The importance of optimizing DNS filtering when under attack

Wes Hardaker

March 2, 2024

1 Common operational responses to attacks

Malicious actors target commonly target an organization's production DNS servers when launching attacks. Because the DNS is the common bootstrap mechanism for all Internet based communication with an organization, successful attacks that take their DNS service offline frequently results in completely disabling the organization as a whole.

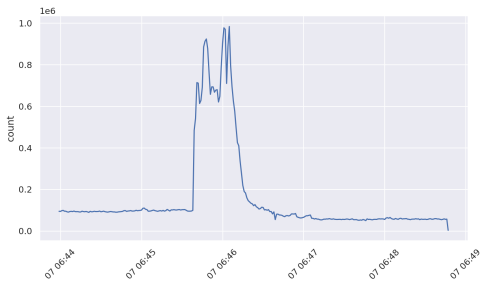

Figure 1: *A small DDoS attack against the USC/ISI root server*

When operators respond to these events, their first course of action is frequently to identify the characteristics of the attack and to apply network filters either on site or upstream. When under pressure, often only the first-order visible elements of attack traffic are placed into these filtering profiles, which may result in both false positives (filtering out legitimate queries) and false negatives (failing to filter out some attack traffic). Figure 1 shows an example DDoS attack against the USC/ISI root server. The predominance of this attack was observed to be random queries with packet lengths of 554 bytes.

2 Creating more specific filters

To combat thees false positives and negatives, we argue that looking for second order effects and other protocol characteristics will help improve the filtering accuracy. Although initial filters may need to be put in place that act as a first order prevention system, we demonstrate the need to examine other common characteristics of attack traffic that allow for more precise updates to existing filters to improve the accuracy in defense responses. Figure 2 shows a missed attack signal (a false negative) that would not be caught by filtering only 554 byte packets.

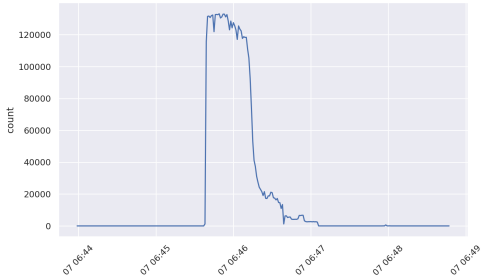

Figure 2: *A false negative with additional observed queries for www.example.com*