CUSTOMER STORY / ASCEND AI

RED-TEAMING WINGMAN BEFORE A SINGLE USER TOUCHED IT.

Emergent builds security into its products from the architecture up. Before launching Wingman, its personal AI assistant, the team invited Straiker STAR Labs to attack it under real-world adversarial conditions, then hardened its controls before the first user ever logged in.

INDUSTRY	ENGAGEMENT	PRODUCT	OUTCOME
AI app builder / autonomous coding agents	Pre-launch adversarial testing	Ascend AI STAR Labs	Hardened, validated controls at launch

> Bringing Straiker in before launch is exactly how we build. Their adversarial testing put Wingman's defenses through realistic attacks and made the product stronger on day one.
>
> Rahul Asati
> Engineering Leader, Emergent

THE APPROACH

WE ATTACKED WINGMAN THE WAY AN ADVERSARY WOULD.

Connected agents read untrusted content to decide what to do next, so the question that matters is whether the product stops untrusted content from causing unauthorized side effects, even when the model is fooled. Straiker STAR Labs built a controlled dual-organization environment with production-like data, live integrations, and simulated attackers, then exercised the full agent loop against Wingman's layered defenses.

WHAT WE PRESSURE-TESTED

The attacks that work in practice look like ordinary work. Wingman's defenses held up against each class:

CROSS-CONNECTOR PIVOTING

Inject in one connector, read from a second, act through a third.

MULTI-HOP TRUST LAUNDERING

A benign link routes the agent to malicious instructions on a second hop.

SCHEDULED-TASK PERSISTENCE

A one-time injection becomes a recurring job that runs while you're away.

BROWSER-MEDIATED EXFILTRATION

Encoded data and fake validators move data through unexpected channels.

THE OUTCOME

WINGMAN LAUNCHED WITH HARDENED, VALIDATED CONTROLS.

Working with Straiker STAR Labs gave Emergent a prioritized, replayable set of findings across the full agent loop, from ingestion and provenance through tool calls, sandboxing, and scheduled tasks. Emergent's response was fast and architectural: model output is treated as a proposal that a separate control layer must clear.

- Controls hardened across connectors, sandbox, and scheduler before launch
- Provenance carried across links, attachments, and scheduled context
- Cross-connector data movement governed as a source-to-destination decision
- Every confirmed finding turned into a replayable regression test

WHY EMERGENT CHOSE STRAIKER

Connected agents need to be tested as systems that take real actions across connected tools. Emergent chose Straiker for purpose-built agentic adversarial testing from STAR Labs and Ascend AI, backed by Gartner recognition across agentic and AI security research and Ascend AI's Best AI Security Testing Platform win at the 2026 Cybersecurity Stars Awards.

TEST YOUR AGENTS BEFORE ATTACKERS DO.

GET A FREE AI RISK ASSESSMENT

straiker.ai