Wednesday, June 3

4:20 pm – 4:40 pm

From Snapshot to Signal: Governed Agentic AI for Continuous ATO

Bobby Tuohy

VP, Product and Platform Strategy

Cav AI

Abstract:

The Department of War is in the midst of the most significant platform modernization cycle in a generation. Programs across the services are racing to field autonomous systems, attritable platforms, and AI-enabled capabilities at unprecedented speed. The defense industrial base is scaling rapidly to meet this demand, with new entrants and commercial technology companies accelerating delivery timelines that traditional defense programs measured in decades.

But there is a bottleneck no one is talking about: every one of these systems must receive an Authority to Operate before it can be deployed. And the ATO process has not kept pace with the speed at which systems are now being built.

The problem compounds further as these modernized platforms increasingly rely on open-source software, third-party AI libraries, and commercial components. Each dependency introduces supply chain risk that must be evaluated against RMF controls — yet the current compliance model has no mechanism to continuously assess third-party risk at the speed these components evolve.

Vulnerabilities appear daily. Dependencies update weekly. The compliance picture remains static.

The consequences are not just operational delay. When Authorizing Officials cannot see the current security posture of a system — because the evidence is months old and manually assembled — they face an impossible choice: slow the mission by withholding authorization, or accept risk they cannot fully quantify. In practice, AOs are routinely forced to accept unnecessary or even unknown risk, not because the underlying system is

insecure, but because the compliance process cannot give them the visibility they need to make an informed decision.

This session examines a fundamentally different model: continuously computing compliance posture from live operational data using a governed agentic AI architecture. We will walk through how a deterministic integration layer can ingest raw security telemetry, resolve it into canonical entities, and emit structured compliance evidence without manual intervention — and how an agentic layer above it can evaluate system state, identify control gaps, and propose compliance claims that human operators review and approve before they take effect.

This is not a theoretical framework. I will share production outcomes from federal deployments where this approach reduced ATO timelines by 75–96%, including environments spanning NIPR, SIPR, and JWICS. We will address the governance challenge that DISA has rightly raised around agentic AI: how to ensure least-privilege access, auditability, and accountability when autonomous systems reason about compliance in high-stakes operational environments.

I hope attendees will take away three things: an understanding of how entity-claim data models enable traceable, machine-verifiable compliance findings; a practical framework for integrating governed AI agents into existing RMF and eMASS workflows; and a clear picture of what Continuous ATO looks like when compliance infrastructure moves at the speed of platform delivery, not the speed of paperwork.