

When the detection fires, AI works the alert against your guidance and shows its reasoning. One large digital commerce company cut mean triage time from 1 hour to under 2 minutes.

What Detection Skills Unlock

Skills On the Detection
Triage, Investigation, and Tuning Skills, Attached to the Detection
The context your engineer had at authoring travels with the alert to every agent, so nothing gets re-investigated from scratch.

Triage returns escalate or suppress, with rationale

Investigation returns a verdict and recommended actions

Tuning returns exclude, include, modify, or fork

Open Skill Library
A Reusable, Versioned, Open Library
Detections inherit expert triage on day one from an open, MITRE-aligned Skill library, managed in the UI and via API. Start with 40+ Skills from Vega Research, then author only what's specific to you.

Bulk Attach
Operationalize Hundreds in One Action
Bulk-assign a Skill by package, source, or label, so expert triage and tuning reach hundreds of detections at once, not one rule at a time.

Transparency & Trust
Whitebox and Human-in-the-Loop
See exactly what the AI did and why. Every run shows its evidence and reasoning, and every tuning change ships only after your engineer approves it, so it stays defensible to an auditor or board.

A Large Digital Commerce Company: One Hour Of Triage To Under Two Minutes

By encoding triage, investigation, and tuning into Detection Skills on Vega, the team moved the work upstream: triage runs the moment a detection fires, and the noise is handled before an alert exists.

<2 min.

Mean triage time
Down from 1 hour. Triage Skills work the alert before it reaches a human.

90%

Investigation time reduction
Hours of manual investigation collapsed to minutes with defenders' judgment and context applied.

Pre-Alert

Where the noise stops
Detection Skills start triage and enrichment before an alert exists, so only what matters escalates.

Same Day

New detections, covered
Each one ships with expert triage and tuning the day it is authored, not weeks later.

Author Once. The Skill Runs Triage, Investigation, And Tuning On Every Fire.

1

Author the Skill
Attach your triage, investigation, and tuning logic to the detection: baselines, escalation criteria, enrichment, exceptions, and correlation. Written once, versioned.

2

Triage the Moment It Fires
The Skill travels with the alert. Triage skills apply your baselines and exceptions, then escalates or suppresses, with its rationale.

3

Investigate to a Verdict
The investigation agent inherits the Skill's context, builds the timeline, entities, and blast radius, and returns a verdict with recommended actions.

4

Tune and Compound
Tuning proposes exclusions and changes grounded in real verdicts. Your engineer approves each one, so every incident leaves the system sharper than it found it.

Architect A Whole World Of Agentic Reasoning Beyond Triage, Investigation, & Tuning

Attackers already share their tooling. Detection Skills let defenders' judgment compound the same way: encode it once, run it on every detection, and improve it in one place. Triage, investigation, and tuning are the first Skills, not the limit. Because the standard is open, your Cyber Defense Engineers can architect new kinds of agentic reasoning on the detection itself, tuned to your environment. It's live at [detectionskills.io](#), built on Anthropic's Agent Skills, and runs deepest in SAM.

Get a Demo

