AUTONOMOUS SECURITY OPERATIONS · PUBLIC SECTOR

Built with their team: 98% of alerts now close on their own.

How Alaris worked hand in hand with a U.S. regional transit agency's team, learning what was benign and what was repetitive, then building the clustering and pre-queue workflows that lifted autonomous close to 98 percent. Told in its own numbers.

OUTCOMES · AT STEADY STATE · PLATFORM-REPORTED

98%

AUTO-CLOSE RATE

Confident benign alerts resolved without a human

~2%

REACH AN ANALYST

Only real suspicious evidence escalates now

0

SOC HEADCOUNT ADDED

The same team, a far shorter queue

20:1

DUPLICATES COLLAPSED

Recurring alerts investigated once, not twenty times

INDUSTRY	ENVIRONMENT	SECURITY TEAM	PRIOR POSTURE	PUBLISHED
Public sector, regional transit	Multi-site transit network, 24/7	Lean in-house SOC	Escalate when unsure	July 2026

IN THIS STUDY

MODULES DEPLOYED

- Detection Engineering
- Alert Triage + Investigation
- Threat Hunting
- Containment + Response
- Reporting
- Workflows

INTEGRATIONS IN SCOPE

Splunk	Microsoft Defender	Microsoft Sentinel	Entra ID	ManageEngine
SIEM	ENDPOINT	CLOUD SIEM	IDENTITY	ITSM · TICKETS

01

The situation

This is a U.S. regional transit agency: a public operator running a multi-site network that has to stay up around the clock, with a hybrid estate of corporate IT and operational systems, Splunk as the primary SIEM, and a lean in-house security team behind all of it.

Alaris was already returning investigated verdicts on every alert. The gap was not accuracy; it was confidence to close. Getting the agent to resolve alerts on its own meant teaching it what this specific team already knew: which alerts were benign here, and which ones repeated. So we worked alongside their analysts to capture that.

02

The pressure

FIELD NOTE

The knowledge existed. It lived in the analysts' heads, not in the system. The work was moving it into policy.

The team could tell at a glance which alerts were noise: the routine service-account activity, the same rule firing on the same host, the known-benign patterns particular to a transit environment. But that judgment was not yet encoded, so those alerts still landed on a person.

~40%

Share of alerts reaching a person before the work, most of them patterns the team already recognized as benign.

The repetition compounded it. The same benign events arrived again and again, and correlated alerts pulled otherwise-clear ones into review, so analysts spent their day re-deciding cases they had effectively settled long ago.

None of this was a detection problem. The team knew what was safe and what was routine. It just needed a way to act on that knowledge without a human doing it by hand every time.

THE ARITHMETIC

At a 60 percent auto-close rate, 4 in 10 alerts reach a person. At 98 percent, fewer than 1 in 40 does.

03

The deployment

HOW WE WORKED

This was hands-on. We sat with their analysts, worked the queue together, and encoded what they knew, one pattern at a time.

We started in the queue with their team. Working through real alerts together, we cataloged what was benign in their environment and what kept repeating, then turned each pattern into something the system could act on. The first lever was pre-queue workflows in Splunk: the alerts the team already knew were benign get filtered at the source, so they never reach an analyst at all.

Next we built out clustering for the repetition. Recurring and duplicate alerts collapse before the queue: a new alert is compared against recent open alerts from the same source over a 60-minute look-back, matched on title similarity and shared entities, and an agent judge confirms a genuine duplicate before collapsing it under a parent, so the agent investigates the cluster once. With the benign and the repetitive handled, we tuned the auto-close policy per tenant: on balanced, the agent closes up to medium severity at 75 percent confidence, judging each alert on its own evidence rather than inheriting suspicion from its neighbors.

AUTO-CLOSE POLICY · BALANCED · TUNABLE PER TENANT

98%	75%+	≤ MED	0
AUTO-CLOSE RATE	CONFIDENCE TO CLOSE	SEVERITY CEILING	CONFIDENT-BENIGN HANDOFFS

THE QUEUE, BEFORE AND AFTER

METRIC	BEFORE	WITH ALARIS
Agent auto-close rate	~60%	98%
Alerts reaching an analyst	~4 in 10	Fewer than 1 in 40
Known-benign alerts	Reached the queue	Filtered before intake
Recurring or duplicate alerts	Each reviewed	Clustered to one
Confident calls	Escalated	Closed on evidence

Detection did not change. What changed is that the team's judgment now lives in the system.

04

The result

METHODOLOGY

Figures reflect a steady-state window, platform-reported and measured against the deployment record.

Steady state looks unremarkable, which is the point. The alerts the team always knew were benign are filtered before intake, the repetitive ones collapse into a single case, and the agent closes what it is confident about. What reaches an analyst now is a short list of alerts that actually merit judgment.

WHAT CHANGED IN PRACTICE

01 Benign, filtered early

The patterns the team knew were safe are removed in Splunk before they ever reach the queue.

02 Repetition, clustered

Recurring and duplicate alerts collapse into one case the agent investigates a single time.

03 Their judgment, encoded

The auto-close policy reflects what the team knows, tuned per tenant and adjusted as trust builds.