Effect of Malicious Traffic on the Network

Kunchan Lan, Alefiya Hussain, Debojyoti Dutta
USC/ISI
4676 Admiralty Way,
Marina Del Rey,
CA 90292
{kclan,hussain,ddutta}@isi.edu

ABSTRACT

The Internet has seen, in recent days, a continuous rise in malicious traffic including DDoS and worm attacks. In this paper, we study the effect of malicious traffic on the background traffic by gathering traces from two different locations. We show that the malicious traffic causes DNS latencies to increase by 230% and web latencies to increase by 30%. Using a packet-level simulations based on an empirically derived model of the worm, we demonstrate that the effect of worm-infected hosts can be disastrous when they trigger a DDoS attack.

1. INTRODUCTION

During the last few years, the Internet has witnessed a surge in *malicious traffic*, such as that generated by denial-of-service (DDoS) attacks and propagation of worm traffic [4]. Most previous work [4, 18, 23, 10, 16, 24, 7] has focused on studying the reasons behind the malicious traffic but not their effects on the normal background traffic. We define normal traffic as network traffic generated due to well-known services and applications, for example, web, ftp, nntp, and smtp.

In this paper, we study of the characteristics of network traffic during phases dominated by malicious behavior of DDoS attacks and worm propagation, and compare it with phases when such activity is negligible. We show that DDoS attacks causes DNS latencies to increase by 230%, and the web latencies to increase by 30%. We find that the attacks do not significantly affect the throughputs of bulk TCP transfer. We also present detailed analysis Linux Slapper Worm and study the worm activity in the network. We then use an empirical simulation model to predict the effect of worm traffic when the worm-infected hosts trigger a DDoS attack.

The main contribution of this paper is to provide a quantitative analysis of the background traffic in the presence of malicious activity. We quantitatively study the effects of DDoS attack and worm traffic on normal background traffic. Currently most backbone links are under-utilized [2]. One would expect that the malicious traffic such as DDoS attacks and worm traffic will not change the background traffic patterns significantly if the links are highly over-provisioned. However, we find that this is not completely true. This work motivates the need to study more closely the reasons behind these observations. We believe that there is a need to do further studies of router mechanisms that can give us better performance in the presence of malicious traffic.

2. RELATED WORK

Several researchers have previously studied DDoS attack detection and response, and worm traffic propagation. In this section we provide a brief overview of DDoS and worm related research and compare how this paper complements previous studies.

2.1 DDoS

DDoS attacks attempt to exhaust the resources of the victim. The resources may be network bandwidth, computing power or operating system data structures. Previous research on DDoS attacks focused on either detecting the attack [9, 18, 23, 10] or responding to the attack [5, 6, 8, 11, 15, 13, 21, 22, 26] by blocking attack packets.

Attack detection techniques can be either based on an *anomaly-detection* approach or a static *signature-scan* technique. A large number of anomaly-detection tools have been designed and implemented previously, such as NIDES [14], Emerald [19] and Bro [18]. Anomaly-detection first establishes a normal behavior pattern for users, programs or resources in the system, and then looks for deviation from this behavior. Some anomaly-detection techniques exploit the absence of correlation between bidirectional traffic to detect an attack [9, 10, 13]. On the other hand, signature-scan techniques passively monitor traffic seen on a network and detect an attack when patterns within the packet match predefined signatures in a database. Snort [20] is a popular signature-scan based attack detection tool. In this paper, we use an anomaly-detection technique that tracks the number of source connecting to a single destination. Traffic is flagged as an attack if there is an abnormally high number of source addresses connecting to a single destination address.

2.2 Worm Traffic

Moore et. al. [16] present analysis of backscatter data gathered during the CodeRed infection last July-August. The data indicates 395,000 computers were infected world-wide with the CodeRed worm and resulted in approximately $2.6 billion in damage. Wang et. al. [24] presents a simulation based study to identify characteristics of worm infection. They study the effect of different factors that can be used to detect and treat infections while they are underway, using hierarchical and clustered network topologies. Zou [7] provides a two-factor worm propagation model that matches well with the observed CodeRed data. It models human countermeasures like patching, filtering and decrease in infection rate as a function of time to explain the decrease in CodeRed scan attempts observed during the last several hours of July 19th. In this paper we attempt to analyze the Apache/mod_ssl worm and use an empirical simulation model to study the effect of a DDoS attack launched from worm-infected hosts.

2.3 Web traffic latency analysis

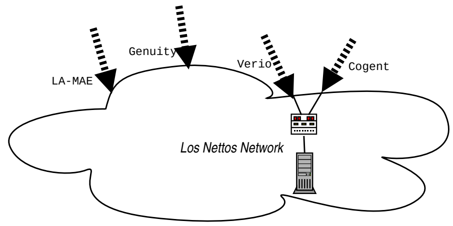

Figure 1: The trace machine monitors two of the four peering links.

Barford et. al. [3] study various factors affecting the performance of HTTP transactions. They show that the server load affects the transfer time for small files, while network load affects the performance of large files. They also show that propagation delay plays a more important role than network variability, such as queuing, in affecting the performance of Web traffic. Our study complements previous work by demonstrating malicious traffic, such as DDoS attack and worm infections, can also significantly increase latency for small and medium web transactions.

3. METHODOLOGY

3.1 Trace collection

We collect traces from two different locations: one at Los Nettos [17], a regional area network in Los Angeles, and the other at the Internet2 [1] peering link at USC. We continously capture detailed packet level traces using tcpdump at both locations and test the presence of attacks or worm infections.

Los Nettos has peering relationships with Verio, Cogent, Genuity, and the LA-Metropolitan Area Exchange as shown in the Figure 1 and serves a diverse clientele including academic institutes and corporations around the Los Angeles area. We monitor the Verio and Cogent peering links that experience an average utilization of 11% at 110Mbps and 38Kpps (packet-per-second). The kernel packet drops are below 0.04% during normal operation. During an attack, if packet rates exceed 100Kpps the drop rate increases to 0.6%. The USC trace machine monitors the Internet2 traffic to and from USC. The average utilization of link monitored by the trace machine is 6% at 60Mbps and 25Kpps.

The captured packet headers are analyzed offline to determine if there was an attack in progress. The detection script flags packets as attack packets if a large number of source IPs connect to the same destination IP within one second. Manual verification is then performed to confirm the presence of an attack. We experience a false positive rate of 25–35%; in other words, those packets have been flagged by the detection script but do not contain an attack after manual examination. A large number of false positives are generated due to network/port scaning and database updates between servers.

3.2 Metrics

We looked at several metrics to understand the impact of malicious traffic such as DDoS and worm on the network.

For web flows, we focus on flows with medium/small size (less than 100KB) to understand the impact of malicious traffic such as DDoS attack on the short-lived transactions. We look at TCP flows larger than 100KB to understand the impact on bulk transfer. We also investigate the impact on the DNS lookup latency. We define DNS lookup latency as the time between the client sends out a request to the DNS sever and the client finally receives a answer from a DNS server that terminates the lookup, by returning either the requested name-to-IP mapping or an error indication. To extract the statistics about lookup latency, we adopt similar approach as used in previous study [12].

Procotols	Los Nettos	USC
TCP	84.243%	95.605%
UDP	13.647%	4.102%
ICMP	1.216%	0.118%
Other	0.894%	0.175%

Table 1: Percentage of packets observed for each protocol at Los Nettos and USC

Service Protocols	Los Nettos	USC
http	39.448%	20.212%
ftp	0.577%	0.116%
dns	11.191%	0.219%
smtp	2.190%	1.075%
nntp	1.584%	10.202%
ssh	0.210%	1.102%
pop3	0.734%	0.118%
P2P	8.220%	15.224%
Games	0.418%	1.637%
Other	35.428%	50.076%

Table 2: Percentage of packets observed for each application at Los Nettos and USC

4. TRAFFIC CHARACTERIZATION

In this section, we briefly characterize the observed traffic. First we show the traffic mix observed in the traces. Table 1 and Table 2 describe the composition of normal traffic seen during peak hours of the day. We observe 13% UDP traffic since Los Nettos hosts a DNS root server. Web traffic constitutes 40% of the observed traffic followed by 11% DNS traffic. At USC Internet2 link, 95% of the network traffic is TCP. We could not classify a large percentage of the traffic since the Internet2 is extensively used for research and most of the packets uses ephemeral ports.

4.1 DDoS traffic

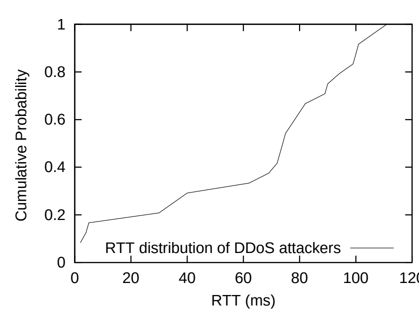

Figure 3: RTT distribution of DDoS attackers

(a) DDoS Traffic volume in bytes

(b) DDoS Traffic volume in packets

Figure 2: The traffic volume generated by DDoS attack in bytes and packets

We have captured 90 DDoS attacks from 15 July to 15 Nov 2002. Most of the attacks have significant impact on the background network traffic. In this paper, we characterize one of the captured attacks and show the effect it had on the background traffic.

Figure 2 shows the average amount of traffic per second as the attack progresses. Twenty eight attackers generate 70Mbps and 90Kpps of attack traffic (a total 11M packets and 8.6Gb of traffic in 192 seconds) directed at a USC host. The attack packets are 60 bytes and have the protocol field in the IP header set to 255. As shown in Figure 2, the magnitude of attack traffic is about three times of normal background traffic in terms of both bytes and packets. Figure 3 shows the distribution of RTT of the DDoS attackers, estimated using the ping utility from the victim network. The attackers have relatively small RTT distribution (less than 120ms) from USC because all attackers are located at different universities in the US. The small RTTs help the attack traffic to quickly reach its peak rate.

4.2 Worm traffic

Worm infection is on the rise. Worms like Code Red and Nimda can infect thousands of hosts within short periods of time and generate significant network traffic [25]. In this paper we study the effect of the Apache mod_ssl worm (aka the Slapper worm) on the network. Our findings suggest that although the Slapper worm did not increase the network traffic at USC or Los Nettos significantly, when the worm-infected hosts trigger a DDoS attack, the effect can be disastrous.

The Slapper worm exploits a bug in Linux-based hosts running Apache web servers with mod_ssl module. During the infection process the worm places source code in the /tmp directory of the target host. The worm then scans for potentially vulnerable systems on port 80 using an invalid HTTP GET request. When a vulnerable Apache host is detected, the worm attempts to connect to the SSL service via port 443 in order to deliver the exploit code. If successful, a copy of the malicious source code is then placed on the victim, where the attacking system tries to compile and run it. Once infected, the victim begins scanning for the other hosts to continue the worm's propagation. The infected system also becomes part of the Apachemod_ssl worm's DDoS network. Infected systems can then share information, including attack instructions, with other infected systems.

Top 10 Top-level Domains

TLD	hosts	hosts(%)
unknown	858	31
net	447	16
com	330	12
us	173	6
ca	126	5
it	106	4
pl	104	4
edu	77	3
tw	70	3
mx	70	3

Table 3: Top ten top-level domains with Linux Slapper Worm infected hosts on Oct

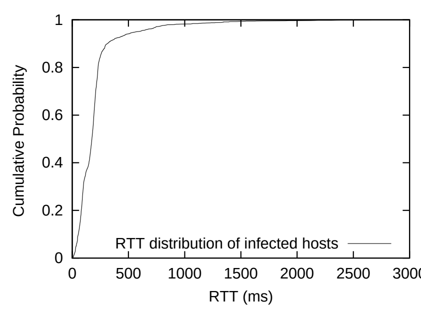

Figure 4: RTT distribution of worm-infected hosts

We observed a total 2727 infected hosts spanning over 39 AS domains distributed all over the world. Table 3 shows the distribution of the number of infected hosts from different domains. We see a large percentages of infected hosts are located in .net and .com domain. Note that we could not determine about 30% hosts due to DNS name resolution failure. Figure 4 shows the distribution of the RTTs of the worm infected hosts. Unlike the RTT distribution of DDoS attack hosts, the RTT distribution of worm-infected hosts shows RTTs of over 1500ms. The huge diversity of RTT distribution suggests that if these worm-infected hosts generate DDoS attacks, they could potentially come from all over the world, making them harder to isolate.

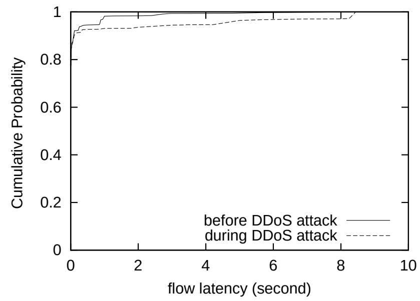

Figure 5: DNS lookup latency increases by 230% during a DDoS attack

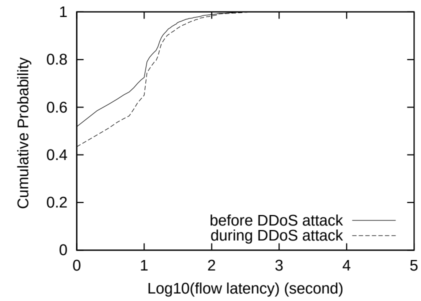

Figure 6: Latency experienced by small and medium web flows increases by 30% during an attack

5. EFFECT OF MALICIOUS TRAFFIC

In this section, we evaluate how malicious traffic changes observed traffic characteristics. Although it is intuitive that traffic characteristics might change on a DDoS attack or a worm infection, we are not aware of any previous work that has quantitatively characterized the effect of such traffic. We study the effect of DDoS traffic on DNS latency and web latency. We observe that DNS latency increases by 230%, and web-latency increases by 30% during a DDoS attack. Finally, based on an empirical simulation model of worm, we predict its effect on the network when the worm-infected hosts trigger DDoS attacks.

5.1 DDoS traffic

DNS latency is defined as the time elapse between the issue of the query to finally the server returns an answer or failure. The effectiveness of DNS strongly affects the performance of many popular network services such as Web traffic and Contents Distributed Networks (CDNs).

As shown in Figure 5, the average latency of DNS lookup has increased from 0.13 seconds to 0.44 seconds during an attack, more than a 230% increase in latency. One possible explanation is that the sudden increase of traffic during an attack leads to higher average buffer occupancies at the routers, resulting in increased queuing delays.

We also look at the effect of DDoS attack on web traffic, We particularly focus on small and medium web flows (which we define as flow size is smaller than 100KB), since such flows are more sensitive to the delay. We define web latency as the time lapse between the issue of HTTP request to the receiving of response data. As shown in Figure 6, the average latency of web flows has increased from 9 seconds to 11.9 seconds, resulting in a 30% increase during the attack. Note that the DNS and web latencies increase even when the link is still under-utilized as shown in Section 3.1.

Even though the DNS and web latencies increase, we noticed that the average throughput of bulk TCP transfers (which we define as flow size larger than 100KB), remains unchanged during the attack as shown in Figure 7. We believe it is because the attack only last for only 192 seconds and has little effect on the long-lived TCP flows.

The above results show that although short duration DDoS attacks might not be disruptive in terms of causing network failures and reducing aggregate throughput, the delay-sensitive traffic such as

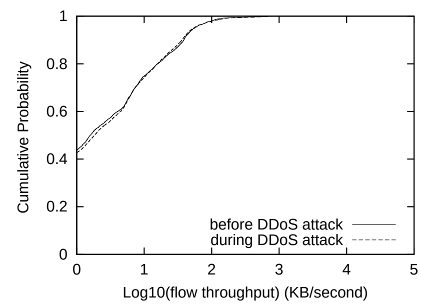

Figure 7: Effect of DDoS attack on throughput of bulk TCP flows

DNS and small/medium web transaction will still be affected by these attacks. Over-provisioning the links on the network does not provide the complete solution, since the short burst of DDos traffic can result in the increases in latency without affecting the throughput. We feel that the above observations can be used as hints to design better AQM mechanisms to provide differential services in order to protect short-lived traffic.

5.2 Worm traffic

The Slapper worm propagation did not generate disruptive amounts of traffic at our data collection point. However, if all the infected machines launched a coordinated DDoS attack, it would have a disastrous effect. In this section, we use hints from the collected Slapper worm data to determine the size of the compromised network. We study its effect on the network when all worm-infected hosts launch a coordinated DDoS attack using a ns-2 simulation.

We derive the topology information of the worm-infected network based on the traces. We simulate its effect on the network when all worm-infected hosts launch a DDoS attack to a victim in the USC campus. We use a simple dumbbell topology with empirical distributions of RTT, flow rates and packet size derived from the traces. The DDoS traffic is modeled as constant bit rate source and currently no background traffic is simulated.

Figure 8 shows the attack intensity when generated by worm-infected hosts. We observed that the different RTT distributions of the attackers cause distinctively different transient ramp-up behavior before the steady state attack rate is achieved. Also when all the worm-infected hosts launch a DDoS attack, the average traffic generated due to the attack is fifty times larger than that generated by

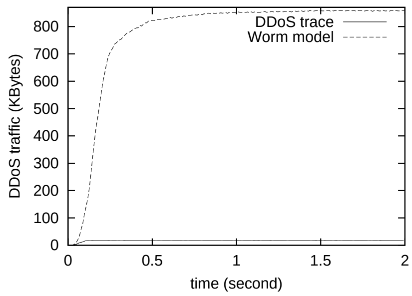

Figure 8: Comparison of DDoS attack intensities; the DDoS attack and when an attack is launched by worm-infected hosts

the DDoS attack that we traced.

6. CONCLUSION AND FUTURE WORK

In this paper, we present a detailed study of how the background traffic changes in the presence of malicious traffic. In particular, we show that the DNS latency increased by 230% and the web latency increased by 30% upon interaction with DDoS traffic. We also analyze the recent Linux Slapper Worm activity. Based on an empirical simulation model of worm, we predict its effect on the network when the worm-infected hosts trigger DDoS attacks.

We have captured 90 DDoS attacks from July 2002 to Nov 2002. This abstract only presents analysis from one attack in the collected traces. We are currently working on a more detailed study of the effect of malicious traffic on background traffic by analyzing more DDoS and worm attacks. In particular, we are studying how different intensities and types of DDoS attacks will change the characteristics of the background traffic. Another aspect of our ongoing effort is to study various worm propagation models in order to predict the overall effect of worm traffic on the network.

7. REFERENCES

[1] Internet 2. http://www.internet2.edu.

[2] Chadi Barakat, Patrick Thiran, Gianluca Iannaccone, Christophe Diot, and Phillipe Owezarski. A fbw-based model for internet backbone traffi c. *Internet Measurement Workshop 2002*, November 2002.

[3] P. Barford and M. E. Crovella. Critical path analysis of TCP transactions. In *SIGCOMM*, Stockholm. Sweden, September 2000.

[4] Paul Barford, Jeffrey Kline, David Plonka, and Amos Ron. A signal analysis of network traffi c anomalies. *Internet Measurement Workshop 2002*, November 2002.

[5] Steven Bellovin. ICMP traceback messages. Internet Drafts: draft-bellovin-itrace-00.txt.

[6] Hal Burch and Bill Cheswick. Tracing anonymous packets to their approximate source. In *Proceedings of the USENIX Large Installation Systems Administration Conference*, pages 319–327, New Orleans, USA, December 2000. USENIX.

[7] Don Towsley Changchun Zou, Weibo Gong. Code read worm propagation modeling and analysis. In *ACM Conference on Computer and Communication Security*, Washington DC, Nov 2002. ACM.

[8] Drew Dean, Matt Franklin, and Adam Stubblefi eld. An algebraic approach to ip traceback. In *In Proceedings of Network and Distributed Systems Security Symposium*, San Diego, CA, February 2001.

[9] Thomer M. Gil and Massimiliano Poletto. MULTOPS: A Data-Structure for bandwidth attack detection. In *Proceedings of the USENIX Security Symposium*, pages 23–38, Washington, DC, USA, July 2001. USENIX.

[10] Danlu Zhang Haining Wang and Kang Shin. Detecting syn fboding attacks. In *Proceedings of the IEEE Infocom*, pages 000–001, New York, NY, June 2002. IEEE.

[11] John Ioannidis and Steven M. Bellovin. Implementing pushback: Router-based defense against DDoS attacks. In *Proceedings of Network and Distributed System Security Symposium*, San Diego, CA, February 2002. The Internet Society.

[12] Hari Balakrishnan Jaeyeon Jung, Emil Sit and Robert Morris. Dns performance and the effectiveness of caching. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop '01*, San Francisco, California, November 2001.

[13] Peter Reiher Jelena Mirkovic, Greg Prier. Attacking ddos at the source. In *10th IEEE International Conference on Network Protocols*, Paris, France, November 2002.

[14] Teresa F. Lunt. Detecting Intruders in Computer Systems. In *Proceedings of the Sixth Annual Symposium and Technical Displays on Physical and Electronic Security*, 1993.

[15] Ratul Mahajan, Steven M. Bellovin, Sally Floyd, John Ioannidis, Vern Paxson, and Scott Shenker. Controlling high bandwidth aggregates in the network. In *ACM Computer Communication Review*, July 2001.

[16] David Moore, Geoffrey Voelker, and Stefan Savage. Inferring Internet denial of service activity. In *Proceedings of the USENIX Security Symposium*, Washington, DC, USA, August 2001. USENIX.

[17] Los Nettos-Passing packets since 1988. http://www.ln.net.

[18] Vern Paxson. Bro: a system for detecting network intruders in real-time. *Computer Networks*, 31(23–24):2435–2463, 1998.

[19] Phillip A. Porras and Peter G. Neumann. EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances. In *Proceedings of the 20th NIS Security Conference*, October 1997.

[20] Martin Roesch. Snort - lightweight intrusion detection for networks. http://www.snort.org.

[21] Alex C. Snoeren, Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio Stephen T. Kent, and W. Timothy Strayer. Hash-based ip traceback. In *Proceedings of the ACM SIGCOMM*, pages 3–14, San Deigo CA, August 2001. ACM.

[22] Dawn X. Song and Adrian Perrig. Advanced and authenticated marking schemes for IP traceback. In *Proceedings IEEE Infocomm*, Anchorage, Alaska, April 2001.

[23] Robert Stone. Centertrack: An IP overlay network for tracking dos fbods. In *Proceedings of the USENIX Security Symposium*, pages 199–212, Denver, CO, USA, jul 2000. USENIX.

[24] Chenxi Wang, J. C. Knight, and M. C. Elder. On computer viral infection and the effect of immunization. In *ACSAC*, pages 246–256, New Orleans, 2000.

[25] Lan Wang, Xiaoliang Zhao, Dan Pei, Randy Bush, Daniel Massey, Allison Mankin, S. Felix Wu, and Lixia Zhang. Observation and analysis of bgp behavior under stress. In *Internet Measurement Workshop*, pages 217–222, Marseille, France, Nov 2002.

[26] E. Zwicky, S. Cooper, D. Chapman, and D.Ru. *Building Internet Firewalls*. 2nd Edition. O'Reilly and Associates, 2000.