ACCELERYNT

ROUTINE ALERTS VS. HUMAN DECISIONS: A CASE STUDY IN SOC AUTOMATION

FROM REPEATED TRIAGE TO DECISIONS WITH THE EVIDENCE ATTACHED

NOVEMBER 2026

©2026 Accelerynt

Contents

Executive Summary 3

The Starting Point:
Commodity Attacks, Human Triage 4

What We Found:
Where the Time Went 4

The Business Risk:
Why It Mattered 4

What We Changed:
Routine Alerts Handled Automatically, Decisions Kept With People 5

Tangible Business Outcomes:
Less Repeat Work, Faster Resolution 6

Accelerynt's Approach:
Automation Built Inside Your Tools 7

Conclusion:
Less Noise In, Better Decisions Out 7

© 2026 Accelerynt

Executive Summary

Accelerynt looks at why alerts keep arriving, so analysts can focus on the decisions that need human judgment.

At a global medical device company running Microsoft Sentinel, alerts reaching an analyst fell 75 percent, median time to resolve routine alerts went from almost 30 days to under 10 minutes, and about 500 analyst hours a week came off repetitive triage.

Those results started with a question about why the same alerts kept arriving. The answer pointed to a queue full of repeat work. Analysts investigated alerts from scratch, even when the same activity had already been reviewed and cleared. Simple fixes like revoking a session or blocking an IP address could wait weeks for a ticket to be approved.

Accelerynt automated the routine work on both sides of the human decision, gathering the evidence before it and carrying out the approved response after it. For routine alerts, that decision is made once, when a person approves how that kind of alert is handled. This case study shows that much of what reaches a SOC is simple, linear, and repetitive work. Finding out why it keeps arriving and automating it leaves analysts with the alerts that are actionable and unique.

Strategic Lessons for Security Leaders

Ask why the alert keeps arriving
Finding out why an alert keeps arriving shows whether you can reduce the noise coming in, so analysts aren't dealing with the same alert over and over.

Start with high-volume, low-risk alerts
Alerts that come up often and can be resolved with a low-risk action are where the value is, so they make the best first candidates for automation.

Name who is accountable, and keep a record
A person signs off on how each kind of alert is handled, and the automation logs what it does, so the team can go back and check.

Measure what your analysts stop doing by hand
That's how you show the change is working, and how you see it if the results slip.

© 2026 Accelerynt

The Starting Point: Commodity Attacks, Human Triage

The customer ran Microsoft Sentinel across its Microsoft environment, and analysts triaged most incoming alerts by hand.

Many of those alerts came from commodity attacks, the standard attacks that happen often in almost any environment, such as password sprays, credential stuffing, phishing and adversary-in-the-middle sign-ins. The 2026 Microsoft Digital Defense Report found that password spray attacks account for more than 99 percent of observed cloud initial access attempts.[1] These attacks follow familiar patterns, and they still need a response, because one left unanswered can succeed. Most of them are simple to resolve with a low-risk action, like ending a session or blocking an IP address. Getting to that action was the hard part.

What We Found: Where the Time Went

Accelerynt looked at what was coming into Sentinel and asked the same questions about each alert type. Why was this alert arriving, and what caused it? Was this work the team had already done? If so, how could it stop landing on analysts over and over?

The answers kept pointing to the same patterns.

Detections hadn't been tuned

Sentinel hadn't been tuned to how the customer's environment normally behaves, so the queue filled with noise. Expected changes looked like new activity, and a person had to look at each one.

Cleared activity came back

Activity the team had already reviewed and cleared kept returning as fresh alerts, including some of the customer's own configuration. Analysts reached the same conclusion about the same activity many times.

Context rebuilt by hand

For each alert, an analyst pieced together the same background again, such as who made the change and whether anyone had already looked at it. Gathering that context was the slowest part of triage. Responding to an attack comes down to how quickly you get to the evidence, and here the evidence was gathered by hand.

Simple fixes waited in the queue

Low-risk fixes like resetting a session token still needed a ticket and an approval. That wait could stretch to weeks or months, and median time to resolve sat at almost 30 days.

© 2026 Accelerynt

The Business Risk: Why It Mattered

The repeat work had two costs. The first was time. The repetitive triage added up to about 500 analyst hours a week, roughly the working week of a dozen full-time analysts, and that time couldn't go to more strategic work.

The second was exposure. While a fix sat in the queue, an attacker could:

- Stay signed in on a session flagged for impossible travel.
- Keep using a token captured in an adversary-in-the-middle sign-in.
- Keep running a password spray against other accounts.
- Log in with a password that worked in a credential stuffing attempt.

The CrowdStrike 2026 Global Threat Report puts the average eCrime breakout time in 2025 at 29 minutes.[2] That's how long it takes an attacker to move beyond the first foothold. A fix that waits weeks leaves that window open far longer.

Analysts were answering the alerts, but the fixes still waited on a ticket.

What We Changed: Routine Alerts Handled Automatically, Decisions Kept With People

Over two years, Accelerynt built about 30 production automations for the customer, starting with the alert types that carried the most repeat work.

Find the repeat work

- Each alert type was traced back to why it kept arriving. Where the cause could be fixed, it was, so analysts didn't keep dealing with it. Some causes were stopped at the source with playbooks. If malware infected a device, a playbook quarantined it so it couldn't reach anything else on the network. If an incident turned up indicators of compromise, dissemination playbooks pushed them to the other security tools that could block them. Either way, the attack was stopped where it started, before it could spread.
- An alert was a candidate for automation when the response was a low-risk decision that could be made from a set of criteria. Anything that needed judgment across systems stayed with an analyst.
- Commodity attacks were the first candidates. They produce high-fidelity alerts, so once a person had approved the response, each alert was checked against the criteria, and a person reviewed the results monthly.

© 2026 Accelerynt

Automate the evidence and the response

- The automations pulled the evidence together and closed out alerts that matched known-good activity.
- Anything left opened as a case with its context already attached, so an analyst could start with the decision.
- For the response, each automation took the smallest action that would solve the problem.

Keep the decisions accountable

- Before anything went live, a person reviewed and approved the response for each kind of alert.
- After that, nobody had to look at each alert of that type. It works like a garage door opener. Nobody approves the door every time you pull into your own driveway, because that decision was made when you got the remote.
- Every automated action was logged, along with who approved it.

Accelerynt started from its own library of integrations, automations, analytics rules and dashboards, customized each piece to what the customer needed, and built new playbooks where it had to. Some of Accelerynt's playbooks are published publicly on GitHub.

Tangible Business Outcomes: Less Repeat Work, Faster Resolution

There was a clear before and after. Because the biggest sources of repeat work came out first, the change showed up early.

75%

fewer alerts reached an analyst

Routine alerts were handled automatically, so more of what reached the team needed a person.

Under 10 min

median time to resolve routine alerts, down from almost 30 days

The fixes that had waited in the queue now ran automatically, and cases opened with the evidence already attached. With a median under 10 minutes, the typical fix now lands well inside the 29-minute average breakout time.

About 500

analyst hours a week off repetitive triage

That time became available for more strategic work.

A record of every automated action

The customer could show what each automation did, when, and on whose authority, straight from Sentinel.

© 2026 Accelerynt

Accelerynt's Approach: Automation Built Inside Your Tools

Everything ran inside the customer's own Sentinel workspace through Azure, with no new platform to buy, and the customer's analysts could read every rule.

How Accelerynt decides what to automate

- Accelerynt uses the same test with every customer: is the response a low-risk decision that can be made from a set of criteria? Each automation is then built for that customer's environment.
- Before building anything, Accelerynt measures the current effort and documents, with the customer's team, what known-good activity looks like. That shared definition is what the automation checks against.
- For each automation, Accelerynt estimates the analyst time saved and names the person accountable for approving the response.
- Incidents closed by automation, and the actions taken on them, are normally reviewed once a month. If a type of alert that usually shows up 50 times suddenly shows up 50,000 times, or a user complains about being logged out, someone takes another look sooner.

Conclusion: Less Noise In, Better Decisions Out

Look upstream at why the SOC is responding at all, and don't stop after the first automation. Teams that keep improving what they already have, and show the numbers, can move their analysts up to the decisions that need them.

IBM's Cost of a Data Breach Report 2026 found that companies using AI and automation in security operations cut breach costs by an average of almost $2 million.[3]

This work also prepares the SOC for AI agents. As AI agents start taking on SOC work, the alerts they handle need to arrive triaged and with the right information, because running agents on untriaged noise gets expensive.

A good place to start is your highest-volume alert types. If your team closes some of them the same way every time, those are the first candidates. If the work is simple, linear, and repetitive, Accelerynt can automate it inside the Microsoft tools you already own.

Automate the routine alerts.
Keep the decisions with people.

(855) 527-2252
info@accelerynt.com
accelerynt.com

Sources

1. Microsoft, *2026 Microsoft Digital Defense Report*, October 2026, p. 73.
2. CrowdStrike, *2026 Global Threat Report*, February 2026, p. 9.
3. IBM, *Cost of a Data Breach Report 2026*, July 2026.

© 2026 Accelerynt

Accelerynt is a Microsoft-native security partner for enterprise organizations. We operate inside the Microsoft tools our customers already own, validating configurations and controls against each organization's approved baseline so problems are caught before an attacker can use them. Through managed detection and response, security assessments and the Accelerynt Security Platform, our practitioners help security teams get stronger over time. Accelerynt is employee owned and based in Plano, Texas.

©2026 Accelerynt

ACCELERYNT

accelerynt.com