2nd Workshop on Steps to Reducing Unwanted Traffic on the Internet (SRUTI '06)

July 7, San Jose, CA

Thanks to sponsorship by AT&T, Google, and Microsoft Research

Summarized by Anirudh Ramachandran and John Bethencourt and edited by Balachander Krishnamurthy

Rob Thomas of Team Cymru began the workshop with a scintillating keynote address on the underground economy. Although much of the research community working on unwanted traffic issues has focused on technical aspects of various subproblems, Rob brought his direct experience with ongoing study of the underground economy dominated by the criminal elements trading in credit cards, passwords, and the like. He painted a grim picture of the underground economy and stressed the need for a closer examination of activities common in that world that are largely unknown to the research community.

The talk was laden with various anecdotes and examples chosen from actual IRC sessions, but the thrust was to convey a sense of the breadth of the activities across the financial underpinnings of the worldwide economy, which is increasingly dependent on the Internet. Periodically, he delved into details of some of the attacks, motivations of the criminals behind them, and their varying technical expertise. Participants often peddle compromised hosts in more valuable domains (such as .mil and .gov), all the way up to gigabit backbone routers. The information gathered from these hosts, such as bank or credit card details and entire identities (Social Security cards, birth certificates, and visa information), are traded.

The offers and exchanges are performed on robust IRC servers. The servers are often, but not always, hosted in countries with lax cybercrime laws. Honor among thieves is missing but attempts to swindle are met with retributions in the form of attacks or, more often, documentation of the fraud and banishment from the trading community. Team Cymru works closely with many partners to reduce the overall threat level. However, there is significant pessimism, given the extent to which criminal elements with significant profit motives are able to stay ahead in the technical arena, where an overwhelming majority of users are less knowledgeable and thus susceptible to social engineering.

- ***The Rising Tide: DDoS from Defective Designs and Defaults***

  Richard Clayton, University of Cambridge, Computer Laboratory

  The first technical paper session focused on flooding and Distributed Denial of Service (DDoS) attacks and its mitigation strategies. Richard Clayton discussed flooding arising from defective software and firmware designs. Defective design in software results in unwanted traffic, such as malformed DNS traffic to root nameservers, Network Time Protocol traffic from hard-coded domain names, etc. Flawed designs of components assumed to be secure, such as wireless routers, cause unwanted attack traffic to hosts on the Internet. The author warns of other possible sources of flooding (i.e., not from compromised hosts) resulting from design flaws. He suggests distributing services, out-of-band authorization, education, and economic disincentives as some of the ways to mitigate flaws in design.

- ***Efficient and Secure Source Authentication with Packet Passports***

  Xin Liu and Xiaowei Yang, University of California, Irvine; David Wetherall and Thomas Anderson, University of Washington

  This paper discusses the design of a packet "passport" system to securely authenticate the source of a packet. The goal is to prevent source address spoofing and help filter unwanted hosts in DDoS attacks. The technique involves generating a Message Authentication Code for each packet over the spoofable fields of a packet header, using preshared keys to perform encryption. The paper also addresses complications with such a scheme, including key distribution, preventing replay attacks (using bloom filters), and secure bootstrapping using shim headers piggybacked on ordinary BGP route announcements.

- ***Cookies Along Trust-Boundaries: Accurate and Deployable Flood Protection***

  Martin Casado, Stanford University; Aditya Akella, University of Wisconsin, Madison; Pei Cao, Stanford University; Niels Provos, Google; Scott Shenker, University of California, Berkeley, and ICSI

  In this paper, the authors propose "flow cookies" for limiting DDoS attacks, building on previous work on capabilities and fil-

tering. The idea is a "hack" to TCP by using the currently unused timestamp field in the header coupled with syn-cookies. It uses cookie middleboxes stationed ahead of susceptible servers to complete TCP handshakes with clients and issue temporary capabilities. The proposal includes the notion of a "trust region"—ISP A trusts ISP B if A is a customer of B—to facilitate broader, incremental deployment of flow cookies along trust boundaries, thus providing benefit for the autonomous systems that choose to deploy it.

In sum, the papers presented in this session addressed both the cause and effect sides of the DDoS problem and proposed novel, incrementally deployable network-based solutions for DDoS mitigation.

The second paper session considered several abuses of resources, along with the difficulties in making accurate measurements of unwanted traffic.

- ***Separating Wheat from the Chaff: A Deployable Approach to Counter Spam***

  Youngsang Shin, Minaxi Gupta, and Rob Henderson, Indiana University; Aaron Emigh, Radix Labs

  The authors proposed two new techniques for spam filtering aimed at reducing Mail Transfer Agent (MTA) processing load. The first technique—token-based authentication—involves adding a cookie-like token as a new header in outgoing messages. Any replies to such messages will naturally include the token header and may be delivered immediately without further spam filtering efforts. The second technique, history-based prioritization, utilizes various characteristics of an incoming SMTP connection from another MTA that may be considered before the message content is processed. Upon receiving a new connection, a receiver that maintains a history of past connections can examine whether that server has contacted it before, how long ago, whether it sent spam earlier, etc. Based on this information, some messages may be deemed unlikely to be spam and are delivered immediately. Other messages may be queued up for processor-intensive filtering. An experimental evaluation of seven months of email data in a university environment revealed that this technique provides 90% accuracy in spam identification without processing the message body. Reducing the number of messages that must undergo expensive spam detection algorithms (which may take as long as 10 seconds per message) reduces delivery latency.

- ***Strider Typo-Patrol: Discovery and Analysis of Systematic Typo-Squatting***

  Yi-Min Wang and Doug Beck, Microsoft Research, Redmond; Jeffrey Wang, PC-Rethinking.com; Chad Verbowski and Brad Daniels, Microsoft Research, Redmond

  The authors examined the practice of registering domains similar (in the sense of string edit distance) to existing, popular domains. Such domains are used in many ways: They may host a page of ads (known as domain parking) or, more maliciously, be used for phishing or distributing malware. A systematic investigation of the problem of typo-squatting as it currently exists on the Internet has been carried out via several automated tools. For each of the most popular 10,000 domains (according to Alexa traffic rankings), many typo-like variations of the domain according to several simple rules are generated. For each typo that resolved in DNS, an HTTP request was sent, and the response was recorded along with any redirection URLs. For one sort of typo (a missing dot after www), 30% of the generated typo domains were found to be registered. The vast majority of typo-squatting was found to be due to a relatively small number of large-scale domain-parking companies, including Applied Semantics (formerly Oingo), which is currently owned by Google.

- ***Tracking the Role of Adversaries in Measuring Unwanted Traffic***

  Mark Allman, ICSI; Paul Barford, University of Wisconsin; Balachander Krishnamurthy and Jia Wang, AT&T Labs—Research

  The authors evaluated a number of techniques used in malicious traffic to foil network measurement systems. They pointed out the importance of network monitoring systems such as intrusion detection systems, firewalls, honeypots, and application-level filters in maintaining awareness of malicious traffic. So far, malicious traffic constructed in an intelligent attempt to bypass or disable such monitoring systems has been given little consideration; this paper attempts to fill that gap. A variety of attacks on monitoring systems were discussed, ranging from direct attacks that attempt to compromise or overload the monitoring system itself to methods for avoiding monitoring systems while compromising other hosts. To better understand these attacks and take a principled approach to countermeasures, two metrics for classifying their effects on measurement systems were proposed: consistency and isolation. These metrics were shown to provide a taxonomy dividing the polluting effects of malicious traffic on measurement systems into four groups.

  The talks of the last paper session concerned detection of botnets used for various purposes

and possibilities for pushback mechanisms that disable or block bot software on the infected host itself.

- ***An Algorithm for Anomaly-Based Botnet Detection***

  James R. Binkley and Suresh Singh, Portland State University

  The authors reviewed some practical botnet detection techniques at a university. The university's dormitory networks are frequently home to botnets, because of the proliferation of poorly secured Windows PCs. These hosts launch DDoS attacks. Since many botnet programs use IRC for coordination, the authors have found that inspecting layer 7 for anomalous IRC traffic is useful in discovering the botnets. The authors define a "TCP work weight" metric for each host on the network as the ratio of TCP control packets sent to total TCP packets sent. The metric is helpful in revealing whether a particular host is participating in DDoS attacks.

- ***Revealing Botnet Membership Using DNSBL Counter-Intelligence***

  Anirudh Ramachandran, Nick Feamster, and David Dagon, Georgia Institute of Technology

  In this paper, the authors propose a different technique for detecting botnets that have yet to be used in malicious activities. DNS-based blackhole lists (DNSBL) are currently used to keep track of hosts that relay large amounts of spam and allow MTAs to easily query this information. Botnets that are not yet listed in DNSBL fetch a much larger price in underground markets such as those discussed in the keynote address. Purchasers of such botnets may first check the IPs in DNSBL to verify that the botnet is still useful for sending spam. Two heuristics were developed to detect such anomalous DNSBL query activity. Running the heuristics against actual DNSBL logs revealed that these query patterns were in fact present and that the hosts queried were later used as spam relays. This discovery provides a method for passive botnet detection before the botnet is used for spam relaying.

- ***Leveraging Good Intentions to Reduce Malicious Network Traffic***

  Marianne Shaw, University of Washington

  Shaw discussed speculative thoughts on a new strategy for combating the compromised machines in botnets. A large majority of such systems are owned by technically ignorant but well-intentioned users ("grandma"). Some of these users may be willing to allow some sort of backdoor system on their host that will allow external systems to block or disable malware on their machine should it begin producing malicious traffic. Since malware on the compromised system will naturally attempt to disable any such mechanism, it would have to be located outside the control of the host operating system. Proposed were several locations, including the firmware of cable modems, NIC firmware, and inside a VM wrapped around the host operating system. This last possibility has been implemented as a research prototype under the Denali VMM. The system developed provides a mechanism for a host under DDoS attack to return a blocking request to the enforcement mechanism in the VMM. Experiments demonstrated an acceptable, yet significant, impact on the host's network performance.

- ***Panel: Real World Experiences***

  Richard Clayton, Sean Donelan, Mark Seiden, Rob Thomas

  The workshop ended with a one-hour panel discussion on real-world problems and the social issues involved in efforts to reduce malicious Internet traffic. The panel opened with the question "If you had a magic wand, what one thing would you change?" Responses focused on increased law enforcement efforts directed at the sorts of communities discussed in the keynote. It was acknowledged that it would be difficult to track down and prosecute all (or even most) perpetrators of online fraud, but any efforts that increased the costs and risks of online fraud would help serve as a deterrent. Further discussion centered on the role of user education in improving the security of end hosts and thereby reducing the numbers involved in malicious activities. The general consensus was that user education in security is at worst a hopeless task and at best of limited utility in certain environments. The lack of faith in end-user education led to examining whether more responsibility for compromised hosts could be placed on the service providers. One problem in this approach is the lack of an incentive for service providers to combat problems affecting other networks by cutting off the access of their own paying customers who have compromised hosts. The panel discussion revealed the inherent difficulty of eliminating malicious Internet traffic without reducing the ability of technically naive users to access network services—the same network services we are ultimately trying to protect.