Light coming from the Dark(nets)

Why would anybody send a message to an address with no occupant? And, why would anyone have interest in receiving such a message? Could one expect anything useful to get delivered to a mailbox installed in front of a vacant lot, much less many mailboxes across a whole city full of vacant lots? Since 2003, CAIDA has collected messages on the Internet in just such a fashion. In this case, the mailbox(es) take the form of a network router instrumented to collect packets destined to a large IP address range that has no devices assigned or attached. Network researchers describe such address space as blackholed or ``dark'', and use network telescopes to observe the substantial amount of traffic destined to various segments of non-existent address space (or darkspace). When collected and analyzed, this data (available through the PREDICT portal[1]) can provide insights to network researchers studying the global Internet.

Since no devices exist in this address space, all messages destined to these addresses are unsolicited and unwanted. This unsolicited traffic, commonly referred to as Internet background radiation (IBR), is useful for the study of malware, botnets, country-level censorship and as potential early warning signals for misconfigurations and Internet outages caused by natural disasters.

The larger image in this poster presents a heat map visualization of our darknet data collection volume. Each vertical bar represents one day of data divided into 24 hourly segments. The color of each segment reflects the size of a compressed file of traffic captured during that hour of the day. We color each data point based on its deviation from the median hourly captured traffic file size, so that hotter colors mean more data.

Data collection on the telescope is a best-effort service -- outages show up as vertical black bars. The early part of the collection, from 2003 into 2008, contains extremely sparse trace data that aligns with CAIDA's Backscatter datasets. Starting in 2008, we retained more complete trace data. Diurnal patterns are cleary visible in the data, with brighter colors dominating the 08:00 to 20:00 (UTC) bands. The images around the heatmap describe research efforts using various samples of this data collection over the last decade.

1.Protected Repository for the Defense of Infrastructure Against Cyber Threats (PREDICT). https://www.predict.org

Conficker

The first serious evidence of the Conficker worm outbreak was reported on November 22, 2008. Conficker engages in different types of observable network scanning via TCP port 445 for additional victims, including random scanning. The UCSD Network Telescope observes a significant fraction of the random scans. This figure plots the unique IP source addresses per day scanning the telescope address space since October 2008, with day-0's marked for the two first variants labeled (several outages of the telescope caused gaps in data). We believe this data[1] is one of the few datasets that provides information on the first few hours of the spread of the Conficker worm.

[1] http://www.caida.org/research/security/ms08-067/conficker.xml

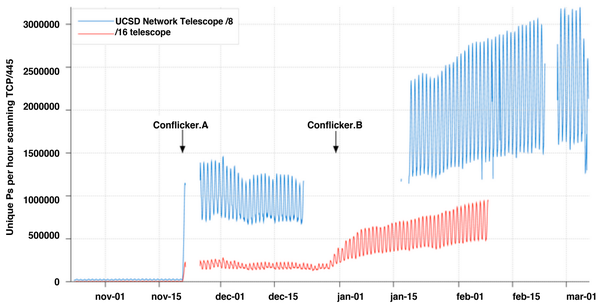

Censorship in Libya

In February 2011 protests erupted in Libya, calling for an end to the Gaddafi regime. On the night of February 18 the government imposed an "Internet curfew", blocking all Internet access until morning (08:01 local time), and repeating it the next day. This figure depicts these two events as observed at the UCSD Network Telescope, showing the packet rate of unsolicited traffic from hosts geolocated to Libya. This analysis is part of a study[1] that combined passive measurements from the Telescope with active measurements and data from BGP updates, which allowed a detailed reconstruction of the country-wide Internet filtering events happened in Lybia and Egypt at the beginning of 2011.

[1]http://www.caida.org/publications/papers/2011/outages_censorship

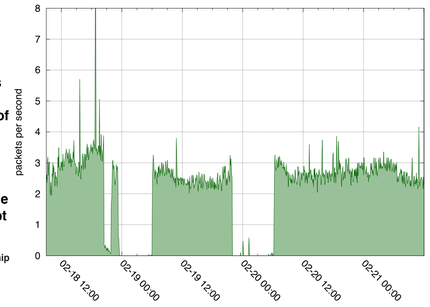

Heat Map

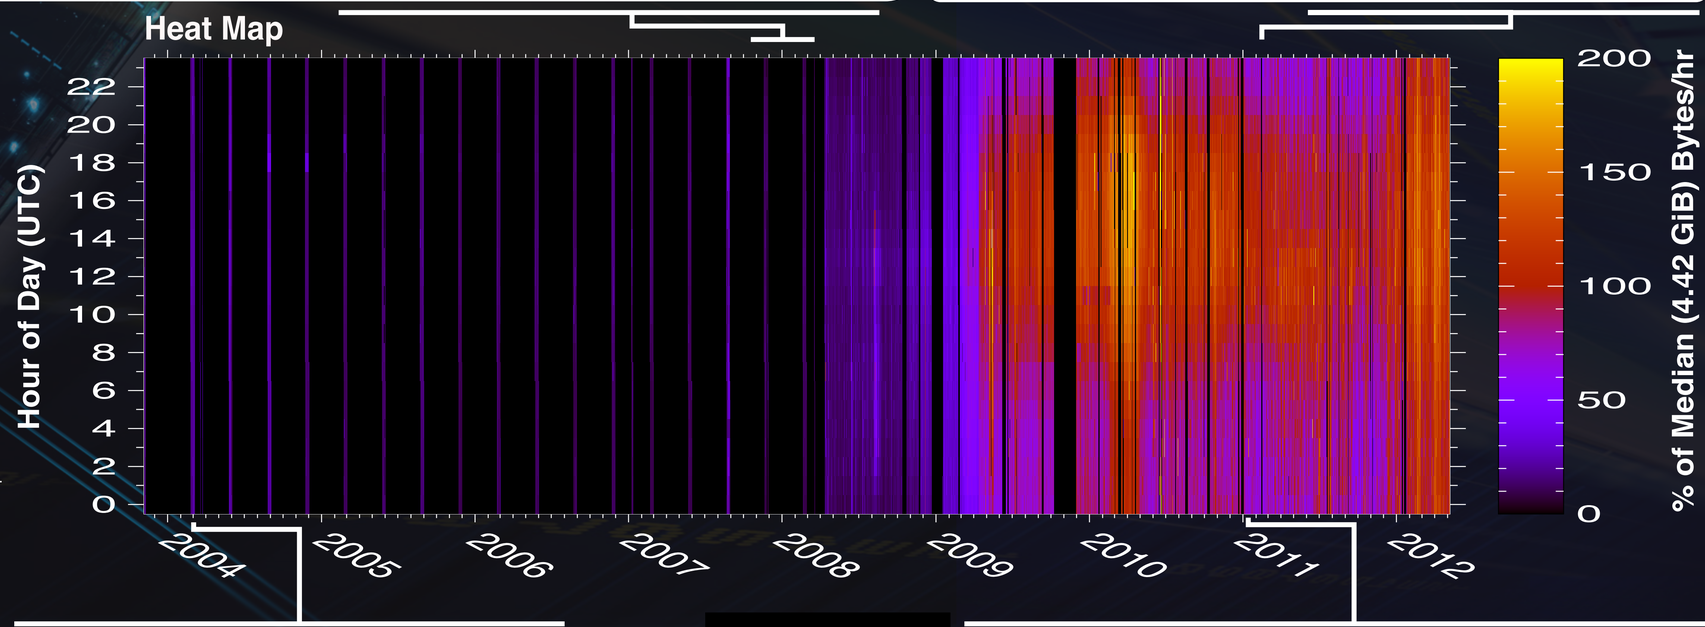

Witty Worm

On Friday March 19, 2004, an Internet worm began to spread, targeting a buffer overflow vulnerability in several Internet Security Systems products with firewall and intrusion detection capabilities. Witty was the first widely propagated Internet worm to carry a destructive payload and spread through a host population in which every compromised host was doing something proactive to secure their computers and networks. Witty spread through a population almost an order of magnitude smaller than that of previous worms, demonstrating the potency of automated worms to rapidly compromise an entire population (of machines running a given piece of software) on the Internet, even in niches without a software monopoly. This image is a snapshot (Fri Mar 19 21:00:36 2004 PDT) from an animation representing the spread of the Witty worm in the United States. The animation is part of a study[1] of this worm based on the analysis of data from the UCSD Network Telescope.

[1]http://www.caida.org/research/security/witty

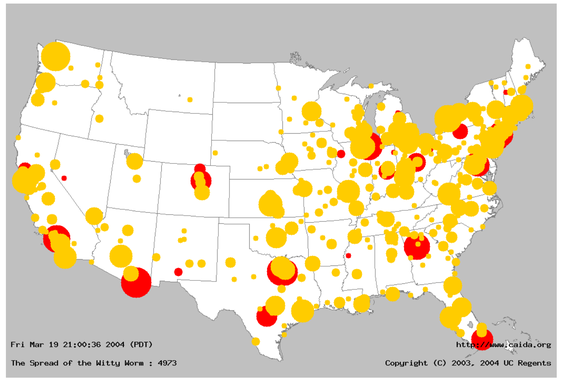

Sipscan

On February 2011 a botnet known as Sality performed a heavily-coordinated scan lasting 12 days and involving about 3 million distinct source IP addresses. The probing activity, looking for SIP servers, probably targeted the entire IPv4 address space with a reverse-byte-order progression of the IP addresses.

This image is a snapshot (Wed Feb 2 09:34:00 2011 UTC) of our "World Map" animation visualizing the spatial and temporal dynamics of this scan. Circles show the geographic coordinates of bots participating in the scanning activity. The size of a circle is proportional to the number of hosts and the color indicates the number of packets sent.

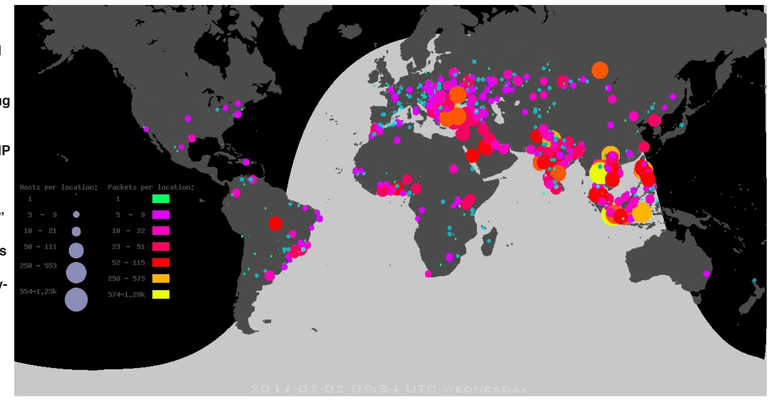

Analysis Team

Emile Aben kc Claffy
Alberto Dainotti Alistair King
Bradley Huffaker David Moore
Colleen Shannon

Poster Design

Justin Cheng Bradley Huffaker

San Diego Supercomputer Center, University of California San Diego

9500 Gilman Drive, mc0505, La Jolla, CA 92093-0505 | 858-534-50000 | www.caida.org/