Realtime publishers

The Essentials Series:
Security Information Management

Foundations of Security Information Management

sponsored by

SecureWorks®

by Dan Sullivan

Copyright Statement

Foundations of Security Information Management

Security information management (SIM) is a multi-step process that, broadly speaking, consists of data collection, data analysis, and incident response. Each of these steps consists of a complex series of operations. In this, the second of three articles about SIM, we examine the operational details of security information management. The article begins with a technical discussion of the operational issues in SIM but addresses organizational issues as well. Specifically, this article covers:

- Data collection and pre-processing in SIM systems
- Data analysis and incident identifications
- The role of SIM in incident response
- Challenges and best practices in SIM

We begin with basic monitoring and data collection.

Data Collection and Pre-processing

Security-related information can be collected from a wide range of sources throughout the network:

- Gateway devices, such as firewalls and routers
- Network appliances, such as Intrusion Prevention Systems (IPSs), antivirus appliances, and content filters
- Application servers, Web servers, database servers, and other core applications that log event information

Before it can be analyzed, this data must be collected in a central staging area and properly normalized.

Data Collection Methods

SIM systems typically use a push data collection method in which log data is sent to a SIM collection appliance using one or more methods. Commonly supported protocols such as Syslog, Simple Network Management Protocol (SNMP), or the Simple Mail Transport Protocol (SMTP) can provide near real-time updates. The key advantage of these protocols is that they are widely supported; unfortunately, they may not provide some of the useful contextual data available to the source system.

Some devices, such as Windows Servers, do not forward event logs natively, and in these situations, software agents must be installed to enable log collection. This introduces another component that must be maintained, adding to the complexity of SIM. In addition, agents are device specific, so organizations with a wide array of devices in use will assume a substantial installation and management challenge.

A third method of data collection is via native application programming interfaces (APIs). These introduce more complexity than the simple log forwarding protocols but provide greater contextual detail than those protocols. They can, for example, provide packet decodes via an API giving details of the security context that triggered the IPS alert.

Pre-Processing and Data Normalization

Pre-processing is a general term used to describe the file manipulation and data normalization that is required to map data into a format the SIM system can manipulate. A SIM, for example, may have data from a firewall, an IPS, and an antivirus appliance all related to a single event; however, because the source systems are manufactured by three different vendors, even common data types may be formatted differently.

Although pre-processing modules extract and reformat data, normalization procedures ensure that data from different systems referencing the same kind of data all use a common frame of reference. For example, timestamps should be comparable between source systems even though each system has its own clock. Similarly, a physical device on the network could be uniquely identified by either a currently assigned IP address or a MAC address; using one identifier scheme makes the analysis operations less complex.

Synchronizing system clocks can be a challenge. One approach to solve this problem is to add a timestamp to an event when it arrives at the SIM appliance so that both the source device timestamp and the SIM appliance timestamp can be used to establish a common network time. Another challenge can arise when a single physical server runs multiple virtual machines. Information tied to a MAC address may need to be correlated with events in one of the virtual machines. Once these challenges are overcome and the raw data is in a standardized format, the data analysis stage can begin.

Data Analysis and Incident Identifications

Data analysis and incident identification is a three-step process:

- Aggregating and filtering events
- Correlating events of interest
- Assessing incidents

Source systems generate a great deal of data. Not all of these log entries and event notifications are relevant to determining whether an incident has occurred. One way SIMs deal with the problem of information overload is through filtering and aggregation.

Filtering is performed with two types of filter sets, positive and negative. Positive filters capture known malicious or suspicious events. These events should not occur under normal circumstances and need to be analyzed further to determine whether there is a threat. Negative rules eliminate known events that are expected under normal operating conditions and have zero chance of being a threat. Of course, events that have never been processed before fall into neither the positive nor negative category; these are considered “anomalous” and should be investigated immediately to determine whether they are threats and update filter sets accordingly. This ability to incrementally improve filter rules enables SIM systems to scale to the breadth of requirements in businesses today.

Aggregation methods are used to further consolidate events. For example, a firewall might generate details about a port scan of the perimeter firewall. Tens or hundreds of individual events documenting the scanning of each port is not useful to network managers, but a single event with an aggregate count and range of those ports scanned is actually useful information.

Once the significant events have been isolated, they are correlated to determine the sequence of events and identify potential attack patterns. For example, a file may be transferred to a server, a configuration file changed, and then a data file copied from the network. Correlation can occur along a number of dimensions in the data:

- Data contained in the event, such as source and destination IP addresses, device, event summary, and event time
- Asset data, such as types of applications running on devices
- Asset classification data, such as the business purpose of the device
- Vulnerability scanning results
- External intelligence data, such as provided by BuqTraq, CERT notices, and subscription feeds
- Time-related attributes, such as time of day and number of events in a given period of time

A series of individual or correlated events might trigger a logic rule indicating a possible breach. This pattern matching by is the last step in the process and the one that triggers an event notification that a significant event has occurred.

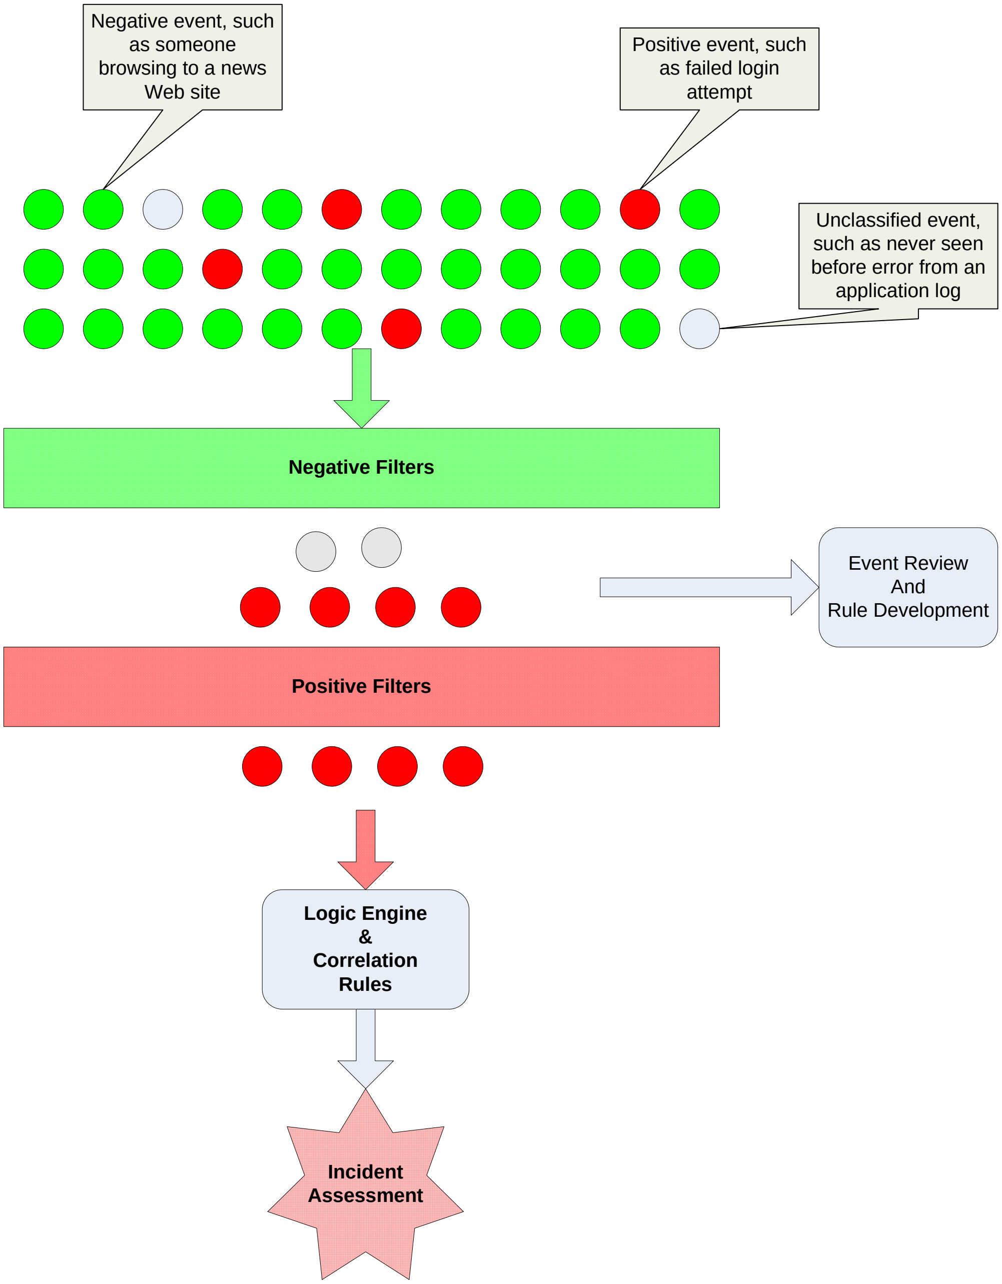
Negative event, such as someone browsing to a news Web site
Positive event, such as failed login attempt
Unclassified event, such as never seen before error from an application log
Negative Filters
Event Review And Rule Development
Positive Filters
Logic Engine & Correlation Rules
Incident Assessment

Figure 1: Large volumes of raw events are filtered to identify those that are significant and then to determine whether an adverse event has occurred.

Once an incident has been identified, IT staff should respond following an incident response plan.

Incident Response

Best practices dictate that you have an incident response plan in place rather than rely on *ad hoc* reactions to individual incidents. A common framework for incident response plans is the five-step cycle depicted in Figure 2. SIM practices and technologies can inform and improve all phases of incident response because a SIM system gives a comprehensive picture of events on the network.

Clearly, SIM's ability to collect data from diverse devices, filter irrelevant events, and correlate data support the detection phase of incident response. This is especially important because the sooner a breach or malicious event is detected and contained, the less chance of a costly loss.

> See the first article in this series, "The Business Case for Information Management" for statistics on the costs of data breaches.

The monitoring services provided by SIMs also aid with the *contain*, *eliminate, and recover* phase to help ensure that all instance of malicious activity have been eliminated. Finally, in the *learn and adapt* phase, SIM data can provide forensic details that illuminate how a breach occurred and where additional security measures are required.

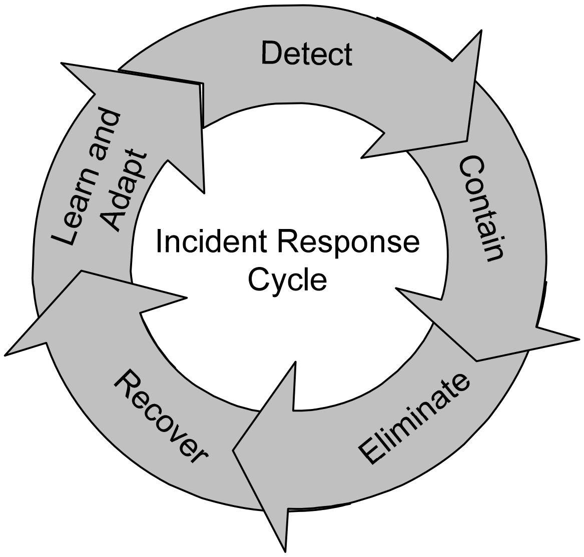

Figure 2: The incident response cycle brings a system back to operational state and leverages lessons learned to keep it secure.

Challenges and Best Practices

SIM is a practice that builds on other security measures and operations. Not surprisingly, the challenges to implementing effective SIM are shared by security practices as well.

SIM systems must scale to meet the volume of data generated on a network. Scalable performance can come from a number of areas, including improved filtering to reduce the number of events that must be analyzed, tuning incident detection rules to reduce the time required to analyze significant event patterns, and adding devices to load balance filtering and analysis operations.

IT staff face the challenge of staying current with the threat landscape. Cybercriminals and others attackers are constantly developing new techniques to avoid detection and exploit vulnerabilities. This situation demands that staff maintain adequate training and levels of expertise.

For additional details, see the third article in this series, "Making Security Information Management Work for Your Organization."

Summary

SIM is a combination of technical and business processes. SIM includes data collection and pre-processing, data analysis, incident identification, and incident response as the major steps to controlling security breaches. It also includes long-term planning, incident response formulation, performance management, and training to ensure that adequate resources are in place to detect, remediate, and recover from security incidents.