COMP3010 Lab

Intrusion Analysis --- Host Infection --- Answers

5. Host Infection: Mass Mailing Malware

1. Is this an attack? If yes, how do you know? Which IDS alerts suggest this is the case? (Hint: do the IDS alerts capture the offending traffic or the responses to it?)

Yes, it is. The dataset contains traffic generated by a mass mailing malware. When analyzing the dataset with Snort in SecurityOnion, it triggers alerts such as:

- "ET SMTP Spamcop.net Block Message" (sid:2012983), and "ET SMTP Abuseat.org Block Message" (sid:2012982). These have been generated by SMTP servers (as the source port of sender is 25/tcp), and contain warning and error messages on the presence of spam traffic that has been blocked.
- "ICMP Destination Unreachable Communication with Destination Host is Administratively Prohibited" (sid: 200000486). These alerts suggest that the recipient has violated an access control policy at the sender's network. Looking at the payload of these error messages, we can see that the original datagram that caused the errors was SMTP traffic with destination port 25/tcp. Therefore, we can assume that the ICMP error messages have been generated as a result of the spam activity.

The IDS alerts from SecurityOnion have not flagged the original spam traffic, but the warning and error messages that have been generated by external hosts/routers as a result of the spam traffic. The recipient of all these error messages is internal host 192.168.0.103

2. Which hosts are affected? It is essential to determine the target. How do you do that?

The recipient of all the error messages that triggered the IDS alerts above is internal host 192.168.0.103. This suggests that it has been sending spam traffic. To verify this, we go back on the dataset and filter all traffic with source IP 192.168.1.103 and destination port 25/tcp ('src host 192.168.0.103 and tcp dst port 25') to verify that spam emails were indeed sent by the infected host.

3. Immediate step to take infected host offline. (No action needed for you)

4. How does the malware work and propagate? (Hint: You might want to open the dataset with Wireshark. You might want to use features, such as 'Statistics -> Protocol Hierarchy').

There are 16 snort alerts at Securityonion entitled 'DNS SPOOF query response with TTL of 1 min. and no authority' (sid:200000254). These alerts indicate DNS responses without authoritative answers and low TTL value of 1 minute. Indeed, looking for dns traffic through wireshark (via display filter udp.port==53, or via Protocol Hierarchy), we can identify several DNS responses, with non authoritative DNS A and MX records, and very low TTL values. This is an indication that the malware could be using fast-flux functionality, which allows it to update DNS records of domains with newly infected IP addresses. Indeed, all the domains that were returned relate to spam and phishing domains.

Snort alerts in VM also pick up suspicious web traffic, which suggests that the infected host is posting zip files on a web server. In order to view objects (data files) exported over http, you can select menu option 'File -> Export -> Objects -> HTTP', where you will see several files being posted to web server go-thailand-now.com (IP Address 64.120.144.69). Retrieving and clearing up these data files would enable further analysis of the malware activity. Short of doing that, we can Follow TCP Stream by selecting any relevant packet, and reveal that the infected host is uploading and receiving data (we can suspect these are updates) to / from the server.

Apart from the http data traffic, there is UDP data being sent from the infected host to 8.3.211.74. You can select and display this traffic from wireshark, via the Protocol Hierarchy. These udp packets are sent at regural intervals (every 20 seconds) and contain the same 20 bytes of data. They all use source port 1025/udp and destination port 8102/udp. This could suggest that a "heartbeat" is sent to a control server, signifying that the infected host is still alive.

The Protocol Hierarchy from wireshark does not suggest any other propagation mechanism, other than email (and possibly http).

5. What is the extent of the damage? Are there any other hosts affected? How would you find out? (Hint: You might want to use BPF tcpdump filters)

To determine if any other internal hosts have been affected, we need to filter for all traffic between 192.168.0.103 and other internal hosts within the same IP range. We could use the following tcpdump filter:

'src host 192.168.0.103 and dst host 192.168.0 and dst host not 192.168.0.1'

which filters traffic from source the infected host and destination within the 192.168.0.0-255 IP range – the IP gateway 192.168.0.1. is excluded.

The results return no packets, which shows that no other internal hosts were involved.

As an additional check, we can search for any traffic involving the 2 suspicious IP addresses, which were discovered earlier, 64.120.144.69 and 8.3.211.74. By applying a tcpdump filter 'host 64.120.144.69 and !192.168.0.103' on the dataset, we can filter any traffic between 64.120.144.69 that does not involve the infected host. A similar filter for 8.3.211.74 reveals that no other host communicated with the 2 suspicious servers.

6. Host Infection: Witty and Slammer Worms

In this section we will examine 2 more datasets, as follows:

Worm Dataset Files	
Worm	**Location (also available at module DLE site)**
Slammer (2003)	`\\seclab-pc30\temp\slammer-short.pcap`
Witty (2004)	`\\seclab-pc30\temp\witty.pcap`

Both files represent "textbook", cleaned up examples of the Slammer and Witty worms. As they do not contain the full packet payload, nor exceed Snort's threshold for portscan traffic, they are not detected by Snort. Nonetheless, it is still useful to examine them using Wireshark, and see how they propagate.

1.Which port number does each worm use to propagate?
Slammer propagates over 1434/udp, whereas Witty uses 4000/udp as source port.

2. Which hosts are infected?

To identify the infected hosts in slammer, you need to identify all source IP addresses sending packets on destination port 1434/udp. You could also view the list of distinct source IPs by viewing Conversations.

The witty dataset contains packets from infected host 208.195.0.5.