

Scan 17

Your challenge is to review and analyze a month's worth of data collected by a Honeynet. All submissions are due no later than 17:00 CST, Friday, 20 July. Results will be released Monday, 22 July.

NOTE: On 1 July we released the challenge with a corrupted README file (don't ask how we were able to corrupt a simple .txt file :). On July 8 we repackaged the download without a README file. This new download contains the exact same data, just minus the README. The README was nothing more than a simple .txt file that restated information already found on this page. The only additional information it had was the firewall logs were not logging NetBIOS scans for that month.

The Challenge:

The past several Scan of the Month challenges have focused on a specific attack or incident. This month is different as we review and analyze blackhat activity over an entire month. What makes Honeynet data unique is reduced number of false positives and false negatives. False positives are when production activity is incorrectly flagged as suspicious. Many organizations are overwhelmed with data, they have difficulty determining what data is production, and what data is suspect. Administrators become overwhelmed with false positives and begin ignoring alerting mechanisms. False negatives are when organizations fail to detect malicious or unauthorized activity. For example, a new attack may be launched, but the organization is not alerted because no signatures exist for such an attack. A Honeynet addresses these two problems by having little or no production traffic, thus all traffic is suspect by nature, captured, and alerted to. This helps create a cleaner and easier to use data set.

Below is the data capture for the month of November, 2000, the same month as from the [Forensic Challenge](#). We provide the data in two layers, [IDS Snort](#) alerts and firewall logs. Snort alerts are generated when suspect traffic matches an existing signature. Firewall logs each *unique* inbound scan in a 24 hour period. The first inbound packet of every unique source is logged. If a system sent one packet or 10,000 packets in a 24 hour period, they would only be logged once. For a better understanding of these data capture methods, we recommend you review the whitepaper [Know Your Enemy: Honeynets](#).

This data was captured from an 8 IP Honeynet connected to a home ISDN line. Nothing was done to attract or lure attackers. Three honeypots were live during this month. We rotated and rebuilt the three boxes on 26 November, at approximately 18:30. Note: no NetBIOS traffic was logged by the firewall at this time. We disabled the logging of NetBIOS traffic on purpose, as we were overwhelmed by the scans, and had already determined the nature of their origin.

Date Brought Online	honeypot operating system	honeypot IP address
November 4	Red Hat 6.2 server	172.16.1.107
November 5	Sun Microsystems 2.6 Sparc	172.16.1.101
October 31	Windows98 Desktop	172.16.1.105
November 8 (<i>same system, rebuilt</i>)	Windows98 Desktop	172.16.1.102

November 26	Sun Microsystems 2.6 Sparc	172.16.1.103
November 26	Red Hat 6.2 server	172.16.1.104
November 26	WindowsNT SP4	172.16.1.106

Data from this month can be downloaded here:

[SotM.tgz](#) MD5 = bf03e549996cc5edac2fda74c54d861d

[SotM.zip](#) MD5 = 3e01eba6b9f6341669568eadcd895fe1

Your challenge is as follows. Please be sure you show your data analysis methods. We want the security community to learn from you and to apply your methods to future data captures. Best of luck!

1. What trends did you identify?
2. What does this activity tell us about the blackhat community?
3. What if anything happened in the firewall and IDS logs that gave us a clue of what was coming? Could any of the attacks been predicted ahead of time. If so, how?
4. What data did you find more valuable, the Snort alerts or the firewall logs of unique scans? Why?
5. What lesson did you learn from this?
6. How long did this challenge take you?

Bonus Question:

Both the Snort alerts and the inbound firewall logs missed a successful attack. The only reason the HoneyNet Project detected this successful attack was because the compromised system attempted an outbound connection (by definition this means a system was compromised). The HoneyNet Project did a writeup on this incident, can you identify the attack and why the HoneyNet failed to alert to the attack (though the attack was captured). Hint: the attack is written up and posted somewhere on our site.

The Results:

Writeup from the HoneyNet Project members.

✍

Writeup from the Security Community

Top Ten

✍

