tierpoint

WHITEPAPER

The True Cost of Alert Fatigue

How AI Reduces 10,000+ Daily Alerts to Actionable Intelligence

Modern SOC teams face an unrelenting flood of security alerts, often exceeding 10,000 per day, with most turning out to be false positives. Each alert still demands time, tools, and analyst focus, driving up investigation workloads and costs while stretching teams thin. The impact extends beyond wasted hours: real threats slip through the noise, analysts burn out, and maintaining around-the-clock coverage becomes increasingly difficult. The result is a widening gap between detection and response capabilities that directly increases an organization's exposure to risk.

This whitepaper quantifies the operational and financial toll of alert fatigue, breaking down the true cost of manual triage across labor, breach recovery, and resource retention. It then examines how AI-powered triage cuts through alert noise by correlating signals across identity, endpoint, and network data to surface only high-fidelity threats. The analysis includes performance benchmarks, ROI estimates, and case studies showing how leading SOCs have reduced investigation times from days to minutes while reclaiming millions in productivity.

The Cost of Alert Fatigue

Security teams are facing an alert overload crisis. Enterprise Security Operations Centers (SOCs) receive tens of thousands of alerts each day, with very high false positive rates. The volume and inaccuracy mean that nearly two-thirds of these alerts are ignored, leaving room for genuine threats to slip through the cracks.

But the impact doesn't stop at the SOC. When alerts and tickets spill beyond the security team, IT service delivery and customer-facing teams absorb the noise. Responding to security alerts is not their primary responsibility; they are already balancing projects, outages, business-as-usual changes, and user support. As a steady stream of moderate and low urgency security tickets arrives, most of which ultimately resolve as blocked threats with no action required, teams quickly learn that the majority are non-events. Over time, this creates alert fatigue outside the SOC, increasing the risk that truly critical security actions are delayed, deprioritized, or approved without proper scrutiny as teams become desensitized.

This overload also creates engineering and operational drag. SOC analysts must constantly retune tools amid the noise, reducing time available for proactive work like detection tuning and threat hunting. Meanwhile, service delivery teams see slower response times, disrupted workflows, and unnecessary ticket volume. This all erodes trust and efficiency across the organization.

The Financial Impact

Alert fatigue creates a significant financial and operational burden for large enterprises. Investigating low-fidelity or false-positive alerts consumes substantial resources through analyst time, tool upkeep, and constant rule tuning. When genuine threats are buried in this volume of noise, the resulting breaches can carry far greater costs, from regulatory repercussions to lasting reputational damage.

The strain on teams compounds the problem. High turnover, long hiring cycles, and chronic burnout leave fewer analysts managing ever-expanding toolsets, each generating its own stream of alerts. The outcome is a cycle of fatigue, attrition, and inefficiency that weakens detection and response capabilities over time.

Breaking that cycle requires making alerts more meaningful and reducing the noise that drives these costs. AI-powered triage helps focus analyst attention on validated threats, improve efficiency, and strengthen both operational performance and security outcomes.

Why AI is a Great Solution

Automated Triage & Intelligent Response

For analysts who are overburdened with alerts, AI provides some much-needed relief. Artificial intelligence can reduce alert volume through:

- **Contextual AI architecture:** Instead of just looking at events in isolation, AI can establish a "normal" baseline for devices, applications, and users. Behaviors are considered in context, so when something happens outside of the norm, it can be flagged. AI tools can also understand which assets are most critical and prioritize alerts based on that criticality.
- **Intelligent filtering and correlation:** Similarly, intelligent filtering can contextualize what is likely to be a true positive versus a false positive and automatically improve accuracy over time. Low-fidelity, noisy alerts can then be suppressed in favor of actual threats. AI can also correlate alerts, connecting what may seem like low-level warnings across systems to uncover larger attack chains. These methods can significantly reduce manual workloads.
- **Alert enrichment:** AI can also automatically enrich alerts with threat intelligence and historical context, not just providing the alert itself but the "why" behind its importance. Systems may even be able to solve low-risk alerts automatically, without the need for manual intervention.
- **Identity-Centric Threat Detection:** A significant share of today's alerts comes from identity systems, and many are low-value authentication events. AI reduces this noise by analyzing user behavior, authentication patterns, and privilege changes to distinguish routine activity from true indicators of compromise. By correlating identity signals across cloud and on-prem environments, AI surfaces the small number of identity-driven threats that matter while suppressing the bulk of benign access alerts.

Integration & Architecture Considerations

AI only meaningfully reduces alert fatigue when it integrates cleanly into the existing security stack. Modern platforms ingest telemetry from SIEM, EDR, identity, and network tools to build a unified view of activity and apply consistent detection logic across systems.

The most effective solutions also provide transparent, explainable reasoning so analysts understand why alerts are suppressed or escalated, preserving trust and auditability. When analysts stay in control of high-impact decisions, AI strengthens the SOC's effectiveness without adding complexity or becoming another source of noise.

AI Governance & Explainability

As AI takes on a larger role in triage and response, security managers and leaders must ensure these systems operate transparently and within clear governance guardrails. Effective platforms explain every suppression, escalation, and correlation decision so teams can validate outcomes, document actions, and stay audit-ready.

Robust governance should cover documented model behavior, data handling practices, and oversight controls that keep AI-driven decisions aligned with regulatory requirements and internal risk policies. With these structures in place, organizations can deploy AI with confidence while maintaining full visibility and accountability for how security decisions are made.

Speed & Efficiency Gains

AI and machine learning (ML) integration into cybersecurity operations can automate the most time-consuming, low-value tasks, giving human analysts more time to hunt the most critical threats and engage in more proactive strategic activities.

As organizations mature, the role of machine learning in security operations follows a clear progression. Early efforts focus on basic filtering and suppression to eliminate redundant or low-fidelity alerts. As maturity increases, platforms advance to a managed stage where AI establishes behavioral baselines, understands asset and user context, and automatically enriches alerts with relevant intelligence to support faster triage.

At more advanced stages, AI introduces inference and dynamic risk scoring, reasoning across multiple signals in real time to distinguish normal behavior from suspicious activity or likely threats. This allows the system to autonomously handle low-level investigations and trigger appropriate response actions, reserving analyst time for high-impact cases. As a result, organizations can shrink response times from hours or days to minutes, driving measurable improvements in metrics like mean time to detect (MTTD) and mean time to respond

What This Actually Means for Businesses

Productivity Gains

First, security teams can become much more productive. Reduced investigation time can lead to organizations being able to reclaim millions annually. This allows existing analysts to focus on high-value projects. Organizations are also avoiding costs associated with breach exposures. Every minute an attacker has infiltrated your systems, costs can increase.

Gains in productivity can also reduce turnover. Analysts who aren't bogged down by alert fatigue are less likely to experience burnout, which can improve satisfaction with their job and lead to lower turnover rates. Instead of paying for all-new tools, the right AI integrations that work with existing stacks can make the most of existing extended detection and response (EDR), firewalls, and security information and event management (SIEM) solutions. AI can serve as the central node and brain for the rest of the stack, making alerts and data more useful.

Operational Efficiency

The continuous, high-quality detection and responses enabled by AI mean that SOCs can meet their service level agreements (SLAs) for handling incidents, which can be a requirement for regulatory bodies and key internal stakeholders. AI-enabled tools also unify multiple feeds, providing better visibility. Managing low-level alerts and prioritizing the biggest threats can make it easier for lean teams to sustainably manage 24/7 operations.

AI's direct impact on alert fatigue can be seen in real KPIs that Chief Information Security Officers (CISOs) already track, including:

- **Mean Time to Detect:** Average time between when the incident begins and when it is detected
- **Mean Time to Respond:** Average time between the detection and containment or remediation
- **False Positive Rate:** Percentage of security alerts that are not threats
- **Cost per Incident:** How much it costs to manage one security incident, which can include downtime, forensic work, and labor costs
- **Coverage-to-Resource Ratio:** The amount of coverage achieved compared to analyst hours

Real World Examples

Many industries can benefit from AI-powered alerts. The following are a few different examples of scenarios where AI can reduce alert fatigue and prevent costly incidents from causing major harm. In the ICU, alarms can be extremely informative, but they can also turn into background noise when they happen too frequently. A recent study suggested the implementation of AI-driven alarm alternatives to emphasize what requires the most urgent action. Many healthcare providers receive dozens of alerts per day. Workers in the Veterans Affairs (VA) system receive over 100 daily alerts. A nationwide initiative with the VA did lead to a small reduction in alerts, but using AI to reduce these signals can reduce burnout and the frequency of errors.

Financial Services High false-positive rates in financial services firms can slow down fraud protection. Traditional rule-based systems in anti-money laundering (AML) can have false-positive rates as high as 95%. Reducing the number of alerts can significantly shorten MTTD and MTTR.

Manufacturing Manufacturers often receive large volumes of operations-related security alerts from sensors, controllers, and connected equipment. Many of these signals reflect normal process variance rather than malicious behavior. AI can filter out routine noise and surface true cyber indicators, such as unauthorized controller changes or abnormal movement between operations networks and IT systems

Energy Energy providers generate constant security alerts from grid systems, pipelines, and remote monitoring equipment, many of which reflect expected fluctuations. AI helps distinguish routine operational activity from genuine threats, including unauthorized access attempts or suspicious changes in control-system data.

How TierPoint Adapt Platform Helps

TierPoint Adapt Managed Detection and Response (MDR) helps minimize alert overload with automated triage and SOAR-driven (Security Orchestration, Automation, and Response) orchestration that cuts through noise and highlights high-impact threats often missed by traditional systems. Within TierPoint's SOC, alerts are continuously tuned and triaged so downstream IT and service delivery teams only receive cases that truly require action. This approach helps prevent alert fatigue, ensuring security issues don't get lost among other operational demands. By combining AI-driven threat detection with human expertise, Adapt MDR helps teams move from reactive firefighting to proactive protection.

Adapt MDR integrates with your stack, enriching it with experts who can turn existing alerts into real action. AI-powered tools can form correlations across identity, endpoint, and network signals, and customers can build trust in their systems through actionable insights, reporting, and regular advisory touchpoints.

Reach out to speak with a TierPoint specialist today!

About TierPoint

TierPoint (tierpoint.com) is a leading provider of secure, connected IT platform solutions that power the digital transformation of thousands of clients, from the public to private sectors, from small businesses to Fortune 500 enterprises. Taking an agnostic approach to helping clients achieve their most pressing business objectives, TierPoint is a champion for untangling the complexity of hybrid, multi-platform approaches to IT infrastructure, drawing on a comprehensive portfolio of services, from public to multitenant and private cloud, from colocation to disaster recovery, security, and more. TierPoint also has one of the largest and most geographically diversified U.S. footprints, with dozens of world-class, cloud-ready data centers in 20 markets, connected by a coast-to-coast network.

Learn More

Call 844.267.3687 or email sales@tierpoint.com to connect with one of our advisors. Or visit us at tierpoint.com to learn more.

tierpoint

Implementation Roadmap Advice

GETTING STARTED

Implementation starts with understanding your baselines. What are your current alert fatigue costs and your starting metrics? This can include alert volume, false positive rates, mean time to respond, and investigation hours, and any other metrics you may be looking to improve. Next, identify any workflow inefficiencies, including tool fragmentation. Which parts of your stack are integrated, and what needs to flow better? Finally, SOCs should assess their AI maturity. How well have they integrated AI into their flows so far, and what can use improvement or support?

SELECTION & DEPLOYMENT

The solutions you choose should meet the criteria that suit your business needs, including:

- ✓ AI/ML capabilities that reduce alert fatigue
- ✓ Integrating with your existing security stack
- ✓ Scalability that allows the solutions to grow with our business
- ✓ Automated response capabilities that lower the frequency of manual intervention

Deploying new AI tools doesn't have to happen all at once. A phased approach can help you gradually integrate new efforts, train team members to become early adopters and advocates, and optimize for improved accuracy along the way, focusing on the KPIs you want to reach.

It's also important to pick vendors with transparent ML models. The vendors you're considering should be able to share documentation on their training data and model architecture. They should also have clear AI governance processes, as well as explainability and security practices for their AI solutions.