Security Management: Making Sense Of Events

Christopher M. King

Security event management promises clarity amid the alarms. But it's an integration challenge.

The typical large enterprise routinely is inundated with security-related alerts from heterogeneous security devices (intrusion detection systems, firewalls, VPN gateways, and platforms). Network security managers are awakened at all hours by various events that seem to demand their immediate attention. These managers find themselves attempting to manually inspect or decipher reports of security anomalies from amid the reams of logs generated by their organization's array of security devices—an impossible task.

To make sense of all this information, security managers need an operational view of the security health of the enterprise. This article will look at strategies to properly alert, categorize and react to security events as they occur.

The Immediate Response

A security "event" is defined as an observable occurrence in a security system or application. If an event is detected that poses a risk to sensitive data or a resource, the organization has two options for its immediate response:

- Block the traffic from the specific individual(s) who are conducting the attack. The traffic could be coming from an attacker's own platform, but this is highly unlikely. More likely, one of the user organization's own assets has been compromised and used as a launch point.
- Shut down the resource that's being attacked, or disconnect it from the network to prevent contamination. This typically involves contacting the owner of the resource.

After the dust settles, a damage assessment needs to be performed before the systems can be restored. This restoration process is the responsibility of the business continuity group. If a legal or criminal prosecution is in order, the local security staff must preserve the evidence for the forensics team—which may be an outside consulting group—and law enforcement authorities.

Security Incident Management

A security incident is an adverse event in a platform or application. The process just described is a high-level view of the incident management process. As Figure 1 shows, the first line of defense is prevention, which requires the proper controls (e.g., a robust security architecture, the latest software patches and rock solid change management processes).

The best source for keeping current on Internet security problems is the Carnegie Mellon CERT Coordination Center (CERT/CC) organization. CERT/CC is the major reporting center for Internet security problems (e.g., the Morris worm, or more recently, the Code Red attacks).

FIGURE 1 Incident Management Process

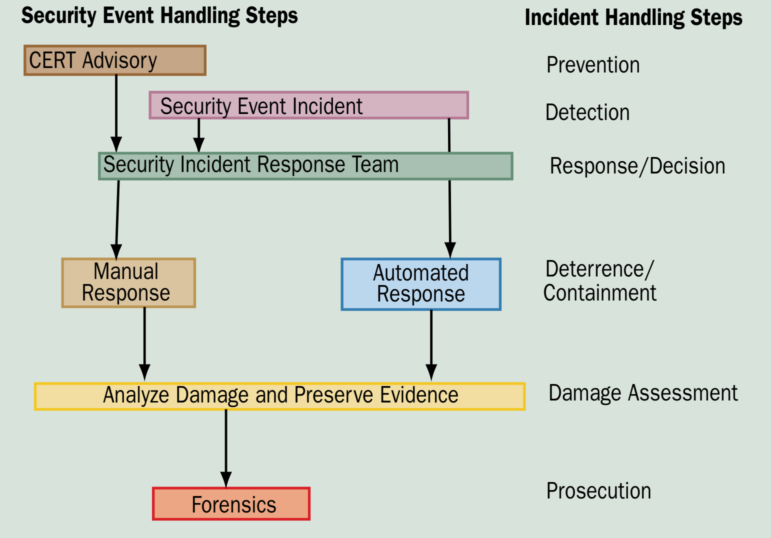

Christopher M. King is the practice director of Greenwich Technology Partners Information Security Practice. He is the lead author of the book, "Security Architecture: Design, Deployment and Operations". He can be reached by email at cking@greenwichtech.com

When an advisory is published by CERT, other security groups or by the vendor, an enterprise's security incident response team must decide how best to deal with the advisory. This could entail applying a patch, or blocking access to particular resources until a patch is available and tested.

Taking such pre-emptive steps will lower the risk of an incident, but you cannot prevent all attacks. When attacks do occur, a management tool like e-security's Open e-security platform or OpenService SystemWatch alerts the security manager that an incident has taken place within the organization; typical examples include online attacks such as the Code Red buffer overflow, a computer virus such as Chernobyl or a malicious email (VbScript). This alerting capability is part of a larger security event management (SEM) strategy, but as we'll see, there's much more to a SEM.

FIGURE 2 Security Architecture

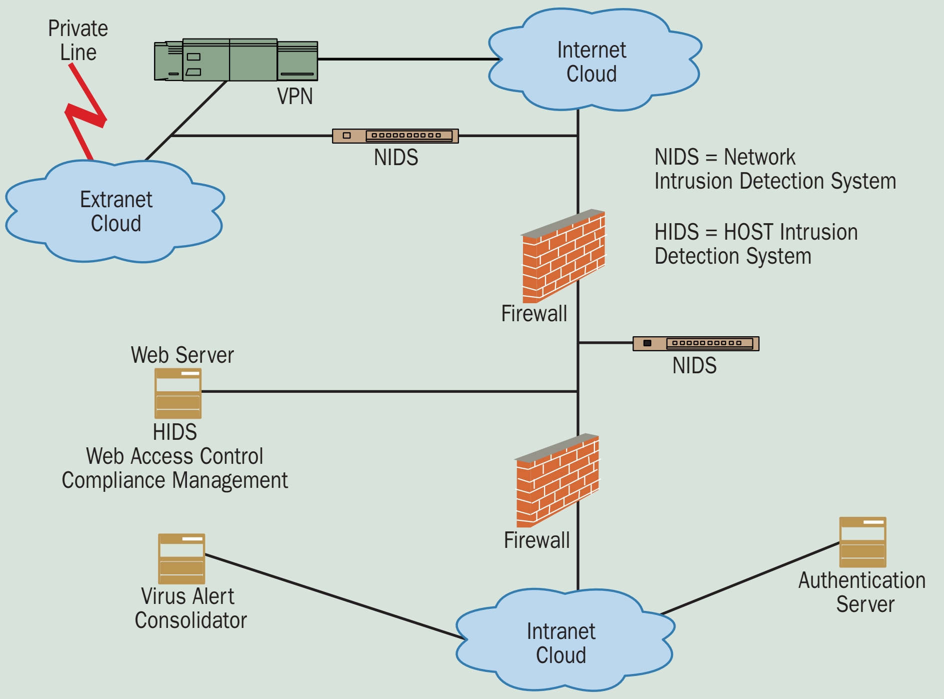

Once an attack has been deemed important enough or has compromised a resource, an escalation process must be initiated. The first system to be contacted is the network management system. This is usually done in a secure fashion using an encrypted and signed message. The new SNMP version, SNMPv3, extends the client/server authentication and encrypts the SNMP communication between server and clients.

In some instances, security devices can respond automatically to an attack, for example by reconfiguring the firewall or sending a reset packet to back to the attacker. Resetting the TCP packet will stop that attack, forcing the attacker to reestablish the session. This ability to respond automatically is called dynamic deterrence.

If automated methods fail, manual intervention is necessary. This involves determining the level of the incident and, if necessary, preserving the evidence and providing containment to prevent spreading. The most effective means of containment is to shut down the platform that has been attacked or pull it from the network, after which time the business continuity plan can be executed or a forensic expert sent in.

Security Architecture Components

The totality of all the security devices, placement and processes within an infrastructure is called the security architecture (Figure 2). To get a handle on how to manage the events based on all the information these systems generate, let's first look at just which systems we're talking about:

■ Network Intrusion Detection Systems (NIDS)—Intrusion detection looks for security anomalies by analyzing network traffic, platforms or application behaviors. The anomaly can be detected either by using known attack signature matching or through behavioral analysis, or by a combination of these two methods.

Attack signature matching is comparable to anti-virus software (i.e., the attack can be detected if the pattern is matched). The shortcoming is it can only identify known attacks; an attack outside the signature list will go undetected.

The ability to add custom signatures is a new feature to the top commercial IDS products this year, and it's a key capability. This solution is imperfect—an attack signature database doesn't know about new attacks until they have begun to appear in networks somewhere in the world, so there will always be some lag time during which an organization is vulnerable to a new attack. But the sooner, more frequently and more easily you can add new attacks to your signature database, the better your security. Many lesser-known IDS products, like Enterasys' Dragon, were built with the ability to add custom attack signatures.

IDS Attack signatures are categorized many ways; one is by the type of attack (e.g., ICMP, flood, Web, TCP, etc.). Another is by severity; at the low end, this might include an "informational" event, such as an outside party scanning or connecting to well-known ports on the network but not taking any other action. The most severe events include attacks that can compromise systems, which are typically buffer-overflow related.

TABLE 1 Typical Firewall Report Entry

Rule #	IP Source	IP Destination	Protocol	Action	Timestamp	Firewall
5	128.127.1.45	www.gtp.com	http	Accept	7/27/2001 9:51:35 AM	Public

Besides running pre-defined signatures, some IDSs can use behavioral analysis to spot anomalies. After a specific amount of time (i.e., a learning period), an alert will be generated if an event is outside the learned parameters. This is very difficult to do at the network level; most pattern analysis is performed at the application and platform level.

Whichever detection method is used, IDSs usually report to a manager device to aggregate the events from all the IDS collectors. The manager serves as a console to the network administrator and is an intermediary device to forward SNMP traps to a network management platform or to act as the alerting facility (i.e., by sending out pages and email).

Other enhancements are available on specific products; for example, Internet Security Systems' (ISS's) Realsecure version 6 has the ability to play back network traffic to help categorize the specific attack.

- **Host Intrusion Detection System (HIDS)—**In addition to being a network-based device, an IDS can also reside on individual hosts. Host-based intrusion detection software performs real-time analysis of system activities, such as accounting, critical processes, and static files. Products like ISS RealSecure agent and Symantec ITA are prominent in this space.
- **Compliance Management Systems—**Host-based compliance systems perform static checks on the platform's baseline policy (e.g., file and directory permissions, a change in the audit subsystem, or a change in executables or network services). Products like ISS System agent and Symantec Enterprise Security Manager (ESM) characterize this space.
- **Web Access Control—**Web access control software, such as OpenNetwork DirectorySmart and Entrust getAccess protect Web-based resources. All access control and authentication information for Web pages and applications is available for monitoring.
- **Host-based Firewalls—**Host-based firewalls (e.g., Checkpoint Firewall-1) are capable of logging security and system health events. When an event occurs, typically the only information reported to the security manager is which rule within the security policy was executed for a given packet.

TABLE 2 Syslog Priority levels

Level	Description
0	System unusable (Not used for PIX)
1	Take immediate action
2	Critical condition
3	Error message
4	Warning message
5	Normal but significant condition
6	Informational
7	Debug message

Table 1 shows such a report. This is a single line entry in a firewall log file. It lists the rule number that was executed, the source IP address that was attempting to go to the hostname (www.gtp.com), the protocol used, the action taken by the firewall (accept, drop, or reject), the time it occurred and the firewall on which it occurred.

The most often-invoked rule is the explicit deny, which is executed if none of the other rules are met. Invoking this rule drops a packet that has been found not to meet the security policy; dropping the packet causes the party requesting access to time out. The explicit deny is the basis for most firewalls' logic (i.e., deny anything that is not explicitly allowed).

Firewalls log many events, but vendors have been unsuccessful in developing a tool that can examine firewall logs in real time, looking for anomalies. Unlike IDSs, firewalls are not designed to look for certain events, but simply to log rules as those rules are invoked. Therefore, an organization can examine firewall logs manually, in search of intrusion attempts or successes, but obviously this will be done after the incident has already occurred.

The network management information provided by firewalls usually includes the results of SNMP calls. For example, an SNMP-get provides the state (up, problem, unknown) current policy loaded, number of drops, reject, accepts and loggged packets.

- **Appliance-based Devices—**Appliance-based devices (e.g., Cisco PIX firewall and VPN gateways) do not have hard disks, and instead log security and system health events using syslog.

Syslog, which traces its roots back to Unix, sends a string message to a system logger server inside the intranet. Each message is tagged with a priority that lists the facility (i.e., the appliance where the log message originates) and level (priority levels range from 0–7—see Table 2).

The Cisco PIX syslog message includes the firewall's host IP address, the facility, severity, timestamp and the log message. The information contained in the syslog message includes the denied packets, connection counts, firewall console access, reboots, and the number of bytes transferred for accounting. Collecting several PIX syslog messages is challenging, because a firewall appliance generates a large amount of data. Syslog was not designed to generate a lot of traffic; it was designed for exception logging. Software packages like Open Systems Private I and netForensics have robust solutions for managing and analyzing events for multiple PIX firewalls.

Security Event Management

Security event management (SEM), then, is the ability to monitor, alert and report on events from security components. The components may be network based (including firewalls, IDSs and VPN gateways), platform based (also includes IDSs, as well as compliance management devices) or application based (e.g., Web access control, DBMS, or SAP).

Figure 3 illustrates how these components provide information to the SEM system, and the subsystems that make up a SEM system are shown in Table 3, p. 38.

No single vendor product can yet provide a true SEM system. Instead, an enterprise must cobble together a combination of vendor-supplied security systems, network management platforms and the enterprise's own middleware and procedures on security and network management.

To be effective, the SEM system requires the following information:

- Inventory of all the sensitive resources under your purview (e.g., database servers, financial processing server, security repositories) and their locations.
- The placement of event sources—i.e., security devices—in the infrastructure, and all the expected outputs.
- A common vulnerability mapping between heterogeneous signature databases.
- The ability to correlate the different security events among the devices.
- The systemic attack signature or pattern matching paradigms to detect anomalies across all the devices.
- The ability to forward alerts to the network management system.
- An adequate escalation procedure and incident response team put in place by the enterprise.

The ability to correlate security events from multiple heterogeneous sources is a very powerful—and elusive—proposition. For starters, there's no consistent treatment of alert levels. Alert levels vary according to the placement of the event source, frequency of events and the resource being monitored.

Furthermore, many vendors' systems that are available to date just aggregate all the events in a single data store for archival purposes, while other systems only provide threshold logic (an account of the event's occurrence). Rarely is there combinational logic (i.e., if-then-else), particularly over a given time period.

What's more, it's almost impossible to stitch multiple manufacturers' products together; the security vendors have done an excellent job in keeping their systems proprietary. The management systems of most security products make it very difficult to forward events from that product—instead, these homogenous systems form "stovepipes" of information.

One of the biggest problems is attack signature names; conventions (security levels, exploit classification) among vendors are all different. For

A SEM system will require products from many vendors that can talk to one another

FIGURE 3 Security Event Management

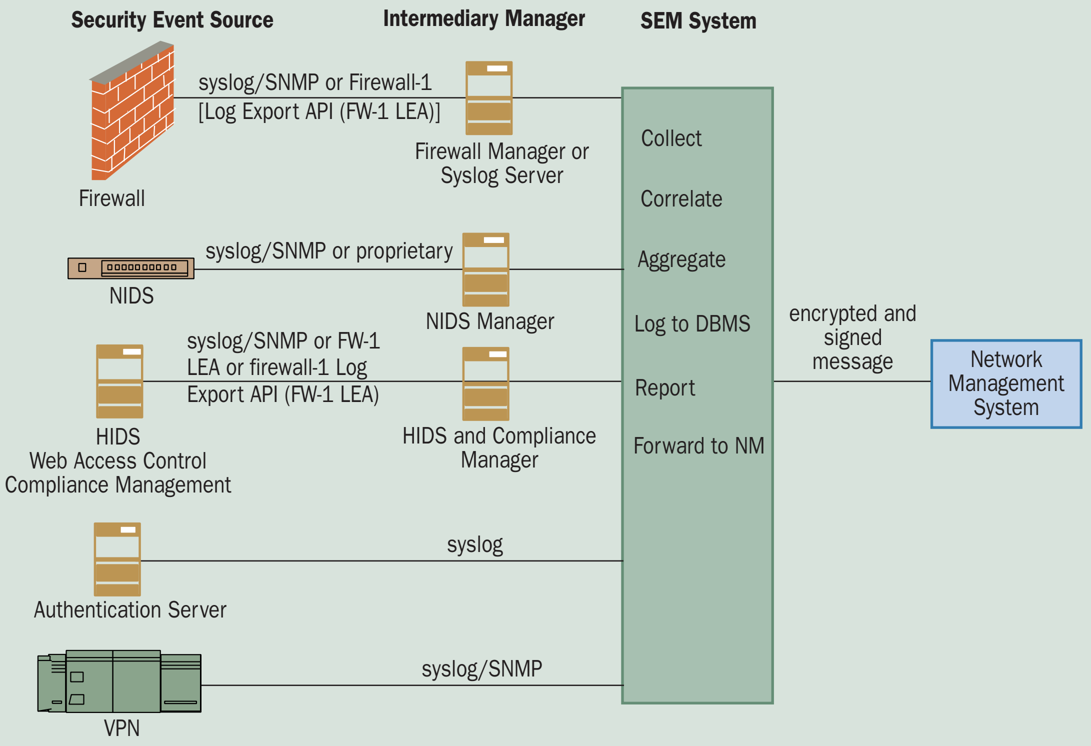

TABLE 3 SEM Subsystems

SEM	Component Description
Triggers	This is the most complex part of the SEM system. The specific alert that is generated from the SEM system has a very high probability of being an intrusion. The logic used to determine whether it's an intrusion is systematic for the entire event source and, based on the corporate policies and procedures, includes escalation, notification and best practices for deployment.
Event Source	Source of the security events, i.e., the security devices. This is typically via an intermediary manager (firewall or IDS manager). It could also be a raw feed via syslog or SNMP trap.
Intermediary manager	Most event sources report to their own manager using a proprietary protocol (e.g., Symantec ESM, ITA, ISS RealSecure). The challenge is in getting this manager to forward information to the SEM system.
Console	A GUI-based and/or Command Line Interface to administer and monitor the security event policy. This console is located at a security operation center (SOC).
Analyzer	A component able to process all the events (correlate, apply security policy). It resides on the corporate intranet.
Agents	Reside on the event source host. They collect data, possibly perform some consolidation, and forward the events to the analyzer.
Database	The database is populated with all the event sources, expected output location, and the resources it's monitoring. The only way to adequately generate reports is from a relational database.
Report Writer	Commercial products provide canned reports (e.g., expired password, failed access attempts, and frequent scans).
Forwarder	The ability to forward events into the existing network management system is key, especially if they are monitoring 7x24x365.
Alerter	Provides paging, email or executes programs.

example a Web PHF attack is called something different by Cisco Secure IDS and ISS RealSecure.

The Mitre Corporation, a federally funded research and development center, is championing an effort to standardize the names for the publicly known vulnerabilities and security exposures, but this initiative, called the Common Vulnerability and Exposures (CVE) is the only one of its kind. Some vendors have begun adopting its recommendations; Version 3 of the Cisco Secure intrusion detection software maps its proprietary network security database (NSDB) to CVE.

Conclusion

Security event management is a combination of the security and network management disciplines. It requires not only the proper infrastructure, but the correct processes. An enterprise trying to implement SEM today faces an impressive integration challenge.

And yet, viewing your security architecture as a collection of piecemeal systems is a recipe for complacency. As an example, one of the major misconceptions in the security industry is in the perception of intrusion detection. The age-old question is: "Have you been broken into?" If the answer is no, they are lying. If the answer is yes, then the next question is: "Did you detect the intrusion in a timely manner?" Rarely is the attack stopped outright by the intrusion detection system—detection is not the same as prevention—but the sooner an attack is detected the better.

The best way to prevent an attack is to provide the adequate controls as close to the resource as possible. For example, it's not enough to deploy IDS, and then believe you're covered. How and where you deploy the IDS is crucial: Network-based IDS would not have caught the Code Red attack, but a host-based IDS that was monitoring static Web pages would have. Another alternative is to use third-party access control software like Okena StormWatch, which controls access between the applications and operating system resources like the file system or network.

Ultimately, the security managers and support teams in the enterprise would want to manage the whole enterprise seamlessly from one location. We may not have reached this goal yet, but better systems and processes can move us closer□

Companies Mentioned In This Article

CERT/CC (www.cert.org)
Checkpoint (www.checkpoint.com)
Cisco (www.cisco.com)
Enterasys (www.enterasys.com)
e-security (www.esecurityinc.com)
Internet Security Systems (www.iss.net)
MicroMuse (www.micromuse.com)
netForensics (www.netforensics.com)
Open Systems (www.opensystems.com)
Okena (www.okena.com)
Symantec (www.symantec.com)
OpenService (www.open.com)
The Mitre Corporation (www.mitre.org)