Using events-per-second as a factor in selecting SEM tools

by Roberto Angelino, M.S.C.S.

Events Per Second, or EPS, as it is commonly referred to in the world of network security, is a measurement that is used to convey how fast a network generates data from its security devices (firewalls, Intrusion Detection Systems (IDS), servers, routers, etc.), and/or how fast an SEM product can correlate data from those devices. A savvy buyer will match the EPS his network is generating to those that can be accommodated by the SEM tool that he is purchasing.

For the purpose of this article, we define the EPS that can be accommodated by an SEM tool more precisely as the number of security-related events a product can receive, normalize, analyze/correlate, and display or act on in the form of results within an acceptable time frame. (For the purpose of this discussion, we will not attempt to rigorously define "an acceptable time frame"; we will instead leave that as the proverbial exercise for the reader.)

Many vendors provide EPS numbers as performance metrics for their SEM tools on their websites or in product brochures. Some vendors do not list EPS numbers, but instead list the number of devices that their products support. This direct or indirect allusion to EPS is intended to impress the prospective buyer with the performance capabilities of the product, and, beyond that, to help buyers make informed decisions that will ultimately lead to satisfaction with their purchase.

However, deriving this level of customer satisfaction requires some effort on the part of the customer. In the following pages we will help to demystify EPS by examining in detail how EPS should be measured, where it should be measured, and how this information is connected to the ultimate goal of achieving a more secure network.

Determining your EPS requirements

The amount of network security device data that is produced by a network is unique to each environment and its security policies. For this reason, in order to determine the EPS being generated by your network, it is important to first make sure that each of the security devices has been configured to generate all the log information that is required, as specified by your security policies – no more and no less. Generating more log information generates unnecessary EPS that will bog down the performance of the SEM devices; generating less EPS leaves the network potentially vulnerable and in violation of the policies.

For example, a very simple part of a security policy, and one that is used by most large organizations may entail logging all successful and unsuccessful login attempts from network devices such as routers, servers, firewalls, switches, etc. So for every one of these devices listed, a log message must be generated and sent to a logging server or SEM product whenever a successful or unsuccessful login attempt is made.

A more complex policy would include the information from the simple example above and in addition, might include logging Network Address Translation (NAT) entries on firewalls and routers. Any user traversing a firewall or router with NAT logging turned on would generate a log message for each packet/session that traverses these devices. This policy would generate significantly more events per second, and, if the information were used correctly, would also provide an additional level of information for event correlation and detection of security threats.

In either case, as soon as each of the security devices is successfully generating the correct number of log events to reflect the policy, you are ready to determine the total EPS generated by your network. The total EPS being generated across all of your devices is the EPS number that the SEM tool you are purchasing must accommodate. You should determine the total EPS being generated by your network both during normal levels of activity and at very high levels of activity, since the SEM tool should handle periodic bursts of high-activity data.

EPS and SEM tools

When a vendor provides an EPS number for its SEM product, this number conveys to the prospective buyer that it can process X number of events.

For example, the security message generated by a Cisco PIX firewall first enters the SEM product with a message such as the one below:

2004 21:19:23: %PIX-2-106001: Inbound TCP connection denied from 10.123.12.146/2493 to 172.169.110.21/113 flags SYN on interface outside
(Cisco PIX "inbound denied" message.)

The SEM device collects this data and normalizes the signature part of this message ("Inbound TCP Denied) into a format that is independent of the vendor originating the message. The signature will tell the correlation process of the SEM product that that an inbound TCP was denied. This information will then be analyzed and correlated with other messages from other devices to determine whether or not there are activities on the network that require action or attention from security personnel.

Some vendors state that they can process, for example, 12,960,000 such Events per Day. This equates to 150 Events per Second (12,960,000 events per day / 86400 seconds per day = 150 events per second). Is this good enough for a real network? That depends. How many security devices are on that network? How many securityrelated events are generated and how much network traffic is there? The frequency of messages from the network security devices needs to be examined in this context if the prospective buyer is to come to a decision regarding the viability of a SEM product on his network.

Measuring EPS on your network

Assume a firewall message is 100 characters/bytes long. On a 100 Mb network with no other traffic on it, this firewall has the potential to generate 128,000 events per second (100Mb or 12,800,000 bytes per second divided by 100 byte messages). If an SEM vendor's product can handle say 10,000 events per second, it would drop 92% of the security messages being sent to it.

Is it possible for a single firewall to generate this much data? The answer is yes. An improperly configured firewall can easily generate this much data. However, let's assume your firewall is correctly configured and, based on your security policy, it generates an average of 500 events per second. Now let's install 10 of these firewalls on a small network, with 10 routers each generating 200 events per second, and 10 IDSs each generating 1000 events per second.

This example small network now generates an average of 17,000 events per second. Calculate the EPS being generated on your own network similarly, using the number of devices and events generated by each. Compare this number with the SEM vendor's specified EPS numbers to determine whether the vendor's product will be adequate to meet your expectations with regard to SEM performance on your network.

Measuring EPS of SEM devices

Where is EPS measured? Most SEM vendors measure EPS at the point where their product accepts data from the network. For SEM vendors that provide security device agents, EPS is measured at the agent. If the SEM device is a hardware appliance, it is measured in the process that accepts the data. Why is this an important place to measure EPS? The input to an SEM device is independent of the architecture of the device. However, this is not the only place where we should measure EPS. A careful buyer should also measure EPS where the user views or responds to the results, because data paths internal to the SEM device and its analysis engines are all different. When an SEM vendor provides you with an EPS performance metric, ask the vendor if this is the rate that the tool can receive data. Then ask at which rate the tool can correlate analysis results. For example, your SEM vendor says its product can correlate 5,000 EPS. If it's measured where the data is accepted, you still need to know the delay (if any) in processing from the point the security message enters the SEM tool through the correlation process, storage (if any), and finally to the display. It may turn out

that the 5,000 EPS measured at the input to the SEM device translates to 4,000 EPS measured subsequent to correlation -- or at the display.

Is this a bad thing? Well -- again, it depends on your network. If your normal network traffic produces 5,000 EPS then this SEM tool may not meet your performance expectations since it will always be lagging 1,000 EPS. If the SEM tool is not scalable (i.e., an incremental rise in frequency and total accumulated data will slow analysis), then it probably does not satisfy the requirements: a serious network event may lag significantly behind the SEM tool's ability to analyze the problem and convey the results to the user in a meaningful amount of time.

Encryption

Most SEM tools are client/server based. This means the vendor provides an appliance and either a web interface or thick client that runs on the user's local machine. The SEM vendor should disclose whether the EPS numbers of the product reflect actual performance with or without encryption (from the appliance/server to the display). This is important because encryption is very CPU intensive and will cause EPS numbers to be lower when it is turned on.

Data filtering

When evaluating SEM tool performance on your network, it is essential to find out if the tool supports data filtering. Data filtering is a double-edged sword. It is a process by which the SEM tool can ignore a burst of device messages once a problem is detected from the originating device. It is fairly easy to use a tool like Nessus in "go-asfast- as-you-can mode" to cause an IDS to produce a lot of output. Some SEM tools have the ability to suppress data from these "noisy" devices (and to then output a message like "1500 bad messages detected from IDS ....).". Although this is a worthwhile feature, the heuristics used to determine whether or not to deploy it need to be intelligent enough to determine when a device is genuinely noisy and when a hacker is just trying to DOS (Denial of Service) the SEM tool by flooding it with IDS messages or causing it to ignore IDS messages in order to mask malicious network activity.

The bottom line is that suppressing noisy devices will improve an SEM tool's EPS numbers, since the tool will not devote CPU cycles to the attempted analysis of noise. However, this feature may increase the chances of successful DOS attacks to your SEM device.

More on hackers & denial of SEM service

The frequency of security event messages is an important factor when evaluating SEM products, not only because of your own performance expectations under normal circumstances, but also because of the potential for security messages to be maliciously generated as part of an external attack for the explicit purpose of exceeding the SEM vendor's abilities to handle them. Why is this important?

Let's assume a hacker has been slowly gathering information about your network over months of probing and has found one or two security-related vulnerabilities that are accessible from outside your network. This hacker may be very smart, and know which SEM product you've purchased. Or, maybe he's not that smart. Maybe he knows this because you've advertised it on your web site. Or perhaps your SEM vendor has advertised that you are his customer on his web site. Or maybe the hacker has access to this knowledge from the inside. (Remember, a very large percentage of network security related incidents come from inside your own company!)

At any rate, the hacker knows that your SEM product can handle 7,000 events per second. The hacker will force your security devices to generate say 12,000 events per second using a tool like Nessus. After only 2 seconds (or 24,000 events) your SEM software will have dropped 10,000 events. The SEM tool (and hence the user) is so overwhelmed by the flood of data, that he misses the real attack that is occurring -- the SEM tool, through the intentional flood of data will have dropped the one piece of security data that contained the information about the hacker's entry into the network. Because of the potential for such an attack, security personnel may never even see the actual compromise if the hacker is DOSing an SEM tool!

Why is this important? Aside from the fact you may have been compromised, which may occur with or without an SEM tool installed on a network, it's obvious that if you do not have a view into your network from the security perspective, you are doing yourself, your company and (depending on the type of business) your users a disservice. Compounding your potential for embarrassment and consequences, many companies may be required by law (Sarbanes-Oxley or Gramm-Leach-Bliley) to disclose the attack, should a compromise result in critical data leaving and/or being destroyed on computers which house this information.

For these reasons it is important that SEM vendors should not make EPS numbers easily available to the general public unless the SEM tool can handle data at wire speeds. Otherwise, the SEM vendor is educating the hackers of the world on how to DOS the SEM tool. However, it's important to convey this information in a trusted form (such as during an in-person sales presentations) since any prospective buyer needs this information to even consider evaluating a SEM product.

Scalability requirements and the resulting increase in complexity

What happens on a large network that generates 400,000 EPS? Or, more appropriately, what needs to happen? The answer: the SEM tool needs to be scalable.

All SEM tools should be scalable. If a SEM tool advertises it can handle 40,000 EPS, then the SEM vendor should provide the ability to deploy 10 SEM devices throughout the network to distribute the workload, correlating events on each device and also across devices. Scalability is a complex topic that requires in-depth discussion that is beyond the scope of this article. Measuring it requires complex testing. SEM vendors should provide tools to performance test their products if none are available.

Although EPS deals with performance, the analysis results should be scrutinized to ensure that in a multi-SEM-device solution, they are still correct.

Some SEM vendors will boast about the type of analysis their tools provide. A good rule of thumb is that the more complex the analysis, the harder it is to test. It is important that, during the evaluation process, the end user test all the analysis that is required to satisfy his company's security policy to ensure the product works as advertised. This is not an easy task even for a single SEM device since there are a lot of vulnerabilities, attacks, compromises, etc., in existence. The correlation of results from multiple SEM devices increases the complexity. Further, once the end-user is satisfied with the functional correctness of the analysis, then the performance aspect of the evaluation should be tackled using the same data but at an accelerated frequency to match that of the user's security network traffic.

Summary

EPS is to security what miles per hour is to a sports car. EPS is an easy concept to grasp since, in the context of SEM devices, it's just a number used to quantify the results that can be produced by a complex real-time correlation process. Networks and their security devices generate a certain number of events per second. In order to assure a satisfying customer experience with an SEM product, it is essential to match the EPS generated by your network with the EPS that can be correlated by your SEM purchase. The bottom line is that SEM products with higher EPS numbers at each of the relevant transition points (reception, normalization, correlation, and display) are more likely to meet the expectations and performance requirements of most networks.

The information in this article has been written for the purpose of educating the SEM tool buyer about the decision-making process that a well-informed buyer uses when evaluating SEM tools. The table below provides a set of summary questions that the prospective SEM buyer can use in determining if a SEM tool will meet his organization's security needs.

Questions to ask

These questions should be asked with respect to a configuration where one SEM tool is used, then applied to a distributed configuration where numerous SEM tools are used together to handle correlation requirements beyond the capability of one SEM tool.

- How many EPS are generated by the security devices on my network?
- What is the EPS of SEM tool I am considering?
- Where is EPS measured?
- What is the EPS to the display?
- What is the EPS of the receive process?
- What was the duration of EPS testing? (i.e., how long was the SEM tool exposed to high frequency of device data?)
- Was filtering on for EPS testing?
- Does the SEM tool support wire speed?
- Was encryption on during EPS testing?

Robert Angelino is an independent consultant who specializes in real-time software and network security.
robertangelino@yahoo.com