


# Trusted AI: Explainable, Auditable, Resilient

Trust is what decides who can truly operationalize AI.

**Dave Hoon**
Chief Technology Officer, Norseman Defense Technologies · February 2026

AI is everywhere. But in national security, "everywhere" isn't good enough. What matters is **trusted AI: explainable, auditable, and resilient under adversarial conditions.**

Black-box algorithms won't cut it when lives or missions are on the line. Leaders need confidence that AI recommendations are traceable, grounded in real data, and able to operate securely at the tactical edge.

## The Three Pillars

### Explainable

An AI system that can't explain its reasoning isn't a decision-support tool—it's a liability. Explainability means an operator can ask "why did the system recommend this?" and receive a traceable, human-readable answer grounded in the actual data and logic the model used. This is increasingly a regulatory requirement under the DoD AI Ethics Principles and NSM-10.

### Auditable

Every AI interaction in a national security context should generate an immutable audit trail: who initiated the query, what data was used, what the model output was, and what action was taken as a result. Without this, there is no accountability and no way to prove compliance during an assessment. Auditability also enables continuous improvement—when you can review every AI decision against actual outcomes, you can identify where models are drifting before they cause mission impact.

### Resilient

National security AI systems operate in contested environments. Adversaries will attempt to manipulate sensor data, poison training sets, and probe model outputs for exploitable patterns. Trusted AI must be designed with adversarial robustness from the start—not as a post-deployment patch. This means adversarial testing during development, input validation at inference time, anomaly detection on model outputs, and fallback modes when confidence thresholds aren't met.

## Odin's Edge: Trust as the Foundation

Norseman's Odin's Edge was built on this principle: trusted AI isn't a feature—it's the foundation. From secure data handling to explainable decision-making, trust is baked in at the architecture level.

The promise of AI is real. But trust is what decides who can truly operationalize it.

Explore **Odin's Edge** or contact **Norseman** to discuss trusted AI architecture for your mission.

Published at norseman.com/insights/trusted-ai-explainable-auditable-resilient. To talk with Norseman Defense Technologies, call 410.579.8600 or email sales@norseman.com.