

Network Working Group
Request for Comments: 4948
Category: Informational

L. Andersson
Acreo AB
E. Davies
Folly Consulting
L. Zhang
UCLA
August 2007

Report from the IAB workshop on Unwanted Traffic March 9-10, 2006

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document reports the outcome of a workshop held by the Internet Architecture Board (IAB) on Unwanted Internet Traffic. The workshop was held on March 9-10, 2006 at USC/ISI in Marina del Rey, CA, USA. The primary goal of the workshop was to foster interchange between the operator, standards, and research communities on the topic of unwanted traffic, as manifested in, for example, Distributed Denial of Service (DDoS) attacks, spam, and phishing, to gain understandings on the ultimate sources of these unwanted traffic, and to assess their impact and the effectiveness of existing solutions. It was also a goal of the workshop to identify engineering and research topics that could be undertaken by the IAB, the IETF, the IRTF, and the network research and development community at large to develop effective countermeasures against the unwanted traffic.

Table of Contents

1.	Introduction	3
2.	The Root of All Evils: An Underground Economy	4
2.1.	The Underground Economy	5
2.2.	Our Enemy Using Our Networks, Our Tools	6
2.3.	Compromised Systems Being a Major Source of Problems	7
2.4.	Lack of Meaningful Deterrence	8
2.5.	Consequences	10
3.	How Bad Is The Problem?	10
3.1.	Backbone Providers	10
3.1.1.	DDoS Traffic	10
3.1.2.	Problem Mitigation	11
3.2.	Access Providers	12
3.3.	Enterprise Networks: Perspective from a Large Enterprise	13
3.4.	Domain Name Services	14
4.	Current Vulnerabilities and Existing Solutions	15
4.1.	Internet Vulnerabilities	15
4.2.	Existing Solutions	16
4.2.1.	Existing Solutions for Backbone Providers	16
4.2.2.	Existing Solutions for Enterprise Networks	17
4.3.	Shortfalls in the Existing Network Protection	18
4.3.1.	Inadequate Tools	18
4.3.2.	Inadequate Deployments	18
4.3.3.	Inadequate Education	19
4.3.4.	Is Closing Down Open Internet Access Necessary?	19
5.	Active and Potential Solutions in the Pipeline	20
5.1.	Central Policy Repository	20
5.2.	Flow Based Tools	21
5.3.	Internet Motion Sensor (IMS)	21
5.4.	BCP 38	22
5.5.	Layer 5 to 7 Awareness	22
5.6.	How To's	22
5.7.	SHRED	23
6.	Research in Progress	23
6.1.	Ongoing Research	23
6.1.1.	Exploited Hosts	23
6.1.2.	Distributed Denial of Service (DDoS) Attacks	25
6.1.3.	Spyware	26
6.1.4.	Forensic Aids	26
6.1.5.	Measurements	27
6.1.6.	Traffic Analysis	27
6.1.7.	Protocol and Software Security	27
6.2.	Research on the Internet	27
6.2.1.	Research and Standards	28
6.2.2.	Research and the Bad Guys	29

7.	Aladdin's Lamp	30
7.1.	Security Improvements	30
7.2.	Unwanted Traffic	31
8.	Workshop Summary	31
8.1.	Hard Questions	31
8.2.	Medium or Long Term Steps	32
8.3.	Immediately Actionable Steps	33
9.	Terminology	33
10.	Security Considerations	38
11.	Acknowledgements	38
12.	Informative References	39
Appendix A.	Participants in the Workshop	40
Appendix B.	Workshop Agenda	41
Appendix C.	Slides	41

1. Introduction

The Internet carries a lot of unwanted traffic today. To gain a better understanding of the driving forces behind such unwanted traffic and to assess existing countermeasures, the IAB organized an "Unwanted Internet Traffic" workshop and invited experts on different aspects of unwanted traffic from operator, vendor, and research communities to the workshop. The intention was to share information among people from different fields and organizations, fostering an interchange of experiences, views, and ideas between the various communities on this important topic. The major goal of this workshop was to stimulate discussion at a deep technical level to assess today's situation in regards to:

- o the kinds of unwanted traffic that are seen on the Internet,
- o how bad the picture looks,
- o who and where are the major sources of the problem,
- o which solutions work and which do not, and
- o what needs to be done.

The workshop was very successful. Over one and half days of intensive discussions, the major sources of the unwanted traffic were identified, and a critical assessment of the existing mitigation tools was conducted. However, due to the limitation of available time, it was impossible to cover the topic of unwanted traffic in its entirety. Thus, for some of the important issues, only the surface was touched. Furthermore, because the primary focus of the workshop was to collect and share information on the current state of affairs, it is left as the next step to attempt to derive solutions to the

issues identified. This will be done in part as activities within the IAB, the IETF, and the IRTF.

During the workshop, a number of product and company names were cited, which are reflected in the report to a certain extent. However, a mention of any product in this report should not be taken as an endorsement of that product; there may well be alternative, equally relevant or efficacious products in the market place.

This report is a summary of the contributions by the workshop participants, and thus it is not an IAB document. The views and positions documented in the report do not necessarily reflect IAB views and positions.

The workshop participant list is attached in Appendix A. The agenda of the workshop can be found in Appendix B. Links to a subset of the presentations are provided in Appendix C; the rest of the presentations are of a sensitive nature, and it has been agreed that they will not be made public. Definitions of the jargon used in describing unwanted traffic can be found in Section 9.

2. The Root of All Evils: An Underground Economy

The first important message this workshop would like to bring to the Internet community's attention is the existence of an underground economy. This underground economy provides an enormous amount of monetary fuel that drives the generation of unwanted traffic. This economic incentive feeds on an Internet that is to a large extent wide open. The open nature of the Internet fosters innovations but offers virtually no defense against abuses. It connects to millions of mostly unprotected hosts owned by millions of mostly naive users. These users explore and benefit from the vast opportunities offered by the new cyberspace, with little understanding of its vulnerability to abuse and the potential danger of their computers being compromised. Moreover, the Internet was designed without built-in auditing trails. This was an appropriate choice at the time, but now the lack of traceability makes it difficult to track down malicious activities. Combined with a legal system that is yet to adapt to the new challenge of regulating the cyberspace, this means the Internet, as of today, has no effective deterrent to miscreants. The unfettered design and freedom from regulation have contributed to the extraordinary success of the Internet. At the same time, the combination of these factors has also led to an increasing volume of unwanted traffic. The rest of this section provides a more detailed account of each of the above factors.

2.1. The Underground Economy

As in any economic system, the underground economy is regulated by a demand and supply chain. The underground economy, which began as a barter system, has evolved into a giant shopping mall, commonly running on IRC (Internet Relay Chat) servers. The IRC servers provide various online stores selling information about stolen credit cards and bank accounts, malware, bot code, botnets, root accesses to compromised hosts and web servers, and much more. There are DDoS attack stores, credit card stores, PayPal and bank account stores, as well as Cisco and Juniper router stores that sell access to compromised routers. Although not everything can be found on every server, most common tools used to operate in the underground economy can be found on almost all the servers.

How do miscreants turn attack tools and compromised machines into real assets? In the simplest case, miscreants electronically transfer money from stolen bank accounts directly to an account that they control, often in another country. In a more sophisticated example, miscreants use stolen credit cards or PayPal accounts for online purchases. To hide their trails, they often find remailers who receive the purchased goods and then repackage them to send to the miscreants for a fee. The miscreants may also sell the goods through online merchandising sites such as eBay. They request the payments be made in cashier checks or money orders to be sent to the people who provide money laundering services for the miscreants by receiving the payments and sending them to banks in a different country, again in exchange for a fee. In either case, the destination bank accounts are used only for a short period and are closed as soon as the money is withdrawn by the miscreants.

The miscreants obtain private and financial information from compromised hosts and install bots (a.k.a. zombies) on them. They can also obtain such information from phishing attacks. Spam messages mislead naive users into accessing spoofed web sites run by the miscreants where their financial information is extracted and collected.

The miscreants in general are not skilled programmers. With money, however, they can hire professional writers to produce well phrased spam messages, and hire coders to develop new viruses, worms, spyware, and botnet control packages, thereby resupplying the underground market with new tools that produce more unwanted traffic on the Internet: spam messages that spread phishing attacks, botnets that are used to launch DDoS attacks, click fraud that "earns" income by deceiving online commercial advertisers, and new viruses and worms that compromise more hosts and steal additional financial information as well as system passwords and personal identity information.

The income gained from the above illegal activities allows miscreants to hire spammers, coders, and IRC server providers. Spammers use botnets. Direct marketing companies set up dirty affiliate programs. Some less than scrupulous banks are also involved to earn transaction fees from moving the dirty money around. In the underground market, everything can be traded, and everything has a value. Thus is spawned unwanted traffic of all kinds.

The underground economy has evolved very rapidly over the past few years. In the early days of bots and botnets, their activities were largely devoted to DDoS attacks and were relatively easy to detect. As the underground economy has evolved, so have the botnets. They have moved from easily detectable behavior to masquerading as normal user network activity to achieve their goals, making detection very difficult even by vigilant system administrators.

The drive for this rapid evolution comes to a large extent from the change in the intention of miscreant activity. Early virus attacks and botnets were largely anarchic activities. Many were done by "script kiddies" to disrupt systems without a real purpose or to demonstrate the prowess of the attacker, for example in compromising systems that were touted as "secure". Mirroring the commercialization of the Internet and its increasing use for e-business, miscreant activity is now mostly focused on conventional criminal lines. Systems are quietly subverted with the goal of obtaining illicit financial gain in the future, rather than causing visible disruptions as was often the aim of the early hackers.

2.2. Our Enemy Using Our Networks, Our Tools

Internet Relay Chat (IRC) servers are commonly used as the command and control channel for the underground market. These servers are paid for by miscreants and are professionally supported. They are advertised widely to attract potential consumers, and thus are easy to find. The miscreants first talk to each other on the servers to find out who is offering what on the market, then exchange encrypted private messages to settle the deals.

The miscreants are not afraid of network operators seeing their actions. If their activities are interrupted, they simply move to another venue. When ISPs take actions to protect their customers, revenge attacks are uncommon as long as the miscreants' cash flow is not disturbed. When a botnet is taken out, they move on to the next one, as there is a plentiful supply. However, if an IRC server is taken out that disturbs their cash flow, miscreants can be ruthless and severe attacks may follow. They currently have no fear, as they know the chances of their being caught are minimal.

Our enemies make good use of the Internet's global connectivity as well as all the tools the Internet has developed. IRC servers provide a job market for the miscreants and shopping malls of attack tools. Networking research has produced abundant results making it easier to build large scale distributed systems, and these have been adopted by miscreants to build large size, well-controlled botnets. Powerful search engines also enable one to quickly find readily available tools and resources. The sophistication of attacks has increased with time, while the skills required to launch effective attacks have become minimal. Attackers can be hiding anywhere in the Internet while attacks get launched on a global scale.

2.3. Compromised Systems Being a Major Source of Problems

The current Internet provides a field ripe for exploitation. The majority of end hosts run vulnerable platforms. People from all walks of life eagerly jump into the newly discovered online world, yet without the proper training needed to understand the full implications. This is at least partially due to most users failing to anticipate how such a great invention can be readily abused. As a result, the Internet has ended up with a huge number of compromised hosts, without their owners being aware that it has happened.

Unprotected hosts can be compromised in multiple ways. Viruses and worms can get into the system through exploiting bugs in the existing operating systems or applications, sometimes even in anti-virus programs. A phishing site may also take the opportunity to install malware on a victim's computer when a user is lured to the site. More recently, viruses have also started being propagated through popular peer-to-peer file sharing applications. With multiple channels of propagation, malware has become wide-spread, and infected machines include not only home PCs (although they do represent a large percentage), but also corporate servers, and even government firewalls.

News of new exploits of vulnerabilities of Microsoft Windows platforms is all too frequent, which leads to a common perception that the Microsoft Windows platform is a major source of vulnerability. One of the reasons for the frequent vulnerability exploits of the Windows system is its popularity in the market place; thus, a miscreant's investment in each exploit can gain big returns from infecting millions of machines. As a result, each incident is also likely to make headlines in the news. In reality, all other platforms such as Linux, Solaris, and MAC OS for example, are also vulnerable to compromises. Routers are not exempt from security break-ins either, and using a high-end router as a DoS launchpad can be a lot more effective than using a bundle of home PCs.

Quietly subverting large numbers of hosts and making them part of a botnet, while leaving their normal functionality and connectivity essentially unimpaired, is now a major aim of miscreants and it appears that they are being all too successful. Bots and the functions they perform are often hard to detect and most of the time their existence are not known to system operators or owners (hence, the alternative name for hosts infected with bots controlled by miscreants - zombies); by the time they are detected, it might very well be too late as they have carried out the intended (mal-)function.

The existence of a large number of compromised hosts is a particularly challenging problem to the Internet's security. Not only does the stolen financial information lead to enormous economic losses, but also there has been no quick fix to the problem. As noted above, in many cases the owners of the compromised computers are unaware of the problem. Even after being notified, some owners still do not care about fixing the problem as long as their own interest, such as playing online games, is not affected, even though the public interest is endangered --- large botnets can use multi-millions of such compromised hosts to launch DDoS attacks, with each host sending an insignificant amount of traffic but the aggregate exceeding the capacity of the best engineered systems.

2.4. Lack of Meaningful Deterrence

One of the Internet's big strengths is its ability to provide seamless interconnection among an effectively unlimited number of parties. However, the other side of the same coin is that there may not be clear ways to assign responsibilities when something goes wrong. Taking DDoS attacks as an example, an attack is normally launched from a large number of compromised hosts, the attack traffic travels across the Internet backbone to the access network(s) linking to the victims. As one can see, there are a number of independent stake-holders involved, and it is not immediately clear which party should take responsibility for resolving the problem.

Furthermore, tracking down an attack is an extremely difficult task. The Internet architecture enables any IP host to communicate with any other hosts, and it provides no audit trails. As a result, not only is there no limit to what a host may do, but also there is no trace after the event of what a host may have done. At this time, there is virtually no effective tool available for problem diagnosis or packet trace back. Thus, tracking down an attack is labor intensive and requires sophisticated skills. As will be mentioned in the next section, there is also a lack of incentive to report security attacks. Compounded with the high cost, these factors make forensic tracing of an attack to its root a rare event.

In human society, the legal systems provide protection against criminals. However, in the cyberspace, the legal systems are lagging behind in establishing regulations. The laws and regulations aim at penalizing the conduct after the fact. If the likelihood of detection is low, the deterrence would be minimal. Many national jurisdictions have regulations about acts of computer fraud and abuse, and they often carry significant criminal penalties. In the US (and many other places), it is illegal to access government computers without authorization, illegal to damage protected government computers, and illegal to access confidential information on protected computers. However, the definition of "access" can be difficult to ascertain. For example, is sending an ICMP (Internet Control Messaging Protocol) packet to a protected computer considered illegal access? There is a lack of technical understanding among lawmakers that would be needed to specify the laws precisely and provide effective targeting limited to undesirable acts. Computer fraud and liabilities laws provide a forum to address illegal access activities and enable prosecution of cybercriminals. However, one difficulty in prosecuting affiliate programs using bot infrastructure is that they are either borderline legal, or there is little evidence. There is also the mentality of taking legal action only when the measurable monetary damage exceeds a high threshold, while it is often difficult to quantify the monetary damage in individual cases of cyberspace crimes.

There is a coalition between countries on collecting cybercriminal evidence across the world, but there is no rigorous way to trace across borders. Laws and rules are mostly local to a country, policies (when they exist) are mostly enacted and enforced locally, while the Internet itself, that carries the unwanted traffic, respects no borders. One estimate suggests that most players in the underground economy are outside the US, yet most IRC servers supporting the underground market may be running in US network providers, enjoying the reliable service and wide connectivity to the rest of the world provided by the networks.

In addition, the definition of "unwanted" traffic also varies between different countries. For example, China bans certain types of network traffic that are considered legitimate elsewhere. Yet another major difficulty is the trade-off and blurred line between having audit trails to facilitate forensic analysis and to enforce censorship. The greater ability we build into the network to control traffic, the stronger would be the monitoring requirements coming from the legislators.

It should be emphasized that, while a legal system is necessary to create effective deterrence and sanctions against miscreants, it is by no means sufficient on its own. Rather, it must be accompanied by

technical solutions to unwanted traffic detection and damage recovery. It is also by no means a substitute for user education. Only a well informed user community can collectively establish an effective defense in the cyberspace.

2.5. Consequences

What we have today is not a rosy picture: there are

- o big economic incentives and a rich environment to exploit,
- o no specific party to carry responsibility,
- o no auditing system to trace back to the sources of attacks, and
- o no well established legal regulations to punish offenders.

The combination of these factors inevitably leads to ever increasing types and volume of unwanted traffic. However, our real threats are not the bots or DDoS attacks, but the criminals behind them. Unwanted traffic is no longer only aiming for maximal disruption; in many cases, it is now a means to illicit ends with the specific purpose of generating financial gains for the miscreants. Their crimes cause huge economic losses, counted in multiple billions of dollars and continuing.

3. How Bad Is The Problem?

There are quite a number of different kinds of unwanted traffic on the Internet today; the discussions at this workshop were mainly around DDoS traffic and spam. The impact of DDoS and spam on different parts of the network differs. Below, we summarize the impact on backbone providers, access providers, and enterprise customers, respectively.

3.1. Backbone Providers

Since backbone providers' main line of business is packet forwarding, the impact of unwanted traffic is mainly measured by whether DDoS traffic affects network availability. Spam or malware is not a major concern because backbone networks do not directly support end users. Router compromises may exist, but they are rare events at this time.

3.1.1. DDoS Traffic

Observation shows that, in the majority of DDoS attacks, attack traffic can originate from almost anywhere in the Internet. In particular, those regions with high speed user connectivity but

poorly managed end hosts are often the originating sites of DDoS attacks. The miscreants tend to find targets that offer maximal returns with minimal efforts.

Backbone networks in general are well-provisioned in regard to traffic capacities. Therefore, core routers and backbone link capacity do not get affected much by most DDoS attacks; a 5Gbps attack could be easily absorbed without causing noticeable impact on the performance of backbone networks. However, DDoS attacks often saturate access networks and make a significant impact on customers. In particular, multihomed customers who have multiple well-provisioned connections for high throughput and performance may suffer from aggregated DDoS traffic coming in from all directions.

3.1.2. Problem Mitigation

Currently, backbone networks do not have effective diagnosis or mitigation tools against DDoS attacks. The foremost problem is a lack of incentives to deploy security solutions. Because IP transit services are a commodity, controlling cost is essential to surviving the competition. Thus, any expenditure tends to require a clearly identified return-on-investment (ROI). Even when new security solutions become available, providers do not necessarily upgrade their infrastructure to deploy the solutions, as security solutions are often prevention mechanisms that may not have an easily quantifiable ROI. To survive in the competitive environment in which they find themselves, backbone providers also try to recruit more customers. Thus, a provider's reputation is important. Due to the large number of attacks and inadequate security solution deployment, effective attacks and security glitches can be expected. However, it is not in a provider's best interest to report all the observed attacks. Instead, the provider's first concern is to minimize the number of publicized security incidents. For example, a "trouble ticket" acknowledging the problem is issued only after a customer complains. An informal estimate suggested that only about 10% of DDoS attacks are actually reported (some other estimates have put the figure as low as 2%). In short, there is a lack of incentives to either report problems or deploy solutions.

Partly as a consequence of the lack of incentive and lack of funding, there exist few DDoS mitigation tools for backbone providers. Network operators often work on their own time to fight the battle against malicious attacks. Their primary mitigation tools today are Access Control Lists (ACL) and BGP (Border Gateway Protocol) null routes to black-hole unwanted traffic. These tools can be turned on locally and do not require coordination across administrative domains. When done at, or near, DDoS victims, these simple tools can have an immediate effect in reducing the DDoS traffic volume.

However, these tools are rather rudimentary and inadequate, as we will elaborate in Section 4.2.1.

3.2. Access Providers

A common issue that access providers share with backbone providers is the lack of incentive and the shortage of funding needed to deploy security solutions. As with the situation with security incidents on the backbone, the number of security incidents reported by access providers is estimated to be significantly lower than the number of the actual incidents that occurred.

Because access providers are directly connected to end customers, they also face unique problems of their own. From the access providers' viewpoint, the most severe impact of unwanted traffic is not the bandwidth exhaustion, but the customer support load it engenders. The primary impact of unwanted traffic is on end users, and access providers must respond to incident reports from their customers. Today, access providers are playing the role of IT help desk for many of their customers, especially residential users. According to some access providers, during the Microsoft Blaster worm attack, the average time taken to handle a customer call was over an hour. Due to the high cost of staffing the help desks, it is believed that, if a customer calls the help desk just once, the provider would lose the profit they would otherwise have otherwise made over the lifetime of that customer account.

To reduce the high customer service cost caused by security breaches, most access providers offer free security software to their customers. It is much cheaper to give the customer "free" security software in the hope of preventing system compromises than handling the system break-ins after the event. However, perhaps due to their lack of understanding of the possible security problems they may face, many customers fail to install security software despite the free offer from their access providers, or even when they do, they may lack the skill needed to configure a complex system correctly.

What factors may influence how quickly customers get the security breaches fixed? Past experience suggests the following observations:

- o Notification has little impact on end user repair behavior.
- o There is no significant difference in terms of repair behavior between different industries or between business and home users.
- o Users' patching behavior follows an exponential decay pattern with a time constant of approximately 40% per month. Thus, about 40% of computers tend to be patched very quickly when a patch is

released, and approximately 40% of the remaining vulnerable computers in each following month will show signs of being patched. This leaves a few percent still unpatched after 6 months. In the very large population of Internet hosts, this results in a significant number of hosts that will be vulnerable for the rest of their life.

- o There is a general lack of user understanding: after being compromised, unmanaged computers may get replaced rather than repaired, and this often results in infections occurring during the installation process on the replacement.

3.3. Enterprise Networks: Perspective from a Large Enterprise

The operators of one big enterprise network reported their experience regarding unwanted traffic to the workshop. Enterprises perceive many forms of bad traffic including worms, malware, spam, spyware, Instant Messaging (IM), peer-to-peer (P2P) traffic, and DoS. Compared to backbone and access providers, enterprise network operators are more willing to investigate security breaches, although they may hesitate to pay a high price for security solutions. False positives are very costly. Most operators prefer false negatives to false positives. In general, enterprises prefer prevention solutions to detection solutions.

Deliberately created unwanted traffic (as opposed to unwanted traffic that might arise from misconfiguration) in enterprise networks can be sorted into three categories. The first is "Nuisance", which includes unwanted traffic such as spam and peer-to-peer file sharing. Although there were different opinions among the workshop participants as to whether P2P traffic should, or should not, be considered as unwanted traffic, enterprise network operators are concerned not only that P2P traffic represents a significant share of the total network load, but they are also sensitive to potential copyright infringement issues that might lead to significant financial and legal impacts on the company as a whole. In addition, P2P file sharing applications have also become a popular channel for malware propagation.

The second category of unwanted traffic is labeled "Malicious", which includes the traffic that spreads malware. This class of traffic can be small in volume but the cost from the resulting damage can be high. The clean up after an incident also requires highly skilled operators.

The third category of unwanted traffic is "Unknown": it is known that there exists a class of traffic in the network that can be best described in this way, as no one knows its purpose or the locations

of the sources. Malicious traffic can be obscured by encryption, encapsulation, or covered up as legitimate traffic. The existing detection tools are ineffective for this type of traffic. Noisy worms are easy to identify, but stealth worms can open a backdoor on hosts and stay dormant for a long time without causing any noticeable detrimental effect. This type of bad traffic has the potential to make the greatest impact on an enterprise from a threat perspective.

There are more mitigation tools available for enterprise networks than for backbone and access network providers; one explanation might be the greater affordability of solutions for enterprise networks. The costs of damage from a security breach can also have a very significant impact on the profits of an enterprise. At the same time, however, the workshop participants also expressed concerns regarding the ongoing arms race between security exploits and patching solutions. Up to now, security efforts have, by and large, been reactive, creating a chain of security exploits and a consequent stream of "fixes". Such a reactive mode has not only created a big security market, but also does not enable us to get ahead of attackers.

3.4. Domain Name Services

Different from backbone and access providers, there also exists a class of Internet service infrastructure providers. Provision of Domain Name System (DNS) services offers an example here. As reported by operators from a major DNS hosting company, over time there have been increasingly significant DDoS attacks on .com, .net and root servers.

DNS service operators have witnessed large scale DDoS attacks. The most recent ones include reflection attacks resulting from queries using spoofed source addresses. The major damage caused by these attacks are bandwidth and resource exhaustion, which led to disruption of critical services. The peak rate of daily DNS transactions has been growing at a much faster rate than the number of Internet users, and this trend is expected to continue. The heavy load on the DNS servers has led to increasing complexity in providing the services.

In addition to intentional DDoS Attacks, some other causes of the heavy DNS load included (1) well known bugs in a small number of DNS servers that still run an old version of the BIND software, causing significant load increase at top level servers; and (2) inappropriately configured firewalls that allow DNS queries to come out but block returning DNS replies, resulting in big adverse impacts on the overall system. Most of such issues have been addressed in the DNS operational guidelines drafted by the IETF DNS Operations

Working Group; however, many DNS operators have not taken appropriate actions.

At this time, the only effective and viable mitigation approach is over-engineering the DNS service infrastructure by increasing link bandwidth, the number of servers, and the server processing power, as well as deploying network anycast. There is a concern about whether the safety margin gained from over-engineering is, or is not, adequate in sustaining DNS services over future attacks. Looking forward, there are also a few new issues looming. Two imminent ones are the expected widespread deployment of IPv6 whose new DNS software would inevitably contain new bugs, and the DNS Security Extensions (DNSSEC), which could potentially be abused to generate DDoS attacks.

4. Current Vulnerabilities and Existing Solutions

This section summarizes three aspects of the workshop discussions. We first collected the major vulnerabilities mentioned in the workshop, then made a summary of the existing solutions, and followed up with an examination of the effectiveness, or lack of it, of the existing solutions.

4.1. Internet Vulnerabilities

Below is a list of known Internet vulnerabilities and issues around unwanted traffic.

- o Packet source address spoofing: there has been speculation that attacks using spoofed source addresses are decreasing, due to the proliferation of botnets, which can be used to launch various attacks without using spoofed source addresses. It is certainly true that not all the attacks use spoofed addresses; however, many attacks, especially reflection attacks, do use spoofed source addresses.
- o BGP route hijacking: in a survey conducted by Arbor Networks, route hijacking together with source address spoofing are listed as the two most critical vulnerabilities on the Internet. It has been observed that miscreants hijack bogon prefixes for spam message injections. Such hijacks do not affect normal packet delivery and thus have a low chance of being noticed.
- o Everything over HTTP: port scan attacks occur frequently in today's Internet, looking for open TCP or UDP ports through which to gain access to computers. The reaction from computer system management has been to close down all the unused ports, especially in firewalls. One result of this reaction is that application designers have moved to transporting all data communications over

HTTP to avoid firewall traversal issues. Transporting "everything over HTTP" does not block attacks but has simply moved the vulnerability from one place to another.

- o Everyone comes from Everywhere: in the earlier life of the Internet it had been possible to get some indication of the authenticity of traffic from a specific sender based for example on the Time To Live (TTL). The TTL would stay almost constant when traffic from a certain sender to a specific host entered an operators network, since the sender will "always" set the TTL to the same value. If a change in the TTL value occurred without an accompanying change in the routing, one could draw the conclusion that this was potential unwanted traffic. However, since hosts have become mobile, they may be roaming within an operator's network and the resulting path changes may put more (or less) hops between the source and the destination. Thus, it is no longer possible to interpret a change in the TTL value, even if it occurs without any corresponding change in routing, as an indication that the traffic has been subverted.
- o Complex Network Authentication: Network authentication as it is used today is far too complex to be feasible for users to use effectively. It will also be difficult to make it work with new wireless access technologies.

A possible scenario envisages a customers handset that is initially on a corporate wireless network. If that customer steps out of the corporate building, the handset may get connected to the corporate network through a GPRS network. The handset may then roam to a wireless LAN network when the user enters a public area with a hotspot. Consequently, we need authentication tools for cases when the underlying data link layer technology changes quickly, possibly during a single application session.

- o Unused Security Tools: Vendors and standards have produced quite a number of useful security tools; however, not all, or even most, of them get used extensively.

4.2. Existing Solutions

4.2.1. Existing Solutions for Backbone Providers

Several engineering solutions exist that operators can deploy to defend the network against unwanted traffic. Adequate provisioning is one commonly used approach that can diminish the impact of DDoS on the Internet backbone. The solution that received most mentions at the workshop was BCP 38 on ingress filtering: universal deployment of

BCP 38 can effectively block DDoS attacks using spoofed source IP addresses. At present, Access Control List (ACL) and BGP null routing are the two tools most commonly used by network operators to mitigate DDoS attacks. They are effective in blocking DDoS attacks, especially when being applied at or near a victim's site.

Unfortunately, BCP 38 is not widely deployed today. BCP 38 may require device upgrades, and is considered tedious to configure and maintain. Although widespread deployment of BCP 38 could benefit the Internet as a whole, deployment by individual sites imposes a certain amount of cost to the site, and does not provide a direct and tangible benefit in return. In other words, BCP 38 suffers from a lack of deployment incentives.

Both BGP null routing and ACL have the drawback of relying on manual configuration and thus are labor intensive. In addition, they also suffer from blocking both attack and legitimate packets. There is also a potential that some tools could back-fire, e.g., an overly long ACL list might significantly slow down packet forwarding in a router.

Unicast Reverse Path Filtering (uRPF), which is available on some routers, provides a means of implementing a restricted form of BCP 38 ingress filtering without the effort of maintaining ACLs. uRPF uses the routing table to check that a valid path back to the source exists. However, its effectiveness depends on the specificity of the routes against which source addresses are compared. The prevalence of asymmetric routing means that the strict uRPF test (where the route to the source must leave from the same interface on which the packet being tested arrived) may have to be replaced by the loose uRPF test (where the route may leave from any interface). The loose uRPF test is not a guarantee against all cases of address spoofing, and it may still be necessary to maintain an ACL to deal with exceptions.

4.2.2. Existing Solutions for Enterprise Networks

A wide variety of commercial products is available for enterprise network protection. Three popular types of protection mechanisms are

- o Firewalls: firewalls are perhaps the most widely deployed protection products. However, the effectiveness of firewalls in protecting enterprise confidential information can be weakened by spyware installed internally, and they are ineffective against attacks carried out from inside the perimeter established by the firewalls. Too often, spyware installation is a byproduct of installing other applications permitted by end users.

- o Application level gateways: these are becoming more widely used. However, because they require application-specific support, and in many cases they cache all the in-flight documents, configuration can be difficult and the costs high. Thus, enterprise network operators prefer network level protections over layer-7 solutions.
- o Anti-spam software: Anti-spam measures consume significant human resources. Current spam mitigation tools include blacklists and content filters. The more recent "learning" filters may help significantly reduce the human effort needed and decrease the number of both false positives and negatives.

A more recent development is computer admission control, where a computer is granted network access if and only if it belongs to a valid user and appears to have the most recent set of security patches installed. It is however a more expensive solution. A major remaining issue facing enterprise network operators is how to solve the user vulnerability problem and reduce reliance on user's understanding of the need for security maintenance.

4.3. Shortfalls in the Existing Network Protection

4.3.1. Inadequate Tools

Generally speaking, network and service operators do not have adequate tools for network problem diagnosis. The current approaches largely rely on the experience and skills of the operators, and on time-consuming manual operations. The same is true for mitigation tools against attacks.

4.3.2. Inadequate Deployments

The limited number of existing Internet protection measures have not been widely deployed. Deployment of security solutions requires resources which may not be available. It also requires education among the operational community to recognize the critical importance of patch installation and software upgrades; for example, a bug in the BIND packet was discovered and fixed in 2003, yet a number of DNS servers still run the old software today. Perhaps most importantly, a security solution must be designed with the right incentives to promote their deployment. Effective protection also requires coordination between competing network providers. For the time being, it is often difficult to even find the contact information for operators of other networks.

A number of workshop participants shared the view that, if all the known engineering approaches and bug fixes were universally deployed, the Internet could have been enjoying a substantially reduced number

of security problems today. In particular, the need for, and lack of, BCP 38 deployment was mentioned numerous times during the workshop. There is also a lack of enthusiasm about the routing security requirements document being developed by the IETF RPSEC (Routing Protocol Security) Working Group, which focuses heavily on cryptographically-based protection requirements. Not only would cryptographically-based solutions face the obstacle of funding for deployment, but also they are likely to bring with them their own set of problems.

4.3.3. Inadequate Education

There exists an educational challenge to disseminate the knowledge needed for secure Internet usage and operations. Easily guessed passwords and plaintext password transmission are still common in many parts of the Internet. One common rumor claims that Cisco routers were shipped with a default password "cisco" and this was used by attackers to break into routers. In reality, operators often configure Cisco routers with that password, perhaps because of the difficulty of disseminating passwords to multiple maintainers. A similar problem exists for Juniper routers and other vendors' products.

How to provide effective education to the Internet user community at large remains a great challenge. As mentioned earlier in this report, the existence of a large number of compromised hosts is one major source of the unwanted traffic problem, and the ultimate solution to this problem is a well-informed, vigilant user community.

4.3.4. Is Closing Down Open Internet Access Necessary?

One position made at the workshop is that, facing the problems of millions of vulnerable computers and lack of effective deterrence, protecting the Internet might require a fundamental change to the current Internet architecture, by replacing unconstrained open access to the Internet with strictly controlled access. Although the participants held different positions on this issue, a rough consensus was reached that, considering the overall picture, enforcing controlled access does not seem the best solution to Internet protection. Instead, the workshop identified a number of needs that should be satisfied to move towards a well protected Internet:

- o the need for risk assessment for service providers; at this time, we lack a commonly agreed bar for security assurance;
- o the need to add traceability to allow tracking of abnormal behavior in the network, and

- o the need for liability if someone fails to follow recommended practices.

Adding traceability has been difficult due to the distributed nature of the Internet. Collaboration among operators is a necessity in fighting cybercrimes. We must also pay attention to preparation for the next cycle of miscreant activity, and not devote all our efforts to fixing the existing problems. As discussed above, the current reactive approach to security problems is not a winning strategy.

5. Active and Potential Solutions in the Pipeline

This section addresses the issues that vendors recognized as important and for which there will be solutions available in the near future.

There are a number of potential solutions that vendors are working on, but are not yet offering as part of their product portfolio, that will allegedly remedy or diagnose the problems described in Section 4.1.

Inevitably, when vendors have or are about to make a decision on implementing new features in their products but have not made any announcement, the vendors are not willing to talk about the new features openly, which limits what can be said in this section.

5.1. Central Policy Repository

One idea is to build a Central Policy Repository that holds policies that are known to work properly, e.g., policies controlling from whom one would accept traffic when under attack. This repository could, for example, keep information on which neighbor router or AS is doing proper ingress address filtering. The repository could also hold the configurations that operators use to upgrade configurations on their routers.

If such a repository is to be a shared resource used by multiple operators, it will necessarily require validation and authentication of the stored policies to ensure that the repository does not become the cause of vulnerabilities. Inevitably, this would mean that the information comes with a cost and it will only be viable if the sum of the reductions in individual operators' costs is greater than the costs of maintaining the repository.

5.2. Flow Based Tools

A set of tools based on flow data is widely used to extract information from both network and data link layers. Tools have been built that can be used to find out the sources of almost any type of traffic, including certain unwanted traffic. These flow-based tools make it possible to do things like DDoS traceback, traffic/peering analyses, and detection of botnets, worms, and spyware.

These tools monitor flows on the network and build baselines for what is the "normal" behavior. Once the baseline is available, it is possible to detect anomalous activity. It is easy to detect variations over time, and decide if the variation is legitimate or not. It is possible to take this approach further, typically involving the identification of signatures of particular types of traffic.

These flow-based tools are analogous to the "sonar" that is used by navies to listen for submarines. Once a particular submarine is identified, it is possible to record its sonar signature to be used to provide rapid identification in the future when the same submarine is encountered again.

Examples of existing tools include Cisco IOS NetFlow <http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html>, sFlow <<http://www.sflow.org/>>, and NetraMet <<http://www.caida.org/tools/measurement/netramet/>> based on the IETF RTFM and IPFIX standards.

There are also tools for working with the output of NetFlow such as jFlow <<http://www.net-track.ch/opensource/jflow/>> and Arbor Networks' Peakflow <http://www.arbor.net/products_platform.php>.

The Cooperative Association for Internet Data Analysis (CAIDA) maintains a taxonomy of available tools on its web site at <<http://www.caida.org/tools/taxonomy/index.xml>>.

5.3. Internet Motion Sensor (IMS)

The Internet Motion Sensor (IMS) [IMS] may be used to watch traffic to or from "Darknets" (routable prefixes that don't have end hosts attached), unassigned address spaces, and unannounced address spaces. By watching activities in these types of address spaces, one can understand and detect, e.g., scanning activities, DDoS worms, worm infected hosts, and misconfigured hosts.

Currently, the IMS is used to monitor approximately 17 million prefixes, about 1.2% of the IPv4 address space. The use of IMS has highlighted two major characteristics of attacks; malicious attacks are more targeted than one might have assumed, and a vulnerability in a system does not necessarily lead to a threat to that system (e.g., the vulnerability may not be exploited to launch attacks if the perceived "benefit" to the attacker appears small). Data from IMS and other sources indicates that attackers are making increased use of information from social networking sites to target their attacks and select perceived easy targets, such as computers running very old versions of systems or new, unpatched vulnerabilities.

This form of passive data collection is also known as a "Network Telescope". Links to similar tools can be found on the CAIDA web site at http://www.caida.org/data/passive/network_telescope.xml.

5.4. BCP 38

In the year 2000, the IETF developed a set of recommendations to limit DOS attacks and Address Spoofing published as BCP 38 [RFC2827], "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing". However, up to now BCP 38 capabilities still have not been widely deployed, perhaps due to the incentive issue discussed earlier.

The IETF has also developed an additional set of recommendations extending BCP 38 to multihomed networks. These recommendations are published as BCP 84 [RFC3704].

5.5. Layer 5 to 7 Awareness

Tools are being developed that will make it possible to perform deep packet inspection at high speed. Some companies are working on hardware implementation to inspect all layers from 2 to 7 (e.g., EZchip http://www.ezchip.com/t_npu_whpaper.htm). A number of other companies, including Cisco and Juniper, offer tools capable of analyzing packets at the transport layer and above.

5.6. How To's

One idea that was discussed at the workshop envisaged operators and standards bodies cooperating to produce a set of "How To" documents as guidelines on how to configure networks. Dissemination and use of these "How To's" should be encouraged by vendors, operators, and standards bodies.

This type of initiative needs a "sponsor" or "champion" that takes the lead and starts collecting a set of "How To's" that could be freely distributed. The workshop did not discuss this further.

5.7. SHRED

Methods to discourage the dissemination of spam by punishing the spammers, such as Spam Harassment Reduction via Economic Disincentive (SHRED) [SHRED], were discussed. The idea is to make it increasingly expensive for spammers to use the email system, while normal users retain what they have come to expect as normal service. There was no agreement on the effectiveness of this type of system.

6. Research in Progress

In preparation for this session, several researchers active in Internet Research were asked two rather open ended questions: "Where is the focus on Internet research today?" and "Where should it be?"

A summary of the answers to these questions is given below. Section 6.2.2 covers part of the relationship between research and miscreants. For example, research activities in each area (please refer to the slide set for Workshop Session 8 which can be found at the link referred to in Appendix C).

6.1. Ongoing Research

Section 6.1 discusses briefly areas where we see active research on unwanted traffic today.

6.1.1. Exploited Hosts

One area where researchers are very active is analyzing situations where hosts are exploited. This has been a major focus for a long time, and an abundance of reports have been published. Current research may be divided into three different categories: prevention, detection, and defense.

6.1.1.1. Prevention

Code quality is crucial when it comes to preventing exploitation of Internet hosts. Quite a bit of research effort has therefore gone into improvement of code quality. Researchers are looking into automated methods for finding bugs and maybe in the end fixes for any bugs detected.

A second approach designed to stop hosts from becoming compromised is to reduce the "attack surface". Researchers are thinking about

changes or extensions to the Internet architecture. The idea is to create a strict client server architecture, where the clients only are allowed to initiate connections, and while servers may only accept connections.

Researchers have put a lot of effort into better scaling of honey pots and honey farms to better understand and neutralize the methods miscreants are using to exploit hosts. Research also goes into developing honey monkeys in order to understand how hosts are vulnerable. Both honey pots/farms and honey monkeys are aimed at taking measures that prevent further (mis-)use of possible exploits.

6.1.1.2. Detection

When an attack is launched against a computer system, the attack typically leaves evidence of the intrusion in the system logs. Each type of intrusion leaves a specific kind of footprint or signature. The signature can be evidence that certain software has been executed, that logins have failed, that administrative privileges have been misused, or that particular files and directories have been accessed. Administrators can document these attack signatures and use them to detect the same type of attack in the future. This process can be automated.

Because each signature is different, it is possible for system administrators to determine by looking at the intrusion signature what the intrusion was, how and when it was perpetrated, and even how skilled the intruder is.

Once an attack signature is available, it can be used to create a vulnerability filter, i.e., the stored attack signature is compared to actual events in real time and an alarm is given when this pattern is repeated.

A further step may be taken with automated vulnerability signatures, i.e., when a new type of attack is found, a vulnerability filter is automatically created. This vulnerability filter can be made available for nodes to defend themselves against this new type of attack. The automated vulnerability signatures may be part of an Intrusion Detection System (IDS).

6.1.1.3. Defense

An IDS can be a part of the defense against actual attacks, e.g., by using vulnerability filters. An Intrusion Detection System (IDS) inspects inbound and outbound network activities and detects signatures that indicate that a system is under attack from someone attempting to break into or compromise the system.

6.1.2. Distributed Denial of Service (DDoS) Attacks

Research on DDoS attacks follows two separate approaches, the first has the application as its focus, while the second focuses on the network.

6.1.2.1. Application Oriented DDoS Research

The key issue with application oriented research is to distinguish between legitimate activities and attacks. Today, several tools exist that can do this and research has moved on to more advanced things.

Research today looks into tools that can detect and filter activities that have been generated by bots and botnets.

One approach is to set up a tool that sends challenges to senders that want to send traffic to a certain node. The potential sender then has to respond correctly to that challenge; otherwise, the traffic will be filtered out.

The alternative is to get more capacity between sender and receiver. This is done primarily by some form of use of peer-to-peer technology.

Today, there is "peer-to-peer hype" in the research community; a sure way of making yourself known as a researcher is to publish something that solves old problems by means of some peer-to-peer technology. Proposals now exist for peer-to-peer DNS, peer-to-peer backup solutions, peer-to-peer web-cast, etc. Whether these proposals can live up to the hype remains to be seen.

6.1.2.2. Network Oriented DDoS Research

Research on DDoS attacks that takes a network oriented focus may be described by the following oversimplified three steps.

1. Find the bad stuff
2. Set the "evil bit" on those packets
3. Filter out the packets with the "evil bit" set

This rather uncomplicated scheme has to be carried out on high-speed links and interfaces. Automation is the only way of achieving this.

One way of indirectly setting the "evil bit" is to use a normalized TTL. The logic goes: the TTL for traffic from this sender has always

been "x", but has now suddenly become "y", without any corresponding change in routing. The conclusion is that someone is masquerading as the legitimate sender. Traffic with the "y" TTL is filtered out.

Another idea is to give traffic received from ISPs that are known to do source address validation the "red carpet treatment", i.e., to set the "good bit". When an attack is detected, traffic from everyone that doesn't have the "good bit" is filtered out. Apart from reacting to the attack, this also give ISPs an incentive to do source address validation. If they don't do it, their peers won't set the "good bit" and the ISP's customers will suffer, dragging down their reputation.

Overlay networks can also be used to stop a DDoS attack. The idea here is that traffic is not routed directly to the destination. Instead, it is hidden behind some entry points in the overlay. The entry points make sure the sender is the host he claims he is, and in that case, marks the packet with a "magic bit". Packets lacking the "magic bit" are not forwarded on the overlay. This has good scaling properties; you only need to have enough capacity to tag the amount of traffic you want to receive, not the amount you actually receive.

6.1.3. Spyware

Current research on spyware and measurements of spyware are aiming to find methods to understand when certain activities associated with spyware happen and to understand the impact of this activity.

There are a number of research activities around spyware, e.g., looking into threats caused by spyware; however, these were only briefly touched upon at the workshop.

6.1.4. Forensic Aids

Lately, research has started to look into tools and support to answer the "What happened here?" question. These tools are called "forensic aids", and can be used to "recreate" an illegal activity just as the police do when working on a crime scene.

The techniques that these forensic aids take as their starting point involve the identification of a process or program that should not be present on a computer. The effort goes into building tools and methods that can trace the intruder back to its origin. Methods to understand how a specific output depends on a particular input also exist.

6.1.5. Measurements

Measurements are always interesting for the research community, because they generate new data. Consequently, lots of effort goes into specifying how measurements should be performed and into development of measurement tools. Measurements have been useful in creating effective counter-measures against worms. Before measurements gave actual data of how worms behave, actions taken against worms were generally ineffective.

6.1.6. Traffic Analysis

One aspect of research that closely relates to measurements is analysis. Earlier, it was common to look for the amount of traffic traversing certain transport ports. Lately, it has become common to tunnel "everything" over something else, and a shift has occurred towards looking for behavior and/or content. When you see a certain behavior or content over a protocol that is not supposed to behave in this way, it is likely that something bad is going on.

Since this is an arms race, the miscreants that use tunneling protocols have started to mimic the pattern of something that is acceptable.

6.1.7. Protocol and Software Security

The general IETF design guidelines for robust Internet protocols says: "Be liberal in what you receive and conservative in what you send". The downside is that most protocols believe what they get and as a consequence also get what they deserve. The IAB is intending to work on new design guidelines, e.g., rules of thumb and things you do and things you don't. This is not ready yet, but will be offered as input to a BCP in due course.

An area where there is a potential overlap between standards people and researchers is protocol analysis languages. The protocol analysis languages could be used, for example, look for vulnerabilities.

6.2. Research on the Internet

The workshop discussed the interface between people working in standardization organizations in general and IETF in particular on the one hand and people working with research on the other. The topic of discussion was broader than just "Unwanted traffic". Three topics were touched on: what motivates researchers, how to attract researchers to problems that are hindering or have been discovered in

the context of standardization, and the sometimes rocky relations between the research community and the "bad boys".

6.2.1. Research and Standards

The workshop discussed how research and standardization could mutually support each other. Quite often there is a commonality of interest between the two groups. The IAB supports the Internet Research Task Force (IRTF) as a venue for Internet research. The delta between what is done and what could be is still substantial. The discussion focused on how standardization in general and the IETF in particular can get help from researchers.

Since standardization organizations don't have the economic strength to simply finance the research they need or want, other means have to be used. One is to correctly and clearly communicate problems, another is to supply adequate and relevant information.

To attract the research community to work with standardization organizations, it is necessary to identify the real problems and state them in such a way that they are amenable to solution. General unspecified problems are of no use, e.g., "This is an impossible problem!" or "All the problems are because my users behave badly!"

Instead, saying "This is an absolutely critical problem, and we have no idea how to solve it!" is much more attractive.

The potential research problem should also be communicated in a way that is public. A researcher that wants to take on a problem is helped if she/he can point at a slide from NANOG or RIPE that identifies this problem.

The way researchers go about solving problems is basically to identify all the existing constraints, and then relax one of the constraints and see what happens. Therefore, rock solid constraints are a show stopper, e.g., "We can't do that, because it has to go into an ASIC!". Real constraints have to be clearly communicated to and understood by the researcher.

One reasonable way of fostering cooperation is to entice two or three people and have them write a paper on the problem. What will happen then is that this paper will be incrementally improved by other researchers. The vast majority of all research goes into improving on someone else's paper.

A second important factor is to supply sufficient relevant information. New information that suggests possible ways to address new problems or improve on old or partial solutions to previously

investigated problems are attractive. Often, understanding of important problems comes from the operator community; when trying to initiate research from a standards perspective, keeping operators in the loop may be beneficial.

Today, the research community is largely left on its own, and consequently tends to generate essentially random, untargeted results. If the right people in the standards community say the right things to the right people in the research community, it can literally focus hundreds of graduate students on a single problem. Problem statements and data are needed.

6.2.2. Research and the Bad Guys

A general problem with all research and development is that what can be used may also be misused. In some cases, miscreants have received help from research that was never intended.

There are several examples of Free Nets, i.e., networks designed to allow end-users to participate without revealing their identity or how and where they are connected to the network. The Free Nets are designed based on technologies such as onion routing or mix networks. Free Nets create anonymity that allows people to express opinions without having to reveal their true identity and thus can be used to promote free speech. However, these are tools that can also work just as well to hide illegal activities in democracies.

Mix networks create hard-to-trace communications by using a chain of proxy servers. A message from a sender to a receiver passes by the chain of proxies. A message is encrypted with a layered encryption where each layer is understood by only one of the proxies in the chain; the actual message is the innermost layer. A mix network will achieve untraceable communication, even if all but one of the proxies are compromised by a potential tracer.

Onion routing is a technique for anonymous communication over a computer network; it is a technique that encodes routing information in a set of encrypted layers. Onion routing is a further development of mix networks.

Research projects have resulted in methods for distributed command and control, e.g., in the form of Distributed Hash Tables (DHT) and gossip protocols. This of course has legitimate uses, e.g., for security and reliability applications, but it also is extremely useful for DDoS attacks and unwanted traffic in general.

A lot of effort has gone into research around worms, the result is that we have a very good understanding of the characteristics of the

technology associated with worms and how they behave. This is a very good basis when we want to protect against worms. The downside is that researchers also understand how to implement future worms, including knowledge on how to design faster worms that won't leave a footprint.

7. Aladdin's Lamp

If we had an Aladdin's Lamp and could be granted anything we wanted in the context of remedying unwanted traffic or effects of such traffic - what would we wish for? The topic of this session was wishes, i.e., loosening the constraints that depend on what we have and focus on what we really want.

There certainly are lots of "wishes" around, not least of which is making things simpler and safer. On the other hand, very few of these wishes are clearly stated. One comment on this lack of clarity was that we are too busy putting out the fires of today and don't have the time to be thinking ahead.

7.1. Security Improvements

Operators at the workshop expressed a number of wishes that, if fulfilled, would help to improve and simplify security. The list below contains a number of examples of actions that ought to improve security. The content is still at the "wish-level", i.e., no effort has gone in to trying to understand the feasibility of realizing these wishes.

Wish: Reliable point of contact in each administrative domain for security coordination.

First and foremost, operators would like to see correct and complete contact information to coordinate security problems across operators.

The "whois" database of registration details for IP addresses and Autonomous System numbers held by Regional Internet Registries (e.g., ARIN, RIPE, APNIC) was intended to be a directory for this type of information, and RFC 2142 [RFC2142] established common mailbox names for certain roles and services. There are several reasons why these tools are largely unused, including unwanted traffic.

Wish: Organized testing for security.

Today, new hardware and software are extensively tested for performance. There is almost no testing of this hardware and software for security.

Wish: Infrastructure or test bed for security.

It would be good to have an organized infrastructure or test bed for testing of security for new products.

Wish: Defaults for security.

Equipment and software should come with a simple and effective default setting for security.

Wish: Shared information regarding attacks.

It would be useful to have an automated sharing mechanism for attacks, vulnerabilities, and sources of threats between network users and providers in order to meet attacks in a more timely and efficient manner.

7.2. Unwanted Traffic

Wish: Automatic filtering of unwanted traffic.

It would be useful, not least for enterprises, to have mechanisms that would automatically filter out the unwanted traffic.

Some filtering of spam, viruses, and malware that is sent by email is already practicable but inevitably is imperfect because it mainly relies on "heuristics" to identify the unwanted traffic. This is another example of the "arms race" between filtering and the ingenuity of spammers trying to evade the filters. This "wish" needs to be further discussed and developed to make it something that could be turned into practical ideas.

Wish: Fix Spam.

A large fraction of the email traffic coming into enterprises today is spam, and consequently any fixes to the spam problem are very high on their priority list.

8. Workshop Summary

The workshop spent its last two hours discussing the following question: What are the engineering (immediate and longer term) and research issues that might be pursued within the IETF and the IRTF, and what actions could the IAB take? The suggested actions can be summarized into three classes.

8.1. Hard Questions

The discussions during this concluding section raised a number of questions that touched upon the overall network architecture designs.

- o What should be the roles of cryptographic mechanisms in the overall Internet architecture? For example, do we need to apply

cryptographic mechanisms to harden the shell, or rely on deep packet inspection to filter out bad traffic?

- o To add effective protection to the Internet, how far are we willing to go in
 - * curtailing its openness, and
 - * increasing the system complexity?

And what architectural principles do we need to preserve as we go along these paths?

- o A simple risk analysis would suggest that an ideal attack target of minimal cost but maximal disruption is the core routing infrastructure. However, do we really need an unlinked and separately managed control plane to secure it? This requires a deep understanding of the architectural design trade-offs.
- o Can we, and how do we, change the economic substructure? A special workshop was suggested as a next step to gain a better understanding of the question.

8.2. Medium or Long Term Steps

While answering the above hard questions may take some time and effort, several specific steps were suggested as medium or long term efforts to add protection to the Internet:

- o Tightening the security of the core routing infrastructure.
- o Cleaning up the Internet Routing Registry repository [IRR], and securing both the database and the access, so that it can be used for routing verifications.
- o Take down botnets.
- o Although we do not have a magic wand to wave all the unwanted traffic off the Internet, we should be able to develop effective measures to reduce the unwanted traffic to a tiny fraction of its current volume and keep it under control.
- o Community education, to try to ensure people *use* updated host, router, and ingress filtering BCPs.

8.3. Immediately Actionable Steps

The IETF is recommended to take steps to carry out the following actions towards enhancing the network protection.

- o Update the host requirements RFC. The Internet host requirements ([RFC1122], [RFC1123]) were developed in 1989. The Internet has gone through fundamental changes since then, including the pervasive security threats. Thus, a new set of requirements is overdue.
- o Update the router requirements. The original router requirements [RFC1812] were developed in 1995. As with the host requirements, it is also overdue for an update.
- o Update ingress filtering (BCP 38 [RFC2827] and BCP 84 [RFC3704]).

One immediate action that the IAB should carry out is to inform the community about the existence of the underground economy.

The IRTF is recommended to take further steps toward understanding the Underground Economy and to initiate research on developing effective countermeasures.

Overall, the workshop attendees wish to raise the community's awareness of the underground economy. The community as a whole should undertake a systematic examination of the current situation and develop both near- and long-term plans.

9. Terminology

This section gives an overview of some of the key concepts and terminology used in this document. It is not intended to be complete, but is offered as a quick reference for the reader of the report.

ACL

Access Control List in the context of Internet networking refers to a set of IP addresses or routing prefixes (layer 3 or Internet layer information), possibly combined with transport protocol port numbers (layer 4 or transport layer information). The layer 3 and/or layer 4 information in the packets making up a flow entering or leaving a device in the Internet is matched against the entries in an ACL to determine whether the packets should, for example, be allowed or denied access to some resources. The ACL effectively specifies a filter to be used on a flow of packets.

BGP route hijacking

Attack in which an inappropriate route is injected into the global routing system with the intent of diverting traffic from its intended recipient either as a DoS attack (q.v.) where the traffic is just dropped or as part of some wider attack on the recipient. Injecting spurious routes specifying addresses used for bogons can, for example, provide bogus assurance to email systems that spam is coming from legitimate addresses.

Bogon

A bogon is an IP packet that has a source address taken for a range of addresses that has not yet been allocated to legitimate users, or is a private [RFC1918] or reserved address [RFC3330].

Bogon prefix

A bogon prefix is a route that should never appear in the Internet routing table, e.g., from the private or unallocated address blocks.

Bot

A bot is common parlance on the Internet for a software program that is a software agent. A Bot interacts with other network services intended for people as if it were a real person. One typical use of bots is to gather information. The term is derived from the word "robot," reflecting the autonomous character in the "virtual robot"-ness of the concept.

The most common bots are those that covertly install themselves on people's computers for malicious purposes, and that have been described as remote attack tools. Bots are sometimes called "zombies".

Botnet

Botnet is a jargon term for a collection of software robots, or bots, which run autonomously. This can also refer to the network of computers using distributed computing software. While the term "botnet" can be used to refer to any group of bots, such as IRC bots, the word is generally used to refer to a collection of compromised machines running programs, usually referred to as worms, Trojan horses, or backdoors, under a common command and control infrastructure.

Click fraud

Click fraud occurs in pay per click (PPC) advertising when a person, automated script, or computer program imitates a legitimate user of a web browser clicking on an ad for the purpose of generating an improper charge per click. Pay per click advertising is when operators of web sites act as publishers and offer clickable links from advertisers in exchange for a charge per click.

Darknet

A Darknet (also known as a Network Telescope, a Blackhole, or an Internet Sink) is a globally routed network that has no "real" machines attached and carries only a very small amount of specially crafted legitimate traffic. It is therefore easily possible to separate out and analyze unwanted traffic that can arise from a wide variety of events including misconfiguration (e.g., a human being mis-typing an IP address), malicious scanning of address space by hackers looking for vulnerable targets, backscatter from random source denial-of-service attacks, and the automated spread of malicious software called Internet worms.

Dirty affiliate program

Affiliate programs are distributed marketing programs that recruit agents to promote a product or service. Affiliates get financially compensated for each sale associated with their unique 'affiliate ID.' Affiliates are normally instructed by the operator of the affiliate program to not break any laws while promoting the product or service. Sanctions (typically loss of unpaid commissions or removal from the affiliate program) are normally applied if the affiliate spams or otherwise violates the affiliate program's policies.

Dirty affiliate programs allow spamming, or if they do nominally prohibit spamming, they don't actually sanction violators. Dirty affiliate programs often promote illegal or deceptive products (prescription drugs distributed without regard to normal dispensing requirements, body part enlargement products, etc.), employ anonymous or untraceable affiliates, offer payment via anonymous online financial channels, and may fail to follow normal tax withholding and reporting practices.

DoS attack

Denial-Of-Service attack, a type of attack on a network that is designed to bring the network to its knees by flooding it with useless traffic or otherwise blocking resources necessary to allow normal traffic flow.

DDoS attack

Distributed Denial of Service, an attack where multiple compromised systems are used to target a single system causing a Denial of Service (DoS) attack.

Honey farm

A honey farm is a set of honey pots working together.

Honey monkey

A honey monkey is a honey pot in reverse; instead of sitting and waiting for miscreants, a honey monkey actively mimics the actions of a user surfing the Web. The honey monkey runs on virtual machines in order to detect exploit sites.

Honey pot

A honey pot is a server attached to the Internet that acts as a decoy, attracting potential miscreants in order to study their activities and monitor how they are able to break into a system. Honeypots are designed to mimic systems that an intruder would like to break into but limit the intruder from having access to an entire network.

IRC

Internet Relay Chat is a form of instant communication over the Internet. It is mainly designed for group (many-to-many) communication in discussion forums called channels, but also allows one-to-one communication, originally standardized by RFC 1459 [RFC1459] but much improved and extended since its original invention. IRC clients rendezvous and exchange messages through IRC servers. IRC servers are run by many organizations for both benign and nefarious purposes.

Malware

Malware is software designed to infiltrate or damage a computer system, without the owner's informed consent. There are disagreements about the etymology of the term itself, the primary uncertainty being whether it is a portmanteau word (of "malicious" and "software") or simply composed of the prefix "mal-" and the morpheme "ware". Malware references the intent of the creator, rather than any particular features. It includes computer viruses, worms, Trojan horses, spyware, adware, and other malicious and unwanted software. In law, malware is sometimes known as a computer contaminant.

Mix networks

Mix networks create hard-to-trace communications by using a chain of proxy servers [MIX]. Each message is encrypted to each proxy; the resulting encryption is layered like a Russian doll with the message as the innermost layer. Even if all but one of the proxies are compromised by a tracer, untraceability is still achieved. More information can be found at <http://www.adastral.ucl.ac.uk/~helger/crypto/link/protocols/mix.php>.

Onion routing

Onion routing is a technique for anonymous communication over a computer network, it is a technique that encodes routing information in a set of encrypted layers. Onion routing is based on mix cascades (see mix networks (q.v.)). More information can be found at <http://www.onion-router.net/>.

Phishing

Phishing is a form of criminal activity using social engineering techniques. It is characterized by attempts to fraudulently acquire sensitive information, such as passwords and credit card details, by masquerading as a trustworthy person or business in an apparently official electronic communication. Phishing is typically carried out using spoofed websites, email, or an instant message. The term phishing derives from password harvesting and the use of increasingly sophisticated lures to "fish" for users' financial information and passwords.

Root access

Access to a system with full administrative privileges bypassing any security restrictions placed on normal users. Derived from the name traditionally used for the 'superuser' on Unix systems.

Script kiddy

Derogatory term for an inexperienced hacker who mindlessly uses scripts and other programs developed by others with the intent of compromising computers or generating DoS attacks.

Spam

Spamming is the abuse of electronic messaging systems to send unsolicited, undesired bulk messages. The individual messages are referred to as spam. The term is frequently used to refer specifically to the electronic mail form of spam.

Spoofing

(IP) spoofing is a technique where the illegitimate source of IP packets is obfuscated by contriving to use IP address(es) that the receiver recognizes as a legitimate source. Spoofing is often used to gain unauthorized access to computers or mislead filtering mechanisms, whereby the intruder sends packets into the network with an IP source address indicating that the message is coming from a legitimate host. To engage in IP spoofing, a hacker must first use a variety of techniques to find an IP address of a valid host and then modify the packet headers so that it appears that the packets are coming from that host.

Spyware

Any software that covertly gathers user information through the user's Internet connection without his or her knowledge, e.g., for spam purposes.

UBE

Unsolicited Bulk Email: an official term for spam.

UCE

Unsolicited Commercial Email: an official term for spam.

Virus

A program or piece of code that is loaded onto a computer without the owner's knowledge and runs without their consent. A virus is self-replicating code that spreads by inserting copies of itself into other executable code or documents, which are then transferred to other machines. Typically, the virus has a payload that causes some harm to the infected machine when the virus code is executed.

Worm

A computer worm is a self-replicating computer program. It uses a network to send copies of itself to other systems and it may do so without any user intervention. Unlike a virus, it does not need to attach itself to an existing program. Worms always harm the network (if only by consuming bandwidth), whereas viruses always infect or corrupt files on a targeted computer.

Zombie

This is another name for a bot.

10. Security Considerations

This document does not specify any protocol or "bits on the wire".

11. Acknowledgements

The IAB would like to thank the University of Southern California Information Sciences Institute (ISI) who hosted the workshop and all those people at ISI and elsewhere who assisted with the organization and logistics of the workshop at ISI.

The IAB would also like to thank the scribes listed in Appendix A who diligently recorded the proceedings during the workshop.

A special thanks to all the participants in the workshop, who took the time, came to the workshop to participate in the discussions, and who put in the effort to make this workshop a success. The IAB

especially appreciates the effort of those that prepared and made presentations at the workshop.

12. Informative References

- [IMS] University of Michigan, "Internet Motion Sensor", 2006, <<http://ims.eecs.umich.edu/>>.
- [IRR] Merit Network Inc, "Internet Routing Registry Routing Assets Database", 2006, <<http://www.irr.net/>>.
- [MIX] Hill, R., Hwang, A., and D. Molnar, "Approaches to Mix Nets", MIT 6.857 Final Project, December 1999, <<http://www.mit.edu/afs/athena/course/6/6.857/OldStuff/Fall99/papers/mixnet.ps.gz>>.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.
- [RFC1459] Oikarinen, J. and D. Reed, "Internet Relay Chat Protocol", RFC 1459, May 1993.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2142] Crocker, D., "MAILBOX NAMES FOR COMMON SERVICES, ROLES AND FUNCTIONS", RFC 2142, May 1997.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC3330] IANA, "Special-Use IPv4 Addresses", RFC 3330, September 2002.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [SHRED] Krishnamurthy, B. and E. Blackmond, "SHRED: Spam Harassment Reduction via Economic Disincentives", 2003, <<http://www.research.att.com/~bala/papers/shred-ext.ps>>.

Appendix A. Participants in the Workshop

Bernard Aboba (IAB)
Loa Andersson (IAB)
Ganesha Bhaskara (scribe)
Bryan Burns
Leslie Daigle (IAB chair)
Sean Donelan
Rich Draves (IAB Executive Director)
Aaron Falk (IAB, IRTF chair)
Robert Geigle
Minas Gjoka (scribe)
Barry Greene
Sam Hartman (IESG, Security Area Director)
Bob Hinden (IAB)
Russ Housely (IESG, Security Area Director)
Craig Huegen
Cullen Jennings
Rodney Joffe
Mark Koster
Bala Krishnamurthy
Gregory Lebovitz
Ryan McDowell
Danny McPherson
Dave Merrill
David Meyer (IAB)
Alan Mitchell
John Morris
Eric Osterweil (scribe)
Eric Rescorla (IAB)
Pete Resnick (IAB)
Stefan Savage
Joe St Sauver
Michael Sirivianos (scribe)
Rob Thomas
Helen Wang
Lixia Zhang (IAB)

Appendix B. Workshop Agenda

Session 1:

How bad is the problem? What are the most important symptoms?

Session 2:

What are the sources of the problem?

Lunch session (session 3):

Solutions in regulatory and societal space

Session 4:

The underground economy

Session 5:

Current countermeasures, what works, what doesn't

Session 6:

If all our wishes could be granted, what would they be?

Session 7:

What's in the pipeline, or should be in the pipeline

Session 8:

What is being actively researched on?

Session 9:

What are the engineering (immediate and longer term) and research issues that might be pursued within the IETF/IAB/IRTF?

Appendix C. Slides

Links to a subset of the presentations given by the participants at the workshop can be found via the IAB Workshops page on the IAB web site at <http://utgard.ietf.org/iab/about/workshops/unwantedtraffic/index.html>. As mentioned in Section 1, this is not a complete set of the presentations because certain of the presentations were of a sensitive nature which it would be inappropriate to make public at this time.

Authors' Addresses

Loa Andersson
Acreo AB

EMail: loa@pi.se

Elwyn Davies
Folly Consulting

EMail: elwynd@dial.pipex.com

Lixia Zhang
UCLA

EMail: lixia@cs.ucla.edu

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

