Keys to Effective Anomaly Detection

Matt Bromiley

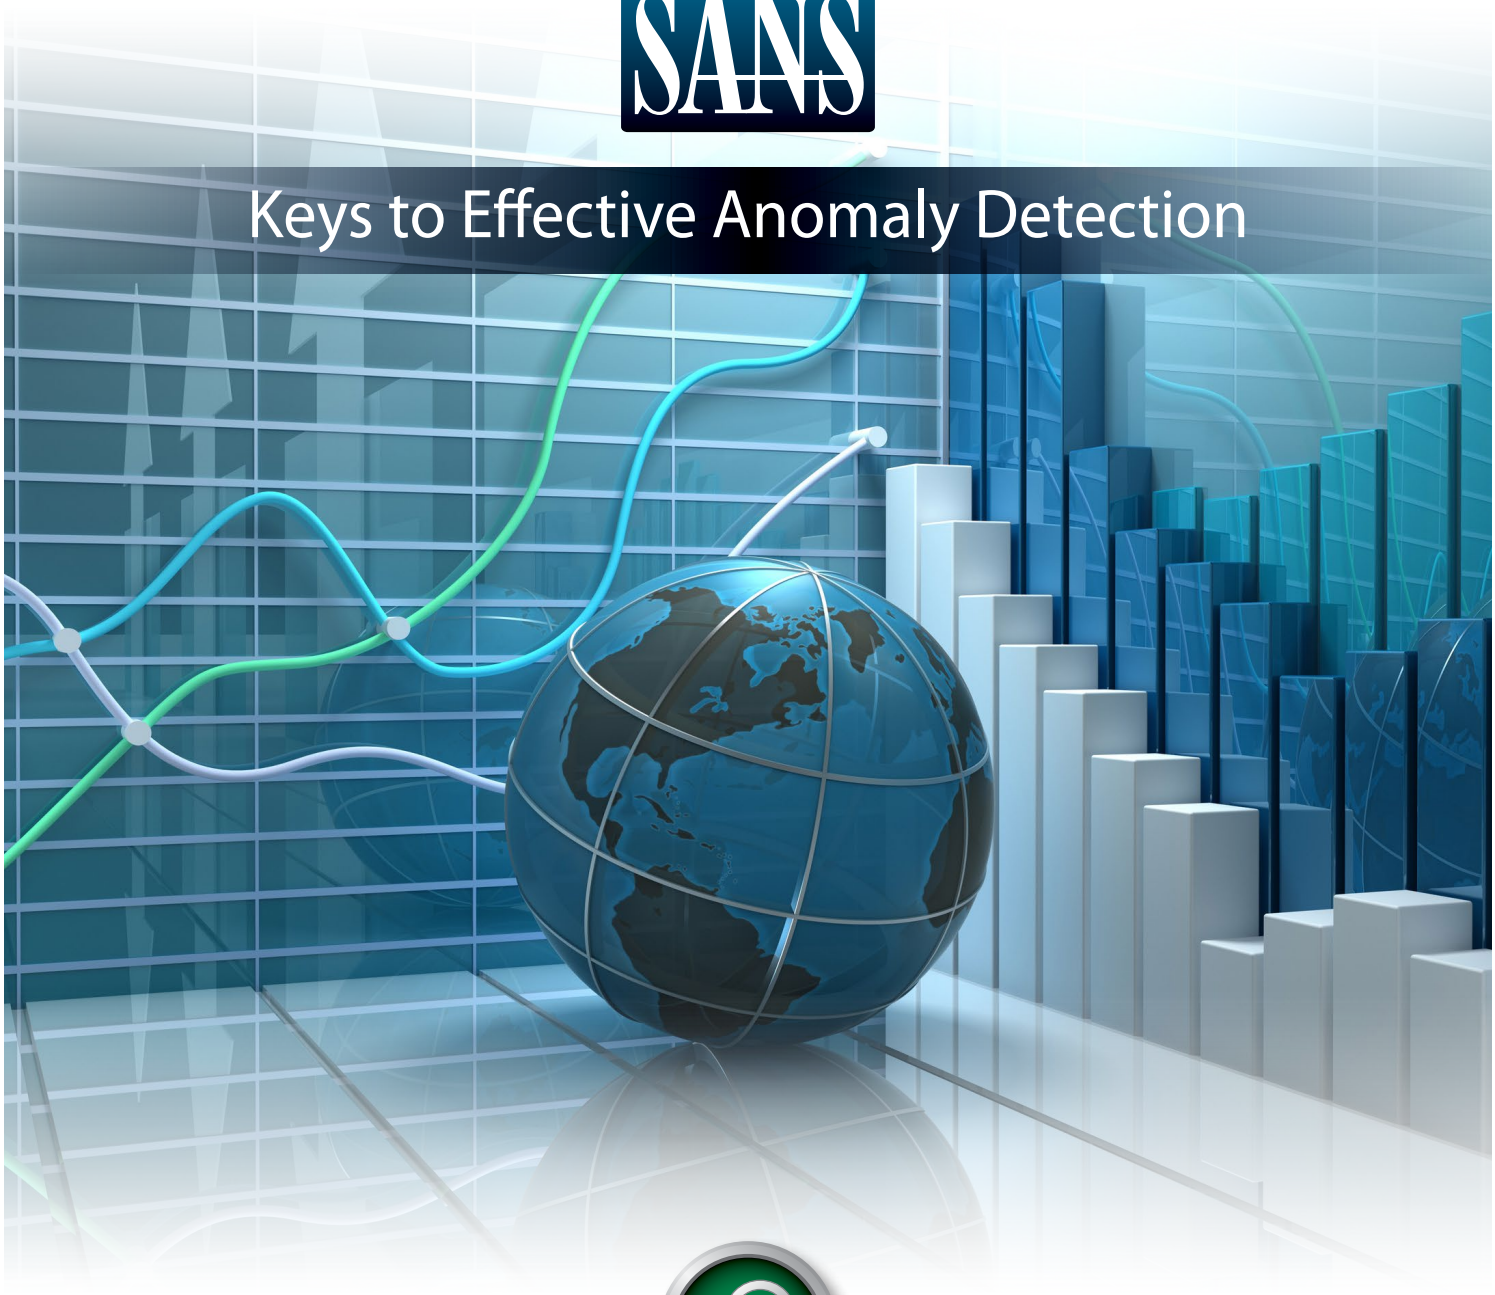

Keys to Effective Anomaly Detection

A SANS Spotlight

Written by Matt Bromiley

October 2016

Sponsored by

Sponsored by Blue Coat Systems

Understanding Anomaly Detection

Simply put, an anomaly is something that seems abnormal or doesn't fit within an environment. A car with five driving wheels would be an anomaly. In the context of an enterprise network, an anomaly is very much the same—something that does not fit or is out of place. While anomalies in an enterprise network may be indicative of a configuration fault, they are often evidence of something much more worrisome: a malicious presence on the network.

The benefits of anomaly detection are easy to see: The organization is taking a step to detect earlier phases of attacks to prevent excessive compromise. In some cases, anomaly detection may be useful in preventing attacks, but that will depend on the organization's ability to quickly distinguish attacker reconnaissance activity from normal Internet scanning. Additionally, anomaly detection can help an organization plan for the unexpected, such as a critical failure or system disruption. Security staff can use the data gathered during a baselining exercise (explained later) in multiple areas of the organization.

But traditional approaches to anomaly detection can be taxing to an organization. If the information security program isn't fully mature, anomaly detection may provide too much data for consumption and overwhelm the team. A smaller team may not be able to effectively baseline the organization, wasting resources on chasing false positives. Furthermore, with manual anomaly detection, the organization will need to dedicate resources to analyzing, interpreting and acting upon detected anomalies. Without a solid understanding of the environment, this can be a daunting task. Before considering implementing anomaly detection, an organization should determine whether its approach to anomaly detection is a cumbersome, manual process or an automated, intelligent system that enables information security teams to gain value from the additional data sources.

The good and bad of anomaly detection programs are summarized in Figure 1.

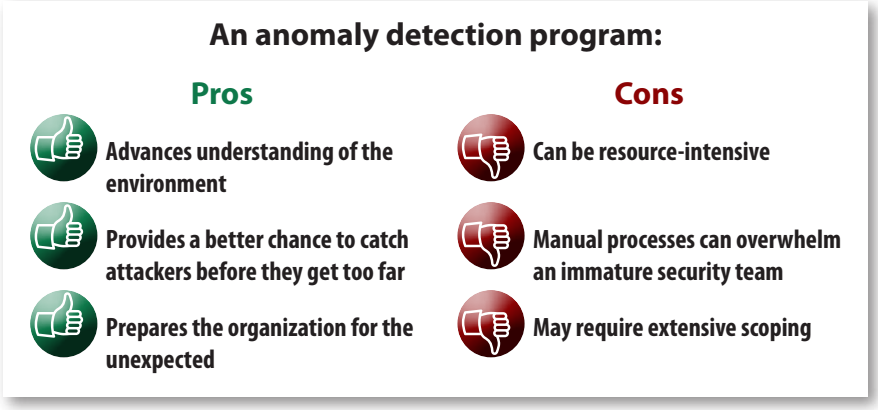

Figure 1. The Pros and Cons of Anomaly Detection

Key Elements of Effective Anomaly Detection

Once an organization decides to implement anomaly detection, it can choose among multiple methods. A crucial first step is to define what is normal. That lets the organization simultaneously define what it expects to see. This technique, called *baselining*, serves to build a profile of what a normal event or time frame, such as user activity on a business day, looks like on the network. However, baselining is easier said than done in many of today's organizations, with their myriad operating systems, user groups, and wired and wireless networks.

Due to the complexities of modern networks, manual anomaly detection struggles to provide results. Many organizations benefit from automated techniques to help baseline their environment, keep it updated and find anomalies in their networks. Automated detection often relies on multiple mechanisms to differentiate between the normal and the abnormal. The following elements, also referred to as *attributes*, are among many often found in automated anomaly solutions and can act as a benchmark when evaluating anomaly-detection options:

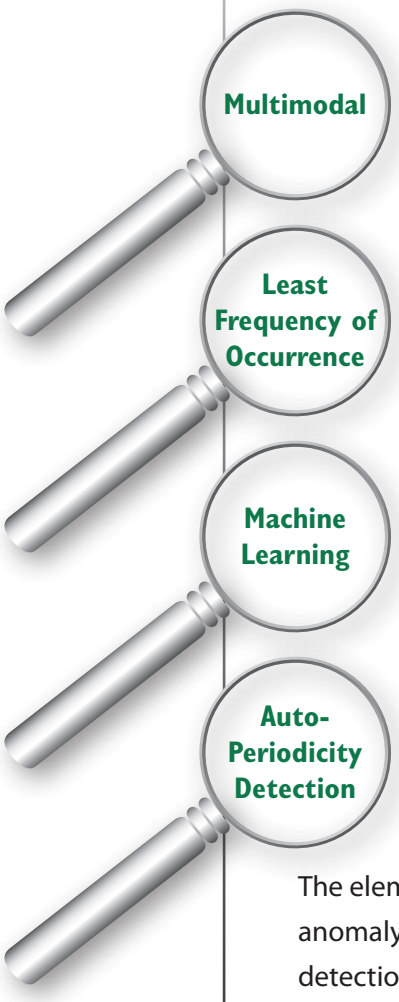

Multimodal anomaly-detection mechanisms take into account the fact that many enterprise networks have multiple modes of activity. For example, we cannot simply say that a network will be active only during business hours. We must also consider other factors, such as remote workers or outsourced operations.

Least frequency of occurrence (LFO) is a method similar to baselining that looks for anomalies based on frequency. Examples of LFO anomalies may include network connections with very low frequency, relative to the network, over an extended period of time. For example, an attacker that communicates over Port 4444 would be an anomaly and not represent a frequent occurrence within the network.

Machine learning, lately a hot topic in the realm of information security, involves complex algorithms that allow machines to grow smarter about the data at hand. Anomaly detection is no different. The success that machine learning has achieved provides a means to overcome much of the manual workload that is often associated with anomaly detection.

Auto-periodicity is centered on an understanding that, although networks may not fit within a bell curve of activity, defined periods of expected, increased activity occur—for example, a business might perform system backups on Saturday evenings. This type of detection identifies these periods and classifies them appropriately to aid in identifying anomalies.

The elements mentioned here are only a few that can be found in many of today's anomaly-detection solutions. Any organization seeking to implement anomaly detection should be sure to ask how detection is conducted and how these elements may contribute to success within the organization's network.

Anomaly Detection as a Part of Network Forensics

Anomaly detection is powerful on its own, but the benefits are enhanced when it is combined with other network security and forensics capabilities. For example, consider an anomaly alert that points an information security team to a particular part of a network. With only an alert, the team is in the right area but has no evidence of what transpired prior to the alert. However, with advanced security features such as full packet capture or efficient network traffic log aggregation, the team has essentially enabled a DVR for the network and created a "system of record." Using the anomaly as a bookend, the team can replay the event to gauge how serious the event is and possibly gain other insights into the attacker's activity and deliver evidence to validate the source.

Incorporating anomaly detection into advanced network forensics also helps the organization determine where it should prioritize its incident response (IR) efforts. In the 2016 SANS Incident Response Survey,[1] 32 percent of respondents reported dwell times, or the time to detection, of over two days. With a higher fidelity of detection, the organization can obtain faster response times and mitigate risk. Anomaly detection can help a team find activity that possibly exposes advanced attackers, who may cause just enough activity to be outside the expected norm.

Lastly, anomaly detection is important for any organization to protect the mission-critical areas of its business. Wrapping detection around the most important areas, often referred to as the "crown jewels defense," helps alert teams to suspicious activity targeting the most sensitive parts of the business. This method can also help the IR team focus on a sensitive area if suspicious activity is detected.

Anomaly detection can help a mature organization enhance its information security team to catch attacks just as they are beginning by isolating the abnormal from the normal. When combined with additional network security features, anomaly detection can be a powerful tool that helps organizations deal with the known and prepare for the unknown.

[1] "Incident Response Capabilities in 2016: The 2016 SANS Incident Response Survey," June 2016, www.sans.org/reading-room/whitepapers/incident/incident-response-capabilities-2016-2016-incident-response-survey-37047

About the Author

Matt Bromiley, a SANS GIAC Advisory Board member who holds the GCFA and GNFA certifications, is an up-and-coming FOR508 instructor. A senior consultant at a major incident response and forensic analysis company, he has experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, and network forensics, as well as memory analysis and network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, sharing with others and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor: