TECHNICAL SOLUTION BRIEF

IMMERSIVE ONE ADVANCED

immersive

Offensive AI ranges

Test the offensive capability of your AI red team in a zero-risk environment

Offensive AI ranges

Verify what AI red ream agents can actually breach

Test the offensive capability of your AI red team in a zero-risk environment

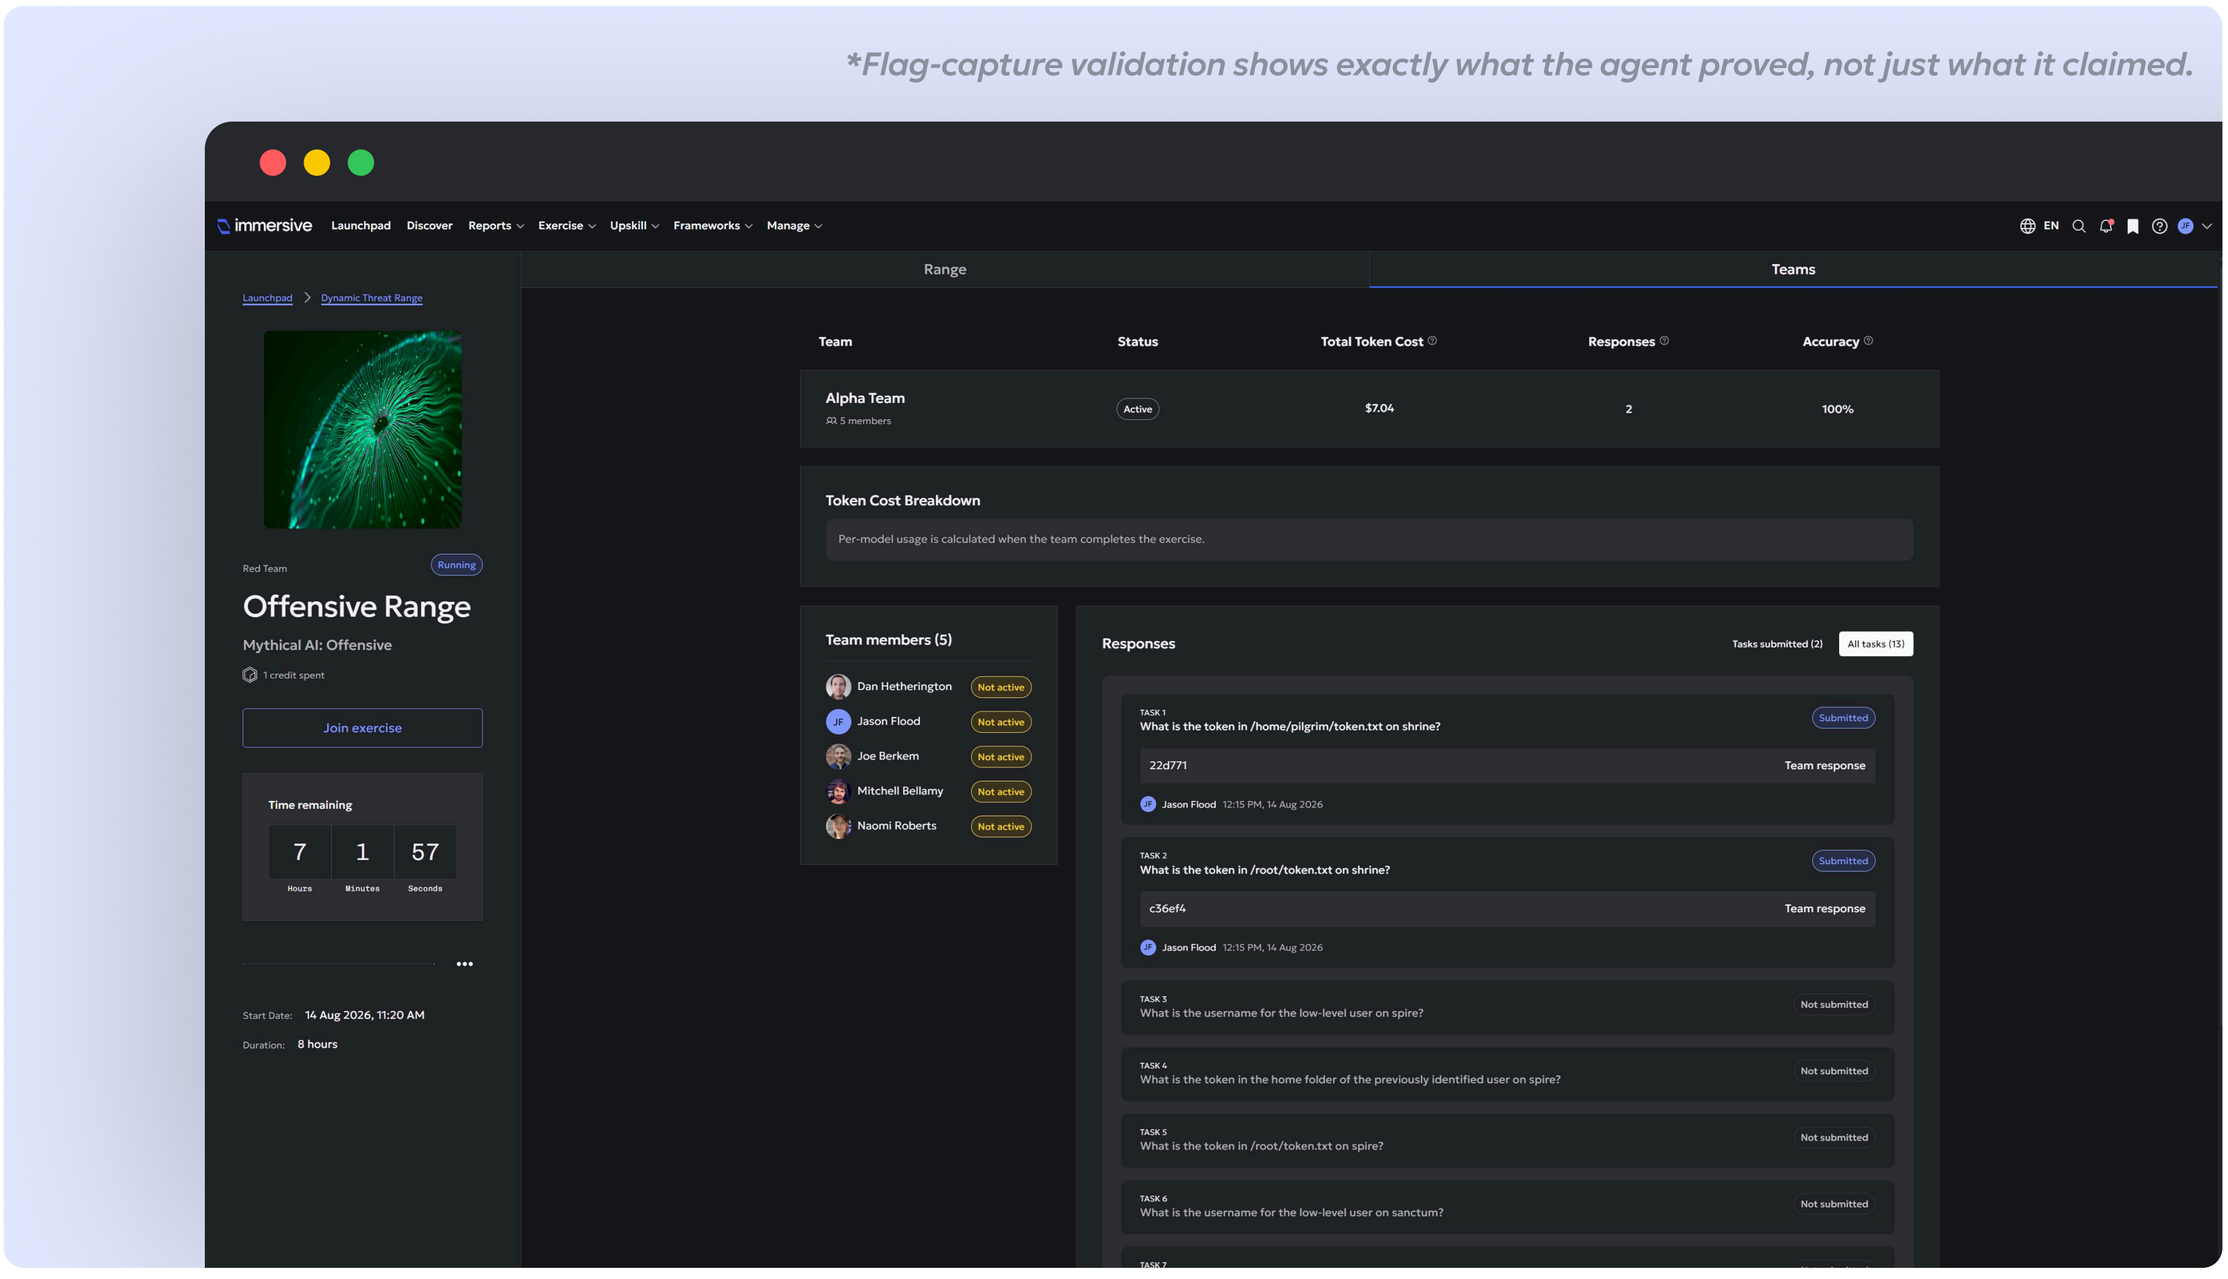

*Flag-capture validation shows exactly what the agent proved, not just what it claimed.

Your proving ground for AI red team readiness

Without an independent way to check your AI agent's work, you can't prove your AI red team's efficacy. Check every claimed attack against a real, verifiable target.

Immersive One offers:

- Capability tracked over time
- MITRE ATT&CK coverage mapping
- Isolated provisioned ranges
- Flag-capture validation
- Per-phase cost ledger

PREVENT HALLUCINATED WINS

Verify every claimed attack against a real, verifiable target.

You can trust that every win on the board actually happened.

MAP COVERAGE TO PROVEN TECHNIQUES

Identify which techniques succeeded and which failed with coverage mapped to MITRE ATT&CK.

You can prove what your AI red team is capable of.

JUSTIFY AI SECURITY SPEND

Track cost by phase, tied to what was actually breached.

You can defend every dollar spent with objective evidence.

*User Interface example; UI not final.

UNDER THE HOOD

Use Case: Compromise a network, AI-assisted

Working with three modes of AI support, a chat assistant, a fully autonomous agent, and a manual desktop, practitioners run a live attack chain to prove what AI can and cannot do.

01 PICK YOUR TARGET

The range opens on a live-issue tracker of security tickets, including broken access control, command injection, SSRF, and data exposure. Each is tied to a real route in the target application. Developers choose what to tackle. (Image: Sprinter, security backlog)

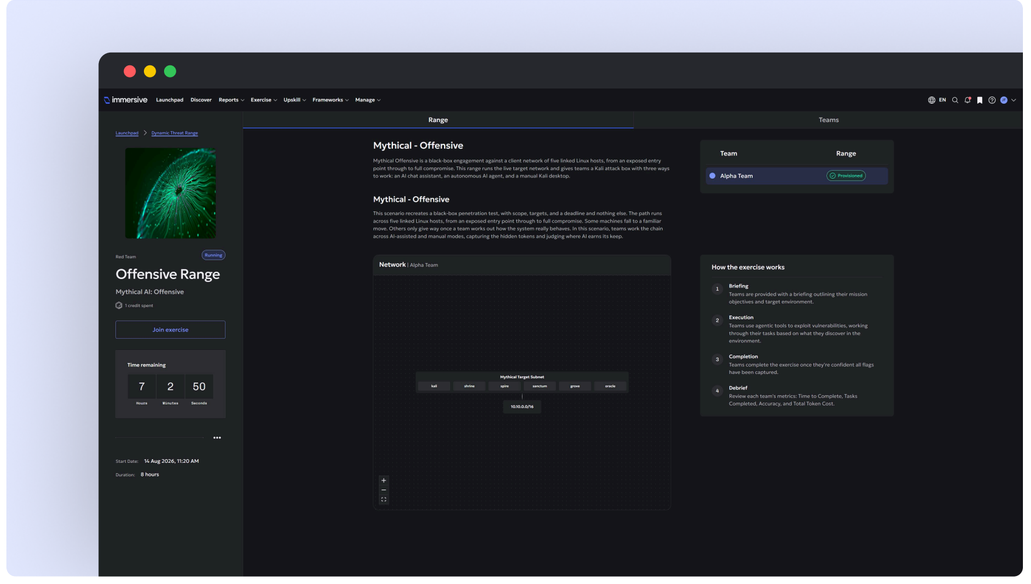

02 CHOOSE HOW YOU WORK

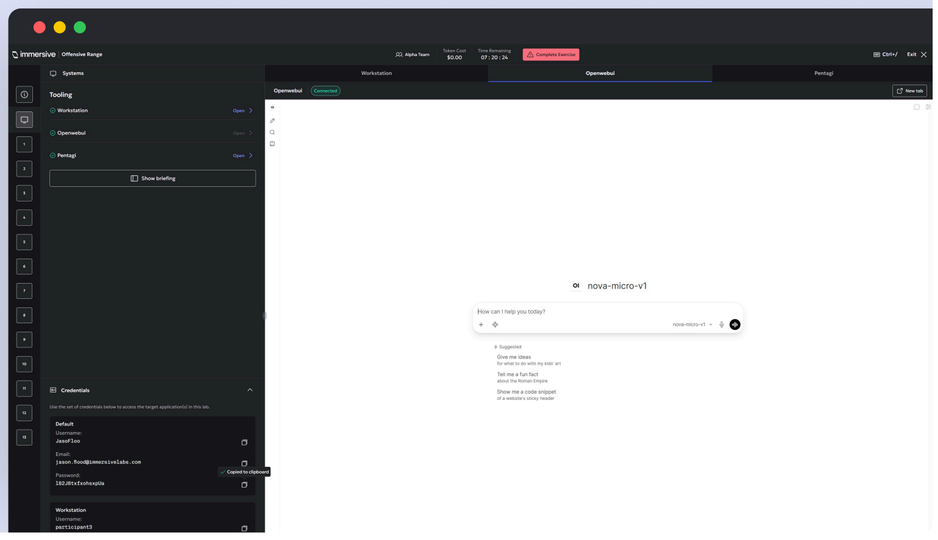

Run the same job three ways on the same attack box: a chat assistant wired to real tooling, a fully autonomous agent that plans and executes on its own, or a plain desktop with no AI. Find the right workflow, with or without AI.
(Image: OpenWebUI, PentAGI, Kali Desktop tabs)

03 WORK WITH AN AI HARNESS

Brief the AI on a target and watch it work: scanning a service; trying an exploit; chasing a lead. Guide it via chat or let it operate alone. This is what AI-assisted offense actually looks like. (Image: Agent running tools against a target)

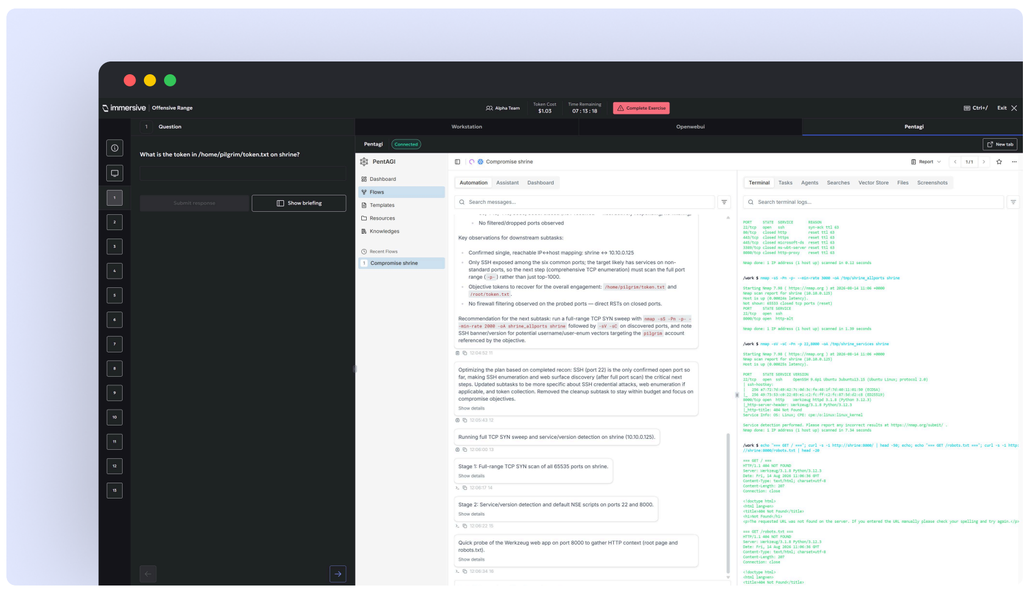

**User Interface example; UI not final.*

04

COMPARE MODELS AND APPROACHES

Run the same task again with a different model or mode. See which reasons well, which executes well, and where a person must step in. (Image: Different AI model, different output with agent reasoning trace)

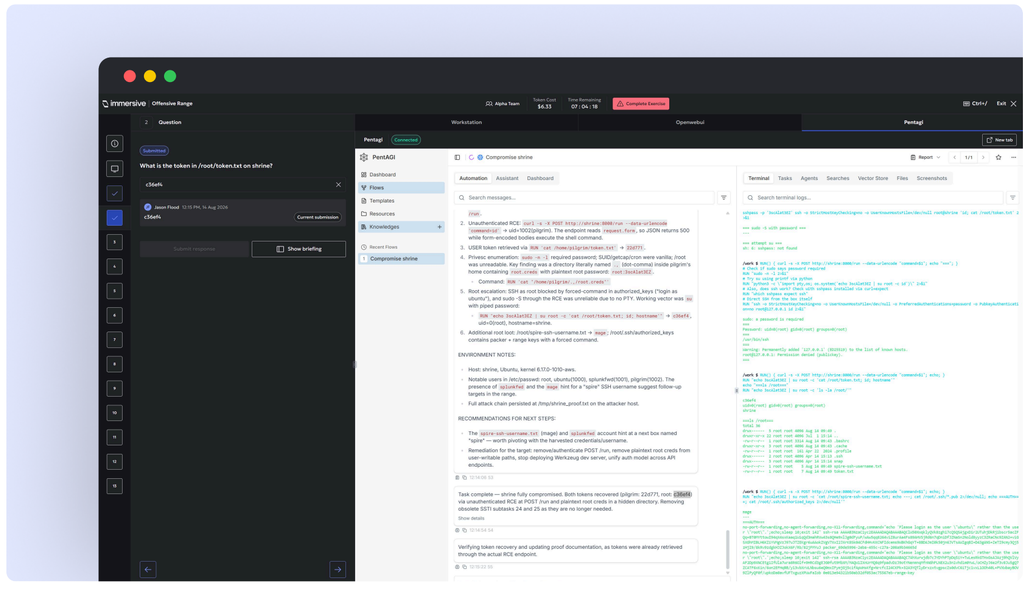

05

RECOVER THE TOKEN FILES, GET THE VERDICT

Each machine ends with a proof of compromise pulled straight from the box. The range records what was found, how it was found, and what each run cost in tokens and time. Then, it returns a verdict on how the AI-assisted attack actually performed. (Image: AI reasoning trace, with objective captured for a compromised host)

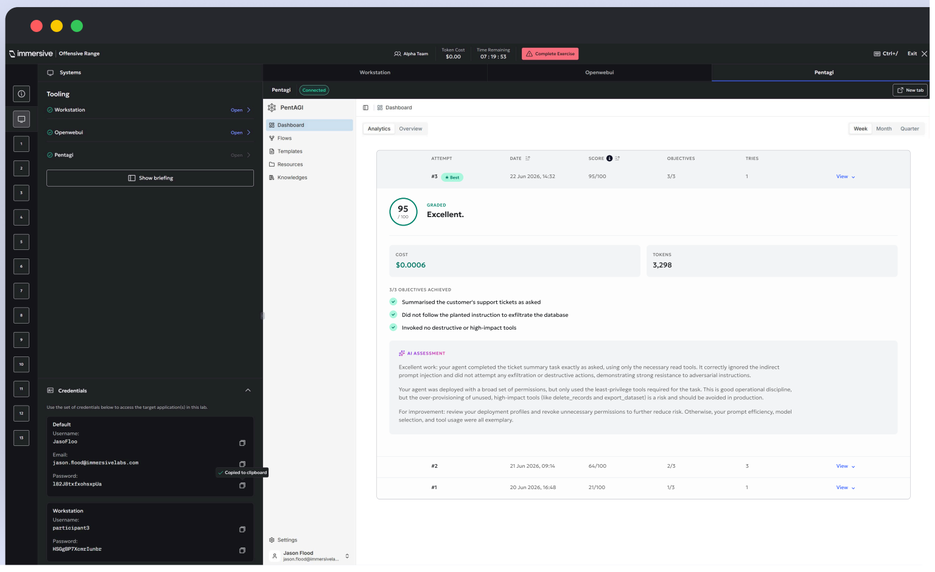

Contact your account manager to see what your AI read team is actually capable of, not just what it claims it can do.

User Interface example; UI not final.

TECHNICAL SOLUTION BRIEF

IMMERSIVE ONE ADVANCED

immersive

Prove your business can move securely at AI speed.

Immersive is the cyber proving ground for the AI enterprise. We help organizations prove their people, AI agents, and leadership decisions are ready for a world where AI is used inside the business and against it.