sandvine

Worm Mitigation on Broadband Networks

Service provider strategies for managing worm attacks

An Industry White Paper

www.sandvine.com

408 Albert Street
Waterloo, Ontario
Canada
N2L 3V3

Executive Summary

Worms are considered more loathsome than viruses to service providers, especially in networked environments, since they use up network bandwidth and often have an evil intent.
Service providers are now bearing the brunt of worm attacks, as the focus has shifted to residential broadband subscribers. These residential subscribers represent the weakest, most uncontrolled point in the Internet, while being very expensive to protect en masse.

Worms continue to improve in terms of their sophistication and detrimental impact. The latest trend in worm creation is the utilization of peer-to-peer (P2P) file-sharing networks, such as KaZaA or Morpheus, as a means to infect innocent victims. By exploiting the benefits of peer-to-peer file sharing, worms spread more efficiently and have a greater potential of exhausting service provider's networks. From a subscriber's perspective, P2P viruses are particularly dangerous because subscribers have no way of determining the integrity of a file until it is downloaded.

Individual protection is not effective given that a very small percentage of unprotected or poorly protected subscriber machines can create untold havoc for the rest of the service providers network. Security professionals know that current network defenses, including anti-viral products and firewalls, are inadequate to prevent the rapid spread of worm infections at the network level. There are thousands of different ways a worm can infiltrate a network, making it nearly impossible to feel completely confident any network is truly protected.

Future worms will without a doubt be more destructive than what has been seen to date. The tools to accomplish attacks across networks are becoming more widespread and easier to use every day. Hyper-powered worms like Slammer and Code Red have created an urgent need for equally powerful network-based responses. In order to better protect against malicious worm attacks, service providers are looking for a more proactive approach to worm mitigation.

Network Worms

An Introduction

Worms are an insidious subset of viruses that crash the performance of broadband networks. Worms exploit common security holes and then reproduce themselves at incredible speed by transmitting their progeny over high-speed Internet connections. Since they take advantage of common weaknesses, almost every computer system is vulnerable, making worms one of the most fearsome security threats.

Considered more loathsome than viruses to service providers, worms consume massive amounts of bandwidth as they replicate. And depending on the number of unsecured servers, a worm can create hundreds of thousands of copies of itself in a matter of hours.

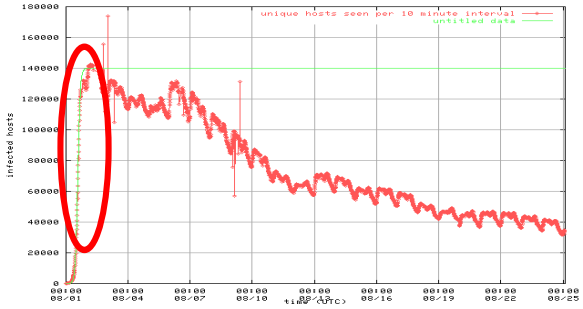

The above figure shows massive Code Red penetration, replicating itself over 250,000 times in approximately nine hours on July 19, 2001.
Source: http://howstuffworks.lycoszone.com/virus3.htm

The latest trend in worm creation is to utilize peer-to-peer (P2P) file-sharing networks, such as KaZaA or Morpheus, as a means to infect innocent victims. By exploiting the benefits of Peer-to-Peer file sharing, worms spread more efficiently and have a greater potential of exhausting service provider's networks.

Worms and Broadband Networks

A Brief History

Worms were first noticed as a potential security threat in the late 1980's. The first worms utilized TCP/IP protocols, common application layer protocols, operating system bugs, and a variety of system administration flaws to propagate. Various problems with worm management resulted, including extremely poor system performance and the complete denial of network service.

More recent worms have attempted to disable anti-virus and security software on infected computers. Some worms steal data by attaching images or document files to the infected messages they send out, while others have destructive payload characteristics that destroy infected systems all together.

Famous Worms

The Code Red attack of July 2001 was the first to gain major publicity. It spread rapidly and globally until almost all vulnerable servers on the Internet had been compromised.

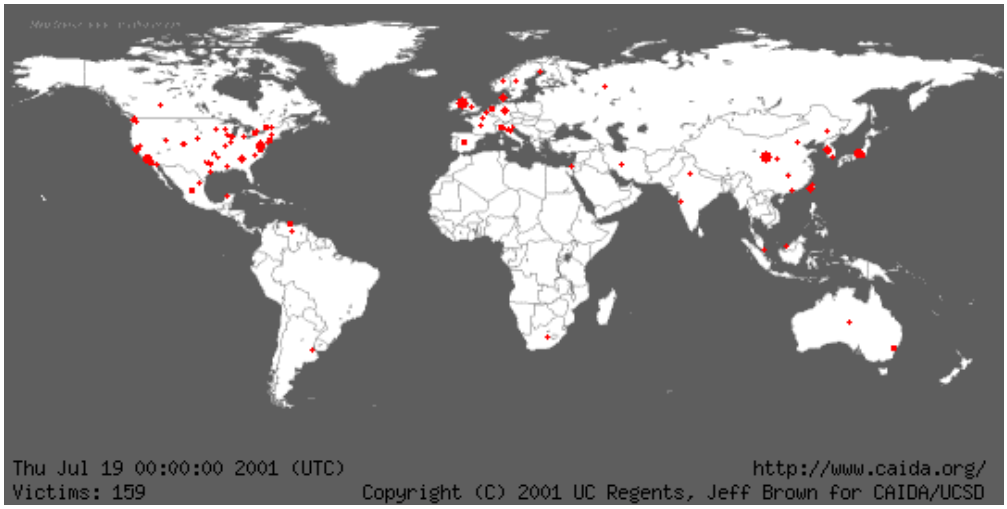

CodeRed, July 19, 2001, infected 359,000 hosts in less than less than 14 hours

The Nimda worm of September 18th 2001 was a "multi-mode" worm capable of infecting from a number of different vectors. It set a new standard of ferocity by spreading very rapidly while penetrating thousands of firewalls. Nimda reached saturation within a few hours and maintained itself on the Internet for months after inception.

The more recent Slammer or Sapphire worm was introduced at 9:30pm on Friday January 24th, 2003, exploiting a vulnerability in Microsoft's SQL server.

Despite being the smallest worm to date, only 376 bytes, it was by far the most network damaging. Scan rates ranged from 3,000 to 30,000 packets per second. This resulted in a very dramatic spread, initially doubling itself in only 8.5 seconds and almost entirely saturating the network in a mere ten minutes. Disruption included blocked networks and infected SQL servers, rendering both unavailable to perform critical tasks. The general public experienced flight cancellations, election interferences, ATM failures and 911 emergency center shut downs.

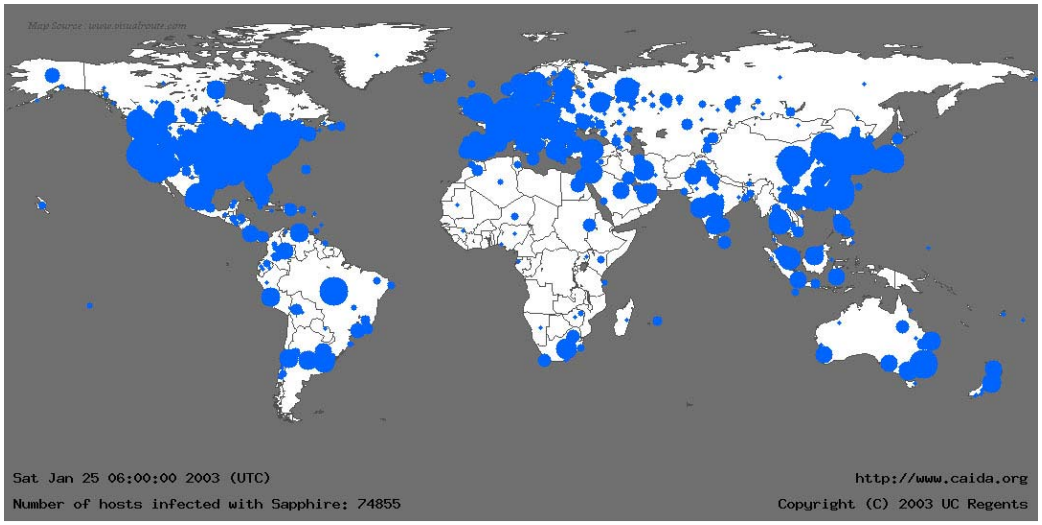

Slammer/Sapphire, January 25, 2003, infected 90% of vulnerable hosts in less than 10 minutes. Within 3 minutes the number of slave servers in Slammer's replicant army was doubling every 8.5 Seconds. (One hundred times faster than Code Red two years earlier.)

P2P and worms – a new vector of infection

The latest trend in worm creation is the utilization of peer-to-peer (P2P) file-sharing networks, such as KaZaa or eDonkey, as vectors of infection. According to anti-virus firm Symantec, the Linux Slapper Worm was the first worm to make use of peer-to-peer networking technology.

By exploiting the benefits of peer-to-peer file sharing, worms spread more efficiently and have a greater potential of exhausting service provider's networks. From a subscriber's perspective, P2P viruses are particularly dangerous because subscribers have no way of determining the integrity of a file until it is downloaded.

Not only do these worms travel faster and with greater ease, they also allow infected servers to maintain contact, potentially providing hackers with control over an entire network.

Virus writers also use P2P technology to endow their worms with networking capabilities. SecurityFocus predicts that "P2P technology might allow conventional Internet worms to update themselves, perhaps making themselves invulnerable to anti-virus software by communicating with other infected systems and receiving new code." Slapper is once again a primary example of a worm exhibiting these sophisticated capabilities.

Worm Traffic & Service Providers

Service providers are now bearing the brunt of worm attacks, as the focus has shifted to residential broadband subscribers. These residential subscribers represent the weakest, most uncontrolled point in the Internet, while being very expensive to protect en masse.

The collective probing and multiplying behavior of worms cause routers and flow-based devices to exhaust resources. Processing this excess traffic degrades subscribers' service levels, ultimately impacting the service provider.

Individual subscriber protection is not effective given that a very small percentage of unprotected or poorly protected subscriber machines can create untold havoc for the rest of the service providers network.

Cleaning network outbreaks can require a complete network upgrade. Disinfecting each computer one by one can be extremely costly and time consuming. The market research firm *Computer Economics* produced widely cited estimates of the total cost of major worm and virus incidents. For recent worms, overall costs range between 0.64 and 2.62 billion worldwide. (See below) It has also been estimated that the public and private sectors combined spend millions of dollars a day to chart and protect themselves against worms. This is an amount that is expected to climb.

Code Red	$2.62 billion
Nimda	$0.64 billion
Slammer	$1.25 billion
Blaster [1]	$2.0 billion

(Blaster cost includes the cost of the near simultaneous Sobig.F virus.)

An additional cost not always accounted for by service providers is that of technical support. Recent findings by major North American service providers report seven in ten broadband subscribers called for technical support during the last year, with 40% having to call twice. At an average cost of $13 a call, and a cost of $150 for site visits, technical support and customer service can cost service providers millions of dollars a year. This cost skyrockets when a worm infects their networks and degrades service levels.

With service providers under intense pressure to cut operational costs, reducing customer support burdens alone stands to significantly improve their bottom line.

Current Solutions

Since the days of Code Red in 2001, worm mitigation strategies have involved network equipment vendors scrambling to find and provide patches, Access Control Lists (ACL's) or other "band-aids" to react to the problem of the day. Often this remedy is more service impacting than the disease itself, as emergency network upgrades and patches can cause outage problems of their own.

The only class of security tools currently employed to proactively protect networks from worms are firewall and anti-virus systems. Firewalls protect organizations and individuals from incidences in the larger network world. An *intelligent* firewall filters all connections between hosts on the organizational network and the world-at-large while a *simple* firewall disallows all connections with the outside world, essentially splitting the network in two. However, both anti-virus and firewall protection assume you know the signature of the virus you are looking for at which time it is often too late.

Security professionals know that current network defenses, including anti-viral products and firewalls, are inadequate to prevent the rapid spread of worm infections at the network level. There are thousands of different ways a worm can infiltrate a network, making it nearly impossible to feel completely confident any network is, in reality, protected.

Worms can easily foil firewall mechanisms by targeting web server hosts or entering the enterprise network via incoming e-mail. Both the Code Red and Sapphire worms treated the Internet as a single, flat address space and spread solely via the worm mode for a single exploit.

Other than virus and firewall protection solutions employed at the residential subscriber level, adequate "layered defenses" at the service provider level have not been developed until recently.

Future Considerations

Future worms will without a doubt be far more destructive than what has been seen to date. The tools to accomplish attacks across networks are becoming more widespread and easier to use every day. Future incidences will exploit more widespread vulnerabilities, will be better tested, work on a broader range of systems, use faster speed algorithms and have more malicious payloads. It is important to stay ahead of the attackers but it is next to impossible to keep a determined attacker out.

Today's large, ambivalent population of multi-user systems is an attractive target for both virus authors and worm developers. Personal computer worms or virus/worm hybrids are becoming a huge threat. With a large homogeneous population of systems available, it is conceivable that authors of malicious code will combine the previously disjoint attacks of viruses and worms. An attack consisting of a worm traversing a network and dropping viruses on the individual hosts becomes a startling possibility.

The worst-case scenario is a flash worm; a worm optimized with knowledge of the Internet's topology. With a flash worm the worm releaser scans the network in advance and develops a complete hit list of vulnerable systems on the network. The worm carries this address list with it, and spreads throughout the list. Flash worms are also hard to contain and have the ability to potentially penetrate the Internet within tens of seconds.

Reflecting the widespread struggle to properly confront the virus threat, the Gartner IT Security Summit, held in London on September 15th, focused much attention on the issue of virus management.

Gartner Inc., a technology research and advisory firm, cautioned companies not to rely solely on the anti-virus solutions Microsoft plans to embed in its Windows operating systems.

Tests conducted by web security specialist Sanctum on behalf of software application testing specialist, Sim Group, show 97% of websites to have significant security flaws.

Fortunately, worms to date have been relatively benevolent in that none have combined extremely damaging capabilities with technological sophistication. Nonetheless, it is only a matter of time.

Future malicious payload scenarios:

- Wipe out hard drives on all infected machines
- Damage hardware by reflashing bios, causing computers to become inoperable
- Perpetrate DDOS attacks on many targets simultaneously
- Search infected machines for intellectual property of a particular or general sort
- Level stealthy attacks by remote and anonymous control via a worm distributor
- Accept new software modules that propagate through the worm and give it new behaviors at run time.
- Corrupt data over time and in a subtle and difficult to detect way

Conclusion

The latest research shows that file-sharing networks have become favored vectors for worm infection, allowing worms to spread more efficiently and have greater potential of exhausting service provider's networks.

Hyper-powered worms like Slammer and Code Red have created an urgent need for equally powerful network-based responses. Besides the obvious threat they pose to subscriber PCs, worms sap processing power from broadband routers and flow-based devices as they struggle to handle excessive, maliciously crafted traffic. This degrades the Internet experience across the service provider's entire network and negatively impacts end users.

In order to better protect against malicious worm attacks, service providers are looking for a more proactive approach to worm mitigation.

Sandvine's award-winning network equipment helps broadband service providers better manage the growing burden of peer-to-peer (P2P) activity while protecting subscribers and preserving their overall Internet experience. Sandvine Peer-to-Peer Policy Management helps service providers take control of P2P traffic, stop the proliferation of destructive worm code and achieve new operational efficiencies. Sandvine products are suitable for all broadband and narrowband networks, cable or DSL. To find out more, contact us at marketing@sandvine.com.

References

- Network World Fusion - http://www.nwfusion.com/news/2002/0913p2pworm.html
- Security Focus - http://www.securityfocus.com/
- NetWorm.org - http://www.networm.org/faq
- Computer Economics - www.computereconomics.com/
- Network World Fusion - http://www.nwfusion.com/news/financial/symantec.html
- S. Stanford, C. Kahn, "Worm Containment in the Internal Network" Silicon Defense Technical Whitepaper, March 2003 - http://www.silicondefense.com/
- Total Telecom - http://www.totaltele.com/view.asp?ArticleID=100848&Pub=tt