# Deciding Build vs Buy for AI in the SOC.

Five questions that decide whether your team, or your team plus a partner, produces the better security outcome.

Crogl investigates alerts where your data already lives: on-premises, in your cloud, or air-gapped. Nothing leaves your environment.

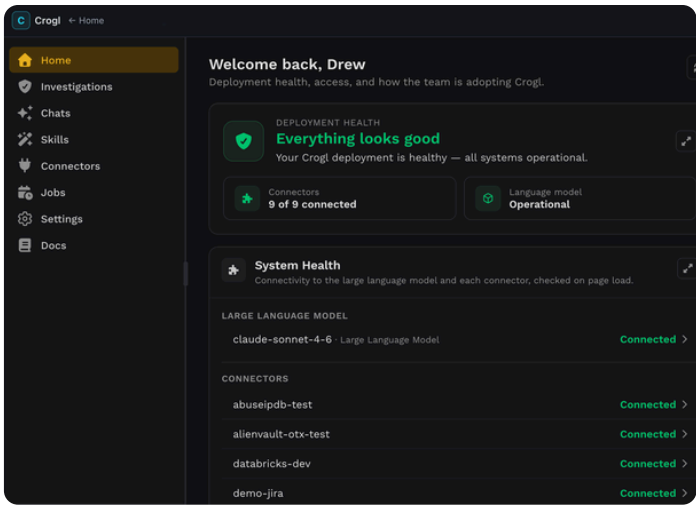


## Building for Operational Outcomes.

An engineer with API access and a few focused weeks can wire a language model to a SIEM and produce a convincing demo. If you are evaluating AI for your SOC, someone on your team has probably already done it, and they were right to. Making it operational requires sustained engineering effort.

Set aside the usual build versus buy argument about maintenance burden and total cost. The question here is narrower: which arrangement, your team alone or your team plus a partner, produces the better security outcome? Both answers are legitimate, and the choice is not all-or-nothing. Who builds what is an outcome decision too: where your team clears the bar, your team should build.

Known campaigns like Volt Typhoon pre-positioning and insider-risk compliance alerts keep piling into your queue. New ones like the July 2026 Hugging Face agent intrusion signal a class nobody has a playbook for yet. Your team's work is multiplying.

An operational outcome has three parts. Coverage: alerts closed with a documented finding. Quality: a finding a SOC manager or auditor can stand behind. Durability: both holding as your environment and the models underneath it change.

**95%**

failure rate for internal enterprise pilots attempting to drive measurable P&L return.

*MIT NANDA Initiative ('25)*

**380%**

cost overrun, due to underestimated data prep pipelines & compounding inference costs

*Industry Research via Uvik Cost Guide ('26)*

**$285,000**

average base salary for a Senior AI Engineer in the U.S.

*Glassdoor via Uvik Software ('26)*

PARTNERING FOR OUTCOMES

# Five Questions About Your Team & Requirements.

## 1 **Attention:** who spends their whole day on this?

The model is the smallest part of the system. Token hygiene, query construction, context management, and knowledge graph optimization are where the engineering hours go, and every model release moves the ground under all of it. Can your engineers keep that harness current?

## 2 **Craft:** can we build enterprise software?

Detection engineering is one discipline. Enterprise software is another: permissioning across humans, agents, and tenants, plus upgrade paths, failure modes, and backward compatibility at scale. Has this team shipped and operated that before?

## 3 **Trust:** who owns trust for model and tool calls?

Agents need lifecycle management, and trust has to hold across models, tool calls, API servers, MCP, and user identities. Guard rails written in a prompt are suggestions, and the model treats them that way. Is your trust in the architecture or in the prompt?

## 4 **Evidence:** how would we prove outcomes?

Change a prompt, a model, or a retrieval path and behavior shifts in ways a spot check will miss. Proving improvement takes held-out investigations, a rubric your SOC manager signed, and a record of what the system concluded. Can your team run that on every model release?

## 5 **Sovereignty:** who else sees our data?

Most AI security demos run against a hosted API. That holds up until the alert carries a patient record, a customer identifier, or CUI on a classified network. Ask what crossed the boundary, what was retained, and for how long. Can the vendor answer that in writing?

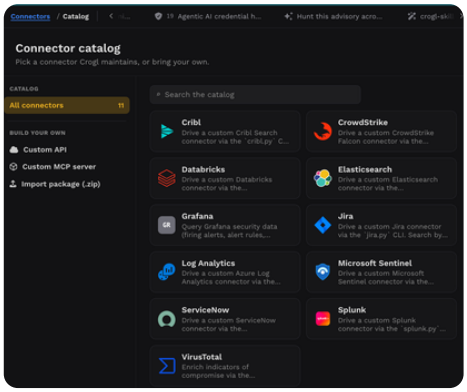


# Own the Problem. Buy the Engine.

### Only you know your workflows & priorities.

The USAF is a well resourced, capable security organization. Partnering with Crogl powers their security outcomes.

### Cost, continuity, & support are only the beginning.

The best outcome is when you have a strong partner who will provide the cost, continuity, and support benefits.

## Build With Crogl.

Score your own team against the five questions and answer: who gets you to operational outcomes?

**Download for free** | crogl.com/download