

Arbiter

The Trust Layer for Agentic AI

Trust Through Verification | Every Voice. One Conversation.


The Challenge: Multi-Agent AI Without Trust

Defense and intelligence organizations are building ecosystems of agentic AI—specialized agents for threat assessment, mission autonomy, and constraint management. But more agents means more complexity: outputs conflict, constraints get misinterpreted, and there's no arbitration when agents disagree.

- **Single Points of Failure:** One model, one opinion, no verification
- **No Arbitration:** When agents conflict, operators have no synthesis layer
- **Invisible Uncertainty:** AI provides answers without confidence levels
- **Hidden Contradictions:** Semantic similarity masks logical disagreements

The Solution: Trust Layer for Agentic AI

Arbiter sits above your AI systems as the verification and arbitration layer. It orchestrates multiple AI models, extracts and compares discrete claims, detects contradictions, validates constraints, and maintains complete audit trails for every AI-assisted decision.



COUNCIL MODE

Orchestrate 3-7 AI models. Extract discrete claims. Detect contradictions. Confidence scoring based on cross-model agreement.

AGENT ARBITRATION

When agents conflict, Arbiter arbitrates. Weighted recommendations with full attribution. Cross-INT fusion.

CONSTRAINT VALIDATION

Upload ROE, ACM, FSCM. Multi-model interpretation catches ambiguities before mission execution.

CONFIDENCE SCORING

Explicit confidence on every output. Degradation-aware. Critical dissent override for safety-critical decisions.

Platform Architecture

TRUST LAYER	ORCHESTRATION	MEMORY LAYER	MODEL LAYER
Audit, confidence, constraint validation	Council Mode, claims, arbitration, synthesis	Constraints, context, mission templates	Claude, GPT, Gemini, Llama, MCP agents

Key Differentiators

- **Trust Layer architecture**—we make agents trustworthy
- **Claim-level verification**—discrete assertion comparison
- **Contradiction detection**—semantic + logical analysis
- **Constraint Validation**—unique to market
- **Critical dissent override**—safety-critical protection

- **Cross-INT fusion**—source diversity weighting
- **Model diversity monitoring**—correlated failure detection
- **Degradation-aware confidence**—explicit uncertainty
- **No vendor lock-in**—model agnostic architecture
- **Edge-deployable**—tactical hardware integration

SOF Mission Applications

Intelligence Analysis

- Multi-model pattern-of-life with consensus
- Target verification before PID/CID
- Cross-INT fusion with source diversity
- Geospatial and temporal claim validation

Mission Planning

- COA development with multi-model risk assessment
- Constraint interpretation validation
- Red team analysis via adversary modeling

Decision Support

- Targeting with ROE constraint validation
- Agent arbitration for conflicting inputs
- Full audit trail for accountability
- Critical dissent override for safety-critical

Edge Operations

- DDIL-optimized tactical edge deployment
- Graceful degradation with confidence flags
- Bandwidth-aware model selection

SOCOM Technology Area Alignment

Agentic Protocols	MCP server exposes Arbiter as arbitration service; ready for A2A
Agentic Workflows	Council Mode + Agent Arbitration for multi-step task execution
Human-Machine Teaming	Confidence scoring, dissent highlighting, operator override authority
Knowledge Representation	Structured constraint storage, mission templates, semantic context
Low SWaP-C	Edge deployment via tactical hardware; quantized local models; graceful degradation
Metrics & AI Accuracy	Consensus as accuracy metric; historical calibration; provenance logs
Collaborative Autonomous	Agent Arbitration synthesizes multi-agent outputs for coordination

Deployment Options

Cloud	Hybrid	Edge Hardware
FedRAMP-ready SaaS. Full Council Mode (5-7 models), complete feature set, automatic updates.	On-premise orchestration with cloud model access. Data stays local, API calls for expanded capability.	Fully disconnected operation. Local quantized models with sync-on-reconnect to cloud.

Technology Readiness: TRL 6 — Core platform operational (MVP), Council Mode synthesis functional, multi-provider integration tested, claim extraction pipeline operational. Constraint Validation and Agent Arbitration in development. Ready for operational demonstration.

Patent pending. Core verification architecture protected by U.S. provisional patent application.

SPARTAN X CORP

peter.galle@spartanxcorp.com

spartanxcorp.com | 1-888-536-1348

SDVOSB Certified • Seaport NxG • Austin, TX • Anchorage, AK • Largo, FL

EXCELLENCE • INTEGRITY • SERVICE