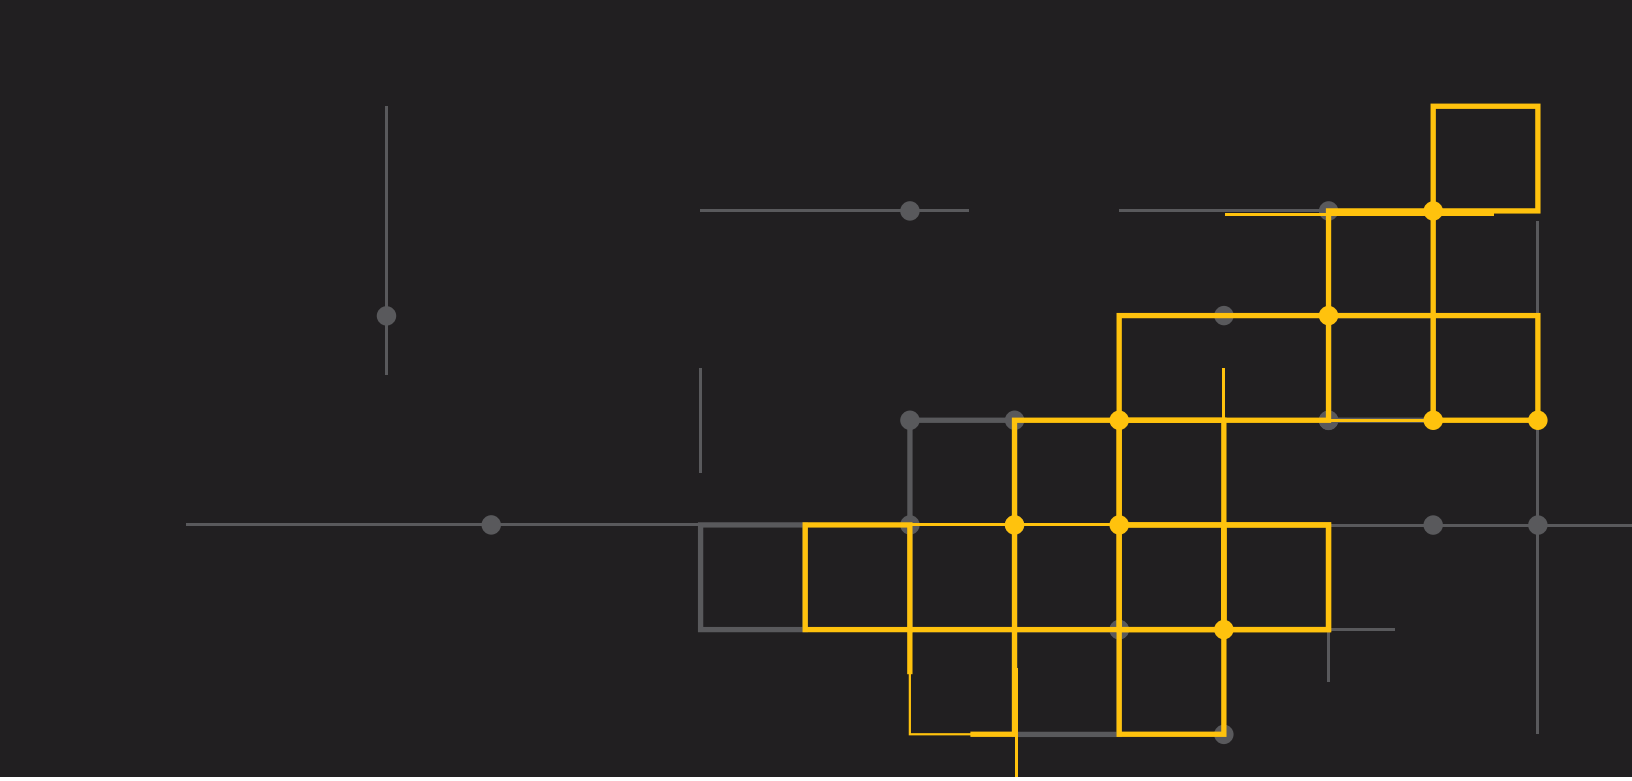

The Value of Enhancing Information Security and Incident Response with Automated Anomaly Detection

Symantec

Executive Overview

To build a truly comprehensive and effective information security and incident response program, organizations need to achieve two key objectives:

- Understand their IT environment and what constitutes normal and abnormal behavior.
- Train the information security and incident response staff and provide them with tools to analyze which behaviors point to real, viable threats.

There's simply too much going on across today's IT networks for information security and incident response teams to manually parse; a month, week or even a day's worth of complete traffic that crosses an organization's network can prove impossible to manually analyze.

The challenge becomes even more complex when considering the combination of individual and distributed services running on the typical IT infrastructure. In addition to industry-standard Windows systems, there may be a mix of mainframe databases and transactional systems, midrange UNIX systems, networking and security technologies, Linux workloads, and virtual machines. Adding to the complexity, individuals tap into these systems using varying endpoint devices with a wide range of operating systems and apps – desktop PCs, Macs, laptop computers, smartphones and tablets. The growing use of cloud services certainly adds to the challenge as well.

There's also the context of the traffic crossing networks that IT organizations need to consider. It's not only about the infrastructure that security teams have to monitor, but also the new behavior of users and the greying line between corporate network and employee access from anywhere using any device. Users' private worlds can encroach into the network. Because organizations don't control users' outside activity, that means potential harm to what the organization does control – which is all the more reason to know what's crossing your network and to retain a historical record of that traffic.

The task of behavioral analysis typically requires a costly investment of human expertise and effort. It's impractical to proactively sift through all the packets crossing the network in a meaningful way. For internal information security and incident response teams and service providers that manage incident response for their customers, keeping on top of what is happening moment-to-moment is literally impossible.

In this white paper, we examine the challenges that information security and incident response teams face when it comes to protecting their IT infrastructure as well as their customer and employee data. It's not enough to simply rely on basic endpoint and/or basic network forensics to detect, block and remediate attacks – your company also needs to proactively recognize evolving new attacks and prepare for the unknown. When used properly, automated anomaly detection closes the window of exposure by alerting you to anomalous, and potentially suspicious, activity much more quickly.

The key to protecting infrastructure and data is to adopt a two-pronged approach where you first establish a baseline against which anomalies can be detected. The baseline must necessarily be revised as behavior gradually changes on a daily, weekly, or monthly basis, either due to changing IT infrastructure or due to changes in user behavior over time.

From there, in-depth, sophisticated analysis is required. Yes, "doing the math" is required to accurately identify abnormal behavior, but the baseline is just the start. User behavior must also be analyzed within the context of what's going on at that point-in-time on your IT network, determine if the behavior matters, and then relate the activity to other time intervals in your data set. Only then can an accurate picture of what's happening be painted.

An apparent anomaly on Monday at 1pm, for example, when investigated further in isolation, may not be an anomaly at all. The same "anomalous activity" could very well occur every Monday at 1pm and therefore could be categorized as normal activity. A baseline could be established that predicts or identifies 1pm on Monday as a normal timeframe for that type of behavior.

However, when using baselines to detect anomalies, analysts and/or the automated systems they use must keep in mind that an activity occurring weekly on Mondays @ 1pm could have an authorized move to Fridays @ 3pm, which would require an adjustment be made to the baseline.

Taking an approach where you dynamically enhance incident response by correlating with automated anomaly detection allows you to weed out false positives – and ensure your information security and incident response team focuses on the major risks that threaten your business and the priorities that can change at any given moment.

Preventing Known Threats Not Enough: Also Detect and Prepare for the Unknown

Most organizations realize prevention alone in today's threat environment is not enough. Yet breaches still go undetected for days, weeks, months or even years. That's because signature-based systems and network management tools used in siloes are not a sufficient means for detecting today's advanced threats. Even when using a security information and event management (SIEM) system, there's simply too much information to sift through, making it humanly impossible to remove the false positives and make sense of the viable threats.

IT security and incident response teams need to know which alarms matter to the organization. Only then can they focus on the true threats facing their critical systems, applications, processes and data. If time is wasted researching numerous alarms that "could" apply to the organization—rather than those alarms that "do" apply – a truly harmful threat might be missed.

The goal is to essentially enable faster responses and reduce the window of exposure to the organization. It's relatively "easy" to manage responses when you know exactly what to look for, because you've seen it before and know how to handle a threat when you see it. The challenge comes when something new infiltrates the network and you have no experience dealing with it. If you don't expect it, you likely won't know the best way to respond.

Anomaly detection is about enabling proactive incident response by giving the incident response team the ability to hunt for potential risks proactively. By utilizing analysis to hunt intelligently, the incident response team can focus attention on anomalies that most likely represent actual threats.

To conduct proper incident response and forensics, you still need to capture "everything" – application data, system log data, and of course, network data – it's a must in today's environment. Until you start investigating a threat, you can't be sure what information you might need.

According to the SANS Institute, capturing "some" packet data is not sufficient. At a minimum, organizations should capture 30 days of packet data, and 60 days of data is even better.[1] The more an organization captures and keeps, the better chance they have of combatting the slow-and-low threats that organizations are experiencing.

The key is to leverage existing threat intelligence against the data you capture. When anomalies are identified in the data, rich metadata, and external threat intelligence, can be leveraged to provide much needed context. Activity that looks normal at 2pm on a weekday – when a batch routine is in process – may represent a major threat if you see the same activity at 2am on a Sunday. Focusing solely on the capture of raw packets crossing the network, for example, is not that useful in detecting anomalous user behavior.

The Challenges: How Previous Anomaly Detection Attempts Fall Short

- **Eyes on Dashboards** – humans are good at spotting simple anomalies in single variables of data on a graph. But this manual approach is lossy and time-consuming, and human errors due to emotions and judgment can produce ineffective results. In addition, no single metric will indicate an advanced attack. While multiple metrics may identify an advanced attack, humans ultimately can't hold enough related items in their memory.
- **Thresholds and Rules** – although this approach may benefit from some level of primitive automation, it also has a lot of drawbacks. For example, thresholds and rules do not work at all on periodic data. Attempts to modify these rules to address periodic data are extremely time-consuming and brittle. The alerts this approach generates can also create a lot of unnecessary noise that distracts the attention of security information and incident response teams.
- **Static Models** – which rely on supervised or trained machine learning usually focus on data that is old, grows stale and thus provides an inaccurate picture of the threat landscape.

Information security and incident response teams also need to consider what happens when the norm is no longer the norm? When an analysis baseline is set, it may get old within one month, one week or even one day, depending on the environment.

Response teams may have some context for an alert, but do they have the 'right' context? Does the context represent a lowest-common-denominator across all organizations, or is it directly relevant to their organization and the threats they face? Is it current – or is it stale?

1 SANS Institute InfoSec Reading Room: A Proactive Approach to Incident Response, by Jake Williams, August 2015: https://www.sans.org/reading-room/whitepapers/analyst/proactive-approach-incident-response-36235

Some additional challenges to consider include:

- Advanced threats may cross from your own network (where you have control) to a hosted network (where you don't have control)
- The line between the corporate network and cloud is also greying
- The classification of what traffic is related to work versus that which isn't related to work is also greying; Google Apps vs Gmail – Skype for business vs. Skype for personal use, social media for marketing vs. personal Facebook

Transactions that are done between a personal device brought into the network and the cloud, and that are not work related, are now part of the network traffic and need to be analyzed. For example, Stacy doesn't typically bring in her laptop, connect it to the wired network, access those systems and that data, and perform those actions...that is not normal.

Response teams must also consider that the value of a metric going over or under <x> will generate an alert, but there may be drawbacks: Do the thresholds and rules work for periodic data, and do they vary by hour of day and day of week? To implement a rule, for example, that if a metric goes over <x> @ 2pm on Monday, you need extreme granularity.

The preferred solution approach to all these challenges would be dynamic, contextual, online and constantly updated models of what's normal in YOUR network.

Achieving Comprehensive Anomaly Detection and Security Analytics

Automated Anomaly Detection

Creating a Baseline of Normal Behavior That Moves with the Organization

An important step towards achieving a comprehensive information security and incident response program is to establish a baseline of normal behavior in the data. This tells you what network activity looks normal so you can identify abnormal activity.

Every organization is unique and constantly changing; moments after you create a baseline, it starts to become inaccurate as the environment changes. So you need to know your own behavior and establish context around the regular activity that occurs and the alerts that are generated. Your system should then automatically adjust the baseline in real time to dynamically adapt to the current environment.

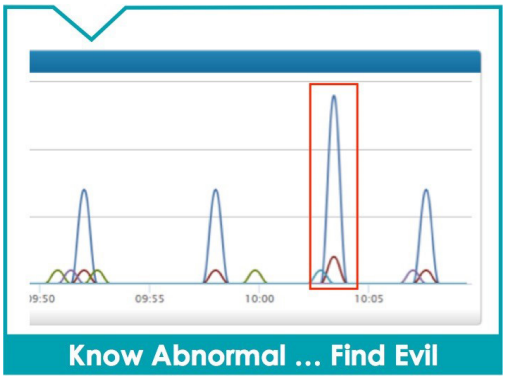

By using multi-modal modeling algorithms, anomaly detection solutions can automatically generate a baseline of "normal" behavior of any time-series data across multiple data dimensions.

Algorithms used against dynamic baselines can also find anomalies related to temporal deviations in values, counts, frequencies, statistical rarities, and population outliers. Automatic periodicity detection and quick adaptation to changing data makes this technology accurate and relevant.

Key Automated Anomaly Detection Attributes

When considering all the capabilities delivered by an automated anomaly detection solution, hiring the necessary resources to manually provide commensurate math skills would be impossible to scale. To illustrate this point, here's a rundown of the key attributes that leading automated anomaly detection solutions provide:

- **Multi-Modal Modeling** – most network data does not fall into a simple "bell curve" distribution. Multi-modal modeling automatically applies a best-fit model to your data, even if different portions of the data behave differently than others.
- **Multi-Dimensional Analytics** – finding anomalies in a single metric is valuable. Finding anomalies in complex behaviors spanning multiple data dimensions (fields) is powerful!
- **Value/Count/Frequency, Rarity, Population** – there are multiple kinds of anomalies involved in advanced attacks. Multiple forms of anomaly detection are thus required to detect them.
- **Quick Adaptation to Rate Changes** – network data changes all the time. Detection technology that relies on thresholds, simple statistics, or even trained machine learning techniques are brittle to changes in data and generate numerous false positive alerts.
- **Auto-Periodicity Detection** – most network data has periodic behaviors. Analytics that automatically detect and model this periodicity are highly accurate and reliable.

- **Calculated Probability** – the mathematically derived probability/improbability of determining whether a given observed data behavior is "normal."
- **Normalized Anomaly Scoring** – lets you know how unusual an event is among other unusual events.
- **Automated Machine Learning** – allows algorithms to analyze data on a continuous basis in order to radically reduce the Mean-Time-to-Detection.
- **Bayesian Analysis** – smart algorithms "know" when they've established a high-fidelity model of normal behavior; they don't generate alerts on anomalies until this stage is reached.

By leveraging these attributes, an automated anomaly detection solution allows your IT security and incident response teams to understand normal baseline activity. The team can then be alerted on abnormal activity that may be a sign of an attack or security breach. Automated anomaly detection also identifies the types of behavior detected and then combines those behaviors to show correlated anomalies.

While one anomaly detection indicator may not be suspicious, multiple indicators could mean your response team needs to quickly investigate. Acquiring these capabilities reduces the required manual effort (perhaps saving hundreds of investigation hours per year) while also decreasing the number of false positives. This leads to faster resolution against any real threats.

Unique for Every Organization

There are no industry standards you can refer to when it comes to anomaly detection because it's all about your data, your activity, your patterns...and your threats. As referenced above, the baseline is dynamic – it needs to adjust to reflect your changing environment. Relying on manual processes has been a challenge with anomaly detection in the past as humans simply can't keep up with the analysis, especially when dynamic baselines are required…automated anomaly detection technology is really the only way to address this, and it does so quite well.

Your unique anomaly detection process pays off even further by finding problems your static signature tools miss. You can thus improve your incident response capabilities and complete investigations much faster and more effectively.

In addition to detecting anomalous behavior, when a comprehensive security analytics solution is at the foundation of the automated anomaly detection process you also acquire all the evidence via full packet capture. You can then conduct swift and focused responses and investigations to the security incidents you have uncovered.

Security Analytics

Exposing Indicators of Compromise and the Full Scope of Attacks

Before diving into the anomaly detection report findings, it's important to supplement your information security and incident response processes by deploying network forensics and security analytics. It's like a security camera for your network – always on, always recording and always providing you with a "system of record" of what happened.

By using a security analytics solution and historical data for analyzing your network activity, accurate and scalable anomaly detection can turn data into real insights. The security analytics approach applies a full-packet capture of all network activity and is a relatively new area of technology that closes the gap where prevention fails. With a full-fidelity record of all network traffic, you can make more informed decisions on whether that anomaly is malicious activity or not. If it warrants further investigation, you have the complete "system of record" at your disposal.

The primary purpose of security analytics is to provide visibility and full evidence into activities within the target environment; this exposes indicators of compromise and the full scope of an attack – thus accelerating the incident response process. When effectively deployed, security analytics provides an early enough warning of malicious activities to stop lateral movement of the threats and to prevent data exfiltration.

Prioritizing Information Security and Incident Response Team Activities

To deliver on the key objectives – increased detection followed by accelerated response and containment – security analytics ingests full packet data and enriches it with metadata for indexing and easy retrieval. The process then creates the baseline for anomaly detection.

Symantec Security Analytics leverages the threat intelligence of the Symantec Global Intelligence Network, which is powered by the collective threat data gathered from over 15,000 customers and millions of end users to provide comprehensive URL and file reputation. Automatic sandboxing of unknown files and URLs by Symantec Malware Analysis delivers risk scoring to Security Analytics. Numerous third-party reputation data providers and outside threat intelligence add context and further enrich the packet data. This process can analyze captured data against this community of threat intelligence to generate insight and identify unknown files.

The process delivers the highest level of visibility possible into activities in the environment and produces the highest-fidelity intelligence for rendering the context of an event. This ultimately enables the best prioritization of activities for the response teams to address.

From there, the solution analyzes the anomaly results. This lets the information security and incident response teams group together anomalies that are related by common influencers in an insightful way that describes the attack progression. Detected anomalies that correspond to elementary attack behaviors are then classified into one or more categories:

- Data exfiltration
- Malware command and control activity
- Compromised endpoints and servers

The Combined Effect of Automated Anomaly Detection and Security Analytics

Security analytics combined with automated anomaly detection enhances historical forensic records and brings traditional post-event analysis closer to real-time detection. The following use cases provide examples of what your response teams may be alerted to:

Determining that Alex in Drafting has a workstation using the DNS protocol in a way that's significantly different from everyone else in the department. In fact, his workstation's unusual behavior indicates a type of data exfiltration called DNS Tunneling.

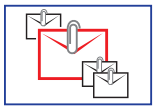

Finding out that Charles in Finance is sending emails with attachments that are significantly larger than the ones he usually sends.

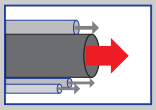

Identifying anomalous/unusual behavior such as Julie in Accounting sending an unusually high volume of data to a file-sharing website.

Discovering that Dan in Sales has suddenly established an unusually large number of network connections to a rare country, such as Belarus.

Determining if something is anomalous or unusual while performing an investigation, such as noticing that Greta in Human Resources is working from the office on the weekend, something she has never done before.

Realizing that Paula from Marketing is signed in from a mobile device and her office PC at the same time.

Any of these unusual occurrences could signal the presence of malicious activity.

The combination of security analytics and anomaly detection can also help response teams quickly identify many other types of activity that indicate a threat may be imminent:

- Unusual counts, signatures, protocols, and destinations
- High traffic to-and-from a restricted country
- Suspicious movement within the network – possible threat recon
- Data exfiltration via a website, email or application – may indicate theft of intellectual property or customer data (credit card, patient records)
- A sudden spike in activity for a specific application type

Protect Your Workforce and Business-Critical Systems

Sophisticated, targeted attacks can take weeks, months or longer to discover and resolve. Information security and incident response teams thus need tools that quickly uncover the full source and scope of an attack to reduce time-to-resolution, mitigate ongoing risk and further fortify the network.

Like a security camera for the network, combining automated anomaly detection with security analytics delivers full network security visibility, advanced network forensics, and real-time content inspection for all network activity. This effectively arms security and response teams with the ability to identify and detect advanced malware crossing the network and to contain zero-day and advanced targeted attacks.

A comprehensive record of all network activity also lets you conduct swift forensic investigations, perform proactive incident response and resolve breaches—in a fraction of the time. Your existing security tools gain the much-needed context and content needed to empower your team to identify and address security issues quickly and thoroughly.

The end result: Better preparation to protect your workforce and business-critical systems against threats that are unique to your environment and often go undetected by traditional security defenses.

Learn More

For more information on how your company can benefit from enhancing information security and incident response with automated anomaly detection as part of an overall more effective detection solution, visit Symantec.com or contact your Symantec representative.

Symantec is a leading provider of advanced web security solutions for global enterprises and governments. Our mission is to protect enterprises and their users from cyber threats – whether they are on the network, on the web, in the cloud or mobile. Our unrivaled capabilities make us uniquely qualified to deliver a security platform to 15,000 organizations worldwide, including over 70% of the Fortune Global 500.

The Six-Step Automated Anomaly Detection/Security Analytics Process

1. Capture Full Packets for Retrospective Analysis
2. Analyze and Enrich Data
3. Find Anomalies
4. Gain Insight
5. Resolve
6. Protect

See All...Know More...Respond Faster – the Business Value of Automated Anomaly Detection and Security Analytics

- Early detection of incidents
- Faster root-cause discovery
- Enable Incident Response Teams to proactively hunt
- Make the Security Operations Team more effective
- Give security analysts new ways to mine network data and protect the organization

About Symantec

Symantec Corporation World Headquarters

350 Ellis Street Mountain View, CA 94043 USA | +1 (650) 527 8000 | 1 (800) 721 3934 | www.symantec.com

Symantec Corporation (NASDAQ: SYMC), the world's leading cyber security company, helps businesses, governments and people secure their most important data wherever it lives. Organizations across the world look to Symantec for strategic, integrated solutions to defend against sophisticated attacks across endpoints, cloud and infrastructure. Likewise, a global community of more than 50 million people and families rely on Symantec's Norton suite of products for protection at home and across all of their devices. Symantec operates one of the world's largest civilian cyber intelligence networks, allowing it to see and protect against the most advanced threats. For additional information, please visit www.symantec.com or connect with us on Facebook, Twitter, and LinkedIn.

SYMC_wp_Security_and_Incident_Response_Anomaly_Detection_EN_v1a