The impact of Microsoft Windows infection vectors on IP network traffic patterns and costs

John Nguyen, Grenville Armitage
Centre for Advanced Internet Architectures
Swinburne University of Technology
Melbourne, Australia
{jnguyen,garmitage}@swin.edu.au

Australian Telecommunications Networks & Applications Conference 2004 (ATNAC2004), Sydney, Australia, December 8-10 2004

***Abstract*- This paper describes a set of tools and techniques to capture and analyse virus-generated IP network traffic. We analyse seven viruses, worms, trojans and spyware that are common in Microsoft Windows environments. We log and analyse the IP traffic generated in the roughly 2 hours after each infection. Based on the resulting IP traffic patterns we estimate the likely financial impact of having an infected PC connected to a consumer-grade, broadband Internet connection.**

Keywords- Viruses, Worms, Trojans, Spyware, Traffic Patterns, Financial Impact

I. INTRODUCTION

Computer virus, worm, trojan and spyware attacks are on a very sharp rise. So-called "malicious programs" are often equipped with sophisticated techniques to trick computer users, and to automatically seek out vulnerable networked hosts. Despite the efforts of anti-virus vendors and the computing security community, virus attacks are still very significant at the global scale.

We all have anecdotal evidence that network viruses are bad, and even some semi-scientific estimates of their financial impact [1][2][3][4]. Similarly, trojan horses and spyware have been rapidly propagating through emails, instant messaging, P2P applications, browser hijacking, etc. Users are bombarded with popups, client browsers redirected, key presses logged and confidential information tracked.

Determining the overall cost impact when a computer system is hit with viruses, worms, trojan horses, spyware etc. is not trivial. A large part of this cost comes from the network traffic generated by an infected host. Therefore we developed a structured process to estimate network damage caused by a variety of infections. Seven well-known Internet viruses, worms, trojan horses and spyware were chosen for our study of their network behaviours and traffic patterns.

Ultimately we aim to answer the following questions from the perspective of a Microsoft Windows machine infected with a typical virus, trojan or worm:

- What type of network attacks, traffic patterns and network loads are caused by each infection?
- How many Mbytes per hour, day or month would be consumed and how much would this cost a typical, broadband-attached 'always on' PC?

Our paper continues with a description of the testbed in section II and summaries of the infections tested in section III. We conclude by estimating the likely financial impact of having an infected PC connected to a consumer-grade, broadband Internet connection.

II. SETUP OF THE CONTROLLED TESTBED

The testbed consists of 2 computers connected via a crossover cable (Figure 1. The victim host runs Windows XP (version 5.1 2600 Service Pack 1) with all the latest patches and security updates at 29th of June 2004. It is injected with a copy of the virus under each experiment.

The sniffing host runs FreeBSD 4.10 with the following components installed and enabled: Bridging and ipfw (firewall), tcpdump packet sniffer, thttpd (web server), sendmail (SMTP server), BIND (DNS server), and tinyproxy (Proxy web server).

Initially, we enable bridging support and firewalling on the sniffing host, blocking everything except DNS traffic (so processes on the victim host can at least resolved targets). We log inbound and outbound Ethernet traffic relating to the victim using tcpdump. This configuration is referred to as "blackhole" case (because all outgoing TCP connections are blocked).

As the experiments evolved, we setup various network services such as DNS, Web and Email on the sniffing host to trick the viruses into thinking this is their ultimate target. A few tiny proxy services are run on some regular ports such as 80, 8000, 8080 in order to log web traffic requests from the victim host. We configured the victim host to send its DNS requests to

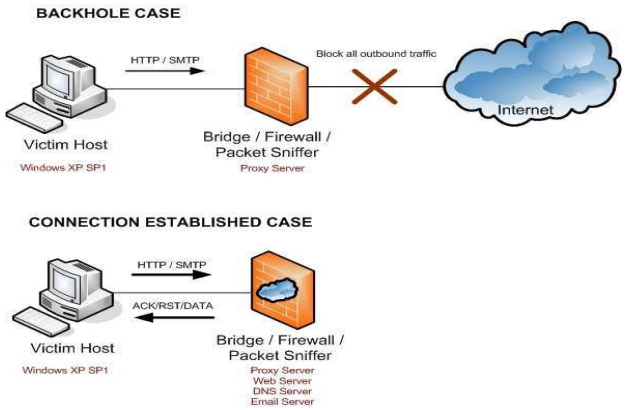

Figure 1 Testbed Configuration

a local DNS server on the sniffing host, which then returned its own address in response to specific DNS requests issued by the infected host. In this manner we tricked the viruses into using

ISBN 0-646-44190-6

the mail and web servers on the sniffing host. This configuration is refered to as the "connection established" case because the victim host can successfully establish http (DoS attack) and smtp (mass mailing) connections to the sniffing host.

III. Selection of viruses, worms, trojan horses and spyware

Table 1 shows the main selection criteria and the list of malicious software chosen for the experiments.

Selection criteria	Virus/Worm/Trojan/Spyware
Popularity	Sasser.A, MyDoom.E
Financial impact	Lovesan, MyDoom.E
Types of propagation and attack	NetSky.R (mass mailing worm), Gator (Spyware), SpyBot (P2P Worm) and SubSeven (Trojan)

Table 1 Selection Criteria

Virus samples were obtained from the following sources:

- Virus Exchange Board (VX Discussion Board)
- Virus Collection Website (e.g.: VX Heavens at http://vx.netlux.org)
- Viri collection hobbyist and trader (many post their email & collection information on the Internet)

A. Sasser.A

Sasser.A is a worm that exploits Windows Directory Service vulnerability on Windows XP and Windows 2000 systems. The worm constantly scans a range of IP addresses on port 445, 50% of them are deduced from the host's own IP address; the others are generated randomly. If a connection to port 445 is successful the worm will send shell code to open a remote shell on TCP port 9996. It then uses the shell on the remote computer to reconnect to the infected computer's FTP server, running on TCP port 5554, and retrieve a copy of the worm.

B. Lovesan

Lovesan is a Blaster worm variant that exploits a Windows' NETBIOS vulnerability. It scans a range of IP addresses on port 135 - two out of five are deduced from the host's address, the others are generated randomly. The worm sends a buffer-overrun request to TCP port 135 of a victim. If successful, the victim starts a command shell on TCP port. The worm runs the thread that opens the connection on port 4444 and waits for FTP "get" request from victim machine. The worm then forces the victim machine to send an "FTP get" request to download and activate a worm copy from the infected machine. The worm can also launch Denial of Service attack against windowsupdate.com. We tested Lovesan in 2 cases: blackhole case and another case where there are ACK/RST packets coming back. We see that when ACK/RST packets are returned the total traffic is five times greater than the blackhole case.

C. MyDoom.E

MyDoom.E is a mass mailing worm capable of carrying out DoS (denial of dervice) attacks against the origin2.microsoft.com site between the 17th and 22nd of each month. It uses its own SMTP engine to send messages with attached copy of viruses directly to the recipient's email server. We tested 4 cases: mass-mailing into blackhole (case 1), mass-mailing successfully (case 2), DoS attack into a blackhole (case 3) and DoS attack successfully (case 4). The most dangerous case is when MyDoom carries out the DoS attack successfully (case 4). The worm spawns multiple "HTTP GET" requests, and gets responses coming back; so both upstream and downstream bandwidth can be totally consumed. If the attack target is down, blocked or not responding (case 3), the worm also tried to propagate itself via email. Although not as intense as a Dos attack (case 4), mass-mailing mode (case 2) can generate many flows of DNS and SMTP traffic, which results in bursts every 10 seconds when emails are sent successfully. The email attachments with worm copy are approximately 20 Kbytes each.

D. Netsky.R

Netsky is another widespread mass mailing worm (written by the same author of the Sasser worm [11][12]). Netsky searches through victim files to obtain valid email addresse and uses its own SMTP engine to send messages with an attached copy of itself (usually with .pif extension). These emails are sent directly to the recipient's email server. Email source spoofing is utilised by this worm to trick users about the origin of the infected emails they receive. The traffic profile shows constant flows of DNS requests from the worm to try resolving MX records of the domains where its victims belong.

E. Gator

Gator is in the adware / spyware category. Gator includes a software component from GAIN (the Gator Advertising Information Network) advertising, which is also bundled with other free software like DivX player; WeatherBug, Kazaa .etc. GAIN displays lots of pop-up advertising and gathers extensive details about user's computer setup and browsing habits [13]. In a period of 1000 seconds (~16 minutes), there are 5 "HTTP GET" requests to pull down data from the servers such as bc2.gator.com, ss.gator.com, etc. There are also occasional "HTTP POST" actions during the experiment.

F. Spybot

Spybot combines characteristics of a virus, a worm (P2P type) and a keylogger program, with more than 1000 variants. Once activated, the worm copies itself into "kazaabackupfiles". Spybot also tries to connect to a few specified IRC servers to report successful infection in order to join a channel to receive commands (DoS attacks, copying itself to hardcode Windows folders.etc.) and logs user's keypresses into a short text file "keylog.txt" stored under Windows system folder. Spybot has a list of IRC server IP addresses that it keeps rotating through in order to establish connections on port 6667.

G. SubSeven

SubSeven is a trojan divided into a client program that the attacker runs on his own machine, and a server run on the victim's computer [14][15]. SubSeven is usually spread via emails, P2P networks, Instant Messaging, etc. SubSeven's control program can instruct the victim to transfer files in and out, therefore the impact of these types of traffic on the network can be quite substantial in those cases

V. Experimental process & tools

We established a testing process (Table 2) and set of tools

ISBN 0-646-44190-6

and to ensure consistency.

Step	Procedure
1. Baseline the test	Re-image the victim host to a clean installation of MS Windows. Measure all traffic, currently running processes, threads and opened ports of the Windows host before any infection
2. Execute & observe behaviours of viruses	Activate virus sample and observe changes done to registry, file system, CPU usage, threads, TCP ports. etc.
3. Sniff traffic from/to the victim host	Run tcpdump from the sniffing host to collect all traffic coming in and out of the victim.
4. Analyse captured traffic	Use Ethereal to analyse traffic patterns, TCP flows, frequency and destination of attacks.
5. Refine the experiment	From results of step 4 refine the experiment: capture for longer period, simulate the target by installing network services such as DNS, Web, Email to respond to virus requests. etc.

Table 2 Experimental Process

We used the following tools to analyse the impact of each virus, trojan or spyware:

- Fport: used to display all victim's opened ports
- Process Explorer: used to display processes & threads under Win32 OS
- tcpdump: used to sniff traffic from the victim's host and write it to a file for later analysis.
- Ethereal: used to analyse traffic patterns and TCP flows
- PacketPlotter: an Excel VBA application to graph exported data from Ethereal.[8]

IV.IMPACT COMPARISON

In order to compress our paper to 4 pages we have elided the results of testing each individual virus or worm. Instead, we jump directly to a quantitative analysis of the financial and link speed impact imposed on victims fortwo scenarios described in Table 3.

The potential cost is calculated based on scenario 1, with an additional assumption that 50% of the user's monthly quota is consumed by 'normal' user activity. Therefore, all virus-generated traffic need to consume the rest of the allowed quota (50%) before the user is charged 15 cents for any extra megabyte. We calculate the actual speed of the plan in scenario 1 as 85% of 256/64 Kbps (217/54 Kbps), roughly taking into account Ethernet framing and ATM overheads used in ADSL links. The percentage utilisation of upstream and downstream bandwidth is then calculated as the percentage of 217/54 Kbps.

Scenario	Plan Details
Typical Home broadband ISP scenario 1	• Used to quantify how much extra dollars to pay a month • Telstra ADSL 500MB Limited Plan • 256/64 Kbps speed (in real life ~ 217/54 Kbps max for 85% efficiency factor) • 15 cent for extra Megabyte upload / download
Typical Home broadband ISP scenario 2	• Used to quantify how many days virus consume all allowed quota • Optus ADSL Value 1GB Plan • 512/128 Kbps speed (~ 435/108 Kbps max for 85% efficiency factor) • Rate limited to 28.8 Kbps until the rest of the month when quota exceeded

Table 3 Assumption scenarios

The results and analysis obtained from all the experiments are summarised in the following subsections.

Figure 2 compares the traffic loads generated by our viruses, worms, trojans, and spyware. Figure 3 extrapolates the experimental results to see the potential impact on a normal Internet link over a month. For example, when myDoom is in its successful DoS attack mode, it can consume all the upstream and downstream bandwidth that is available to the user (~ 90.52 GB/month if allowed to run un-checked). In its mass-mailing mode myDoom floods the link with DNS and SMTP traffic. This can add an extra of 11.77 Gigabytes of traffic load. Thirdly, Lovesan IP address scans with returned acknowledgement can also bring in to the network an addition of 3.91 Gigabytes of traffic.

Assuming that an Internet user is on the Telstra 500 MB limited ADSL (256 Kbps downstream / 64 Kbps upstream) plan and each extra megabyte of traffic is charged at 15 cents, Figure 4 shows a comparison of the estimated amount of money the users have to pay. The graph also compares three different scenarios. The worst-case scenario is when the infected computer is left online 24 hours a day for an entire month and 50% of monthly allocated quota has already been used. The

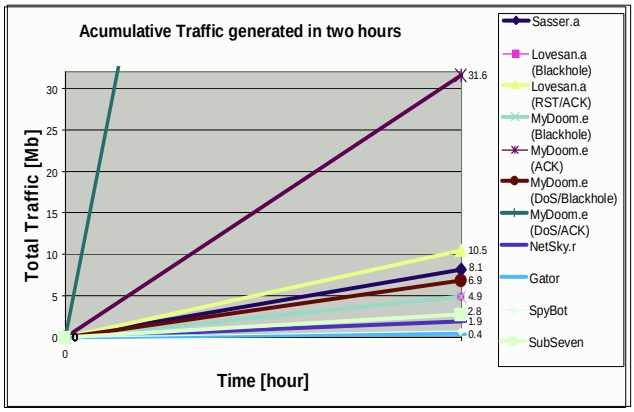

Figure 2 Impact of viruses on network traffic load (Number of Gb(s)hourly)

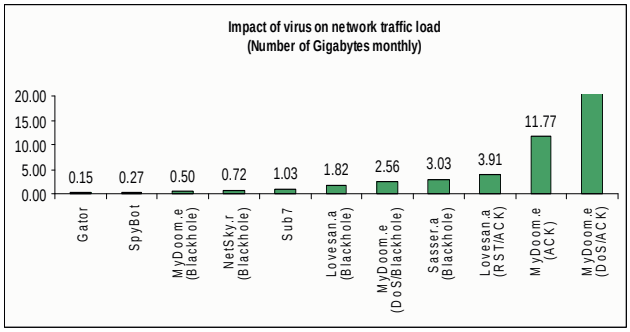

Figure 3 Impact of viruses on network traffic load (Number of Gb(s) monthly)

Australian Telecommunications Networks & Applications Conference 2004 (ATNAC2004), Sydney, Australia, December 8-10 2004

ISBN 0-646-44190-6

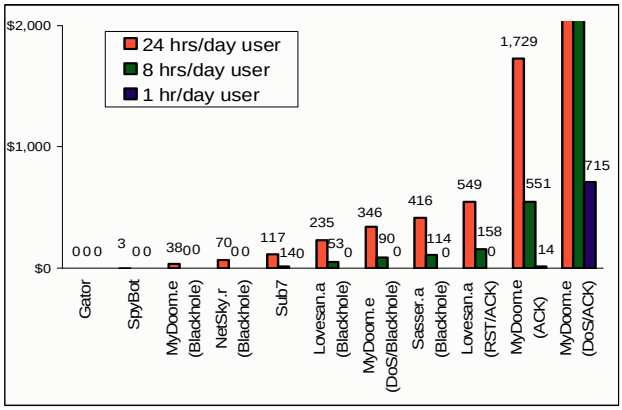

Figure 4 Financial impact of viruses on Internet users (based on Telstra ADSL 256/64 500MB Limited Plan)

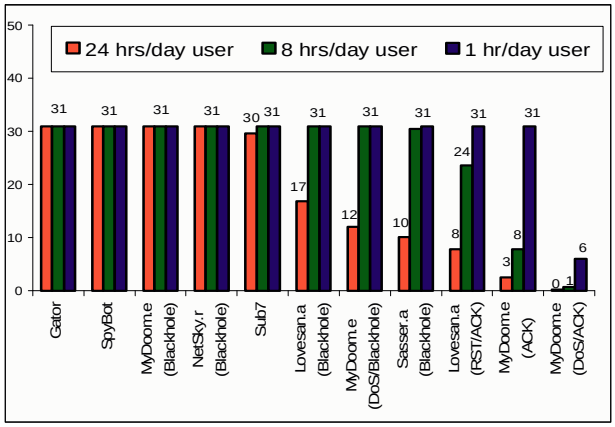

Figure 5 How many days to use up my monthly quota? (based on Optus ADSL512/128 1GB Limited Plan)

average-case scenario is when the infected computer is left online 8 hours a day for an entire month and 50% of monthly allocated quota has already been used. Finally, the best-case scenario is when the infected computer is left online 1 hour a day for an entire month and 20% of monthly allocated quota has already been used.

Clearly an infected machine can incur substantial extra charges on one's monthly Internet bill. Although an adware like Gator seems to cost nothing for the user, nevertheless if many of the same type programs are installed, the cost can add up very quickly.

Figure 5 shows how quickly an infected customer's entire monthly quota would be consumed (based on the Optus 1 GB limited ADSL 512/128 plan). A continuous and successful DoS attack, mass-mailing or IP/port scanning can use all allocated quota within one to three days. The user's Internet link speed is then capped at 28.8 Kbps until the end of the month.

V. Conclusion

We have evaluated the characteristic traffic profiles and network load that a selection of seven viruses, worms, trojans and spyware can generate. Our experimental trials were short-lived, roughly two hours each, but we believe the results to be a resonable predictor of traffic loads over hours, days an weeks.

We have also estimated the likely financial impact for infected Internet users by calculating the total traffic load generated by viruses in a period of time. Our analysis and comparisons reflect the variations due to each infection's modes of attack and propagation. If users are charged by their ISP on the amount of traffic a virus generates, there can be a bill-shock for him or her at the end of the month. We also note that trojans and spyware such as Spybot or SubSeven can create additional damages if they open up backdoors for unauthorised access to the victim's computers.

Studying previous viruses is one of the important steps to improve our ability to deal with the virus problems of the near future. The idea of our research was to address the needs to understand threats and consequences imposed on the network by virus attacks. Our hope is to use this as a stepping-stone for our future research.

Acknowledgments

We would like to thank Warren Harrop and Lawrence Stewart for their technical assistance.

References

(all web references are as of the date of publication of this paper)

[1]"Cost Impact of Major Virus Attacks Since 1995"
http://www.computereconomics.com/images/default/cmr/IT%20Bytes%20April%202004.pdf
[2]"Mydoom virus biggest in months"
http://news.bbc.co.uk/1/hi/technology/3432639.stm
[3]"MyDoom is most expensive virus yet"
http://www.vnunet.com/news/1152514
[4]"MyDoom.0 Hammers Search Sites"
http://www.pcworld.com/news/article/0,aid,117066,pg,1,RSS,RSS,00.asp
[5]"Email Virus Propagation Modeling and Analysis"
http://tennis.ecs.umass.edu/~czou/research/emailvirus-techreport.pdf
[6]"Evaluation of a Pentium PC for use as an Ethernet Bridge"
http://caia.swin.edu.au/reports/030326A/CAIA-TR-030326A.pdf
[7]"Developing an Effective Incident Cost Analysis Mechanism"
http://www.securityfocus.com/infocus/1592
[8]"Packet Plotter"
http:// home.intergga.ch/kummerj/packetplotter/
[9]"Consumers and ISPs go head-to-head on bill-shock"
http://www.zdnet.com.au/news/communications/print.htm?TYPE=story&AT=39148078-2000061791t-10000003c
[10]"Reverse Engineering Malware"
http://www.zeltser.com/sans/gcih-practical/revmalw.html
[11]"70% of viruses written by one man"
http://itvibe.com/default.aspx?NewsID=2769
[12]"Virus writing on the increases"
http://www.sophos.com/pressoffice/pressrel/uk/20040728topten.html
[13]"Gator eWallet"
http://www.scumware.com/apps/scumware.php/action::view_article/article_id::1068605442/topic::Scumware,-Spyware,-Adware-&-Malware-Applications/
[14]"SubSeven Official Site"
http://www.subseven.ws/
[15]"Backdoor.SubSeven"
http://www.symantec.com/avcenter/venc/data/backdoor.subseven.html
[16]"Email Virus Propagation Modeling and Analysis"
http://tennis.ecs.umass.edu/~czou/research/emailvirus-techreport.pdf

ISBN 0-646-44190-6