

>>> A new approach to network disaster protection

Generating fine-grained signatures in real-time

>>> Contents

Introduction	3
Anomaly detection	4
The need for fine-grained signatures	5
The failure of flow-based anomaly detection	6
More insight with packet-based anomaly detection	8
Removing the anomaly	10
Conclusion	11

>>> Introduction

Traditional approaches to business continuity and network disaster protection are out-matched by modern attacks and security threats. A new approach is needed.

The availability and security of today's mission critical networks is more important than ever. Corporations, government organizations and service providers rely heavily on those networks, and the ability of the attached hosts and servers to run applications and allow access to vital data.

Yet, while the intrinsic value of these networks has increased greatly, the typical approach to network security has very much remained steeped in tradition. Even in today's environment of rapidly evolving threats and custom made attacks, such as worm outbreaks or distributed denial of service (DDoS) attacks, network operations and security organizations still rely heavily on prior-knowledge of what should and should not be allowed on the network. Firewalls are set up statically, intrusion detection systems (IDS) and intrusion prevention systems (IPS) are configured with occasionally updated sets of static signatures, routers and switches are configured with static sets of access control lists (ACLs).

The consequence can be seen whenever a new worm, or a site-specific DDoS attack makes headlines - organizations are still vulnerable. Traditional network security approaches, based on signatures of known attacks or static rule sets (prior knowledge) are out-matched by modern attacks and are not sufficient.

This whitepaper explains how a dynamic and intelligent approach to network security, coupled with the built-in capabilities of modern network infrastructure devices, can be used to greatly improve the security and availability of networks, even when faced with zero-day anomalies, for which no prior knowledge existed.

>>> Anomaly detection

Traditional approaches to securing the perimeter of networks are outdated. The ubiquity of wireless access points, mobile devices and extranet links mean that there is no single demarcation line between trusted and untrusted address spaces anymore. Worm-infected devices may be unwittingly carried into the organization and connected to the network, for example.

A first step in the fight against these types of anomalies is to utilize a system for behavioral anomaly detection. These systems do not rely on signatures or prior-knowledge, but instead learn about the normal behavior and profile of the network traffic on their own, and can then notify network operators or security staff, if the network traffic shows signs of unusual activity. Typically, these systems are deployed within the network, not at its edges, and do not sit inline, but instead passively listen to network information.




Figure 1: An anomaly detection solution (green) sitting in a network – showing Espion’s solution as an example

Primitive anomaly detection solutions may observe a network for a while, and then set thresholds for various traffic types, which should not be exceeded. This approach may easily result in false positives due to the fact that network usage changes over time, new applications are added, the user-base expands, etc.

More intelligent solutions, such as Espion’s neural network powered EX3000 and SP5000 appliances, look into underlying characteristics of network traffic, in order to detect subtle shifts in how the different types of traffic behave and relate to each other. Such a modern anomaly detection solution is typically capable of detecting an anomaly within seconds.

>>> The need for fine-grained signatures

One of the key requirements in the fight against dynamic network threats, such as worm outbreaks and DDoS attacks, is the ability to extract fine-grained signatures of the anomalous traffic. ‘Fine-grained’ in this context means: Sufficient information to surgically remove the anomalous traffic, with no or only minimal impact on regular traffic.

Interestingly, many of the devices which today make up an organization’s network infrastructure, are already capable of more or less fine-grained filtering. Many routers, for example, cannot only filter IP addresses, but also protocols, ports, and even more detailed pieces of information, such as type of service fields, packet lengths, and so on.

Therefore, once an anomaly detection solution discovers an anomaly on the network, it will perform an analysis of the observed traffic, and produce a fine-grained signature of the anomaly. This signature can then be used to accurately filter the anomaly, using the existing network infrastructure elements.

The spread SQL/Slammer worm illustrates how this might have worked. This worm was spreading extremely fast, taking advantage of a vulnerability in Microsoft’s SQL server. The entire worm was contained in a single UDP packet, allowing the worm to send the scan and worm-body in that one packet. Individual worm instances generated those packets to random addresses at the rate of thousands per second. Very quickly, the network impact of this worm became so tremendous that different organizations around the world had to suspend operations, resulting in millions of dollars of loss to these businesses.

Which signature did this worm have on the network? The worm sent its packets to UDP port 1434. Therefore, a first and obvious step would be to simply put ACLs for this protocol and port on the routers and switches in the network. However, since this port is used by a legitimate service, an ACL like this would then also prevent any legitimate traffic. This, in turn, could effectively shut down the business as well, in which case the cure may be worse than the initial problem.

Adding source addresses to the ACLs may also not be sufficient, since many machines were infected and scanning at the same time. Filtering them out would require a very long, and constantly changing list of ACLs. Clearly, this is infeasible.

However, one characteristic about the worm which stood out and that was also accessible to many network infrastructure devices, was the packet length. The worm’s packets consisted of exactly 404 bytes (including IP and UDP headers). The percentage of legitimate packets, on UDP port 1434 with a length of exactly 404 bytes, is very small. Therefore, filtering not only on port and protocol, but also on length, constitutes surgical removal of the SQL/Slammer worm: The anomalous packets disappear, while the legitimate traffic continues, even though it may be very similar to the anomalous traffic.

>>> The failure of flow-based anomaly detection

Unfortunately, many anomaly detection solutions only identify very coarse information about the anomaly, such as IP addresses, ports and protocols. As a result, any removal attempt of the traffic is more akin to the use of a chain-saw when instead a scalpel is needed. Applying filtering recommendations based on information presented by those devices can possibly cause significant impact on the network's normal operation.

Let us consider another example: An e-commerce web-site is under a TCP-Syn attack on port 80. So, what traffic should be removed now? All the traffic to the web-site, as identified by IP address and/or port? The service provider or web-hoster may consider this to be a good idea, since this would keep their network operational. However, the owner of the attacked site clearly has different priorities, since such action would essentially shut down their business.

Should all TCP-Syn packets to that site's IP address and port be filtered out? Since TCP-Syn packets are needed to establish any new connection, this would not be ideal either, since any legitimate connection to the site would be prevented.

To understand why many anomaly detection systems provide only such limited information, we have to consider that most of those solutions are based on the processing of flow records, such as Cisco's Netflow. Flows are useful for many things. However, it is not possible to gain enough information from flow records to provide truly fine-grained filter suggestions and signatures for network anomalies.

A flow is a summary record about one direction of a connection between two hosts. Therefore, it contains information as to when the connection started, when it stopped, which IP addresses were involved, which protocol and port, as well as byte and packet counts. Once the connection is closed, or timed out, the flow-generating device, typically a router, will pack the flow into a packet, along with other recently generated flow records, and send it to a registered flow consumer. The flow consumer typically is a server, which receives the flow records and performs some processing on them. >>

We can see that the anomaly detection server is quite far removed from the actual anomalous traffic. It only receives an abstraction, or summary, of what really happened on the wire. It has no insight into the actual traffic. It has no means to collect packet samples. All it can base its 'filter' suggestions on is the information contained in the flow records. This information only provides IP addresses, ports and protocols. Sometimes some flag information can be gleaned from the record as well, but as we have seen, that alone does not help much. In the case of the TCP-Syn flood, a filter recommendation from those solutions may look like this:

TCP-Syn flood

Filter all TCP-Syn packets in port 80, to address a.b.c.d

More information cannot be provided resulting only in a coarse-grained recommendation. Obviously, implementing this filter would stop all new connections from reaching the web-site, essentially shutting it down, and thereby completing the mission of the attacker.

The big constraint, which is inherent in all flow-based anomaly detection solutions, therefore is that they are incapable of providing fine-grained signatures for the anomalies they detect. Only meta information about the traffic is accessible. But any filtering of anomalous traffic needs to be performed by the existing inline devices, which see individual packets. Obviously, there is a mis-match between the information that the flow-based anomaly detection solutions operate on, and the information that is available to the inline devices.

Due to this mismatch, flow-based analysis of network anomalies can only deliver part of the solution. Much more insight into the actual traffic is needed. An alternative exists in packet-based anomaly detection solutions.

>>> More insight with packet-based anomaly detection

In a packet-based approach, the anomaly detection solution does not rely on third-party summary information, such as flow records, but looks at the actual, offending packets, directly and immediately. A packet-based anomaly detection solution will be deployed where it matters - in the network, listening to the traffic on a spanning port or even via a network tap.

With complete insight into traffic, down to the individual bits in the various packets, a packet-based anomaly detection solution, such as those delivered by Espion, can discern characteristics which further identify and differentiate the offending packets from the legitimate packets.

In the case of Espion EX3000 and SP5000 appliances, it is a variety of patent-pending advanced algorithms, in conjunction with specialized neural-networks that extract common properties of the anomalous packets. These commonalities can be translated, in real-time, into signatures, which accurately describe only those packets that are part of the anomaly. Any data element in the network header, and soon also in the packet body (payload) can appear in the anomaly signature. >>



Figure X An additional differentiating characteristic of the SQL/Slammer worm was that all packets had the same length. Legitimate packets were unlikely to match this size, and therefore, a fine-grained signature that took the length into account was able to identify only the worm packets with high accuracy.

Revisiting the previous TCP-Syn flood example, and introducing all the additional information we can see through the packet-based analysis, the filter recommendation may now look like this:

TCP-Syn flood

Filter all packets that:

- are TCP-Syn AND
- have the destination IP address a.b.c.d AND
- have the destination port 80 AND
- have a TCP window size of 16000 AND
- use an initial sequence number of 1234567 AND
- have an overall length of 40 bytes

With much more information, in particular the window size and sequence number, is available, identifying the offending packets takes place with vastly improved accuracy.

This approach is not only limited to DDoS attacks, but also works very well for worm outbreaks. We discussed earlier, the SQL/Slammer worm, for example, had a very distinct packet length, which can be used to great effect. An intelligent, packet-based anomaly detection solution will make this information available in seconds, allowing worm outbreaks to be stopped in an instance.

>>> Removing the anomaly

Fortunately, more inline devices are being deployed, which can filter packets based on fine-grained signatures. Many intrusion prevention systems (IPSs) or deep packet inspection (DPI) firewalls allow for signatures of greater detail. Traditionally, these solutions have been hampered by the historical, and therefore by definition out-dated nature of the downloaded, static signature sets they are required to work with.

But now, with the arrival of intelligent, packet-based anomaly detection from Espion, it is possible to extend the usefulness of previous investments into network infrastructure, and breathe new life into the various inline systems, by supplying them with real-time, relevant signatures, even to zero-day attacks or anomalies that are unique to the target network.

Espion translate any anomaly signatures it has identified into the specific rules for a variety of well-known inline devices. The network operator can then easily select which devices they would like to see the specific signature. Once displayed, it can then be readily applied to whichever inline device is best suited to stop the anomaly.

A packet-based anomaly detection system can produce signatures based on the same kind of data, which is available to the actual inline network infrastructure elements. There is no mismatch between the types of data, as is the case with flow-based solutions. This natural fit results in much more efficient signatures, which enable the surgical removal of anomalies from the network.

>>> Conclusion

Using intelligent packet-based anomaly detection, in conjunction with already existing inline devices, such as IPSs, firewalls, but also ordinary routers and switches, allows for the surgical removal of anomalous traffic. A flow-based solution simply cannot offer the same amount of detail and insight as a packet-based anomaly detection system. Filter-recommendations that are based on flow-information are bound to be coarse, and potentially can impact the normal operation of the network to an extent that can result in complete loss of connectivity.

An intelligent and packet-based anomaly detection solution, however, can improve the ROI of existing network infrastructure elements by utilizing them to their fullest in the ongoing protection of mission critical networks.

> About Espion

Esphion protects enterprises, service providers and Governments from network disaster. Esphion's breakthrough technology uses neural, behavior-based, real-time analysis to detect known and unknown threats within seconds. It then generates fine-grained signatures, allowing you to stop even the most aggressive attacks and eliminate internal network threats and insider misuse. Launched in 2002 by industry veterans, Esphion is backed by leading venture capital firms. Its customers include Fortune 100 enterprises and major service providers and financial institutions. For more information visit www.esphion.com

www.esphion.com

sales@esphion.com

Esphion Corporate Headquarters

20 William Pickering Drive, Albany
Auckland
New Zealand
Ph: +64 9 414 2060
Fx: +64 9 415 0228

Esphion Australia

Level 6, 90 Mount Street
North Sydney
NSW 2165
Australia
Ph: +61 2 9955 3611
Fx: +61 2 9959 5760