# How MeshMesh cleared every stage of Salesforce's AI security review.

CLIENT
**Mesh Mesh, Inc.**
AI-native platform

SECTOR
**SaaS · AI Platform**
Salesforce AppExchange

ENGAGEMENT
**Embedded Security Function**
vCISO · Arch · Red Team · GRC

OUTCOME
**4 / 4 Stages Cleared**
No conditional approvals

PREPARED BY
Zivis · Jim Goldman & Jake Miller

THE ENGAGEMENT

# How MeshMesh cleared every stage of Salesforce's AI security review.

A multi-layered Salesforce security review — a sequence of deep technical evaluations, testing cycles, and stakeholder reviews, where AI platforms face the highest scrutiny. Timelines don't flex. Reviewers change between stages. The evidence bar only rises. Traditional penetration testing would have satisfied a checklist. It would not have answered the questions Salesforce security was actually asking.

| CLIENT | ENGAGEMENT | OUTCOME |
|---|---|---|
| **Mesh Mesh, Inc.**<br>AI-native platform | **Embedded AI Security Function**<br>vCISO · Architecture · Red Team · GRC | **Cleared Every Stage**<br>Multi-stage Salesforce review |

## Three pressures stacked at once

01 / TIME

### Timelines don't flex for security reviews.

As scrutiny around AI intensified, the bar kept rising — requiring consistent, high-confidence responses across repeated evaluations and testing cycles.

02 / SURFACE

### The new security perimeter isn't fixed.

With LLMs and autonomous agents, every API call, tool invocation, and downstream service becomes part of a shifting trust boundary most teams can't fully see.

03 / SCRUTINY

### Every stage brought new reviewers.

Initial security review, a Salesforce-appointed third-party assessment, and AppExchange review — each stage introduced new reviewers and raised the evidence bar.

## Why a pen test alone doesn't pass this review

A pen test answers a static checklist. An AI security review tests how your platform behaves when attackers treat the model, its agents, and their tools as the attack surface — a shifting perimeter no checklist or static threat model covers.

TRADITIONAL CHECKLIST

**A pen test answers**

- Can logins be bypassed?
- Are APIs authenticated?
- Are roles properly scoped?
- Can SQL be injected?

COVERED + EXTENDED

ADVERSARIAL AI ATTACK SCENARIOS

**Continuous adversarial testing**

PROMPT INJECTION
Adversarial prompting, instruction override.

TOOL IMPERSONATION
Tools and APIs weaponized through agents.

MULTI-HOP INJECTION
Injection chained across agents and systems.

TRUST BOUNDARY BREACH
LLM between authenticated traffic and data.

EXCESSIVE AGENT AUTONOMY
Actions beyond user intent or authorization.

DATA EXFILTRATION
Model outputs leaking protected data.

THE SOLUTION

# We didn't deliver a test. We became a security function.

Most AI security vendors ship one of four playbooks — Scanners, Pen Test firms, vCISO firms, or GRC firms. Alone, none clears a multi-stage Salesforce security review. **Zivis is all four, in one team** — in the room, in the code, in the policies, and in the reviewer's meeting.

## WHAT ZIVIS DID

01
### vCISO in every Salesforce review

Jim Goldman — Salesforce's first VP of Global Security GRC, now Zivis co-founder — joined every review call. Reviewers weren't being managed; they were speaking directly to someone who had built the function they operate within.

02
### Continuous, shift-left security

Same team reviewed each new feature's architecture before release, then ran adversarial testing after. Most engagements break down where these don't reconcile. We closed the loop inside one team.

03
### Pen testing + proprietary adversarial AI

Web, API, and LLM OWASP Top 10s executed in parallel — plus Zivis's proprietary adversarial AI taxonomy: context manipulation, tool impersonation, multi-hop prompt injection. Built on a living threat model of adversarial AI attack surfaces reviewers are now asking about.

04
### GRC-as-a-Service in procurement's language

Compliance controls documented and evidenced as fast as findings were remediated. No separate GRC vendor needed to translate pen test results. Every finding tracked to verified closure with retest evidence.

### Multi-stage Salesforce review — traced end to end

NO CONDITIONAL APPROVALS · NO REWORK


TIME TO SIGNAL
**< 48h**
Engagement to initial AI Security Blitz results

REPORT IN HAND
**~2 wks**
Across application and API layers

DEEP AI COVERAGE
**~4 wks**
Full proprietary adversarial scenario library

IN AN ACTIVE SECURITY REVIEW RIGHT NOW?

**Talk to Jim & Jake directly.**

zivis.ai/talk-to-us

Book a Call →