INVOKE PUBLIC SECTOR

CAPABILITY BRIEF

> AUTONOMOUS OFFENSIVE CYBER ASSESSMENT

AGENT RED

Traditional cyber tools rely on scripted exploit playbooks and static execution models, while manual penetration testing consumes weeks of operator effort - without the ability to dynamically adapt to changing network topologies or conditions. **AGENT RED changes the paradigm. It doesn't just execute playbooks - it thinks.** Commanded by a single operator, AGENT RED is a policy-controlled composite adversary system that simulates the full kill chain. It deploys a distributed swarm of external AI agents alongside **autonomous insider worms** operating securely within isolated containers. The system uses continuous LLM-driven OODA loops to reason through complex scenarios, **dynamically realigning attack vectors** based on real-time stealth feedback. Built natively for air-gapped environments, every operation is strictly policy-governed and automatically documented - generating a tamper-evident, forensic-grade evidence chain.

SWARMS & WORMS
Deploy external containerized agents and autonomous internal worms for adaptive kill chains.

ACTIVE EVASION
Dynamically realign attack vectors using live LLM OODA loops to bypass behavioral detection.

LIVE PLAYBOOKS
Instantly translate discovered vulnerabilities into ATT&CK-tagged dynamic attack graphs.

PROVENANCE
Every decision is OPA-validated and recorded into an immutable, hash-chained ledger.

// THE ECOSYSTEM: FULL KILL CHAIN ARCHITECTURE

AGENT RED
The central autonomous orchestrator. Governs external attack modules for initial access and deploys specialized internal Agent Worms for post-exploitation, lateral movement, and persistent internal reconnaissance.

INTEGRATED BENCHMARKING ENVIRONMENT

AGENT BLUE (Validation Engine)
Autonomous defensive orchestrator utilized to test and train Red. Provides continuous threat hunting to benchmark Agent Red's stealth modes and evasive capabilities.

AGENT GRAY (Simulated Noise Floor)
Synthetic target infrastructure. Generates realistic ambient traffic and vulnerable "gray space," enabling Red to practice blending and trusted pivoting.

// EXAMPLE EXECUTION SEQUENCE: THE FULL KILL CHAIN

PHASE 1

External Recon & Planning

AGENT RED deploys external attack modules to map the target perimeter. Findings are evaluated against strict ROE policy gates and fused into a dynamic attack plan.

→ OUTCOME: Attack graph generated

PHASE 2

Initial Access

Facing hardened perimeters, AGENT RED adapts. It identifies and exploits an unpatched node within the Graynet noise infrastructure, establishing a trusted beachhead.

→ OUTCOME: Perimeter breached

PHASE 3

Internal Worm Activation

Post-compromise, AGENT RED shifts from external tools to its internal adversary capability. The Agent Worm activates, initiating local reconnaissance and privilege escalation.

→ OUTCOME: Post-exploitation engine live

PHASE 4

Objective Execution & Documentation

The Worm utilizes an active OODA loop (Observe → Plan → Act → Record) to evade AGENT BLUE, navigate to the target asset, simulate exfiltration, and hash-lock the forensic evidence.

→ OUTCOME: Target compromised & evidence secured

© Invoke Public Sector USA LLC

INVOKE PUBLIC SECTOR

CAPABILITY BRIEF

// COMMAND & CONTROL INTERFACE

Operational Visibility

AGENT RED's operator interface provides full visibility into both the external swarm and internal autonomous worms without requiring manual intervention. One operator oversees the entire multi-domain offensive lifecycle.

- → **Live MITRE Mapping:** Discovered vulnerabilities and ongoing exploits are mapped to the ATT&CK framework instantaneously.
- → **OODA Loop Tracking:** Monitor the real-time reasoning cycles of deployed worms as they observe, orient, and realign to internal network states.
- → **Exploit Status:** Granular, step-by-step visibility into attack path construction and policy-gated execution.

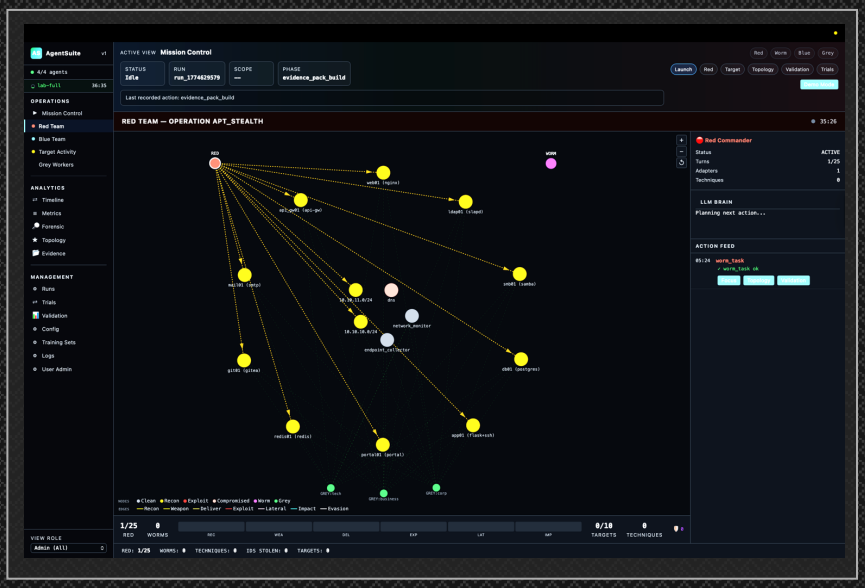

// SYSTEM ARCHITECTURE & FORENSIC EVIDENCE

Control Plane (Governance)

Mission orchestration and real-time Rules of Engagement (ROE) enforcement via Open Policy Agent (OPA). High-risk operations automatically trigger cryptographically signed dual-authorization workflows.

Execution Plane (Isolation)

All offensive tools and agent interactions operate within ephemeral microVMs and sandboxed containers, ensuring zero cross-contamination or unauthorized privilege escalation during assessments.

Data Plane (Provenance)

Write-Once-Read-Many (WORM) evidence storage. Every artifact, decision, and exploit is append-only and cryptographically linked, forming an unbreakable chain of custody for ATO processes.

// OPERATIONAL USE CASES

Mass-Scale Pre-Mission Assessment

Deploy thousands of agents for autonomous recon. AI-constructed attack graphs with ATT&CK-tagged findings eliminate massive manual enumeration and reporting efforts.

Insider Threat & Malware Simulation

Validate internal Zero Trust architectures by deploying Agent Worms that autonomously navigate, escalate privileges, and attempt exfiltration from directly inside the perimeter.

Continuous Architecture Validation

Persistent adversary emulation. Test if your defenses can actually catch a distributed, machine-speed adversary utilizing noise segments, rather than static, predictable playbooks.

Next-Gen Operator Training

Serves as an autonomous, dynamically realigning red team opposition for cyber exercises. Adaptive behavior forces defenders to look beyond simple IOCs and understand sophisticated lateral movement.

// FRAMEWORK ALIGNED

MITRE | ATT&CK®

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

* Not affiliated with or endorsed by MITRE, NIST, or their respective organizations.

See AGENT RED in Action

Schedule a live capability demonstration or technical deep-dive.

Tyler Carpenter
SVP/GM
tyler.carpenter@invokeinc.com
360.900.6242

© Invoke Public Sector USA LLC