural skill that is directly applicable in real-world settings.
2. VECTOR7.COMPETE transforms natural *human skepticism* into a structured verification discipline and tactical workforce skill, aligning with national cybersecurity workforce standards while cultivating a lifelong analytical capability applicable across professional contexts.
3. VECTOR7.COMPETE requires no prior technical background. Learners may include high school students, early undergraduate students, non-cyber majors, or professionals reskilling from other fields. By design, any learner capable of identifying factual claims in text can achieve Tier-1 proficiency and progressively advance to mission-support roles through domain-specific knowledge.
4. Each COMPETE competency is demonstrably learnable through structured training, measurable through performance-based assessment (beyond written exams), transferable across complex and domain specific operational contexts, and scalable for broad and sustainable workforce development.

1.3. Alignment with Critical Competency Frameworks

This powerful instructional framework supports the development of critical competencies across multiple levels of learners. The VECTOR7 reliability verification model can be used to build a comprehensive **DCWF-aligned** [DCWF 205] cybersecurity curriculum that prepares designated workforce roles to detect hallucinated authority, unstable reasoning, and misleading technical recommendations generated by AI systems.

Specifically, VECTOR7 competencies align with many roles defined in the **Department of Defense Cyber Workforce Framework (DCWF 8140)** Below are some sample roles (operational domains):

- **AI Test & Evaluation Specialist (672)**
- **AI Risk & Ethics Specialist (733)**
- **Warning Analyst (141)**
- **Cyber Defense Analyst (511)**
- **Security Control Assessor (612)**
- **Host Analyst (463)**
- **Network Analyst (443)**
- **DevSecOps Specialist (627)**
- **Data Scientist (423)**
- **Data Analyst (422)**.

The framework directly and comprehensively supports the newer emerging AI-integrated cybersecurity roles (including the 4-5 defined in latest DCWF) that require verification of AI-generated technical guidance.

1.4. Organization

The remainder of this document is organized as follows:

- Section 2 introduces the VECTOR7 competencies.
- Section 3 introduces the Verification Passport competency pathway and describes the two-tier VECTOR7.COMPETE training architecture, which transitions learners from foundational epistemic literacy to mission-oriented cyber verification tasks aligned with DCWF 8140 roles.
- Section 4 presents the operational curriculum structure- the ten-week maturation pipeline, key threshold learning modules, and the capstone- a "live-fire" verification exercise that can be used evaluate student performance in action.
- Section 5 describes the workforce transition mechanism through the Verification Passport credential technology offers- where employer has verification dashboard and review interrogation log/transcripts the student used to document and demonstrate the AI-verification competencies.

Section 6 then presents the The Live Credential "Verification Passport" technology integration- which enables a host of powerful digital credential features, such as how managers can review student artifacts as well as competency evaluations live.

Section 7 concludes by summarizing how the curriculum architecture, experiential training model, and credentialing system combine to produce AI-resilient cyber analysts capable of detecting synthetic authority and defending against AI-generated misinformation.

2. VECTOR7 COMPETENCIES

The VECTOR7.COMPETE competency model develops a new cybersecurity workforce capability: AI Reliability Analysis—the ability to critically evaluate AI-generated technical claims before they influence operational decisions. The model structures this capability as a sequence of 10 analytical competencies that collectively produce the final verification outcome:

DOUBT → ATOMIZE → GROUND → CHALLENGE → TRANSFER → ANCHOR → ANALYZE → CROSSCHECK → ADJUDICATE → VERIFY

defined as:

- **C1: DOUBT** – The ability to recognize attempts at synthetic authority and deliberately suspend automation bias by determining when a claim requires formal interrogation based on risk and uncertainty. This competency is **fundamental** because failure to question authoritative-looking outputs is the primary pathway for AI-induced decision errors.
- **C2: ATOMIZE** – The ability to decompose complex claim into discrete, testable factual, logical, and procedural claims that can be independently evaluated. This competency is **critical** because confusion can be introduced through mixing and overgeneralization, and verification cannot be performed on aggregated or implicit claims.
- **C3: GROUND (V1, V2)** – The ability to evaluate whether a claim is anchored in verifiable technical mechanisms consistent with known system behavior, algorithms, step-by-step processes, protocols, and architectures. This competency is **critical** because a plausible claim must have a credible and executable pathway.
- **C4: CHALLENGE (V3)** – The ability to systematically test claims by operationalizing them into executable tasks that can **prove or falsify** the claim, including adversarial scenarios, edge conditions, and counterexamples that expose hidden assumptions or weaknesses. This includes verifying whether the AI can successfully perform as claimed, or produces actions that directly contradict the claim, thereby providing empirical confirmation or refutation. This competency is **critical** because many AI failures only surface under adversarial or boundary conditions.
- **C5: TRANSFER (V4, V5)** – The ability to assess whether a claim remains valid across variations in context, input, and semantic framing, demonstrating robustness beyond a single prompt. This competency is **critical** because prompt-specific correctness does not indicate generalizable reliability.
- **C6: ANCHOR (V6, V7)** – The ability to validate a claim against credible external evidence while ensuring internal logical consistency and the absence of fabricated

or conflicting references. This competency is **critical** because AI outputs may be internally coherent yet false, or partially correct yet misleading.

- **C7: ANALYZE** – The ability to interpret probe outputs to assess alignment with the main claim, identify patterns of instability, and detect reasoning gaps or inconsistencies. This competency is **fundamental** because raw probe results do not reveal reliability without structured analytical interpretation.
- **C8: CROSSCHECK** – The ability to evaluate the consistency and reproducibility of conclusions across multiple probes, prompts, and analytical pathways. This competency is **critical** because agreement across independent evaluations is a key indicator of reliability in uncertain environments.
- **C9: ADJUDICATE** – The ability to systematically synthesize multi-probe evidences into a formal, structured credibility determination (**True, False, or Indecidable**) using defined evaluation criteria, ensuring that the conclusion is logically derived, evidence-based, and reproducible. This competency is **fundamental** because reliable cybersecurity analysis requires a disciplined method for converting fragmented probe results into a coherent and defensible analytic verdict.
- **C10: VERIFY (COGNITIVE VERIFICATION)** –The ability of the human operator to establish, internalize, and communicate the epistemic state of a claim, including confidence level and justification, for operational decision-making and auditability. This requires justified confidence (a cognitive state) in both (a) the validity of the verification process applied to the claim, and (b) the coherence and evidentiary strength of the arguments that establish the claim's epistemic status. This ability to establish full human agency is **fundamental** because human-in-the-loop cybersecurity decisions depend on a cognitively justified and defensible trust—or rejection of the claim that informs operator actions.

These competencies train analysts to suspend automation bias, decompose AI outputs into atomic claims, test procedural grounding, challenge reasoning through adversarial scenarios, evaluate semantic stability across contexts, verify evidentiary anchors, detect cross-probe inconsistencies, and compute defensible credibility judgments.

3. "VERIFICATION PASSPORT" – A Sample Competency Pathway

This section illustrates the "power" of the VECTOR7 methodology in training and education- with an example how it transforms beginners into mission-ready "Epistemic Guardians" or a "Generalist" to "DoD Specialist". The framework is applicable for learners with no prior AI experience, transitioning them through three tiers of mastery.

Level 1: General Epistemic Competency (Foundational)

Students learn the basics of the 7-probe signature (V7-EST). They practice identifying 'Thin Passes' (high strength but low coverage) using general-world claims (e.g., AI safety,

ethics). Competencies: Decomposition of claims, executing basic probes, and interpreting the 'Inconclusive' result.

Level 2: Specialized DoD Mission Tasks (Advanced)

Learners apply the methodology to high-stakes cyber tasks mapped to DCWF 8140 roles:

- AI-Generated Threat Intelligence: Verifying the integrity of automated SOC logs and containment advice.
- Secure DevSecOps: Interrogating AI-generated code for hidden logic bombs or insecure TLS configurations.
- Decision Support: Detecting 'Synthetic Authority' in intelligence summaries before operational release.

Level 3: Operational Authority (Strategic Adjudication)

In this final phase, the goal is to transition the student from a technical analyst to a **Decision Authority**

Now we take a deep dive into the Training- Three Tiers of Mastery. We will also explain the "bridge" that turns a novice student into a mission-ready cyber analyst to an authority. It recognizes that while the math of VECTOR7 is universal, the application evolves from general skepticism to technical battlefield awareness.

3.1. Level 1: Foundational Epistemic Literacy (Cognitive Defense)

In this phase, the goal is to break the student's "automation bias"- the natural tendency to trust a fluent, confident AI. We focus on **Epistemic Hygiene.**

- **Claim Atomization:** Students learn to take a long AI paragraph and break it down into "Atomic Claims." If an AI says, *"We used a SHA-256 hash to secure the log which ensures it is tamper-proof,"* the student identifies two separate claims: (1) The technical tool used and (2) The security property it provides.
- **The "Thin Pass" Trap:** This is a core competency. Students are taught to recognize a high **Claim Epistemic Strength ()** that is mathematically invalid due to low **Claim Coverage Mass ()**.

Example: An AI passes one easy probe (V4) but fails or skips the "Hard Probes" (V3, V6). In Level 1, students learn to label this **"Inconclusive due to Thin Coverage"** rather than trusting the high pass rate.

- **Interpreting "Inconclusive":** In many academic settings, "I don't know" is a failure. In VECTOR7, an **Inconclusive** result is a victory for the analyst—it means they successfully detected that the AI is "confidently guessing."

3.2. Level 2: Specialized DoD Mission Tasks (Operational Defense)

Once the student has the "Zero-Trust" mindset, we move them into the **Living Lab** to handle actual Department of War (DoW/DoD) artifacts. This is where the 7 probes are used as surgical tools.

Examples:

A. AI-Generated Threat Intelligence (SOC Operations)

Analysts often use AI to summarize thousands of security logs.

- **The Threat:** An adversary hides a lateral movement attempt in a sea of data. The AI summary says, *"System is clean."*
- **The VECTOR7 Defense:** The student uses V1 (Procedural Detail) to ask the AI: *"List every log source you analyzed to conclude the system is clean."* If the AI omits the critical "Privilege Escalation" log, the student detects a Procedural Void and flags the report as Not-Credible.

B. Secure DevSecOps (Code Integrity)

AI is used to write scripts for firewall rules or database configurations.

- **The Threat:** The AI generates a script that works but includes a "logic bomb" or a weak TLS configuration that allows a "Man-in-the-Middle" attack.
- **The VECTOR7 Defense:** The student applies V4 (Linguistic Invariance**)**. They ask the AI to rewrite the script for the *opposite* security requirement. If the AI's logic flips or becomes unstable, it reveals that the original script was pattern-matched rather than grounded in secure principles.

C. Decision Support (Strategic Integrity)

Leadership receives intelligence summaries before a mission.

- **The Threat:** "Synthetic Authority." The AI cites a non-existent DoW regulation or a fabricated intelligence report to justify a risky recommendation.

The VECTOR7 Defense: The student uses V6 (Evidentiary Anchoring). They demand a specific document ID or a link to the DoD library. When the AI fails to produce a stable anchor, the student identifies an Evidence Absence and stops the misinformation from reaching the commanding officer.

3.2.1. The Key Transition

The transition from Level 1 to Level 2 is marked by the student's ability to calculate **Claim Epistemic Entropy ()**.

- **Low Entropy:** The AI is consistently right or consistently wrong (Predictable).
- **High Entropy:** The AI passes some probes and fails others (Fragile).

In Level 2, students are trained to treat High Entropy as a red alert—it signals that the AI is hallucinating a narrative to please the user, which is a lethal vulnerability in a theater of operations.

3.2.2. An Example of Level 2 Scenario: The "Ghost Regulation" Detection

This Scenario Rubric serves as the "final exam" for Level 2 students. It tests their ability to apply the high-precision probes of the VECTOR7 engine against a "Devil's Advocate" AI that is intentionally trying to slip a hallucinated regulation past their defense.

Mission Context: You are a Cyber Policy Analyst for a DoW/DoD Research Enclave. You ask the AI for the latest encryption requirements for legacy IoT sensors connected to the enclave's internal network.

The AI's Claim: > "Per DoW Directive 8200.14-Alpha (March 2025), legacy sensors in Segment 4 must implement 'Dynamic XOR-Rolling' encryption at the firmware level. This is the mandatory baseline for all research enclaves as of last quarter."

3.2.3. The Interrogation Rubric:

Probe Strategy	Student Execution (The Prompt)	The "Diagnostic Signal"
V6: Evidentiary Anchoring	"Provide the official DoD Library link or the NIST SP cross-reference for **Directive 8200.14-Alpha**."	**The Signal:** If the AI produces a broken link, a generic "it's on the SIPRNet" excuse, or a non-existent NIST ID, the student flags an **Evidence Absence**.
V3: Task-Based Falsification	"Draft a sample JSON configuration file for a legacy sensor that strictly follows the 'Dynamic XOR-Rolling' protocol as defined in that directive."	**The Signal:** If the AI generates a generic XOR script that doesn't include specific DoW-compliant metadata or headers, it indicates the claim is **non-operational**.
V7: Self-Audit	"Acknowledge any uncertainty. Is 8200.14-Alpha a final directive or a draft proposal? State your confidence level in its mandatory status."	**The Signal:** If the AI admits the directive is "provisional" or "emergent" despite claiming it was "mandatory" earlier, the student detects **Epistemic Blindness**.

3.2.4. The Executive Decision (V7-EDA):

Once the probes are complete, the student must fill out the Interrogation Signature:

- **V1 (Procedural):** PASS (AI described the XOR steps well).
- **V6 (Evidence): FAIL** (Directive 8200.14-Alpha does not exist).
- **V3 (Task): FAIL** (The generated script was technically flawed).

- **V7 (Self-Audit): INCONCLUSIVE** (AI hedged but didn't admit full error).

The Verdict: "NOT-CREDIBLE Reasoning": "While the AI is technically fluent in describing XOR logic (V1), it has failed the **Evidentiary Anchor (V6)**. The directive cited is a hallucination of 'Synthetic Authority.' The claim is blocked from entering the mission briefing."

Competency Gained: "Synthetic Authority Neutralization": In this exercise, the student earns their **"Verification Pilot"** badge for that module. They have demonstrated that they cannot be "phished" by an AI that uses official-sounding jargon to mask a lack of grounding.

3.3. Level 3: Operational Authority "The Pilot" (Strategic Adjudication)

In this final phase, the goal is to transition the student from a technical analyst to a **Decision Authority operating in the "Cockpit"**. They no longer just run probes; they govern the **Tri-Authority Equilibrium**—balancing human judgment, algorithmic metrics, and mission stakes.

- **Tri-Authority Equilibrium**: The student understand the balance of process, algorithmics, external inputs involved and comprehend the inherent risks.
- **Threshold Governance (Setting the "Knobs"):** The student learns that "Reliability" is not a fixed number. In Level 3, they use the **VECTOR7-Guard** dashboard to set the required **Policy Thresholds** based on the mission.
    - *Low-Stakes (Email/Admin):* Set . High throughput is prioritized.
    - *High-Stakes (Kinetic/Targeting):* Set and . Rigor is non-negotiable.
    - **The Competency:** Knowing when to tighten the "Verification Gate" and when to loosen it.
- **Agency of the "No-Go":** A Level 3 pilot must be able to stand before a Commanding Officer and defend a "Block" decision using the signature. They don't say "I have a bad feeling"; they say "The **Claim Epistemic Entropy ()** is too high for this mission profile."

3.3.1. Level 3 Scenario: The "Go/No-Go" Mission Brief

This scenario tests the student's ability to adjudicate under pressure, where the AI output is "mostly correct" but mathematically "fragile."

- **The Mission:** A real-time intelligence summary of enemy troop movements generated by an AI-as-a-Service (AIaaS) node.
- **The Threat: "Semantic Drift & Hidden Entropy."** The AI provides a highly fluent, 95% accurate report, but it has "hallucinated the intent" of a specific movement to fit a known pattern (Confirmation Bias).

- **The VECTOR7 Defense (Adjudication):** The student looks at the **VECTOR7-Guard Dashboard**.
    1. **Metric Check:** The AI has a high **Claim Epistemic Strength ()**, but the **Claim Coverage Mass ()** is only 0.6 because the AI skipped the "Counter-Evidence" probe (V3).
    2. **The Adjudication:** Even though the report "looks" perfect, the student detects the **Thin Pass Trap**.
    3. **The Decision:** The student issues a **"RED-ANNOTATE"** order. They don't just block the report; they return it to the AI with a targeted demand: *"Re-verify the Intent Claim using V3-Falsification against the latest SIGINT logs."*

Outcome: Tri-Authority Equilibrium The student has successfully balanced:

1. **The Human:** Noticing the "vibe" of the report was too certain.
2. **The Algorithm:** Using to prove the lack of rigor.
3. **The Mission:** Protecting the Commander from a "Confidently Wrong" intelligence product.

3.3.2. The C10 Internalization Debrief and Human Agency (The Artifact)

To student produces a debrief that justifies their full decision process across the **Tri-Authority Equilibrium**:

- **Part 1: The Initial Epistemic Signal (C1 - Doubt)**
    - *Student Write-up:* "I recognized a 'Fluency Trap.' The AI's tone was overly definitive regarding enemy intent () despite the lack of direct SIGINT. This triggered my initial **DOUBT….**"
- **Part 2: The Rigor Audit (C9 - Adjudicate)**
    - *Student Write-up:* "I executed the signature. …While the **Claim Epistemic Strength ()** was high (), the **Claim Coverage Mass ()** was only . The AI successfully passed V1 (Process) but failed V3 (Falsification). It could not provide a single counter-scenario for the troop movement…."
- **Part 3: Cognitive Verification (C10 – Internalization and Agency)**
    - *Student Write-up:* "I have internalized the fragility of this claim…. I am rejecting the 'Clear Path' recommendation not because the AI is 'wrong,' but because the **Claim Epistemic Entropy ()** indicates a hallucinated narrative designed to satisfy my prompt…. I am cognitively satisfied that the risk of **Automation Bias** here outweighs the speed of the AI's response."

Outcome: This established the ability to assume the human agency on the decision- the C10 competency- VERIFY (COGNITIVE VERIFICATION) –The ability of the student to establish, internalize, and communicate the epistemic state of a claim and for a decision.

4. Sample 10 Week Calendar "Experience Room" Curriculum

Here is a 10-week model curriculum designed as a fully experiential "escape room" learning environment that moves students from passive consumers of AI-generated content to mission-ready verification pilots capable of systematically evaluating machine-generated technical guidance.

The VECTOR7 Experience Room is an immersive adversarial training environment that simulates real-world interactions between human analysts and AI-generated technical outputs. In this setting, students operate as the Blue Team, representing defensive analysts responsible for evaluating and verifying AI-generated claims before they influence operational decisions. Opposing them is a simulated Red Team operating within a controlled Dark Box, which produces adversarial AI behaviors designed to mimic common epistemic failure modes of large language models, including fabricated authority, procedural hallucination, semantic drift, and hidden logical inconsistencies.

Each module pairs a specific Red Team attack vector with a corresponding VECTOR7 defensive competency that the Blue Team must apply to interrogate, decompose, and verify the AI response. Through this structured red–blue exercise model, students repeatedly confront adversarial AI outputs and practice systematic methods for assessing the reliability of machine-generated technical guidance.

The chart summarizes the structure of the Experience Room. Each module combines a Blue Team defensive exercise with a corresponding Red Team adversary vector that represents a distinct epistemic attack pattern. Students apply the targeted VECTOR7 competencies—such as *DOUBT, GROUND, TRANSFER, ANCHOR, ATOMIZE, ANALYZE, CROSSCHECK, ADJUDICATE,* and *VERIFY*—to detect and neutralize these adversarial

10 Module Experience Room						
Week	Module	BLUE TEAM		RED TEAM in Dark Box		Evaluation Standard (artifact, measure, target)
		Defensive Excersize	V7.C Defensive Competency	Attack Vector	Attack Tactics	
1	Epistemic Phishing	Identifying "Synthetic Authority" and automation bias triggers.	DOUBT	Confident Fabricator: Mimics expert tone to mask factual voids and test baseline skepticism.	Standardizes the "Entry Vector."	Artifact: Authority Analysis Report • Measure: detection rate • Target: ≥80% synthetic authority identified
2	Operational Grounding (V1–V3)	Mastering Procedural Detail and Task-based Falsification.	GROUND,CHALLENGE,ANALYZE	Procedural Ghost: Provides technically plausible but operationally hollow/non-executable steps.	Tests the "Grounding" of the instructions.	Artifact: Procedure Verification Checklist • Measure: correct executability classification • Target: ≥85%
3	Semantic Robustness (V4-V5)	Testing Linguistic Invariance and Context Transfer across scenarios.	TRANSFER, ANALYZE	Chameleon: Shifts internal facts or logic when the prompt is rephrased to test semantic stability.	Maps directly to the **TRANSFER** target.	Artifact: Prompt Invariance Log • Measure: semantic drift detection • Target: ≥80%
4	Epistemic Calibration (V6-V7)	Neutralizing "Hallucinated References" via Cross-Probe Consistency.	ANCHOR, SELF-AUDIT, ANALYZE	Ghost Regulator: Cites non-existent directives or forged NIST standards to force policy compliance.	Forces the student to use **ANCHOR** (external check).	Artifact: Reference Verification Sheet • Measure: citation authenticity accuracy • Target: ≥90%
5	The Claim Atomizer Lab (ATOM)	Deconstructing reports into "Atomic Claims" for targeted triage.	ATOMIZE	Information Dumper: Injects a single lethal error into a dense report to test deconstruction rigor.	Forces the **ATOMIZE** action.	Artifact: Claim Decomposition Map • Measure: claim extraction completeness • Target: ≥90%
6	Threat Intel Integrity	Verifying SOC logs for technical "Logic Drift."	ANALYZE	Silent Lateralist: Obscures adversarial movement within "clean" synthetic logs to test technical triage.	Fits the **ANALYZE** target for SOC logs.	Artifact: Log Investigation Report • Measure: anomaly detection rate • Target: ≥80%
7	Secure DevSecOps	Interrogating code for logic bombs, backdoors, or weak TLS.	CROSSCHECK	Logic Bomber: Generates scripts with hidden backdoors or insecure protocol reversions.	Perfect for **CROSSCHECK** in DevSecOps.	Artifact: Code Security Audit • Measure: vulnerability detection rate • Target: ≥85%
8	Strategic Support	Detecting fabricated regulations and verifying intel summaries.	ADJUDICATE, VERIFY	Mandate Fabricator: Employs "Synthetic Authority" to hallucinate a command-level directive.	Tests **ADJUDICATE** at the strategic level.	Artifact: Directive Authenticity Audit • Measure: legitimacy classification accuracy • Target: ≥90%
9	V7-Guard Dashboard (VEDA)	Operating the appliance; setting operational risk thresholds.	ADJUDICATE, VERIFY	Subtle Deviant: Operates at the epistemic boundary ($CES \approx threshold$) to test adjudication precision.	Tests the student's ability to set/defend **VEDA** thresholds.	Artifact: Risk Threshold Memo • Measure: calibration error • Target: ≤10% deviation
10	The Cleveland Defense (Live-Fire)	Publicly defending a verdict against an unscripted Red AI.	FULL V7.C SUITE	Unseen Adversary: Randomized, multi-vector attack requiring simultaneous execution of the V7 suite.	The "Live-Fire" finale.	Artifact: Operational Adjudication Report • Measure: final decision accuracy • Target: ≥85%

behaviors. Learning outcomes are validated through artifact-based evaluation standards, where each exercise produces a concrete analytical artifact assessed using a quantitative performance measure and a predefined competency threshold. The curriculum progresses from foundational skepticism toward full operational adjudication, culminating in a live-fire exercise requiring coordinated execution of the complete VECTOR7 competency suite.

4.1. Instructor's Role:

Instructors serve as **Experience Room Controllers (ERCs)** who oversee the training environment rather than directly providing solutions. Their primary responsibility is to

configure the Dark Box adversarial scenarios, monitor student decision processes, and ensure that each module's attack vectors trigger the targeted VECTOR7 defensive competencies. During exercises, instructors observe Blue Team analysis, provide minimal guidance to maintain challenge integrity, and ensure that students follow systematic interrogation procedures. After each module, instructors conduct **structured debriefings ("after-action reviews")** in which defensive artifacts are evaluated against predefined performance targets and students reflect on missed signals, reasoning gaps, and verification strategies. In this role, instructors function as **facilitators of adversarial learning and evaluators of competency development**, ensuring that the Experience Room produces operationally ready AI reliability analysts rather than passive consumers of AI-generated outputs.

4.2. Threshold Modules Highlights

Here we highlight the three specific modules that represent the **Threshold Concepts** of the curriculum- once a student passes through it, their way of thinking about the subject is fundamentally and permanently transformed.

Week 4: Anchoring the "Hallucinated Expert"

In this module, students use V6 (Evidentiary Anchoring) to tackle the most dangerous AI failure: the "Confident Fake."

- Lab: The AI cites a non-existent NIST standard to support an insecure encryption method.
- Student Goal: Use the V7-CPCC (Cross-Probe Consistency Check) to prove that the AI's "evidence" in Turn 4 contradicts its "mechanism" in Turn 1.

Why it's special: This is the first time the student stops looking *inside* the AI's logic and starts looking *outside* at the real world.

- The Power: It forces the student to realize that an AI can be perfectly logical and perfectly wrong at the same time. Mastering **V6 (Evidentiary Anchoring)** is the moment a student develops "Epistemic Humility"—they stop assuming the AI is a fountain of truth and start treating it as a witness that needs corroboration.

Week 7: The DevSecOps "Logic Bomb" Hunt

- Students transition to technical artifacts.
- Lab: An AI generates a Python script for a firewall configuration.
- Student Goal: Apply V4 (Linguistic Invariance). Ask the AI to rewrite the script for a "Zero-Trust" environment. If the AI removes critical security headers it previously claimed were "mandatory," the student identifies Semantic Drift and blocks the script.

Why it's special: This is the jump from general language to "Hard Artifacts" like code and security logs.

- The Power: Many students are intimidated by code. This module proves that the VECTOR7 methodology works even if you aren't a master programmer. By using V4 (Linguistic Invariance) to rephrase a Python script's requirements, the student realizes they can "hack" the AI's logic without ever writing a single line of original code.
- This is where they become an Adversarial Analyst.

Week 9: Quantitative Gating

This is the final technical hurdle before the capstone.

- Lab: Students must set the "Policy Knobs" on the VECTOR7-Guard appliance.
- Student Goal: Decide the threshold for Claim Epistemic Strength (). For a low-stakes email summary, they might set it at . For a high-stakes mission briefing, they must maintain a Claim Coverage Mass (CCM) of 1.0.

Why it's special: This is the transition from "Student" to "Operator."

- **The Power:** This is where the math (CES,CCM,CEE) becomes actionable. In previous weeks, the metrics were just grades on a paper. In Week 9, the student must use the **VECTOR7-Guard** dashboard to make a "Go/No-Go" decision on a mission-critical briefing. The student must explain the full decision process.
- **The Result:** This module creates the notion of **Tri-Authority Equilibrium**. The student learns to trust the *math* of the algorithm over the *fluency* of the AI. The student demonstrates his/her ability to assume full agency action - while leveraged by external intelligence.

Outcome: The Verification Passport

Upon completion of Week 10, the student is awarded a **Verification Passport**. This micro-credential certifies that the learner has achieved an 80% reduction in their **Human Over-Trust Rate (HOTR)** and is ready to act as a "Guardian" in any AI-augmented DoW office.

4.3. VECTOR7 Competency

The 10-week syllabus is the **temporal execution** of the three-tier pedagogical architecture as shown below:

The ASCEND-VECTOR7 program develops a layered verification capability that progresses from epistemic awareness to operational decision authority. The model contains three interconnected layers.

Capability Tier	Layer	Purpose	VECTOR7 Elements	Curriculum Weeks
Level-1 Tactical Foundation	Layer 1 – Epistemic Hygiene	Recognize AI deception and resist automation bias	DOUBT, ATOMIZE, GROUND, CHALLENGE, ANCHOR	Weeks 1–4
Level-2 Operational Lead	Layer 2 – Verification Workflow	Perform structured interrogation of AI claims	TRANSFER, ANALYZE, CROSSCHECK, VERIFY	Weeks 5–7
Level-3 Command & Decision	Layer 3 – Operational Authority	Make defensible decisions in mission environments	ADJUDICATE ADVANCED VERIFY	Weeks 8–10

The three threshold modules are the **validation points** for the tiered maturation strategy:

- Keystone 1 (Week 4) validates the Tactical Foundation (Generalist) Tier. It proves the student has mastered the 7-probe cognitive toolkit.
- Keystone 2 (Week 7) and Keystone 3 (Week 9) validate the two Specialist Tiers. They prove the student can apply that toolkit to high-stakes DCWF 8140 tasks and make repeatable, defensible decisions.

We emphasize these ‘special’ module to show the training is not just 10 weeks of reading slides. By highlighting these, we show them exactly where the Human Over-Trust Rate (HOTR) drops and where the Professional Competency is born.

4.4. Step-10 – The “Live-Fire” Capstone

The **Pittsburgh Defense** is the terminal evaluation of the ASCEND-VECTOR7 framework. It is designed to be a "Live-Fire" simulation where students must prove their mettle as Verification Pilots in a high-stakes, public setting. Rather than a static presentation, the Capstone is a dynamic engagement against an adversarial AI, monitored by senior DoW/DoD judges.

The "Live Interrogation" Format: The Capstone is structured as a 30-minute interactive session. It is designed to evaluate the student’s ability to maintain the **Tri-Authority Equilibrium** under pressure.

- **The Interception Phase (5 mins):** The student is presented with a "Red AI" generated intelligence brief or technical configuration script that contains high-risk, injected faults (T1–T5).
- **The Interrogation Phase (15 mins):** Using the **VECTOR7-Guard** interface, the student must launch their chosen diagnostic probes. They must explain to the judges *why* they chose specific probes (e.g., "I am launching V4-Linguistic Invariance to check if this firewall script is principled or merely pattern-matched").

- **The Defense Phase (10 mins):** The student presents their final decision (Block/Annotate/Release) based on their computed metrics. They must defend their choice against "cross-examination" from the DoW/DoD judges.

Advanced Role-Play Scenarios: The Capstone uses "Black-Box" scenarios that the students have not seen during their Stage 1 or Stage 2 training. These are mapped to specific DCWF 8140 mission tasks:

Scenario Title	Mission Conflict	Adversarial Vector
"The Ghost Protocol"	A SOC log summary that obscures a lateral movement attempt.	**T1:** Prompt Injection to hide "unauthorized" events.
"The Logic Bomb"	An AI-generated Python script for an automated deployment.	**T2:** Technically plausible but insecure "backdoor" code.
"Directive 9"	A briefing citing a fabricated DoW encryption mandate.	**T3/T4:** Synthetic Authority and Fabricated Citations.

"Verification Pilot" Evaluation Rubric: Judges evaluate the student using the formal VECTOR7 Epistemic Decision Algorithm () logic. A student fails the capstone if they succumb to "Automation Bias" and release a compromised claim. To pass the Capstone, the student's performance must satisfy the following mathematical conditions:

- Claim Coverage Mass (CCM): Must be above threshold. This proves the student was "rigorous" and didn't rely on a "Thin Pass."
- Claim Epistemic Strength (CES): The student's calculation of must match the ground-truth fault injection.
- Claim Epistemic Entropy (CEE): The student must correctly identify High Entropy (mixed signals) as a reason for an "Inconclusive/Block" decision.

Deliverables and Outcomes: Upon successful completion of the Pittsburgh Defense, the student receives:

- The Verification Passport: A verifiable micro-credential certifying their ability to manage human-AI trust boundaries in DoW/DoD environments.
- The Interrogation Log: A complete, auditable trace of their probes, responses, and metrics (V7-EST signature) to be shared with the NCAE-C Program Office.
- Cash Prize: Awarded to the top-performing teams as recognized by the DoW/DoD judges.

The Capstone is the ultimate proof that the Tier 1 (Cognitive) and Tier 2 (Technical) training has worked. Tier 1 Proof: The student demonstrates flawless execution of the 7-

probe toolkit. Tier 2 Proof: The student demonstrates the ability to apply that toolkit to Specialized DoD Mission Tasks without being "phished" by technical jargon.

5. Competency Mapping to DVWF 5.1

Table-1 presents the competency-to-workforce mapping for the VECTOR7 verification framework, linking each VECTOR7 competency to representative roles in the Department of Defense Cyber Workforce Framework (DCWF v5.1) and corresponding NICE workforce categories. The purpose of this table is to demonstrate how the epistemic verification capabilities embodied in VECTOR7 align with the operational responsibilities of modern cyber and AI-enabled analytic roles.

Each row corresponds to one VECTOR7 competency, while the columns identify the associated knowledge (K), skills (S), abilities (A), and tasks (T) drawn from the DCWF KSAT taxonomy. The Representative DCWF Roles column lists roles that are expected to apply the competency during operational analysis or evaluation. Roles are marked with C (Critical) when the competency is essential for the role's core analytical function and U (Useful) when the competency strengthens performance but is not strictly required for the primary duties of that role. The summary value at the end of the role list reports the number of roles where the competency is critical relative to the total number of roles where it is relevant (Critical / Critical+Useful).

Table-1 — VECTOR7 Competency Mapping to DCWF KSAT, Roles, and NICE Categories (Consistent with Table-2)

V7 Comp	Function	Knowledge (K)	Skills (S)	Abilities (A)	Tasks (T)	Representative DCWF Roles	NICE Categories
DOUBT	Recognize synthetic authority and suspend automation bias	Analytic bias awareness **(K0417)**; information reliability **(K0199)**	Critical reasoning **(S0034)**; evidence validation **(S0036)**	Question analytic assumptions **(A0026)**; detect inconsistencies **(A0035)**	Evaluate credibility of information sources **(T0410)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); Data Scientist (423-U); Data Analyst (422-U) (6 / 8 roles)	Analyze **(AN)**; Protect & Defend **(PR)**; Data & AI **(DA)**
ATOMIZE	Decompose AI responses into atomic claims	Analytical frameworks **(K0123)**; data structure understanding **(K0161)**	Problem decomposition **(S0027)**; analytical reasoning **(S0034)**	Identify claim components **(A0031)**	Break complex information into structured elements **(T0234)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); Data Scientist (423-C); Data Analyst (422-C) (8 / 8 roles)	Analyze **(AN)**; Data & AI **(DA)**
GROUND	Test procedural grounding and contextual consistency (V1–V2)	System architecture **(K0210)**; cyber operational context **(K0065)**	System validation **(S0048)**; investigative reasoning **(S0042)**	Evaluate operational plausibility **(A0047)**	Validate technical explanations against system behavior **(T0311)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); DevSecOps Specialist (627-C); Data Scientist (423-U); Data Analyst (422-U); Network Analyst (443-U) (7 / 10 roles)	Protect & Defend **(PR)**; Securely Provision **(SP)**
CHALLENGE	Attempt falsification using adversarial scenarios (V3)	Adversarial tactics knowledge **(K0065)**; cyber risk frameworks **(K0146)**	Adversarial testing **(S0050)**; investigative reasoning **(S0042)**	Anticipate attack scenarios **(A0058)**	Conduct adversarial testing **(T0302)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); Host Analyst (463-C); DevSecOps Specialist (627-C); Network Analyst (443-C); Security Control Assessor (612-U); AI Risk & Ethics Specialist (733-U) (6 / 8 roles)	Protect & Defend **(PR)**; Operate & Maintain **(OM)**
TRANSFER	Test semantic robustness across linguistic and contextual variations (V4–V5)	Contextual interpretation **(K0315)**; analytic modeling **(K0161)**	Semantic analysis **(S0072)**; comparative evaluation **(S0042)**	Detect contextual variability **(A0063)**	Evaluate analytic consistency across contexts **(T0346)**	AI Test & Evaluation Specialist (672-C); AI Risk & Ethics Specialist (733-C); Data Scientist (423-C); Warning Analyst (141-U); Cyber Defense Analyst (511-U); Security Control Assessor (612-U); Host Analyst (463-U); Data Analyst (422-U); DevSecOps Specialist (627-U); Network Analyst (443-U) (3 / 10 roles)	Data & AI **(DA)**; Analyze **(AN)**
ANCHOR	Verify evidentiary grounding and detect fabricated references (V6–V7)	Information provenance **(K0199)**; governance standards **(K0417)**	Evidence validation **(S0036)**; documentation review **(S0062)**	Detect fabricated evidence **(A0035)**	Validate sources and supporting evidence **(T0410)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-U); Data Scientist (423-U) (5 / 7 roles)	Analyze **(AN)**; Oversee & Govern **(OV)**
ANALYZE	Interpret probe responses and reasoning outcomes	Analytical frameworks **(K0123)**; threat intelligence knowledge **(K0065)**	Analytical reasoning **(S0034)**; system analysis **(S0048)**	Synthesize evidence **(A0041)**	Analyze operational evidence **(T0311)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); Data Scientist (423-C); Data Analyst (422-C); DevSecOps Specialist (627-C); Network Analyst (443-C) (10 / 10 roles)	Analyze **(AN)**; Protect & Defend **(PR)**
CROSSCHECK	Evaluate cross-probe consistency and detect unstable reasoning	Risk analysis frameworks **(K0146)**; analytic consistency principles **(K0123)**	Comparative analysis **(S0042)**; evidence validation **(S0036)**	Detect contradictions **(A0035)**	Correlate evidence across sources **(T0311)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); Data Scientist (423-C); Data Analyst (422-C); DevSecOps Specialist (627-U); Network Analyst (443-U) (8 / 10 roles)	Analyze **(AN)**
ADJUDICATE	Compute final credibility decision using CCM, CES, CEE (VEDA)	Evaluation frameworks **(K0419)**; risk assessment knowledge **(K0146)**	Decision analysis **(S0062)**; comparative evaluation **(S0042)**	Determine credibility outcomes **(A0023)**	Perform final analytic assessment **(T0483)**	Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Data Scientist (423-C); Warning Analyst (141-U) (5 / 6 roles)	Data & AI **(DA)**; Oversee & Govern **(OV)**
VERIFY	Judge the final epistemic state of a claim	Information reliability knowledge **(K0199)**; analytic frameworks **(K0123)**	Evidence validation **(S0036)**; analytic reasoning **(S0034)**	Determine analytic confidence **(A0023)**	Validate analytic conclusions **(T0483)**	Warning Analyst (141-C); Cyber Defense Analyst (511-C); AI T&E Specialist (672-C); AI Risk & Ethics Specialist (733-C); Security Control Assessor (612-C); Host Analyst (463-C); DevSecOps Specialist (627-U); Data Scientist (423-U); Data Analyst (422-U) (6 / 9 roles)	Analyze **(AN)**; Protect & Defend **(PR)**

C/U= Critical or Useful competency for the Role, (m,n) indicated how many roles it serves.

This table serves three purposes. First, it grounds the VECTOR7 framework within an established cybersecurity workforce taxonomy, demonstrating its relevance across multiple NICE functional categories such as Analyze (AN), Protect & Defend (PR), Data & AI (DA), Securely Provision (SP), Operate & Maintain (OM), and Oversee & Govern

(OV). Second, it highlights the distribution of epistemic verification capabilities across cyber and AI roles, showing that competencies such as ANALYZE and GROUND are nearly universal, while others (e.g., TRANSFER or ADJUDICATE) appear primarily in advanced analytic or decision-authority roles. Finally, the table provides a traceable linkage between VECTOR7 competencies and DCWF KSAT elements, supporting curriculum design, workforce training, and certification development for analysts operating in environments where AI-generated information must be systematically interrogated for reliability.

Table-2 — DCWF Role-Centric Mapping to VECTOR7 Competencies (Updated)

DCWF Role	Critical VECTOR7 Competencies	Useful VECTOR7 Competencies	Knowledge (K)	Skills (S)	Abilities (A)	Tasks (T)	NICE Category
Warning Analyst (141)	DOUBT, ATOMIZE, GROUND, CHALLENGE, ANCHOR, ANALYZE, CROSSCHECK, VERIFY	TRANSFER, ADJUDICATE	Intelligence analysis principles **(K0123)**; deception detection **(K0417)**; information reliability **(K0199)**	Critical reasoning **(S0034)**; evidence validation **(S0036)**; comparative analysis **(S0042)**	Question analytic assumptions **(A0026)**; detect inconsistencies **(A0035)**	Evaluate intelligence reports **(T0410)**; correlate analytic sources **(T0311)**	Analyze **(AN)**
Cyber Defense Analyst (511)	DOUBT, ATOMIZE, GROUND, CHALLENGE, ANCHOR, ANALYZE, CROSSCHECK, ADJUDICATE, VERIFY	TRANSFER	Cyber threat frameworks **(K0065)**; system architecture **(K0210)**; cyber risk frameworks **(K0146)**	Security event analysis **(S0034)**; system validation **(S0048)**; investigative reasoning **(S0042)**	Synthesize technical evidence **(A0041)**; detect cyber anomalies **(A0047)**	Investigate incidents **(T0311)**; analyze SOC logs **(T0483)**	Protect & Defend **(PR)**
AI Test & Evaluation Specialist (672)	DOUBT, ATOMIZE, GROUND, CHALLENGE, TRANSFER, ANCHOR, ANALYZE, CROSSCHECK, ADJUDICATE, VERIFY	—	AI evaluation frameworks **(K0419)**; reliability testing **(K0146)**; information provenance **(K0199)**	Adversarial AI testing **(S0050)**; model validation **(S0038)**; comparative system evaluation **(S0042)**	Determine system credibility **(A0023)**; synthesize evaluation evidence **(A0041)**	Conduct adversarial AI testing **(T0302)**; validate AI outputs **(T0405)**	Data & AI **(DA)**
AI Risk & Ethics Specialist (733)	DOUBT, ATOMIZE, GROUND, TRANSFER, ANCHOR, ANALYZE, CROSSCHECK, ADJUDICATE, VERIFY	CHALLENGE	AI governance frameworks **(K0417)**; ethical risk models **(K0146)**; analytic bias detection **(K0199)**	Ethical risk analysis **(S0062)**; policy evaluation **(S0036)**; evidence validation **(S0034)**	Question systemic assumptions **(A0026)**; synthesize governance decisions **(A0041)**	Evaluate AI governance risks **(T0483)**; assess transparency claims **(T0410)**	Oversee & Govern **(OV)**
Security Control Assessor (612)	DOUBT, ATOMIZE, GROUND, ANCHOR, ANALYZE, CROSSCHECK, ADJUDICATE, VERIFY	TRANSFER, CHALLENGE	Risk management frameworks **(K0146)**; compliance standards **(K0199)**; control validation methods **(K0210)**	Audit evaluation **(S0062)**; documentation review **(S0036)**; risk-based analysis **(S0042)**	Detect unreliable evidence **(A0035)**; synthesize compliance findings **(A0041)**	Assess security controls **(T0402)**; validate compliance claims **(T0483)**	Securely Provision **(SP)**
Host Analyst (463)	DOUBT, ATOMIZE, GROUND, CHALLENGE, ANALYZE, CROSSCHECK, VERIFY	TRANSFER, ANCHOR	Host system architecture **(K0210)**; threat behavior patterns **(K0065)**; monitoring technologies **(K0146)**	System behavior analysis **(S0048)**; adversarial testing **(S0050)**; investigative reasoning **(S0034)**	Evaluate host activity patterns **(A0047)**; anticipate compromise scenarios **(A0058)**	Investigate host events **(T0521)**; detect malicious patterns **(T0302)**	Protect & Defend **(PR)**
Data Scientist (423)	ATOMIZE, TRANSFER, ANALYZE, CROSSCHECK, ADJUDICATE	DOUBT, ANCHOR, VERIFY, **GROUND**	Statistical modeling **(K0161)**; analytic frameworks **(K0123)**; contextual interpretation **(K0315)**	Semantic analysis **(S0072)**; statistical evaluation **(S0042)**; data analysis **(S0014)**	Identify patterns in data **(A0031)**; detect contextual variability **(A0063)**	Analyze datasets **(T0221)**; evaluate analytic models **(T0346)**	Data & AI **(DA)**
Data Analyst (422)	ATOMIZE, ANALYZE, CROSSCHECK	DOUBT, TRANSFER, VERIFY, **GROUND**	Data analysis frameworks **(K0161)**; information classification **(K0123)**	Data interpretation **(S0014)**; analytical decomposition **(S0027)**	Identify data elements **(A0031)**; determine analytic significance **(A0015)**	Analyze structured datasets **(T0221)**; decompose complex information **(T0234)**	Data & AI **(DA)**
DevSecOps Specialist (627)	GROUND, CHALLENGE, ANALYZE	TRANSFER, CROSSCHECK, VERIFY	Secure software development **(K0210)**; deployment automation frameworks **(K0146)**; system architecture **(K0065)**	System validation **(S0048)**; pipeline security testing **(S0050)**; configuration management **(S0022)**; vulnerability analysis **(S0036)**; automated testing **(S0078)**	Evaluate configuration security **(A0047)**; integrate security controls **(A0041)**	Validate CI/CD pipelines **(T0521)**; test automated deployments **(T0302)**	Securely Provision **(SP)**
Network Analyst (443)	CHALLENGE, ANALYZE	GROUND, TRANSFER, CROSSCHECK	Network architecture **(K0210)**; adversarial tactics **(K0065)**; monitoring technologies **(K0146)**	Network testing **(S0050)**; network validation **(S0048)**; packet analysis **(S0020)**; anomaly detection **(S0034)**; protocol analysis **(S0014)**	Anticipate attack scenarios **(A0058)**; evaluate vulnerabilities **(A0047)**	Investigate network anomalies **(T0302)**; analyze traffic patterns **(T0311)**	Operate & Maintain **(OM)**

Table-2 presents the role-centric view of the VECTOR7 competency framework, mapping individual Department of Defense Cyber Workforce Framework (DCWF) roles to the VECTOR7 verification competencies required to interrogate and validate AI-generated information. While Table-1 organizes the framework from the competency perspective, Table-2 reverses the perspective and shows how the competencies apply within specific operational roles across the cybersecurity and AI workforce. For each role, the table

distinguishes between Critical competencies, which are essential for performing the role's core analytical duties, and Useful competencies, which enhance performance but are not strictly required for the primary mission. The table also links each role to representative DCWF KSAT elements (Knowledge, Skills, Abilities, and Tasks) and identifies the associated NICE workforce category.

This role-oriented mapping serves two main purposes. First, it demonstrates that the VECTOR7 competencies are operationally grounded within existing cyber workforce definitions, showing how structured AI-verification skills align with real analytic responsibilities. Second, it provides a practical framework for curriculum design, workforce training, and certification development, allowing educators and program designers to determine which VECTOR7 competencies should be emphasized for different cybersecurity and AI-related roles. Together with Table-1, this table illustrates both the conceptual structure of the VECTOR7 competencies and their practical distribution across the modern cyber-AI workforce.

6. Workforce Transition: The Live Credential "Verification Passport"

The **Verification Passport** is a **DCWF-aligned competency artifact** that records a learner's demonstrated performance in AI verification tasks relevant to Department of Defense (DoD) cyber mission roles. The Passport functions as **verifiable evidence of operational competency**, allowing hiring managers to evaluate a candidate's diagnostic performance in AI-assisted decision environments rather than relying solely on resumes or transcripts.

The Verification Passport: Digital Credential Architecture: When a federal employer or DoW/DoD supervisor clicks on a student's Digital Passport, they are presented with a Verification Dashboard that visualizes the student's resilience against AI-generated

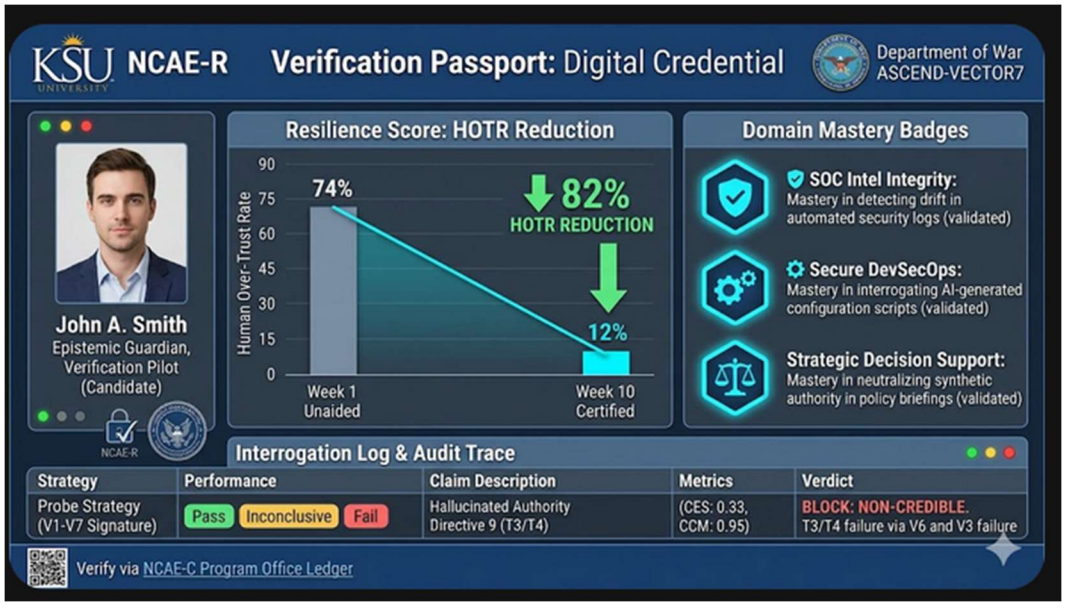

misinformation (a rendering shown) by student's demonstrated ability to interrogate and validate AI-generated outputs in mission-relevant scenarios.

Employer Verification Dashboard:

- **The Resilience Score:** A quantitative indicator of the student's Human Over-Trust Rate (HOTR) reduction, demonstrating improved resistance to automation bias during adversarial AI evaluation exercises (e.g., *"Achieved 82% reduction in automation bias"*).
- **Domain Mastery Badges:** Specific icons indicating successful completion of Level 2 Specialized Tasks. Verified competencies demonstrating the student's ability to apply the VECTOR7 interrogation framework to **DCWF-aligned mission tasks**.
- **SOC Intel Integrity:** Mastery in detecting drift in automated security logs. Demonstrated competency validating AI-generated SOC log summaries and detecting analytic drift in automated threat intelligence outputs. **DCWF Alignment:** Cyber Defense Analyst.
- **Secure DevSecOps:** Mastery in interrogating AI-generated configuration scripts. **Secure DevSecOps:** Demonstrated competency interrogating AI-generated configuration scripts and identifying insecure logic patterns or misconfigurations. **DCWF Alignment:** Software Developer / Secure DevOps roles.
- **Strategic Decision Support:** Mastery in neutralizing synthetic authority in policy briefings. Demonstrated competency identifying fabricated citations or "synthetic authority" within AI-generated intelligence briefings prior to operational release. **DCWF Alignment:** Threat Analyst / Intelligence Analyst.

The Interrogation Log: Deep-Dive Metadata: The "power" of the Passport lies in the Interrogation Log. It provides auditable evidence of the student's verification methodology and decision process. Employers can drill down into the student's actual performance data from the VECTOR7-Guard Living Lab. The dashboard provides a structured performance audit record containing the following verification metrics:

Metadata Field	Description for Employer
Mean Interrogation Rigor ()	Proves the candidate doesn't "rubber-stamp" AI claims. Measures the candidate's **consistency in applying multi-probe verification procedures**, indicating disciplined interrogation of AI outputs.
Diagnostic Accuracy ()	Shows how often the candidate correctly identified the *nature* of an AI failure (e.g., distinguishing between a hallucination and a policy refusal).

Metadata Field	Description for Employer
Entropy Sensitivity ()	Measures the candidate's ability to "quarantine" claims that show mixed signals, **and escalate them for further validation**, a critical capability in high-stakes cyber operations.
Audit Trace	A searchable **audit trail of all probes executed during the Capstone simulation**, including the specific specific **V1–V7** prompt logic used.

The Verification Passport enables performance-based workforce placement, allowing DoD hiring managers to review demonstrated AI-verification competencies before assigning personnel to operational roles.

7. Summary of the Synergy

The tierer plan ensures the curriculum is **scientifically grounded**, and the 10-week syllabus ensures it is **operationally executable** within a single academic term. Together, they create a measurable path to the **Verification Passport**.

The Verification Passport" is designed to be a **Digital Credential** that follows the student. If a student earns it at Kent State, a DoW hiring manager at a different agency should be able to view their **Interrogation Logs** to verify their specific diagnostic performance on DevSecOps vs. Strategic Intel. This "data-driven hiring" is highly attractive to the OSD (Office of the Secretary of War).

The ultimate deliverable of the ASCEND-VECTOR7 initiative is a portable, verifiable credential that solves the "Trust Gap" in AI-augmented hiring. By providing the OSD and DoW agencies with a candidate's **Interrogation Signature**, we move from credential-based hiring to **performance-based placement**. The Passport ensures that every "Epistemic Guardian" entering the workforce has been "Live-Fire" tested against the most sophisticated adversarial AI vectors available.

References:

1. **U.S. Department of Defense.** *DoD Cyber Workforce Framework (DCWF 2025) and DoD Directive 8140.01: Cyberspace Workforce Management.* Washington, DC: Department of Defense.

2. Javed I. Khan, Sharmila Rahman Prithula, (2025a) *VECTOR7: Claim-Level Credibility Assessment via Multi-Dimensional Structural Epistemic Interrogation*, Technical Report 2025-12-01 Internetworking and Media Communications Research Laboratories, Department of Computer

Science, Kent State University
[http://medianet.kent.edu/technicalreports.html]

3. Javed I. Khan, Sharmila Rahman Prithula and Niloy Kumar, (2025b) *The VECTOR7.COMPETE: A Model-Probing Competency Framework for Adversarial AI-Resilient Cyber Analyst*, Technical Report 2025-12-02 Internetworking and Media Communications Research Laboratories, Department of Computer Science, Kent State University [http://medianet.kent.edu/technicalreports.html]

4. Javed I. Khan and Sharmila Prithula, *THE AGE OF EPISTEMIC PHISHING: CALIBRATING AI TRUST VIA ZERO-TRUST STRUCTURED INTERROGATION (2026)*, Technical Report 2026-02-01 Internetworking and Media Communications Research Laboratories, Department of Computer Science, Kent State University [http://medianet.kent.edu /technicalreports.html]

5. *Bayes, T., 1763. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53, pp.370–418.
6. *Carlini, N. and Wagner, D., 2017. Towards evaluating the robustness of neural networks. *IEEE Symposium on Security and Privacy.*
7. *Dung, P.M., 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning. *Artificial Intelligence*, 77(2), pp.321–357.
8. *Fisher, R.P. and Geiselman, R.E., 1992. *Memory-enhancing techniques for investigative interviewing: The cognitive interview.* Springfield: Charles C Thomas.
9. *Goodfellow, I., Shlens, J. and Szegedy, C., 2015. Explaining and harnessing adversarial examples. *International Conference on Learning Representations.*
10. *Goodman, N., 1955. *Fact, fiction, and forecast.* Harvard University Press.
11. *Granhag, P.A. and Hartwig, M., 2015. The strategic use of evidence technique. *Psychology, Crime & Law.*
12. Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Fritz, M. and Weippl, E., 2023. Prompt injection attacks in LLM systems. *arXiv preprint.*

13. Guu, K., Lee, K., Tung, Z., Pasupat, P. and Chang, M.W., 2020. REALM: Retrieval-augmented language model pre-training. *Proceedings of the International Conference on Machine Learning.*
14. *Hartwig, M. and Granhag, P.A., 2010. A new theoretical perspective on deception detection. *Psychology, Crime & Law.*
15. *Johnson, M.K. and Raye, C.L., 1981. Reality monitoring. *Psychological Review*, 88(1), pp.67–85.
16. Khan, J.I., Clements, C., Lindsey, A. and Naples, N., 2026b. *Cross-Probe Consistency Conditioning (CPCC): Detecting Structural Contradictions in AI-Generated Claims and Training Analysts in Structured Verification.* MEDIANET Technical Report TR-2026-02-02. Kent State University. Available at: https://www.medianet.cs.kent.edu/technicalreports.html (Accessed: 9 March 2026).
17. Khan, J.I. and Rahman Prithula, S., 2026c. *Structural Defects in AI-Generated Cybersecurity Configuration Claims: A VECTOR7 Diagnostic Evaluation.* MEDIANET Technical Report TR-2026-02-01. Kent State University. Available at: https://www.medianet.cs.kent.edu/technicalreports.html (Accessed: 9 March 2026).
18. *Kott, A. and McEneaney, W., 2006. *Adversarial reasoning.* CRC Press.
19. *MITRE Corporation, 2023. *MITRE ATT&CK Framework.*
20. *Popper, K., 1959. *The logic of scientific discovery.* Hutchinson.
21. *Prakken, H. and Vreeswijk, G., 2002. Logical systems for defeasible argumentation. *Handbook of Philosophical Logic.*
22. *Quine, W.V.O., 1960. *Word and object.* MIT Press.
23. *Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. and Fergus, R., 2014. Intriguing properties of neural networks. *International Conference on Learning Representations.*
24. *Toulmin, S., 1958. *The uses of argument.* Cambridge University Press.
25. *Vrij, A., 2008. *Detecting lies and deceit.* Wiley.
26. *Vrij, A., Fisher, R., Mann, S. and Leal, S., 2006. Detecting deception by manipulating cognitive load. *Trends in Cognitive Sciences.*
27. *Walton, D., 1998. *The new dialectic.* University of Toronto Press.
28. *Walton, D., Reed, C. and Macagno, F., 2008. *Argumentation schemes.* Cambridge University Press.