



SANS Institute

Information Security Reading Room

Large Scale Network Incidents - What Can We Do?

Jay Garden

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Large Scale Network Incidents – What Can We Do? Jay Garden - 17 November 2002

1. Abstract

Network incidents such as worms and distributed denial of service attacks are a problem not only for the systems targeted by them but also for the Internet in general, as the attacks also involve other systems and consume the available network resources. This assignment looks into the similarities between the two types of attacks and discusses ways to mitigate the risk from an Internet-wide perspective.

2. Background

In February of 2000 a young hacker going by the handle of *Mafiaboy* opened the world's eyes to Distributed Denial of Service (DDoS) attacks when he bombarded eBay, Amazon, CNN and a number of other high profile e-commerce companies with a catastrophic volume of IP packets from multiple sources [Vert]. While his name may be the one associated with this type of attack, the mechanisms - such as Stacheldraht, TrinOO, and TribeFloodNet - were developed by others well before the attacks [Ditt]. Since then, several other DDoS systems have been developed and used on the Internet, generally building on the functionality of the previous ones.

More recently, the Internet's 13 Root DNS servers were attacked with a DDoS attack. According to Computerworld the impact was minimal and not particularly sophisticated [Vija], but with a more sophisticated or longer duration the results could have been much more severe, disrupting all Internet services that depend on domain name resolution. The 2002 Australian Computer Crime and Security Survey [AUS] noted that 41% of the respondents had encountered a denial of service attack in the last twelve months. Unfortunately there is no reason to believe that these forms of attacks have reached a peak either in frequency or sophistication.

Back in 1988 the Morris Worm [Spaf] caused thousands of VAX and SUN-3 computers, a large percentage of the Internet at the time, to seize up as they attempted to run multiple copies of the worm. The worm exploited vulnerabilities in *Sendmail*, *fingerd* and the trust relationships and passwords on the target systems in order to gain system privileges to run a version of itself on the compromised host. That process would then search for other computers to infect. While the action of compromising the host was damaging in itself, the flooding activity on the network due to the many copies of the worm all attempting to infect other computers effectively disrupted even computers that were not vulnerable to infection by the worm.

More recent worms such as *Nimda* and *Code Red* have also successfully propagated through the Internet. Almost 360,000 hosts were estimated to be infected by *Code Red II* in the first fourteen hours [CAIDA]. As well as infecting many computers, both had the effect of consuming countless processor cycles and bandwidth as they attempted to infect other computers. The statistics provided by the Messagelabs Viruseye service [Mess] indicate that even now, over a year after *Nimda* was first discovered, it is still generating significant unwanted traffic.

Worms and DDoS have completely different mechanisms but the load on the Internet and the outcomes for the majority of users can be quite similar. Both can affect systems wherever they are the direct target or not. They both clog up networks, and other peoples unprotected systems can result in denial of service to otherwise protected systems.

While a simple solution against these attacks appears to be to ensure that systems are hardened, patches are up to date, and the processor power and memory is adequate to handle the maximum bandwidth available to it, that would be an expensive solution over the long term and does not solve the issue of the lack of guaranteed Internet connectivity. Methods to make worms propagate faster, for example the - to date, theoretical, *Warhol* worm [Stan] - indicate that these attacks are going to get faster and more resource hungry. The continuing stream of varieties and sophistication of computer viruses and hacking attacks also suggest that DDoS suites and computer worms will continue to evolve for at least the next few years. To date, attacks such as these appear to be intended as proof-of-concepts or showcasing their skills; the effect from a truly motivated attacker such as from a nation state at war or a terrorist/activist organization could have a much greater effect on the target and the compromised systems, and result in more extensive collateral damage to the Internet.

The potential for successful attack through either of these vectors is also increasing through evolution of the huge number of small systems on the Internet. Advancements in peer-to-peer networking, affordable always-on high-speed connections, and convergence of other technologies such as cell-phones and WebTVs has increased the volume, availability and bandwidth of systems that could potentially be compromised for a DDoS attack or worm infestation.

3. Phases of a DDoS attack

In general, preparation and conducting a DDoS attack will have the following phases.

The DDoS system will be **developed or an existing system adapted**. The exploit or exploits used to compromise the zombie computers will then be selected and coded into distribution software or manual processes developed. The attack type will then also be selected, whether it is a basic flood attack or something more efficient such as a flood combined with crafted packets. The system will probably require testing, perhaps on the Internet or on the perpetrator's own systems.

Once the system has been developed the perpetrator can commence **building the zombie network**. One or more computers will first be compromised to provide a remote login for the perpetrator, and the control and zombie tools will be downloaded to them. From these 'master' system(s) the perpetrator then performs a scan of a specific address and port range to determine potentially vulnerable hosts to be used as zombies. Built into this or as a discrete step following it, the zombie exploit developed or selected earlier would then be used to attempt to initialize the backdoor zombie process onto the shortlist of hosts identified in the scan phase. The IP address database of the zombie computers will then be populated either through the host compromise phase, or through them reporting their existence subsequently. The network will then operate in the **standby phase**, potentially for months. However, the longer the zombies are left in this phase the more chance there is of them being detected, removed or altered, renegeing their use to the perpetrator.

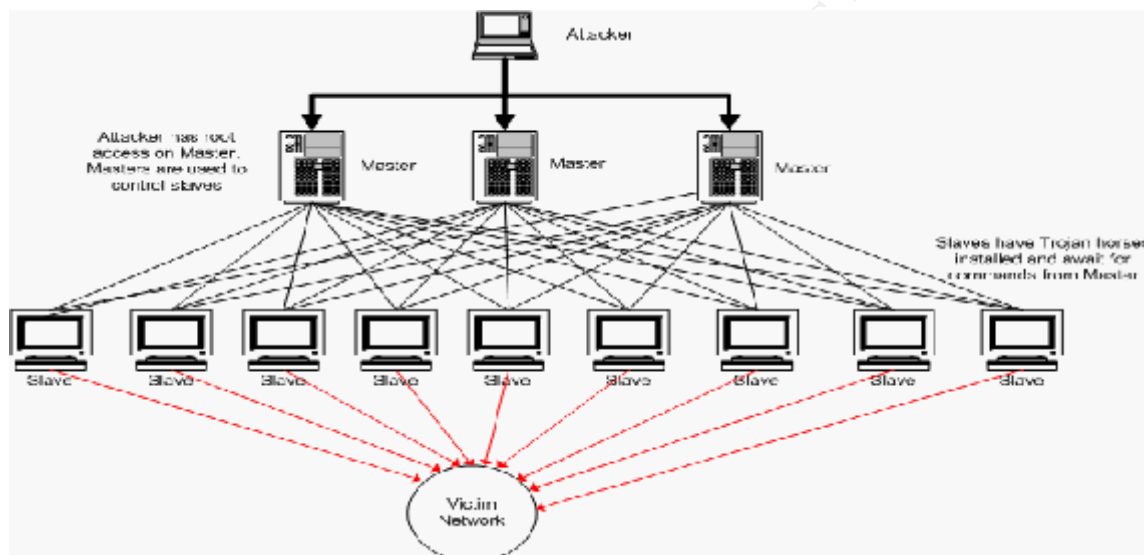


Figure 1: A DDoS Network (Val Popenko [Pipe])

When the attack is mounted the first thing normally to occur is the perpetrator will login to the master system and send the command sequence. This will result in an '**attack command**' being sent to each of the zombies. This command will contain the name or IP address of the target computer or network and any specifics of the attack. It will generally not require acknowledgement so will either have no return address or a spoofed one.

The zombie computers will then attempt to **flood the target** with packets as defined by the attack type. As with the master to zombie communications, the packets will generally have no or spoofed addresses. The success of the attack will depend on the ability of the packet flood getting to the target and having sufficient volume and duration to overwhelm its processing and/or memory capability. However, as the packets flow from the zombies to the target they will also affect the systems around and between those systems. Thus a **secondary denial of service** effect can also result on those systems. If the attack traffic can be blocked the only impact on the target will be a

possible denial or reduction of service from that communications path for the duration of the attack.

Following the attack, an **investigation** may be conducted to trace the zombie computers and the master back to the perpetrator. Because of the lack of true return addresses this can be an extremely difficult, if not impossible, task. Owners of the zombie computers might then be warned of their compromise, or may have detected the spurious traffic themselves, so the zombie process can be removed and the compromise vector sealed. Otherwise the zombies will revert to **standby** mode in anticipation of their next instructions.

4. Phases of a worm incident

While worms are an entirely different type of attack than a DDoS attack there are a number of similarities in the phases and secondary impacts from them, for example, a multitude of compromised systems and a reduction of service for computers targeted by it or around infected systems.

As with a DDoS attack, the worm system will be **developed or an existing system adapted**. The exploit or exploits used to compromise the computers to be infected will then be selected and coded into the worm. Manual or automated process might also be developed to compromise computers to be used as the initial launch points. Whereas with a DDoS attack an attack type will be selected, with a worm a payload will be selected. Many of the worms to date have not included a payload. The system will probably require testing, possibly only on the perpetrator's own systems.

Once the system has been developed the perpetrator can commence **setting up the launching base** by compromising one or more computers to provide a remote login for them to install the worm and its initialization program.

The worm could be activated then or set up to **standby** for a specific time, event or command to activate. As with the DDoS zombie and master computers, the longer they are left in this phase the more chance there is of them being detected, removed or altered, reneging their use to the perpetrator.

When the worm is **initialized into the wild** it will begin to **propagate** by scanning for computers that could be vulnerable to its exploits. This scan may operate in a similar fashion to building a zombie network, with scanning within a range of IP addresses on a specific port or ports. Once a potential host is discovered, the worm will attempt the exploit to gain system privileges so a copy of the worm process can be installed and activated. The worm infection communications will generally be two-way and use the genuine IP addresses but this does not have to be the case if it can attempt the exploit 'blind'. However, there is not a lot of benefit to the perpetrator in doing this, since there is little risk to them of being tracked through the many infection cycles back to them once the worm has gained momentum. The success of the worm to infect the most number

of hosts will depend on its ability to locate and exploit potential hosts with as minimal duplication of attempts as possible against the ability of the network to detect and block the malicious communications. The worm might then be programmed to go into another standby phase.

While the intention of a worm is not generally denial of service, as an infected host attempts to infect other systems it will not only affect the hosts it is attempting to infect, but also those systems around and between those systems regardless of whether they are vulnerable. Thus a **collateral denial of service** effect can also result on those systems. If the infection traffic can be blocked the only impact on those systems will be a possible denial or reduction of service from that communications path until the worm goes into a standby phase or the infection process is stopped or blocked. The further upstream (closer to the infect host) the blocking is the more probability it will have of limiting the overall success of the worm.




Figure 2: Onset of Code Red I v2, Code Red II, and Nimda: Number of remote hosts launching confirmed attacks corresponding to different worms, as seen at the Lawrence Berkeley National Laboratory (Staniford, Paxson, Weaver [Stan]). (Note: Hosts are detected by the distinct URLs they attempt to retrieve, corresponding to the IIS exploits and attack strings. Since Nimda spreads by multiple vectors, the counts shown for it may be an underestimate)

Following the infestation, an **investigation** may be conducted to trace the infected computers and back to the perpetrator. With a worm that continues to attempt to propagate, tracing will be relatively simple. Worms that infect in bursts and then sleep

for a period may be much more difficult to trace, especially if spoofed or no IP addresses are used. Owners of the infected computers should then be notified of their compromise, or may have detected the spurious traffic themselves, so the worm can be removed and the compromise vector sealed.

With a super-worm like the theoretical type described in “Owning the Internet in your Spare Time” by Stuart Staniford, Vern Paxson, and Nicholas Weaver [Stan] there is a greater emphasis on the set-up phase, starting the propagation phase further along the propagation curve (see Figure 3) resulting in a much shorter period between initial detection and saturation.




Figure 3: Spread of Fast Propagation Worms Compared to Current Varieties (Staniford, Paxson, Weaver [Stan])

5. What is being done now?

Currently the primary method to limit the vulnerability and effects of DDoS and worms is to ensure that your own systems are secure from compromise of the most common attacks and are robust enough to resume operation once the barrage ceases [CER1]. One of the difficulties with DDoS attacks is that it is very difficult to trace the source of the attack. In the case of the February 2000 attacks, Mafiaboy was discovered primarily through his discussing his exploits on an IRC channel being monitored by the Royal

Canadian Mounted Police [Vert], not through tracing the attacks themselves even though the attacks were spread over several days giving much more opportunity to trace the attack than will usually be available.

General compliance to best practices such as RFC2267 Network Ingress Filtering [RFC] also means that attacks utilizing IP address spoofing will have a much lower chance of success, and those that do use their real addresses will be easier to trace.

Worms on the other hand can often be traced back to the infecting computer relatively easily, however determining where it originated from is often a much harder task since the initialization host or hosts are possibly many iterations back. By that time the worm has already run its course and the damage has been done.

It is apparent that we cannot stop many of the effects of these attacks solely by getting more powerful computers and keeping them patched. While this is a good start, as the Internet expands and bandwidth to endpoints is increased, management of the problem this way will be forever in catch-up mode, and a computer attached to a network that is blocked up with such spurious activity will still be denied service even if the computer remains operational.

One of the other problems with depending on relying on tighter security of host computers to limit these attacks is that the Internet is global and is not under the control of any one country or group. It is unlikely that a large enough percentage of users will have the motivation and ability to secure their systems, or that every country will enforce rules on unsociable Internet users.

In a presentation to the North American Network Operators Group (NANOG) in May 2001, Massimiliano Poletto noted the "DDoS attacks are really network operations and performance problems, rather than strictly security events. To effectively address these attacks, the network infrastructure must be able to quickly identify unusual packet streams at high rates and help operators move closer to the packets' sources. "[Pole]

Current network and host based intrusion detection systems can often detect DDoS attacks and worm propagation, however they will generally report after the fact and cannot usually stop the event from occurring. They can however help to trace it and to minimize the impact.

The primary emerging technologies to manage this issue are network intelligence and filtering systems such as Esphion's NetDeflect product [Esph]. Designed for placement on access routers or service provider feeds, the hardware software combination will detect attacks through signature-based detection or through assessment of traffic characteristics and statistics. They can be configured to dynamically adjust the protocol and address filters of the local router or of a remote router to drop any packets associated with a perceived attack. While these types of systems are useful when they control an organization's Internet entry point, they are most useful when the filtering is performed in an 'upstream' location such as at an ISP or regional Internet nodes. Such

systems are an important element, but only one component, of an effective mitigation strategy.

Another important component is coordination of national and global detection and reaction to such attacks. *How to Own the Internet in Your Spare Time* [Stan] proposes formation of a "Cyber-Center for Disease Control" (CDC) with the six roles of: Identifying outbreaks; rapidly analyzing pathogens; fighting infections; anticipating new vectors; proactively devising detectors for new vectors; and resisting future threats.

In 1999, a Distributed-systems Intruder Tools Workshop hosted by CERT/CC in Pittsburgh [CER2] proposed a number of responsibilities and actions that should be taken by managers, system administrators, ISPs, and incident response teams. Their recommendations stress the need for a wide-ranging strategy aimed at reducing the risk from various directions not just reacting to attacks as they occur.

The recent US Strategy to Secure Cyberspace [Pres] includes several recommendations that relate directly to protection from these types of attacks, placing responsibilities on home users and small businesses, large enterprises, government, universities, network providers, Computer Emergency Response Teams (CERTs), law enforcement authorities, and global security groups. The Strategy recommends (R4-39) formation of a Cyberspace Network Operations Center "to share information and ensure coordination to support the health and reliability of Internet operations in the United States." It also recommends establishment of a Cyber Warning Information Network (R4-40) to key government and non-government cybersecurity related network operation centers, to disseminate analysis and warning information and perform crisis coordination.

6. A Strategy for Limiting the Risk of These Network Hogs on a National and International Scale

6.1 General Risk Mitigation

The *New Zealand Security of IT publication 104: Risk Analysis* [GCSB] lists the methods of reducing risk as:

- a. *Risk Avoidance. Risk may be avoided by eliminating or relocating an asset. For example, a decision may be made to discontinue processing a particular class of information on a system.*
- b. *Transfer of Risk. Risk may be transferred where an asset is moved to a different security domain. For example, processing may be moved to a different site or outsourced. Insurance is both a transfer of financial risk and an impact reduction measure.*
- c. *Reduction of Vulnerability. Countermeasures to reduce vulnerability can be viewed as barriers which increase the effort an attacker would have to expend to achieve a successful attack. In the case of risks arising from error, negligence, or accident vulnerabilities may be reduced by countermeasures such as improved staff training.*
- d. *Reduction of Impact. It may be cost-effective to accept a certain level of risk where the impact of an incident can be reduced, for example through insurance. Countermeasures may reduce impact by reducing the cost of recovery through, for example, disaster recovery planning."*
- e. *Detection. Detection countermeasures may reduce risk or facilitate the early detection of an incident and therefore reduce its impact. Examples include error logs, access control logs or journals, and recordings from security cameras.*

In some circumstances the risk may also be reduced through reducing the level of threat, for instance countering the potential perpetrator's motivation to carry out the incident.

Since the intention of this paper is to look into a global solution to be able to use the Internet as we do now, I will not review methods for avoiding or transferring the risk. The remainder of this paper will explore each of the other four methods of risk reduction for the generic phases of DDoS and worms to reduce the individual and global risk from these types of attacks.

6.2 The Proposed Approach

A multiple pronged approach will be important if we are not only to protect against current attacks but also to build up a culture and infrastructure that is resilient to new threats and vulnerabilities. This section takes each of the DDoS and worm phases and looks at methods to detect the attack or the build-up to the attack, to limit the vulnerability to the attack, to reduce the threat, or to minimize the impact

The generic phases of worms and DDoS attacks discussed in Sections 3 and 4 are:

- (1) Initial development, selection and/or adaptation of the worm or DDoS system.
- (2) Building the zombie network or setting up the launch base
- (3) Standby phase
- (4) Start command or condition
- (5) D-day. Attack or propagation phase
- (6) Collateral denial of service effect
- (7) Tracing and retribution
- (8) Repeat attack or re-infection

Phase (1) Initial development, selection and/or adaptation of the worm or DDoS system.

This is one of the more difficult phases to reduce the risk, but would also be the best, before any damage is done or systems compromised. There are several ways this phase could be detected, with the emphasis or responsibility on CERTs and government law enforcement and security agencies to monitor hacker news and discussion sources. Many potential perpetrators could also be dissuaded from commencing on a worm or DDoS campaign by authorities successfully tracing and prosecuting similar incidents, and ensuring the outcomes are well-known amongst the potential attacker communities. Hacker and programmer groups could potentially make the biggest difference here, by showing young hackers that these types of attacks are not acceptable behavior. As with graffiti, if the developers of these tools are given an opportunity to showcase their system in a controlled environment they may also be less inclined to take them 'into the wild'.

Phase (2) Building the zombie network or setting up the launch base

This phase should be easier to detect, trace and prosecute than the first phase. Traditional IDS and Honeypots may detect the perpetrator as they attempt to compromise the master, zombie or launch systems. Network anomaly detection by gateways and ISPs may be able to pick up the attempts on their networks. However, obviously, this would not detect attempts on networks that did not have this capability enabled, so it would need to be widespread and impossible for perpetrators to differentiate the systems with detectors in advance. Canary systems could be used to capture a working example of the zombie or worm code.

The vulnerability of individual networks and computers to infection or compromise will be reduced by ensuring best practices for servers, clients and gateways are widely available to a global audience. The guidelines should include maintenance and configuration management. The SANS [SANS] and Center for Internet Security [CIS] tools and guides provide such guidance to a global audience. System administrators, including home users, must be encouraged and helped to keep patches up to date, and to detect and remove malicious software on their computers. This is most easily done through integration of the malware signatures into anti-virus systems. On a national scale, government, CERTs and security experts should work with vendors on the core vulnerabilities being exploited. Common Criteria protection profiles may also require adjustment as attacks such as these evolve.

The threat of systems being targeted in this phase will also be reduced by increasing the effectiveness and reputation of cross-border prosecution.

Managers of systems that are compromised can minimize the impact to them by the use of host based integrity checking and logging. This will allow them to assess the scope of compromise. The use of appropriate and tested back-up and recovery processes will minimise the impact from the event.

Phase (3) Standby phase,

The standby phase may not involve any activity by the perpetrator or the compromised systems. The most effective means of detection will generally be a host-based integrity or anti-virus system. Full file system anti-virus scans that include worm and zombie signatures will detect the known systems and may be effective against some new varieties.

Phase (4) Start command or condition

The worm or DDoS will be activated by a command or when a specific condition is satisfied. The command will generally be very difficult to detect, being short and usually without a known signature. However, if the perpetrator is covering his tracks it may not

use the genuine source IP address, so may be blocked through wide-spread compliance with ingress and egress filtering at gateways and major nodes.

This is the moment of no return for the perpetrator, so an expectation of being traced and prosecuted could stop the majority of them at this point.

Phase (5) D-day. Attack or propagation phase

The trigger has been pulled and the propagation or flood commences. In the case of DDoS, this will involve the master sending command signals to the zombies, then the zombies commencing the attack. For a worm, it will involve scanning for other systems to infect and then attempting to infect them. With both types this period of multiple identical communications should be relatively easy to recognize with a network intelligence system such as NetDeflect or a Network IDS with an appropriately configured statistical rule-base. Coordinated Honeypots and Canary systems scattered around the Internet may also be useful at this point if the attack is new or unusual and needs to be analyzed before it can be effectively defended against.

The vulnerability at this phase can be reduced by having built-in redundancy and controlled bottlenecks for critical systems, along with the measures described for Phase 3. Dynamically controlling the unwanted communications will help to protect the target systems, especially if the filtering can be done well upstream from the target, the closer to the source the better. For many systems some of this type of 'firewalling' could be set in well in advance. For instance, ISPs could offer an option to block service requests to client workstations such as home users. Such a service would have protected them from the unwanted Code Red worm requests to port 80 (HTTP) and the exploits associated with it. General implementation of egress and ingress filtering [RFC] will also either block the traffic or make it more difficult for the perpetrators to cover their tracks.

The impact at this phase can be reduced both to the individual and the Internet in general by ensuring that technical outlines of the attack modus-operandi and removal/recovery procedures are available worldwide as soon as possible. CERTs and anti-virus vendors have traditionally been very good at providing this information to the public. To minimize the impact on the individual systems that are either directly or indirectly affected by the attack, owners of critical systems should ensure they have documented and tested back-up and recovery processes. They should also consider planning for a contingency transmission path for crucial communications.

For new varieties of attacks, coordination between CSIRTs, CERTs, network operators, security experts, security vendors and government security agencies will be crucial if the Internet is to be restored in a timely manner. The contact points will need to be established in advance, and the communications may have to be performed through a pre-established secondary communications path, since the Internet may not be in a usable form at the time.

Once the incident has passed it will be important to minimize the possibility of a repeat, so an awareness campaign may be required to encourage users and administrators to check their systems are not compromised, are patched, and have the current anti-virus, IDS and DoS signatures installed.

Phase (6) Collateral denial of service effect

The collateral effect of worms like *Code Red* and *Nimda* is still being felt, with some systems that have been infected for many months still attempting to infect others. These systems can often be detected and traced from the networks they are attempting to infect. The owners can be informed so that they will, hopefully, disable the instance of the worm and patch the vulnerability the worm or perpetrator used to gain entry.

The impact of the unwanted traffic can be minimized by filtering that traffic, preferably as far upstream and close to the source as possible.

As with the systems directly affected as described in the previous phase, contingency communication paths may be required for critical systems in situations where the Internet becomes clogged up with attack traffic.

Phase (7) Tracing and retribution

Zombie and worm infected computers should be relatively easy to trace if adequate logging has been configured and kept. The perpetrator will generally be more difficult to trace but this may be facilitated by police and intelligence services via traditional intelligence gathering mechanisms, and network forensics.

If the perpetrator is traced, they should be prosecuted and the results shared with the community to ensure awareness and consistency. *UN General Assembly Resolution 55/63 Combating the Criminal Use of IT* [APEC] is a step towards ensuring international consistency of these law enforcement functions and frameworks.

Where appropriate, others should also be noted for their non-action, for instance network providers/ISPs for not blocking obviously malicious flood traffic, and owners of zombie owners for their computer's participation.

Phase (8) Repeat attack or re-infection

Once an incident occurs it is important to learn from it so we are not vulnerable to the same attack again. Computers that have been infected or compromised should be traced and the owners informed if they are not aware already. It is important that they are provided with the tools and understanding so that they can remove the malicious

software and prevent re infection/compromise. The IDS/virus/attack detection systems may have to be adjusted or even replaced.

In some circumstances governments and CERTs should work with the system developers, vendors and/or network operators on the core vulnerabilities being exploited. Maintenance of Common Criteria Protection Profiles and evaluations may also be warranted to ensure new systems are hardened against similar vulnerabilities.

7. Conclusion

Effective protection against evolving attacks such as DDoS and worms will depend on individual system security and filtering the more obvious and wide ranging attacks at major nodes. However, trying to protect against these attacks solely on a real-time basis will be expensive and one step behind the attacks. While in theory, filtering and strengthening systems could work to stop these attacks, in practicality, with the shape and nature of the Internet, those steps alone are unlikely to be particularly effective for making the Internet as a whole a safe and reliable place to operate.

We need to also deter the perpetrators by building up education, response coordination, prosecution and international agreements in response to these threats. While this may not stop a committed attacker, it should deter those who are tempted to do it for the technical challenge.

Governments, CERTs and Internet community need to jointly develop capability to detect these attacks at the build up phase through intelligence gathering as well as host and network based anomaly and signature based detection systems. Detected activity needs to be coordinated and shared quickly and in a trusted fashion.

Internet security groups need to ensure that best practices, patches and incident support are readily available to the owners of target systems and the intermediary systems.

Operators of gateway and backbone routers should be encouraged to configure their gateways and networks along the recommendation in RFC2267. This will not only help to limit the success of many attacks, but may also hinder the DDoS set-up phase and worm propagation. It should also aid in mapping the compromised systems. They should also be encouraged to provide active filtering of known attack signatures wherever possible without blocking legitimate traffic.

As a last resort, traffic restrictions may need to be enforced on networks and countries that continue to provide weak defenses against attacks originating or being hosted on them. It may be time to stop treating the Internet as one big fully-interconnected cloud.

8. References

[APEC] *Report of Economy Implementations of the Ten Measures Included in UN General Assembly Resolution 55/63 Combating the Criminal Misuse of IT*
APEC Telecommunications and Information WG :
www.apectelwg.org/apec/are/telmin5sub08.htm

[CAIDA] *CAIDA Analysis of Code-Red*, Cooperative Association for Internet Data Analysis
www.caida.org/analysis/security/code-red/

[CER1] *Tech Tips of Denial of Service Attacks*, CERT/CC
www.cert.org/tech_tips/denial_of_service.html

[CER2] *Results of the Distributed-system Intruder Tools Workshop*, multiple participants with CERT/CC
www.cert.org/reports/dsit_workshop-final.html

[CIS] Centre for Internet Security operating system security benchmarks and scoring tools
www.cisecurity.org/

[Ditt]. *The Dos Project's "Trinoo" Distributed Denial of Service Attack Tool*
David Dittrich. staff.washington.edu/dittrich/misc/trinoo.analysis

[EspH] *Overview of NetDeflect*, EspHion <http://www.esphion.com/products1.htm>

[GCSB] *NZSIT104: Risk Analysis*, NZ Government Communications Security Bureau.
www.gcsb.govt.nz/pubs/nzsit/104index.htm

[Mess] Messagelabs Virus Eye Threat list
<http://www.message-labs.com/viruseye/threatlist.asp>

[Pipe] *SANS GSEC Practical: Analyzing Distributed Denial of Service Attacks*
Val Pipenko, 23 Aug 02 www.giac.org/practical/Val_Pipenko_GSEC.doc

[Pole] *Practical Approaches to Dealing with DDoS Attacks*. Presentation to NANOG by Massimiliano Poletto, May 01 www.nanog.org/mtg-0105/poletto.html

[Pres] *US strategy to Secure Cyberspace*, The President's Critical Infrastructure Protection Board. www.whitehouse.gov/pcipb/

[RFC] *RFC2267: Network Ingress Filtering: Defeating DoS Attacks Which Employ IP Source Address Spoofing*, IETF. www.ietf.org/rfc/rfc2267.txt?number=2267

[SANS] SANS Consensus Guides
store.sans.org/store_category.php?category=consguides&sans_store=a60ecde9efb35efdc138ced771214835

[Spaf] *The Morris Worm*. Gene Spafford. "Spaf-lworm-paper-CCR.ps" on ftp://coast.cs.purdue.edu/pub/doc/morris_worm

[Stan] *How to Own the Internet in your spare time* by Stuart Staniford, Vern Paxson and Nicholas Weaver. www.icir.org/vern/papers/cdc-usenix-sec02/

[Vert] *The Hacker Diaries: Confessions of Teenage Hackers*, Dan Verton
ISBN 0-07-222552-1

[Vija] *Attack on root servers resulted in moderate damage - this time*
Jaikumar Vijayan and Patrick Thibodeau, Computerworld, 28 Oct 02.
www.computerworld.com/securitytopics/security/story/0,10801,75454,00.html

© SANS Institute 2003, Author retains full rights.