CS 261 Computer Security
David Wagner

Attacks

Scribe: Erika Chin
October 30, 2007

1 Notable Worms

Worm	**Date**	**# Infected**	**Time to Infect**
Code Red	7/2001	360,000	6 days
Code Red II	8/2001	?	?
Nimda	9/2001	? (on the order of millions)	22 minutes
Slammer	1/2003	75,000	10 minutes
Blaster	8/2003	8,000,000	2 days

*The first three worms were crude attacks. The Slammer worm distinguished itself by its rapid spread due to its single packet infection capabilities.

These worms resulted in billions of dollars in financial loss. Networks went down, people lost email access, and companies lost productivity. Eventually, CEOs and CFOs went to Microsoft to tell them that it was unacceptable and that they had to improve the quality of their software. With each successive worm propagating faster and faster, the potential for damage, disruption, and financial loss increased. Furthermore, these worms showed the potential to spread faster than system administrators could respond. This sparked an academic interest in worms and other attacks.

2 How Bad Can it Get?

In terms of payload and economic and financial loss, we examine how harmful these attacks can be. Malicious attacks can impact the economy through:

- Loss of data
  - Example: One could gain access to critical data, encrypt it, and demand ransom.
  - Cost: Unknown. Most companies have backups.
- Attack to critical infrastructure
  - Example: One could take down power grids and water systems.
  - Cost: It is difficult to evaluate the feasibility of attack or estimate the cost as there is not much public data on this subject.
- Loss of hardware
  - Example: One could overwrite the BIOS.
  - Cost: $1,000-2,000 per machine
- Unauthorized transactions
  - Example: One could create a worm to sell U.S. bonds.

- Loss (and cost) of a system administrator’s time
  - Cost: $50-100 per hour per system administrator and it could take up to an hour per machine to restore data from backups.
- Loss of productivity
  - Example: Without working machines, employees may have nothing to do but sit and wait for their computers to come back online. On a larger scale, many e-commerce companies lose revenue when their sites go down or when users cannot gain access to the network.
  - Cost: This can vary greatly depending on the company. Some businesses are based on e-commerce. Amazon makes about $100,000-200,000 per hour. Airlines, likewise, make about $100,000 in online ticket sales. Other companies lose employee productivity waiting on the system administrator. An average of $50 is lost per hour per employee.

In total, costs per machine could be $500-$1,000. With over 200,000,000 machines in the U.S. alone, an attack could cost the nation $5-50 billion.

3 Building a High-speed Worm

If one wanted to build a high-speed worm without worrying about noise (detection), one could seed the worm with a hit list (a list of known vulnerable hosts). This would ensure rapid, exponential infection. One could also use multiple propagation mechanisms. For example, Nimda used 5-6 techniques. It used email with malicious attachments (to pass through firewalls), random scanning in octets (to spread quickly), and infected web servers so that anyone who viewed the HTML would also get infected (to spread widely).

4 Building a Stealthy Worm

There are many ways to create a stealthy worm. One could:

- Rate limit the propagation. Make sure that the worm spreads slowly.
- Wait some time after infection/download. This way if the user notices the infection, they will not be able to figure out what it is from immediately.
- Avoid creating noticeable changes in the network. Network monitors log information and create statistics on traffic. A stealthy worm must not trigger any alarms or increase network traffic.
- Create a worm that piggybacks on existing network connections. For example, one could exploit a vulnerability in a P2P program, ex. bit torrenting programs. Then, when the infected user opens a connection to another client, the worm can piggyback across that connection. Another example is the client-server-based worm. If one could find vulnerabilities in a web server and browser, then the worm could propagate across servers and browsers stealthily. This method is rare, however, as it would require exploiting two vulnerabilities, one in the server and one in the client.

5 In Recent Years

In recent years, a market for vulnerabilities has developed. These vulnerabilities can sell for $1,000 to $50,000. Purchasers include large companies, like Cisco, that want to patch their vulnerabilities before the public is aware of it. It is also suspected that the U.S. Government purchases these vulnerabilities (although it is not known whether they do this to attack or defend systems).

Most worms no longer make headlines anymore. The Blaster worm was the last of the major, publicized worms. This is because the motivation for spreading worms has shifted. In the early 2000s, young hackers exploited systems to gain fame and bragging rights. These hackers sought to create noisy, disruptive worms. The more damage they caused, the more attention they got. Now, malware writers are motivated by the financial prospects of worm creation. These writers would rather receive money than gain fame. Flashy worms do not make money. They are often fixed very quickly, and thus the attacker "wastes" the vulnerability. Instead, attackers look for ways to create stealthy worms, so that they can stay under the radar and make money. This is discussed further in Section 7.

From 2001 to 2003, we saw the damage due to worms escalate. The number of infected machines increased while the time to infect decreased. These attacks have shown us that quickly propagating attacks are a dangerous threat. The good news is that we have not yet seen any worst-case scenarios. Still, these automated attacks propagate faster than humans can respond, and for that reason, we need a defense.

6 Network Telescopes

In order to do a longitudinal study of worms in the wild, we use network telescopes. To create a network telescope, one must go to the ISP and request a large address space of unassigned IP addresses. Traffic to these addresses are then directed to a small set of machines (a darknet). Because these IP addresses are unassigned and no legitimate computer is attached to any of the darknet's IP addresses, we know that any incoming traffic is illegitimate.

Assume we have a darknet consisting of all 1.2.x.y addresses. By setting up a honeynet and running a VM with an unpatched version of Windows, we make it easy for the darknet computer to get infected. Then by monitoring the network traffic (observing packets with the same payload, connections going to the same port, or just a general spike in network), we can then observe infection and subsequently the malware's scanning and probing techniques.

The network telescope may also see a lot of backscatter. When spammer's spoof IP addresses (and that IP belongs to the telescope), we can monitor the spoofed traffic (by the SYN/ACKs that get sent to our IPs).

The telescoping technique only allows us to see a small portion of the Internet. (In the case of the example, we would potentially see $1/ 2^{16}$ of the Internet.) Still, this is enough to observe large events, including the Slammer worm.

For the telescope to be effective, it is important to keep the addresses secret and cycle them frequently. Otherwise, attackers can specifically design their worm to avoid those IP addresses to evade analysis.

7 The Evolution of Attacks on the Internet

In recent years, large underground markets for illegal activities have formed. Attackers and fraudsters have been able to make a living by specializing in an exploit area, using an IRC to meet like-minded people, and selling goods and services. Some things that have been posted for sale on these IRCs:

- Credit card numbers
- Authentication information (Ex. mother's maiden name)
- Checking and savings account numbers
- Login username and password
- Hacked machines
- Hacking software to set up phishing sites, exploit vulnerabilities, etc.
- Spamming services
- Lists of email addresses
- Goods bought using someone else's credit card

Note: To facilitate sales, the salesperson may show the goods to the IRC administrator. The administrator will then verify the goods and mark the person's account as "verified." This mark will show others that the seller is not a ripper (a person who scams the scam artists).

Services are also advertised on the IRC. People act as:

- Receivers - people who go to the bank and pick up money
- Confirmers - people who will pose as the victim when a bank calls to confirm a wire transfer. Often, confirmers will advertise their ability to do male or female voices, accents, location, etc.
- Mules - people (who have an account at the victim's bank) that will log into a victim's account and transfer money to their own accounts
- Drops - people who offer a physical location where goods can be shipped to
- Callers for a change of billing address (COBs) - As a protective measure, some merchants may only ship to the billing address. In order to circumvent this, COBs call the victim's credit card company and pose as the victim in order to change the billing address. Because credit card companies send a change of address confirmation to both the old and new address at the next billing cycle, often, the COB will call the day after the bill goes out. This gives the fraudster a month to use the stolen credit card.

Note: With these types of service transactions, the seller of the information and the receiver/confirmer/mule will negotiate what percentage of money the person makes.

Note: In other cases, fraudsters may advertise a job position looking for someone to do things such as transfer money into a bank account. Thinking that s/he is making authorized transactions, the person who responds to the ad may unwittingly transfer a victim's money into his/her own account and then to the fraudster's account. This innocent middleman will then get prosecuted for fraud once the victim notices their loss.

In a study conducted where an IRC was monitored over seven months, the researchers observed:

- Over 100,000 credit card numbers being offered
- Balance claims on checking, savings, and mortgage accounts totaling over $50 billion
- $5-$15 being offered for a compromised host. An attacker could easily get $5,000-15,000 if they sold a 1,000 machine botnet. With these bots, someone could:
    - Conduct DDoS
    - Threaten DDoS for extortion, targeting grey area sites: porn sites, lottery sites
    - Spam to generate impressions for an ad
    - Manipulate stock prices
        - $10 billion lost per year
    - Conduct phishing attacks – use the botnet to run the webserver
        - Identity theft has cost $12 billion per year
    - Conduct click fraud – use the botnet to generate fake clicks on the fraudster's hosted ads

It is suspected that $20 billion per year is lost due to Internet related fraud.