Trellix

WHITE PAPER

No Alert Left Behind — Conquering Alert Fatigue with GenAI

Win the Challenge to Get to 100%

Take the No Alert Left Behind Challenge

- 100% of alerts investigated
- In under 3 minutes
- With prebuilt playbooks
- With a full audit trail
- From all your security data sources

Take the challenge on Trellix.com

Cybersecurity faces a scalability crisis when it comes to investigating alerts, one that's largely of its own making.

To catch threat actors, it's essential to collect huge amounts of data from as many sources as possible, and then alert on that data. This approach has seen great success at uncovering significant cyber incidents, such as the infamous SolarWinds breach. At the same time, business stakeholders are asking the security organization for assurances that they are responsibly addressing the threat landscape.

Yet, as data volume increases, the solutions that have brought success in the past begin to falter. An organization may have a billion events come through. Using today's tooling, that organization may be able to whittle that number down to 1,000 alerts—an impressive decrease. But, those 1,000 alerts are still too many for a human SOC team to investigate. Consider informational alerts, such as password resets. A password reset may be a completely normal business event, but in the right circumstances it might be really important. You simply don't know.

Understanding if that password reset is benign or malicious takes many additional data points and context. The problem is that when you have 1,000 potentially interesting alerts, you don't have enough time in the day to investigate all of them. One alert may have 60 events that go along with it. To understand the context, you must read everything associated with that alert. Very quickly, it becomes impossible to scale. On average, SOC teams investigate a small proportion of their total alerts, leaving 80%-90% unseen. The dots that could tell you whether that password reset is simply a forgetful executive trying to log in or the start of a stealthy attack never connect.

Getting to 100% with GenAI

Artificial intelligence (AI) and machine learning (ML) have been present in cybersecurity for decades. Recent advances in AI and generative AI (GenAI) have reached the point where AI capabilities can read every alert and tell humans if there's something important they need to look at.

By deploying AI in the background to do work on your behalf, you are actually getting hours of human work accomplished. It's no longer necessary for a human to make 20 clicks just to figure out what happened around one alert. AI handles this machine-level work for you. The work that is data-intensive, fatiguing, and error-prone from a human perspective is exactly what AI excels at—leaving the interesting investigative work for the human SOC team.

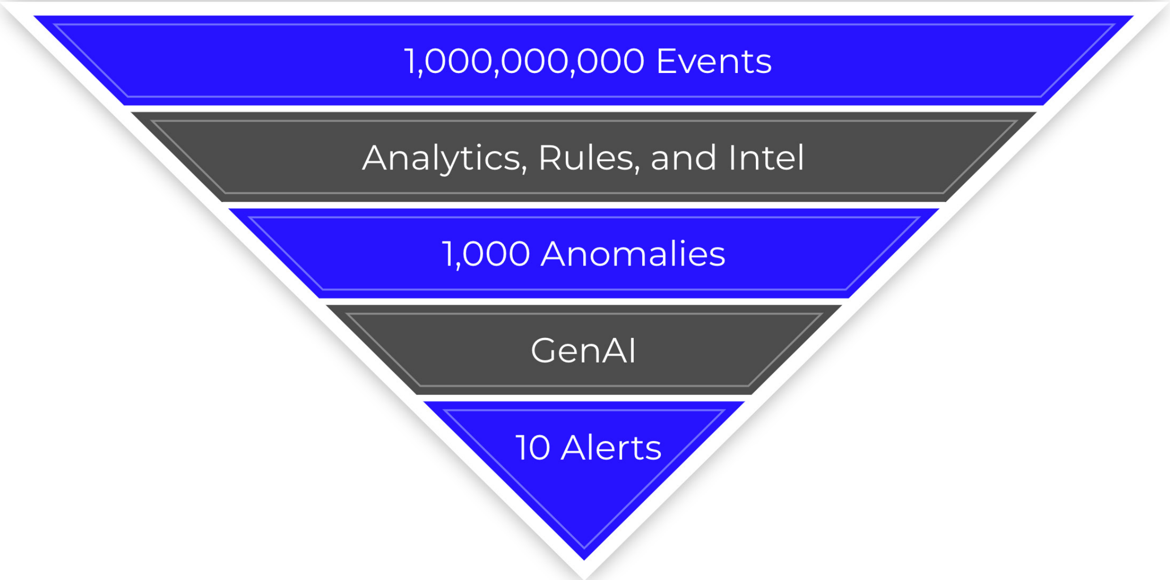

Figure 1: How Trellix Wise GenAI capabilities reduce alert fatigue so SOC teams can focus on what matters.

Protecting your organization becomes an exercise in shifting increasing amounts of work to AI-guided investigations so that your security analysts can focus on the most complex threats. The challenge is to increase the percentage of alerts your team is able to handle. From whatever your starting point, the goal is to get to a point where no alert is left behind.

We've gathered five key questions to guide you on the journey, and we'll look at each in detail below.

- Do you have 100% of alerts covered today?
- Can you investigate every alert in under 3 minutes?
- Do you have prebuilt playbooks for every alert?
- Do you have a full audit trail of investigations?
- Do your investigations cover all your security telemetry?

Do you have 100% of alerts covered today?

We get it, there are far too many alerts. That's why many SOC teams only investigate approximately 10% of their alerts, and those may not even be the most critical. You could spend a lot of time tuning out false positives and writing better rules, but that is a losing battle as new behaviors trigger new avalanches of alerts. But GenAI makes investigating 100% of your alerts possible.

GenAI can answer the questions traditionally answered manually by a SOC analyst and prioritize what matters most for faster investigations.

Pairing extended detection and response (XDR) with GenAI lets you pack a one-two punch and get even more from your alerts. XDR stems the tide of alerts by suppressing false positives and correlating events across security tools; GenAI can then work faster, with the right data, to visualize the complete threat story.

Can you investigate every alert in under 3 minutes?

Be wary of GenAI solutions that claim broad coverage of alerts with no timeline. We are still in the early days of GenAI in cybersecurity and there are numerous vendors jumping on the bandwagon who show outcomes without measurements.

Scaling is an issue for some vendors as the backend processing of thousands of alerts and the ability to ingest them rapidly requires massive computing power and the right AI models. For the biggest impact, your GenAI should give every alert a full investigation within minutes. With the right models, threat intelligence, analytics, and backend processing power this is possible in as little as three minutes.

Do you have prebuilt playbooks for every alert?

GenAI needs guidance based on leading threat-hunting and incident response (IR) practices. Your GenAI should know the right questions to ask because it's infused with playbooks and analytics backed by more than a decade of experience in the complete IR workflow. Ideally you would like to have the wisdom of the world's leading IR experts built into those workflows.

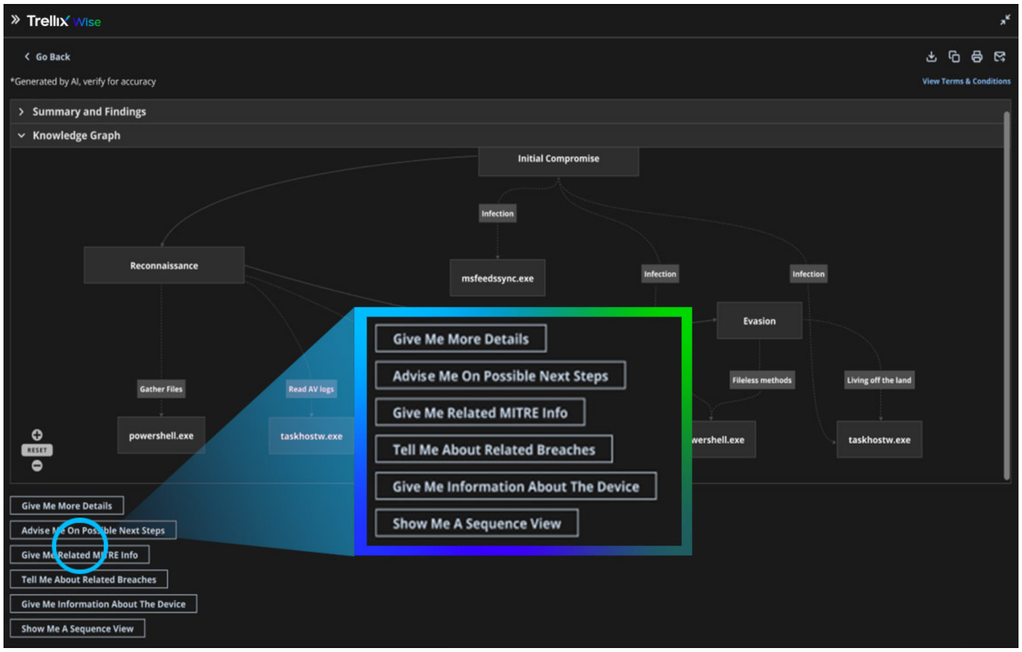

*Generated by AI, verify for accuracy

Figure 2: Trellix Wise offers AI-powered guidance.

For over 10 years, Trellix® has set the stage with tips for alert investigations that follow the workflows a human analyst would perform. These tips ensure that analysts have guidance built on best practices from experts who have performed thousands of threat hunting and incident response scenarios, providing an ideal foundation for automating workflows and guiding analysts who may not have experienced a specific attack.

Do you have a full audit trail of investigations?

Trust is a critical piece in getting to 100% alert investigation. You should be able to audit how GenAI is using your data, so you can trust its analysis.

Can you say with confidence that the AI is making the decisions that you would make if you had time to look at every single alert? Having a full, transparent audit trail that reveals exactly what information the AI was considering and what decision it made is how you build that trust.

An audit trail helps you verify the chain of command on your sensitive data, ensures data and applications access only what is required, and ensures transparency into the prompts, responses, and actions taken to provide accountability for the business.

Your data isn't just valuable to you; it is valuable to GenAI service providers seeking to build better models. Without a complete audit trail, you risk lawsuits and privacy violations.

Protecting data isn't the only value a complete audit trail provides. Seeing what decisions GenAI made and why helps upskill your team. Tracking the speed at which investigations were performed is also beneficial .Budget holders can report on time and efficiency gains to prove a return on their investment.

Do your investigations cover all your security telemetry?

For the broadest coverage, have GenAI use all your telemetry and data sources. Every attack includes multiple vectors, such as inbound traffic at your firewall/cloud perimeter, network, email, endpoint, identity, and more. While investigating telemetry at each one of those points is useful it does not provide the complete story or timeline of an attack scenario. Endpoint telemetry may provide a wealth of knowledge, but how does that correlate with the phishing email one of your users clicked or the exfiltration of data from a cloud storage repository? You need to be able to stitch all of these telemetry sources into one complete story.

When you have line of sight to all the steps an adversary took, analysts benefit from the complete story of the attack. Building a complete picture of the attack lifecycle based on the industry-standard MITRE framework ensures that analysts can understand how an attack unfolded and be able to determine steps for remediation.

Trellix Wise: Ensuring that no alert is left behind

Ensuring that every alert is investigated can seem like a daunting task. It’s possible today thanks to Trellix Wise, our capability for advanced analytics, includes all our ML, AI, and GenAI capabilities that our customers have leveraged for over a decade.

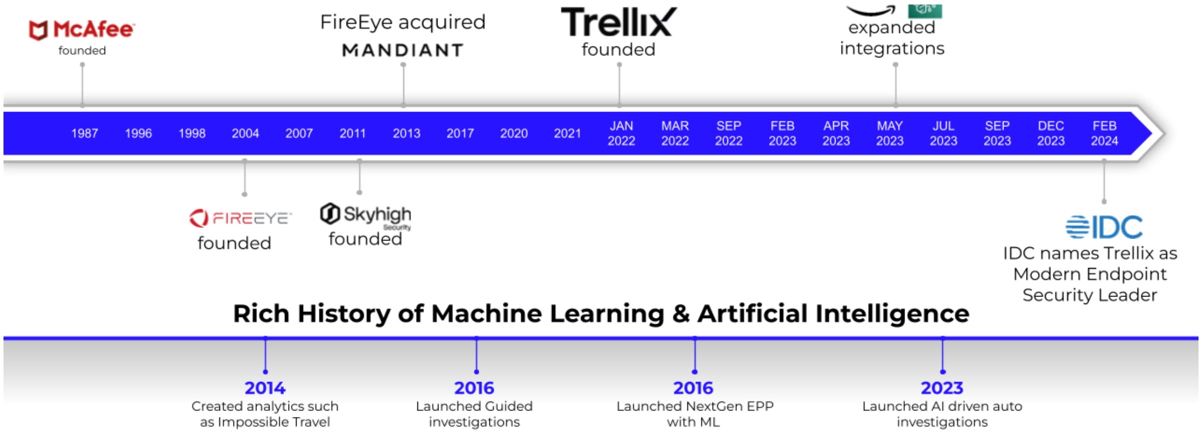

Figure 3: Trellix Wise is built on over a decade of AI modeling and 25 years in analytics and machine learning.

Built on the principles of transparency, choice, and responsibility with over a decade of AI modeling and 25 years in analytics and machine learning, Trellix Wise capabilities relieve alert fatigue and surface stealthy threats. Wise leverages ML, AI, and GenAI-powered investigations to deliver a revolutionary increase in efficacy and coverage while lowering the cost and skill set required to stop attacks. Wise leverages 3x more third-party integrations than competing solutions and delivers real-time threat intelligence leveraging 68 billion queries a day from more than 100 million endpoints.

Wise is embedded in every component of the Trellix portfolio to ensure that:

- 100% of your security alerts can be fully investigated
- Each investigation is performed in under 3 minutes
- Every investigation leverages prebuilt playbooks imbued with the wisdom of the world’s leading threat intelligence and IR experts
- You can trust the results of Wise investigations and use them for upskilling analysts with a full audit trail of the actions performed by Wise
- Each security alert you receive, regardless of whether it is from a Trellix native control or a third party source is investigated

With the power of Trellix Wise and the comprehensive portfolio of endpoint, network, data security, email, and XDR you’re able to overcome the challenges of alert fatigue, surface stealthy attacks, and ensure that No Alert is Left Behind.

Take the Next Step

- Request an assessment of your environment
- Learn more about Trellix Wise