

WHITE PAPER

HIGH-FIDELITY THREAT INTELLIGENCE: UNDERSTANDING FALSE POSITIVES IN A MULTI-LAYER SECURITY STRATEGY

Dave Dubois, Global Security Product Management

Version: 1.0, Jan 2018



CenturyLink®

A Multi-Layer Approach to Cybersecurity

Implementing a multi-layer approach to cybersecurity has been an established best practice for years. Combining elements of several network layers can create an effective protection scheme that allows access to legitimate sites, devices and applications by authorized users — while blocking access to addresses that may cause harm to the enterprise. Some examples include:

Layer 3, IP Addresses: Blocking access to an IP address that has been identified as hosting “potentially malicious” content is a simple and effective way to protect the enterprise from harmful content. However, the enterprise may also have legitimate business with other sites hosted by the same IP address, so incorporating information from additional layers will prove to be more accurate and effective.

Layer 4, Ports and Protocol Types: Incorporating protocol types and port numbers into filters can create a more precise and effective filtering scheme. For instance: If it is established that our users have legitimate business with a specific IP address (say, 101.102.103.104), which always occurs over the TLS protocol (used for encrypting data) on ports 22011 and 22012, then we can block access all to 101.102.103.104, except traffic that is using a TLS protocol on ports 22011 and 22012.

Layer 7, Domains and URLs: Incorporating domain names and URLs into your overall security strategy can also improve accuracy and effectiveness.

1. There are two major approaches at this level. One is during the Domain Name Resolution phase of the communication, where if a user attempts to get an IP address for a known malicious domain, the request is “sinkholed” and the user is given an error message indicating the requested domain is not reachable due to policy decisions.
2. The second approach is to use web content filtering (WCF), a capability found within next-generation firewalls (NGFW). This approach can be deployed using either a “whitelist” or “blacklist” method. A blacklisted domain would simply block access to a specific domain (i.e. “mailware.com”). The whitelist method is used in combination with Layer 3, for instance, “block all access to IP address 101.102.103.104 except for salesforce.com.”

Understanding False Positives

While implementing a multi-layer approach is the most effective way to develop a security strategy, it takes careful coordination between devices to prevent confusion between policies that are implemented at different layers. One typical scenario is confusion resulting from deploying web content filtering with CenturyLink® Adaptive Threat Intelligence.

At the time of this writing, Adaptive Threat Intelligence works primarily at Layer 3 by understanding which IP addresses may be hosting malicious sites. While future releases of Adaptive Threat Intelligence will incorporate domain analytics (reputation at the domain name level), current products are focused at the IP level. However, web content filtering primarily works at the Layer 7 (the “application level,” where domain names are used). This can lead to Adaptive Threat Intelligence reporting that the customer has attempted to communicate with a known malicious site, even though it was blocked by the WCF application. Customers may report this as a “false positive,” when it is more precisely a “true positive that has been mitigated.”

Let’s examine how that can happen. Before any communication can occur at the HTTP level, there needs to be a TCP connection established between the client and the server. Connection establishment takes a three-packet handshake. It’s possible that during the handshake, the Adaptive Threat Intelligence service will observe this communication and create a “threat event” for the customer that shows up in the portal reports and the syslog feed. Subsequently, the WCF application will block further communication to the potentially malicious server. Even though this occurs after the TCP handshake, it happens long before any potential damage can occur.




Figure 1 shows a simplified data flow diagram of the communication getting blocked post TCP handshake. Here we observe the end user attempting to contact a compromised server over an internet connection. The three packets that make up the TCP handshake are labeled SYN, SYN/ACK and ACK, which correspond to the TCP flags that are set in each of these packets. Once these packets are sent and received, a connection is established and HTTP traffic to the server can ensue. It's during this handshake that Adaptive Threat Intelligence will properly identify contact with the compromised server and report it as a threat event (depicted by the red star in the diagram).

Next, we observe that the end user's browser creates and sends an HTTP GET packet, which is used to request information from the server. However, the WCF application identifies this attempt to contact the potentially malicious server by its domain name and blocks the communication from proceeding further (depicted by the red X in the diagram).

In this case, while both the WCF and Adaptive Threat Intelligence services are working as designed, the customer may report this as a "false positive" because when they assign an analyst to investigate the threat event, they conclude there is nothing further that needs to be done. Security management personnel may identify this as "lost time" that the analyst could have spent pursuing unaddressed threats.

Solution: High-Fidelity Threat Intelligence

In the current release of Adaptive Threat Intelligence, we developed the ability to deploy filters to reduce the reporting of threat events that are identified solely through the TCP handshake. The screenshot in Figure 2 depicts the relevant filters.




Figure 2

The filter mode must be set to "All Options" to activate the relevant filters. The "Packet Size" filter is key, which defaults to "*" — a wildcard selection meaning "any value." It may be desirable to select specific threat events within the TCP handshake for filtering.

Figure 3 depicts the TCP Flags pulldown. The “TCP Flags” filter is used to select which parts of the handshake should be filtered. The default value here is also “*”; which will filter all TCP Packets that match any of the TCP Flags.

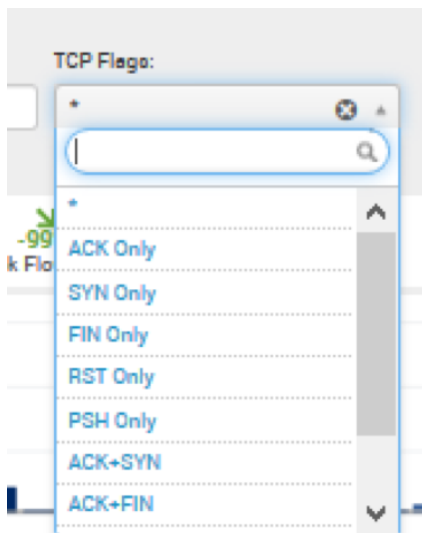


Figure 3 - TCP Flags Filters

Once you have made your packet filtering selection, use the “Packet/TCP Flags Filter Handling” pulldown to select the disposition of the threat events that match the filter. Choices are as follows:

- Include: Use the filter parameters to select which threat events should be INCLUDED in the report
- Exclude if either matches: Exclude any threat event that matches EITHER the “Packet Size” filter or the “TCP Flags” filter.
- Exclude if both matches: Exclude any threat event that matches BOTH the “Packet Size” filter and the “TCP Flags” filter.

All the packets in the TCP handshake are 40 bytes in length and there are no other TCP packets of interest of this length. One quick and easy way to filter out TCP handshakes is to enter “40” in this field, and set the value of “Packet/TCP Flags Filter Handling” to “Exclude if either match.” This will filter out all packets of 40 bytes and below, causing all threat events that were tied to TCP handshake packets to be eliminated from the report. Security teams can use this capability to identify and focus on high-priority work — threats that are not yet mitigated.

Please note, this approach will filter out ALL threat events that are based on packets in the TCP handshake — even ones that are NOT mitigated for in WCF or other applications. As such, it’s advisable to occasionally run without these filters to ensure no actual threats are being missed.

Further Considerations Concerning Threat Events

Adaptive Threat Intelligence attempts to give as much real-time information about threats as possible.

Once an IP address is entered into the threat flow, Adaptive Threat Intelligence will report all observed interactions with that address. Included in each threat event are the source port, destination port and service fields — which will reflect the information in the observed packets and may differ from the ports and services typically used by the malware.

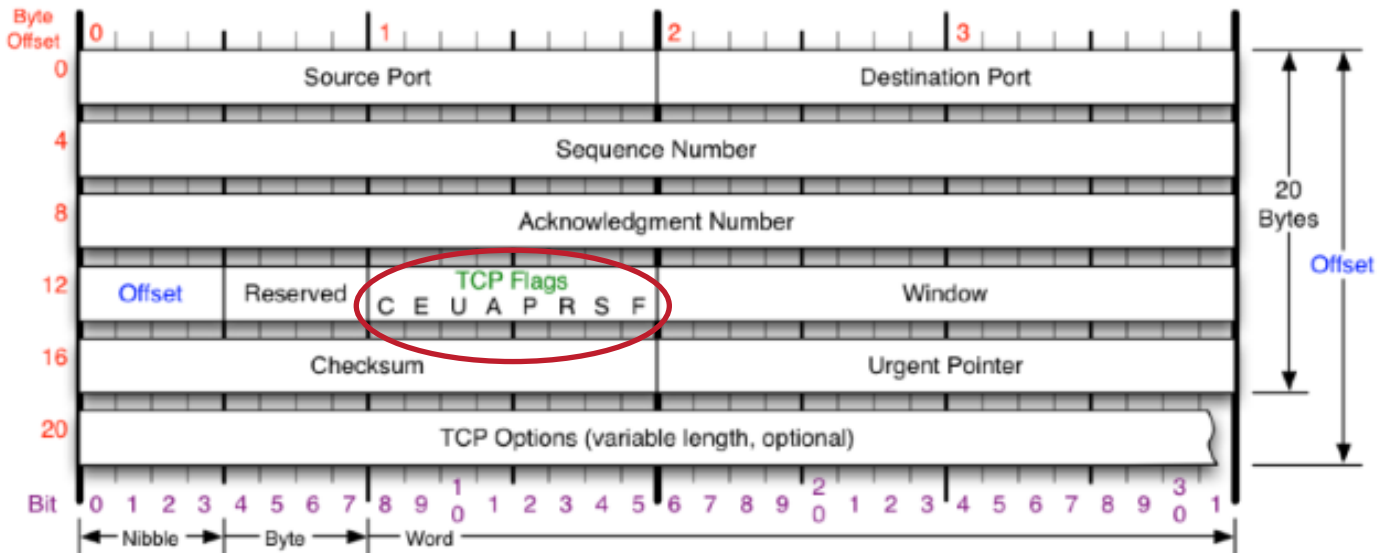
An example scenario would be if IP address 101.102.103.104 is identified as hosting a server that distributes Gafgyt malware, which carries a payload that infects Linux systems to launch DDOS attacks. Once the IP address has been identified to Adaptive Threat Intelligence as potentially malicious and until the threat is remediated, every observed interaction an enterprise has with 101.102.103.104 will result in a threat event being entered into the threat report. Subsequently, there may be threat events that will indicate interaction over port 25 (SMTP), which is not characteristic of the Gafgyt variants.

Conclusion

Using a multi-layer approach for deploying a security strategy is an industry best practice. However, not all information from these layers easily correlates into a coherent story. Having a flexible filtering strategy can help create high-fidelity threat intelligence that aids security personnel in optimizing incident response and investigation.

Reference –TCP Flags

This paper makes liberal use of the term “TCP Flag.” There is plenty of literature that explains the TCP protocol in all its glory, but for a quick reference, a diagram of the TCP packet header is depicted in Figure 4 below. The TCP Flags are circled in red.



TCP Flags	Congestion Notification	TCP Options	Offset																											
C E U A P R S F	ECN (Explicit Congestion Notification). See RFC 3168 for full details, valid states below.	0 End of Options List 1 No Operation (NOP, Pad) 2 Maximum segment size 3 Window Scale 4 Selective ACK ok 8 Timestamp	Number of 32-bit words in TCP header, minimum value of 5. Multiply by 4 to get byte count.																											
Congestion Window C 0x80 Reduced (CWR) E 0x40 ECN Echo (ECE) U 0x20 Urgent A 0x10 Ack P 0x08 Push R 0x04 Reset S 0x02 Syn F 0x01 Fin	<table border="1"> <thead> <tr> <th>Packet State</th> <th>DSB</th> <th>ECN bits</th> </tr> </thead> <tbody> <tr> <td>Syn</td> <td>0 0</td> <td>1 1</td> </tr> <tr> <td>Syn-Ack</td> <td>0 0</td> <td>0 1</td> </tr> <tr> <td>Ack</td> <td>0 1</td> <td>0 0</td> </tr> <tr> <td>No Congestion</td> <td>0 1</td> <td>0 0</td> </tr> <tr> <td>No Congestion</td> <td>1 0</td> <td>0 0</td> </tr> <tr> <td>Congestion</td> <td>1 1</td> <td>0 0</td> </tr> <tr> <td>Receiver Response</td> <td>1 1</td> <td>0 1</td> </tr> <tr> <td>Sender Response</td> <td>1 1</td> <td>1 1</td> </tr> </tbody> </table>	Packet State	DSB	ECN bits	Syn	0 0	1 1	Syn-Ack	0 0	0 1	Ack	0 1	0 0	No Congestion	0 1	0 0	No Congestion	1 0	0 0	Congestion	1 1	0 0	Receiver Response	1 1	0 1	Sender Response	1 1	1 1	Checksum Checksum of entire TCP segment and pseudo header (parts of IP header)	RFC 793 Please refer to RFC 793 for the complete Transmission Control Protocol (TCP) Specification.
Packet State	DSB	ECN bits																												
Syn	0 0	1 1																												
Syn-Ack	0 0	0 1																												
Ack	0 1	0 0																												
No Congestion	0 1	0 0																												
No Congestion	1 0	0 0																												
Congestion	1 1	0 0																												
Receiver Response	1 1	0 1																												
Sender Response	1 1	1 1																												

Call 1.877.453.8353 | Click centurylink.com | Email info@centurylink.com

This document is provided for informational purposes only and may require additional research and substantiation by the end user. In addition, the information is provided “as is” without any warranty or condition of any kind, either express or implied. Use of this information is at the end-user’s own risk. CenturyLink does not warrant that the information will meet the end-user’s requirements or that the implementation or usage of this information will result in the desired outcome of the end user. This document represents CenturyLink’s products and offerings as of the date of issue.



Services not available everywhere. CenturyLink may change or cancel products and services or substitute similar products and services at its sole discretion without notice. ©2018 CenturyLink. All Rights Reserved. The CenturyLink mark, pathways logo and certain CenturyLink product names are the property of CenturyLink. All other marks are the property of their respective owners. 18552647