Gaining Insight into AS-level Outages through Analysis of Internet Background Radiation

Karyn Benson, Alberto Dainotti, kc Claffy*
CAIDA, UCSD
{karyn,alberto,kc}@caida.org

Emile Aben
RIPE NCC
emile.aben@ripe.net

Categories and Subject Descriptors
C.2.3 [**Network Operations**]: Network Monitoring

Keywords
Outages, Packet Loss, Network Telescope, Malware

1. INTRODUCTION

Internet Background Radiation (IBR) is a mix of unsolicited network traffic mostly generated by malicious software, e.g., worms, scans, which is captured by darknets [7]. In previous work, we extracted signal from IBR traffic that allowed us to *identify* large-scale disruptions of connectivity at an Autonomous System (AS) level due to government censorship or natural disasters [4, 5]. Here we explore another IBR-derived metric which may *provide insights into the causes* of macroscopic connectivity disruptions. Specifically, we propose a metric indicating packet loss (e.g., because of link congestion) from a specific AS.

2. ANALYSIS

We analyze traffic captured at the UCSD Network Telescope, a /8 darknet [1]. Because a darknet receives but does not respond to traffic, when an external host attempts to open a TCP connection, the corresponding flow (defined as the traditional 5-tuple with a timeout) is comprised solely of SYN retransmits. It is expected that the number of SYN retransmits is consistent across hosts using the same application and underlying OS. Consequently, drastic changes in the number of packets per flow associated with a specified application and OS may reflect packet loss.

Based on this observation, we attempt to extract a strong (statistically significant) and stable (low noise) signal from IBR indicative of the pattern of SYN retransmits. We selected Conficker-like traffic (i.e., IBR to TCP port 445) as a candidate for signal extraction because it constitutes a large percentage of the packets collected at the UCSD telescope – more than 40% – and is a globally pervasive component of IBR [3].

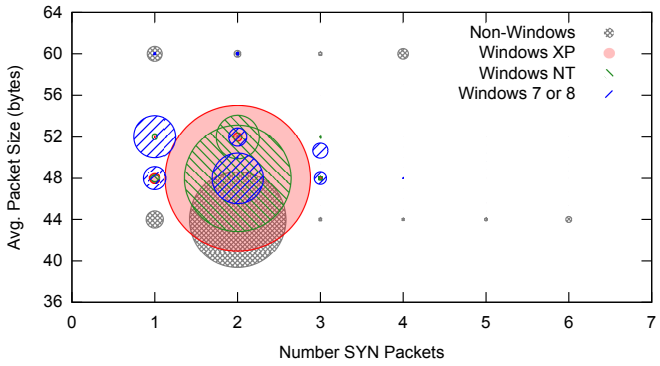

Figure 1: Packet size vs. number of packets by OS (Jan. 2012) for Conficker-like traffic. The radius of each circle is proportional to the percentage of SYN flows for that OS, packet size and number of packets. Flows of more than 7 packets (very infrequent) were excluded from this graph.

Figure 1 illustrates the distribution of SYN flows destined to TCP port 445 as a function of packet size, number of retransmits, and operating system (as identified by *p0f* signatures [8]). To obtain a strong and stable signal, we select only flows (i) from Windows XP and Windows NT (respectively 90% and 6.6% of the total number of flows) and (ii) with packet size of either 48 or 52 bytes. Most of the selected flows contain two SYN packets, consistent with the behavior of Conficker-infected hosts [3].

Using these selected flows, the first metric we tried was the average number of SYN packets per flow. However, when applying this metric to flows sent by a set of source IPs, e.g., computing SYN packets per flow from all IP addresses that map to a specific AS, a single host or flow could skew the average. For example, when a single host conducts a horizontal scan by sending one packet to every IP address in the darknet, the AS-level average is approximately 1 packet per flow regardless of other host activity in the AS. Conversely, a single flow consisting of many SYN packets could distort the average in the other (increasing) direction.

The following improvements reduce the impact of such anomalies: (i) we exclude all flows with more than three SYN packets (97% of Conficker-like flows had three or fewer SYN packets); (ii) we calculate the average number of packets per flow for each distinct source IP, and compute the mean of this distribution, thus limiting the influence of a single source IP. In the following, we call the resulting metric γ.

Figure 2 shows our metric across all ASes during January 2012[1] calculated in time bins of one hour. Both the number of source IPs

*This material is based upon work supported by the National Science Foundation under Grant No. 1059439 Support for the UCSD network telescope operations and data collection, curation, analysis, and sharing is provided by DHS S&T NBCHC070133 and UCSD.

CoNEXT Student'12, December 10, 2012, Nice, France.

[1]The telescope was down for about 40 hours starting on 2012-01-14 and for about 120 hours starting on 2012-01-19.

and γ approximately follow a sinusoidal pattern with a phase of one day. The value of γ is always between 1.98 and 2.02.

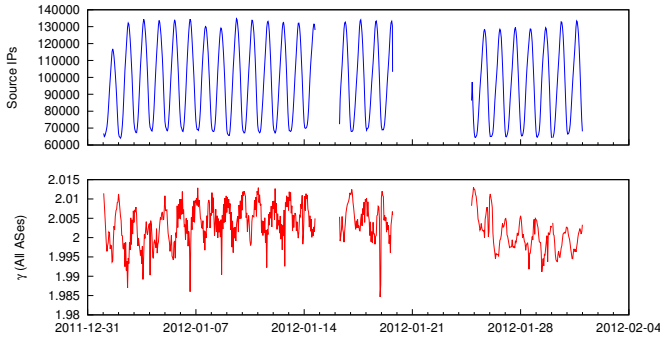

Figure 2: γ and number of source IPs for all ASes (Jan. 2012)

To evaluate the ability of our metric to distinguish service disruptions due to packet loss from those due to complete disconnection of infrastructure, we apply it to two case studies: the "Dodo-Telstra" routing leakage [2,6], and the Libyan Internet blackout [5].

On February 23, 2012, a multi-homed network operator Dodo announced internal BGP routes to its provider Telstra, a major ISP in Australia, which erroneously accepted them. As a result, Telstra sent all of its traffic to the small network, Dodo, instead of a large transit provider, inducing a bottleneck that resulted in a complete outage [2,6]. Figure 3 shows that immediately before the outage, γ dropped from about 1.82 to 1.25 (in the previous month $\mu = 1.88$ with $\sigma = 0.0874$) even though the number of source IPs remained around 12. Such a significant drop in γ is a direct consequence of the congestion on the affected links. Routers started dropping packets, including some of the SYN packets sent by conficker-infected hosts. Once the congestion deteriorated to a complete outage, the telescope did not observe any sources sending SYN packets.

Our second case study applies this metric to the Libyan Internet blackout in February and March 2011, when the Libyan government used BGP disconnection and later packet filtering to implement nationwide censorship [5]. Figure 4 shows that when a subset of hosts are allowed communication through the filtering system, the level of γ returns to values near to pre-censorship (despite having fewer source IPs sending traffic), showing that the outage was not caused by an event inducing packet loss.

In both case studies, the metric γ provided insight into the nature of the outage. In the Dodo-Telstra case, network congestion preceded the outage. In response to congestion, the network dropped packets, which decreased the number of packets in each flow, lowering the value of γ. In the Libya case, although the traffic volume was smaller, γ stayed approximately the same as pre-censorship levels whenever Conficker-like traffic was observed. This is consistent with filtering packets by subnet: the number of sources decreases but nature of the allowed communication, such as packets per flow, does not change.

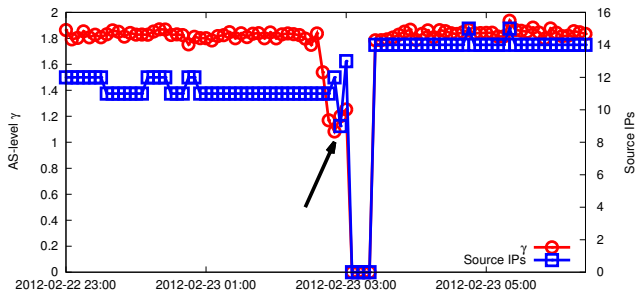

Figure 3: In the Dodo-Telstra BGP incident, γ decreased when there was a bottleneck although the number of Source IPs stayed the same.

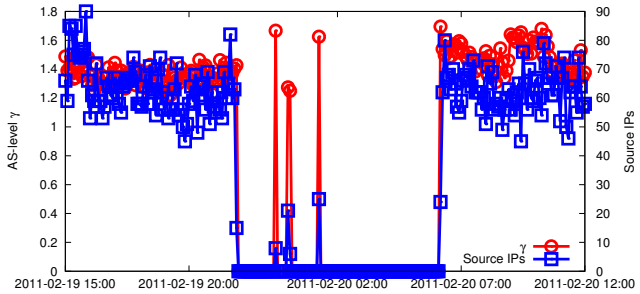

Figure 4: Libya's second outage was caused by packet filtering. γ returns to pre-censorship levels when some hosts are allowed through the firewall.

3. CONCLUSION

We are exploring an IBR-derived metric to help characterize connectivity disruptions due to transit bottlenecks that induce packet loss, e.g., link congestion. Our metric is based on SYN retransmits, which are often associated with packet loss. Here, however, the retransmits are caused by the darknet not responding to the SYNs, so we actually infer packet loss if some of them are missing.

There are several limitations to our metric: it only measures the packet loss between a given source and our darknet. It also relies on the presence of Conficker-like traffic. But our simple metric applied to a large darknet enables us to monitor one aspect of network connectivity (i.e., reachability to our darknet) from all over the world.

We used two case studies to show that this metric can distinguish a transit bottleneck-induced outage from an intentional nation-wide disconnection caused by packet filtering. In the future we would like to test this metric on other connectivity scenarios and other darknet traffic, explore other similar IBR-related metrics to characterize network disruptions, and integrate such metrics into a system for comprehensive detection and diagnosis of such disruptions.

4. REFERENCES

[1] UCSD Network Telescope, 2010. `http://www.caida.org/data/passive/network_telescope.xml`.

[2] How the internet in australia went down under. `http://bgpmon.net/blog/?p=554`, Feb. 2012.

[3] E. Aben. Conficker/Conflicker/Downadup as seen from the UCSD Network Telescope. `http://www.caida.org/research/security/ms08-067/conficker.xml`, 2008.

[4] A. Dainotti, R. Amman, E. Aben, and K. C. Claffy. Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet. *SIGCOMM Comput. Commun. Rev.*, 42(1):31–39, 2012.

[5] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of Country-wide Internet Outages Caused by Censorship. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement*, IMC '11, pages 1–18, New York, NY, USA, 2011. ACM.

[6] G. Huston. Leaking routes. `www.potaroo.net/ispcol/2012-03/leaks.html`, 2012.

[7] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston. Internet background radiation revisited. In *Proceedings of the 10th annual conference on Internet measurement*, New York, NY, USA, 2010. ACM.

[8] M. Zalewski. p0f v3. `lcamtuf.coredump.cx/p0f3/`, 2012.