

Event Data versus Log Data, and the Difference between IT Security and Breach

If you are a CISO, chances are you are not getting a lot of sleep these days. A small excerpt from your job description looks something like this:

- Develop Organizational Information Technology Security Strategic Plan and Program for the entire enterprise
- Determine the most critical areas to address
- Design and update information security and privacy policies and standards
- Provide interpretations of current policies related to specific situations as they arise
- Develop business cases for security initiatives
- Plan, execute and evaluate security programs

The above excerpt entails only about one quarter of the bullets affixed to this particular job description, pulled from www.cisecurity.com. From securing the enterprise to audit trails and compliance, the CISO of tomorrow will have to manage a complex and ever-changing IT environment that without process, she/he can never catch up to. Pinching down on the CISO are c-level peers, internal compliance auditors, industry standards groups and government entities with acronym-laden names like PCI SSC, FISMA, FERC, NERC, SOX, HIPAA, and many others. Many moving pieces need to come together in orchestrated chaos to keep the bad guys out, all while maintaining compliance. Staying ahead of the game with an automated alerts-based, systematic approach is the key to survival and the difference between CISO job security or becoming a cyber-breach headline.

A critical component of alerts-based IT security is advanced “event” notifications. It is so important to one industry consortium – the Payment Card Industry Security Standards Council – that they list event management as one of their standards requirements. The goal of the PCI Data Security Standard (PCI DSS) is to protect credit

cardholder data wherever it is processed, stored or transmitted. And it is the

PCI Security Standards Council that has issued 12 security controls and processes for PCI DSS requirements. These 12 requirements pertain to every merchant – no matter how large or small – that processes a credit card number as well as the service providers within the ecosystem of payment devices, applications, infrastructures and users.

According to the PCI Security Standards Council, the standard was developed to “encourage and enhance cardholder data security and facilitate broad adoption of consistent data security measures globally.” Anyone who has worked in IT and understands data and its expanding complexity understands the difficulties ahead. But the penalties are harsh and credit card transactions via electronic means have become so standardized with retailers, that much of the battle for PCI DSS compliance is well under way. There still, however, is a long way to go for PCI DSS compliance across the board.

The 12 requirements that the Council outlines are fairly



obvious; things like fortify fire wall, secure and encrypt the CC data, tracking and monitoring, and other policy suggestions, with the common thread being *monitor all user data*. In particular, PCI DSS requirement 10.2 states: “Implement automated audit trails for all system components to reconstruct the following events,” and then it goes on to list seven of these “events.” Arguably the most compelling requirement for the software vendor, 10.2 speaks to automated event notification, archival of the event, and replication of the event for forensics. Just like with your friendly CSI team on TV, after an IT breach, auditors want forensic data to recreate the crime scene and they are going to be looking to you, Mr. CISO, to get this data.

Somehow, you are going to have to initiate a string of processes that occurs when a security event takes place. This means you are going to need be well versed in 1) event log management, and 2) automated response mechanisms that kick in when an event occurs.

So what do we mean by an “event?” To answer that question, we first need to have a short discussion on “event log management.”

A Discussion of Event Log Management

PCI-DSS regulations, as well as many other security standards often refer to the management of “events” and “event data.” As with many other aspects of Security Information & Event Management (SIEM), it is important to identify what is meant by this term in order to be compliant with the objectives of these standards.

At first inspection, it may seem that a discussion of the term “event” can be had simply: The meaning of “event” requires very little elaboration. This is actually *not the case at all*.

Any discussion regarding the meaning of “event” can become quite deep, and is subject to a lot of debate and argument. And, the semantics of “event management” is actually fairly profound and complex.

What Does the Word “Event” Actually Mean?

In IT, the term “event” is often used casually, synonymously with any log entry (i.e. anything that occurs at a particular

time, and has a time stamp associated with it). But based on this definition, an event may be something as simple as accepting a single arbitrary packet into the network. It may even be a “heartbeat,” issued by some program at some arbitrary interval ranging from once each day to once each microsecond. Consider the extreme case, where an event is defined as the high-speed system clock of a computer ticking one count. It is certainly classifiable as an “event,” but is anyone really interested in logging this for each device, one million or more times each second from every managed device?

Given that, when the PCI-DSS standard refers to the requirement of logging all events on the network, it cannot possibly be referring to the logging of *each and every conceivable event*. We must interpret that the standard be referring only to “significant events”. As any operator of a SIEM system will tell you that most of what you are logging is pretty much junk anyway. And, you will be required to keep this data for several years. The data will be as voluminous as possible (to satisfy auditors, but also forensics) and each passing year will make that operator more appreciative of how low-cost disk storage has become.

Within the terabytes worth of log data you have collected, there may be a few “significant events.” Hopefully, that number will be low, optimally zero.

The other data is okay to log (actually required by auditors) but this data is not necessarily “event data.” It is simply data to support and clarify the more particular and important “event data,” which is what you need for SIEM to be at its most effective state.

What We Mean by ‘Event Management’

CorreLog approaches the concept of “events” in a fashion slightly different from the typical understanding of the term. The philosophical perspective of CorreLog is that “events” are separate from normal log data. Events are GENERATED from log data, and serve as “significant markers” in time that complement the log data.

Various applications and devices can make a determination of what constitutes an “event.” For example an IDS system may detect a possible intrusion, and log that information.



But an event is separate from the data upon which this determination is based. The invalid login is not significant. However, the fact that it has occurred a dozen times on a secure platform is quite significant. In other words, the invalid logon attempt is just data. The event is revealed when multiple other data items are considered (*correlated*) to make the judgment call that a brute force attack has occurred.

As we have seen over the years, events are not so much things you log as they are *things you generate* that attempt to define user behavior. Events are “judgment calls.”

We see the event as something detected by the SIEM, based upon the type of data you are interested in looking at for threat detection. This perspective is slightly different than other ways of looking at SIEM. And CorreLog supports this perspective in various different ways, such as through its integrated “ticket” facility, which detects an event from the raw log data (through match patterns, triggers, thresholds, and many other integrated tools.) A particular log message can signify an event; more likely the event is related to many different log messages, possibly received from many different sources. So now we have event data and log data being monitored and archived. What’s the difference between the two?

Differentiating “Event Data” from “Log Data”

Event data is a highly specific subclass of log data. It has various characteristics that distinguish it from other log data. Specifically:

1. **Pertinence.** An event is distinguished from other log data based upon its relevance. How interesting is this event? How often does it occur? (It is important to note that, according to the basic tenets of information theory, the relevance of an event goes down the more often it occurs. An event that occurs regularly is not very relevant.)
2. **Context.** The event must have enough of a context to make it understandable. For example, the statement “something very important happened at noon today” has no specific context. It could not be classified as an event by our definition. It may be pertinent (possibly based upon the source of this data), but there is not enough information to call this an event. Contrast the statement with something more specific, such as “You won the state lottery today at noon.” The latter statement is definitely an event, because you now know the context.

3. **Timeliness.** Information cannot really be classified as an event unless it can be associated with a fairly precise time. The more timely the message, the more likely it can be considered an event. For example, stating that some long expected movie release will occur in the fall does not make that film into an event. It is simply information about an event. The actual event occurs at the precise day of the release, which in this case is the eventful movie premier. It has a specific time associated with it. In the best case, the event occurs in real-time.
4. **Actionability.** If information is pertinent, and timely, it must still be actionable to be useful. The event requires something else that logically follows – perhaps further investigation, or corrective action. If a log message is not actionable, it can't really be classified as an event. Note that the event may not actually precipitate an action, but should still be actionable. For example, your birthday is an actionable event, even if you choose not to throw a party.

If a log message has the above characteristics, it is easy to classify this data as an “event” as opposed to simply log data. What we are doing is correlating log messages to define probable user behavior. Remember that an event that occurs regularly (user is logged in every day from 8:30 a.m. to 5:30 p.m., and goes to the same server locations during that time only) is not pertinent. When we see the same user have a login at 2:00 a.m. from a different IP address and accessing different server directories, this is anomalous user behavior and needs an immediate action in the form of a ticket alert.

In our case, the CorreLog SIEM would identify the “event” in a special navigation tab in the software and then access an internal “ticket” facility, which identifies that these log messages may be event data with an action to be taken. If true, the system would generate further events based upon rules and functions (defined by you) that reside in other parts of the software. The “ticket” facility includes specific functions that identify pertinence, context, timeliness, and actionability. In this scenario, CorreLog correlates seemingly random log messages into actionable information that may



www.correlog.com

indicate probable user behavior that is threatening to your enterprise security.

Once the tickets are generated, they are managed either internally (through various management facilitates that are built-in to CorreLog) or can be managed externally by a separate incident management system.

Conclusions

Without the basic understandings of what constitutes an “event,” successful implementation of automated and systematic threat detection, or SIEM, is unlikely. A good understanding of the semantics associated with the term “event” is not difficult to achieve, and can greatly clarify the objectives of the SIEM and security strategy of the organization.

The word “event,” referenced by many standards and vendors, can be interpreted in multiple ways. So a clear understanding of what an auditor or vendor means with regard to “event data” is extremely important.

From the CorreLog perspective, “event data” is different from “log data.” Event data has different properties, and different data sources from more common and random log data.

In particular, an “event” is not typically logged from a device, but is actually a log message internally generated

by the CorreLog Server, which serves as a significant marker to the log data. An event is the result of correlation. It may correspond to a single log message, but more likely it relates to multiple log messages that have been analyzed, correlated, and reduced to actionable tickets in an incident management system. A SIEM’s role is to assist in the understanding of the log data reducing management time and resources (*dollars!*) and as much as possible eliminating the risk of human error or intrusion.

The event data clarifies what is actually going on within the managed enterprise. The event is pertinent, has context, is timely, and is actionable. Any discussion of “Event Management” (as opposed to “Log Management”) must take into account the four things that define the two. The discussion has to be two completely separate, albeit related topics.



About CorreLog, Inc.

CorreLog, Inc. delivers security information and event management (SIEM) combined with deep correlation functions. CorreLog’s flagship product, the CorreLog Security Correlation Server, combines log management, Syslog, Syslog-NG, SNMP, auto-learning functions, neural network technology, proprietary semantic correlation techniques and highly interoperable ticketing and reporting functions into a unique security solution. CorreLog furnishes an essential viewpoint on the activity of users, devices, and applications to proactively meet regulatory requirements, and provide verifiable information security. CorreLog automatically identifies and responds to network attacks, suspicious behavior and policy violations by collecting, indexing and correlating user activity and event data to pinpoint security threats, allowing organizations to respond quickly to compliance violations, policy breaches, cyber attacks and insider threats. CorreLog provides auditing and forensic capabilities for organizations concerned with meeting SIEM requirements set forth by PCI/DSS, HIPAA, SOX, FISMA, GLBA, NCUA, and others. Maximize the efficiency of existing compliance tools through CorreLog’s investigative prowess and detailed, automated compliance reporting. CorreLog markets its solutions directly and through partners. Visit www.correlog.com for more information.

1004 Collier Center Way, Suite 103 · Naples, Florida 34110 · 1-877-CorreLog · 239-514-3331 · info@correlog.com