A Security Study of the Internet: An Analysis of Firewall Behavior and Anonymous DNS

Hal Burch and Dawn Song

July 20, 2004
CMU-CS-04-141

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Hosts connected to the Internet are exposed to a wide array of attacks. Multiple methods are used to limit and impede attacks. This paper looks at how and if some of these methods are deployed on the Internet. The most common method employed is to limit network access to hosts using firewalls. What percentage of IP addresses are behind firewalls? What do these firewalls block and allow? What common policies are installed in firewalls? These questions are extremely important for understanding how firewalls are used as a security defense mechanism on the Internet and were previously unaddressed. In this paper, we first set off to answer these questions by performing a systematic study of firewall behavior on the Internet. Another well-adopted method to limit information about hosts is to give IP addresses anonymous hostnames based on their IP addresses on the public Internet, called *anonymous DNS*. This makes the function and even existence of such machine difficult to determine. In this paper, we then analyze the behavior of anonymous DNS on the Internet, e.g., what fraction of hosts have anonymous names and how much information is contained in Internet hostnames. To the best of our knowledge, we are the first ones to systematically study the behavior of firewalls and anonymous DNS on the Internet. In this paper, we propose a methodology for such a study and describe our measurement results.

Keywords: network measurement, security, census, firewall, firewall behavior, server correlation, split DNS, anonymous DNS

1 Introduction

From 1980 to 2000, the Internet grew from around 200 hosts to over 100 million [14], as more and more corporations connected their internal networks to the Internet. However, corporations discovered that their internal networks were not secure. While this may pose limited risk if only employees are able to access the network, once a company connects its network to the Internet, it becomes a large problem. Rather than protecting themselves by completely disconnecting themselves from the Internet, corporations chose instead to limit access by deploying secure Internet gateways [4], which came to be known as "firewalls". Firewalls examine the traffic between two or more networks and selectively drop packets based on a set of rules.

Firewalls are an important first-line defense mechanism to limit network access to hosts, and arguably the simplest form of network defense. How prevalent are firewalls? How securely are they configured? These questions provide important information about firewall deployment on the Internet and could unveil potential security concerns. In this paper, we set off to find answers to these questions and many other questions related to the deployment of these security defense mechanisms.

For example, there are two basic ways to configure a firewall: block by default and block by exception. In block by default, the default behavior is to not allow access to IP addresses behind the firewall. The firewall may allow access to particular services on particular hosts. In block by exception, only access to certain services is disabled. This is more common in lower-security networks, such as universities and residential ISPs. In these cases, the network operator is not expected to block services. In fact, even limited blocking may be considered unwanted by the network users. In block by default, the firewall administrator decides which services are useful enough to specifically allow through the firewall. In block by exception, the firewall administrator decides which services pose a large enough security risk to block them at the firewall. Block by default is preferred by security experts and is more common in practice. One interesting question one may ask is what percentage of IP addresses are beyond firewalls that are block by default.

For another example, firewall configurations are also often correlated between different services. A firewall administrator that decides to block one service may likely block another service which is similar. This paper examines these correlations to look at how accesibilities of different services are related.

Another question related to the firewall analysis is to study the behavior of publicly accessible services. How many machines are running public servers for different services? Moreover, do servers often run multiple services? To increase security, it is recommended that administrators split public services across machines. If a machine runs many services, compromising any of them likely compromises the entire machine. By splitting the services between multiple machines, the likelihood that any one machine is compromised is decreased and the impact of gaining access to a machine is decreased, since taking the machine offline affects only the services running on that machine and services that depend on those services. Economic realities, however, are such that operating more machines increases hardware and administrative costs. Thus, many installations run multiple services on machines to reduce costs.

In addition to firewalls, corporate networks are often also divided from the Internet in naming. That is, hostnames valid within a corporate network may not resolve using domain name service (DNS) outside that corporation. This is known as "split DNS". This paper will focus on split reverse DNS, where, for example, 192.0.2.12 may reverse resolve to "inside.example.com" within a corporation, but reverse resolve to "host-192.0.2.12.example.com", or some equally uninformative hostname, on the Internet. This provides a layer of obfuscation, as it makes it more difficult for an outside attacker to determine whether 192.0.2.12 is in use, much less that it is a server. Split DNS is akin to firewalls for DNS: the Internet is allowed only partial access to internal DNS, usually limited to resolving external servers.

Split DNS is commonly implemented as assigning mechanically-generated hostnames to IP addresses. This paper focuses on mechanically-generated hostnames derived from the IP addresses, such as the examples above. Such hostnames will be referred to as "anonymous DNS" because the hostname reveals little about the IP address. Split DNS is not the only reason for anonymous DNS. For example, many ISPs do not give meaningful hostnames to the IP addresses of their residential customers. These IP addresses do not meaningfully resolve anywhere. This paper examines the prevalence of anonymous DNS, and how much information is actually contained in reverse lookups. Most companies using anonymous DNS generate hostnames for all IP addresses they own, regardless of whether or not that IP address is in use.

Anonymous DNS affects attempts that use DNS to estimate the number of hosts on the Internet. As firewalls grew in popularity, attempts to measure the number of hosts on the Internet using ping became inaccurate, so projects like ISC Internet Domain Survey [14] use DNS. Anonymous DNS make such projects inaccurate. This paper quantitates the error resulting from using DNS to count hosts.

In summary, this paper specifically addresses the following seven questions:

1. How prevalent are firewalls? What percentage of IP addresses are behind firewalls?

2. What do firewalls block? What services are commonly allowed and blocked by firewalls?

3. How much do firewalls block? What percentage of IP addresses are beyond firewalls that are block by default?

4. How is firewall behavior correlated? Given a firewall blocks or allows a service, what other services is that firewall likely to block or allow?

5. What accessible services do hosts run?

6. How are running servers correlated? Given a server running a publicly available service, what other publicly available services is it likely to run?

7. How prevalent are anonymous hostnames? What fraction of hosts and IP addresses have anonymous hostnames?

To the best of our knowledge, we are the first ones to systematically study the behavior of firewalls and anonymous DNS on the Internet. In this paper, we propose methodologies and describe our findings to address these questions. Many results are surprising and provide new insights for understanding the deployment of security defense mechanisms on the Internet.

The rest of the paper is organized as follows. Section 2 reviews previous work. Section 3 details the methodology used. Section 4 presents the results, including answers to the questions posed above. Section 5 gives the conclusions.

2 Previous Work

Most previous work which examined behavior of hosts on the Internet focused on version of services running on systems. Early previous work includes multiple RFCs: 669[7], 679[9], 701[6], 702[8], 703[10], 751[20], 844[5], 847[35], and 876[33]. RFC 669, 679, 701, 702, and 703 are periodic surveys looking at telnet server behavior to observe the deployment of the new telnet protocol. RFC 751 is a survey that looks at how many mail daemons allow e-mail to unknown recipients and how often a particular method to send e-mail is available. messages is available. RFC 844 looks at the deployment of ICMP. RFC 846 looks at the deployment of TCP. With the exception of RFC 751, all of these exclusively dealt with the rate of the deployment of a new service or protocol. None of these surveys included more than 400 hosts, because the Internet did not have more than 400 hosts. That is small-scale when looked at in modern terms.

More recent work by Provos and Honeyman [31] studied the deployment of a new version of SSH [36] after the discovery of a vulnerability in one version of the protocol. They scanned 400,000 IP addresses at University of Michigan daily to determine the version of SSH server running on the systems that responded (around 2,300 of the IP addresses). Again, they were focused on the version of servers running, rather than looking at responsiveness of systems.

Several projects have used DNS to estimate the number of hosts on the Internet. They include Netsizer [24], ISC Internet Domain Survey [14], and RIPE Region Hostcount [32].

3 Methodology

To analyze firewall behavior, this paper will probe IP addresses for their responsiveness to particular services. Some firewalls contain specific rules that allow particular external hosts special access to their corporate network. This paper focuses on the behavior of a request from a general host, as that is the common case and it is not clear how to reasonably detect these special hosts. Two types of of responsiveness are considered. The first category is a host responding to the service request. This does not mean that the host necessarily runs the service, simply that no firewall blocks the request. The second category is a host listening for a service. This means that not only did no firewall block the request, but also that the host offered access to the service at some level. The specific packet responses implying these two response types are service-specific and detailed below.

Testing firewall behavior is composed of four pieces: selecting the services to test, selecting the IP addresses to test, testing the IP addresses, and separating firewall responses from end-host responses. There are too many services to test all of them, so a representative subset must be selected. There are too many IP addresses to test all of them, even if the set of limited to routed or assigned IP addresses. Thus, the IP addresses must be sampled using a representative subset of sufficient size to limit the error. Once the services and IP addresses are selected, the selected IP addresses must be tested for responsiveness to those selected services. This requires defining what responsive means in the context of each service.

The last piece may be less obvious. The goal is to test if IP addresses are responsive to the service. Some firewalls respond to services for IP addresses beyond them. Because the firewall acts on behalf of the IP address, the IP address appears responsive, even though the firewall does not allow the packet to reach the IP address. To remove these false positives, firewall responses must be detected and removed.

Service	Use	Type	Protocol/Port
Chargen	Debugging	Historical	TCP/19
DNS	Name Lookup	Current	UDP/53
Echo	Debugging	Historical	TCP/7
Finger	Identification	Historical	TCP/79
FTP	File Transfer	Current	TCP/21
Gnutella	File Sharing	P2P	TCP/6346
HTTP	WWW	Current	TCP/80
Ident	Identification	Historical	TCP/113
Kazaa	File Sharing	P2P	TCP/1214
NetBIOS	File Server	Local	TCP/137
NNTP	USENET	Current	TCP/119
Ping	Measurement	Current	ICMP
POP3	E-mail	Current	TCP/110
SMTP	E-mail	Current	TCP/25
SSH	Terminal	Current	TCP/22
Telnet	Terminal	Current	TCP/23
UDP	Measurement	Current	UDP/33437

Figure 1: Services selected to measure. “Historical” means that although the service is still in use, it is much less popular than it used to be.

3.1 Service Selection

It is not possible to test all possible services. There are 65,536 TCP ports and 65,536 UDP ports, too many to test with any level of detail in a reasonable time. Even limiting the search to ports with published services, there are almost 2000 of those listed in FreeBSD’s services file. This is still too many to test responsiveness of all these services for a reasonable number of IP addresses. Thus, a cross-section of services were selected. The seventeen services selected to test are listed in Figure 1. These servers were selected for a variety of reasons:

1. Chargen (character generator) [27] is almost exclusively used for bad purposes. Security experts suggest blocking it and there are few reasons not to do so.
2. DNS (domain name service) [22] is required for most other services to be useful. It is also a UDP service, unlike most of the services selected.
3. Echo [28] is similar to Chargen, but not as useful for bad purposes. Like Chargen, there are few reasons not to block it.
4. Finger [37] is considered bad as it reveals user names. security experts suggest blocking it, but, unlike Chargen, it provides a service useful for end-users.
5. FTP (file transfer protocol) [30] is a popular service used primarily to publish files for public download. FTP was selected for its popularity.

6. Gnutella [12] is a peer-to-peer (P2P) file sharing service. Given the potential for legal problems due to some uses of P2P services, Gnutella is expected to be blocked more often than other services.
7. HTTP (hypertext transfer protocol) [11] is used to transfer web pages. For some users, the Internet is web pages. HTTP was selected for its popularity, as almost every companies runs at least one HTTP server.
8. Ident (identification) [16] identifies the local user owning a particular TCP port. The most common use is by FTP servers. Some FTP servers connect back to the client using the Ident service to learn the identity of the user attempting to access the server. If firewalls simply block the packets, users are forced to wait for these Ident connections to timeout before using such FTP servers. For this reason, its behavior is expected to be unusual.
9. Kazaa [18] is another peer-to-peer (P2P) file sharing service. It is included for the same reasons as Gnutella. Kazaa provides a service to compare with Gnutella.
10. NetBIOS [13] underlies Microsoft Windows networking. It uses multiple ports. Port 137 is the naming service, selected because it is expected to be the least intrusive. Microsoft Windows networking is primarily used for local file sharing, so remote access is expected to be blocked by most firewalls.
11. NNTP (network news transfer protocol) [17] is used to transport USENET articles. It was selected as an established service run by many companies, but less popular than FTP or HTTP.
12. Ping is a service popularly used for measurements and debugging. It is considered relatively low-risk to allow this to pass a firewall. It runs on ICMP [25], unlike any of the other services. For these reasons, its behavior is expected to be different.
13. POP3 (post office protocol v3) [23] is used by users to retrieve their e-mail. The POP3 service that runs on port 110 is usually not encrypted, so it is considered more of a security problem than alternatives, such as IMAPS and POP3S.
14. SMTP (simple mail transfer protocol) [19] is used to deliver electronic mail between mail servers. Like HTTP, almost every company operates at least one publicly-accessible mail server.
15. SSH (secure shell) [36] provides remote terminal access to a host. It provides encryption and key-based authentication. It is expected that firewalls will block most SSH accesses so that users wanting terminal access are forced to use a VPN or go through a small set of SSH servers to gain access to internal machines.
16. Telnet [29], like SSH, provides remote terminal access. Although telnet can be encrypted, it is usually not. For this reason, security experts suggest using SSH instead. It is expected to be blocked more than SSH.
17. The UDP “service” is not a true service. If a UDP packet is sent to a non-listening port, most hosts will respond with an ICMP port unreachable message. This is commonly used by

traceroute [15] and related technologies for measurement and debugging. Allowing traceroute past a firewall is considered relatively low-risk and is useful for debugging network problems.

One service not selected is SNMP [3]. SNMP was excluded for two reasons: SNMP is intrusive and an SNMP server should not respond to unauthorized packets. SNMP is a network management service, so attempting access is often viewed as an attack. SNMP packets contain a community string, which acts somewhat like a password. If the community string does not match any community string configured to have access, the SNMP server does not respond to the request. In such cases, it is difficult to distinguish a firewall blocking the packets and not having a proper SNMP community string. For these reasons, SNMP was not tested.

3.2 Sample Selection

It is not reasonable to test all four billion IP addresses. Even limiting oneself to routed or assigned IP addresses leaves too many IP addresses to test. Thus, a random set of IP addresses must be selected.

The primary goal of the measurements is to analyze how prevalent different services are. Thus, IP addresses are randomly selected from the entire IP address space. Large parts of the address space are not used. A local BGP feed indicates that less than 30% of the IP address space is globally routed. However, aggressively limiting the IP addresses to active address space is prone to error. Although IP addresses are selected from the IP address range, only IP addresses within /8 CIDR blocks indicated as allocated according to ARIN's reverse lookup databases are tested. All other IP addresses are considered to be non-responsive. ARIN lists 124 /8 blocks as in use, after excluding multicast and private address space. It is more appealing that testing takes slightly longer than to introduce additional error by excluding active IP address space by filtering aggressively.

Two million IP addresses are selected, of which 968,482 IP addresses are in the CIDR blocks indicated as used by ARIN. The number of IP addresses selected determines the accuracy of the result. Because responsiveness is unusual, Poisson confidence intervals [1] are used. Consider a Poisson test that yields x positives. From statistics, the $1-\alpha$ confidence interval for $E[x]$ is given by $(\frac{1}{2}\chi^2_{\frac{\alpha}{2}}(2x), \frac{1}{2}\chi^2_{1-\frac{\alpha}{2}}(2x+2))$, where χ^2 is the chi-square distribution. For 423 positive responses, the 95% confidence interval is approximately $(383.6, 465.3)$, an error of less than 10%. If at least 908,400 IP addresses would respond positively to some test, then a sample of two million IP addresses is expected to contain 423 such IP addresses. Therefore, given that two million IP addresses will be tested, any test where at least 908,400 IP addresses would respond positively will, with high probability, estimate the correct number within 10%.

Selecting IP addresses from the space uniformly does not sample firewall behaviors uniformly, as one firewall may represent a large fraction of the IP addresses. However, this paper is more interested in the sampling blocking done for IP addresses than sampling blocking done by firewalls. This is exactly what is measured by selecting random IP addresses. Part of the reason for this focus is that firewalls may have different rules for different IP addresses. For example, a firewall may allow public access to the network with the web servers, at least for HTTP, but block HTTP access to the rest of the network. It is not clear what the behavior of the firewall would be considered is such a case. By sampling IP addresses, this ambiguity does not arise.

3.3 Testing Method

Most of the services selected are TCP-based. To test a TCP service, a TCP SYN packet is sent to an IP address to test if that host is responsive to that service. Any TCP response is considered responsive. Any TCP response that has the SYN and ACK flags set is considered a listening IP addresses. A TCP RST packets is sent for each TCP packet received that is not itself a RST, telling the tested host to discard the connection attempt.

In the case of a listening host, no attempt is made to determine what service is actually running on the port, simply that some service is listening to the port. In addition, no testing is done on other ports to determine if a service is running there. If a host used TCP wrappers, the connection would be established and then immediately closed. The testing records such hosts as listening, even though such machines provides no real service.

An alternative method to test responsiveness is to send a TCP ACK or TCP SYN/ACK packet instead of the TCP SYN packet. For firewalls that do not maintain state but allow outbound connection, TCP ACK packets may be allowed by the firewall. For firewalls that do maintain state, this packet should be dropped, as there would be no associated connection. Thus, testing using a TCP ACK may detect stateful firewalls and may find more hosts. However, a TCP service is only accessible if a TCP connection can be established. It is not possible to establish a connection with a host if TCP ACK packets elicit responses but TCP SYN packets do not. Thus, using TCP SYN packets better measures accessibility.

Not all of the services selected are TCP. In particular, DNS and UDP use UDP and Ping uses ICMP. Each of these services was tested in a different way.

The DNS service is tested by sending a PTR request for 1.0.0.127.in-addr.arpa. This is a request for the hostname of 127.0.0.1, which is the localhost address. Most DNS systems include this entry. Requesting this hostname has no security implications, as IP packets sent to this IP address are destined for the host that sent them. Any IP address that responds with a DNS packet (including a DNS error response) is considered responsive and listening. Any IP address that responds with an ICMP port unreachable is considered responsive but not listening.

The UDP service specifies its testing method. A UDP packet is sent to a high-numbered port. An IP address responding with an ICMP port unreachable message is considered responsive but not listening. An IP address cannot be considered listening to the UDP service. The port 33437 is one of the ports used by traceroute.

Ping is tested by sending an ICMP echo request packet (type 8, code 0). An IP address responding with an ICMP echo reply (type 0, code 0) is considered responsive and listening. An IP address responding with an ICMP port unreachable is considered responsive but not listening. Although an ICMP port unreachable may seem impossible, a few IP addresses respond in this way.

Sending a packet to test for a service may also induce an ICMP error response. Common ICMP error responses include ICMP time exceeded, ICMP host unreachable, ICMP filtered, and ICMP network unreachable. These may be responses from firewalls or may be responses from other network elements. Such error messages mean the host did not respond because the packet did not reach its destination. Therefore, with the exception of ICMP port unreachable, all ICMP error responses are ignored.

These techniques are effectively attempting an inbound connection. Depending on the definition of "connected", inbound connection attempts cannot find all the hosts connected to the Internet. Hosts behind firewalls that block all incoming connections or other devices that obscure its existence, such as a network-address translation (NAT) boxes, cannot be detected. For firewalls that block all incoming connections, it is not possible to use inbound connection attempts to determine which of the IP addresses beyond a firewall are active and which are not. It is not possible to determine how many hosts are beyond NAT boxes using inbound connection attempts. Other techniques, such as looking at inbound or outbound traffic, may be able to detect active hosts beyond such devices, but they still cannot detect hosts which do not access the Internet. Inbound connection attempts identify the subset of the hosts connected to the Internet that are publicly accessible to some degree, providing a subset of all hosts connected to the Internet.

Because it is impossible to detect a host behind a firewall that blocks all incoming connections using these methods, it is not possible to analyze what percentage of hosts are behind firewalls that block a given service. Thus, the analysis will focus on hosts either not behind a firewall or protected by a firewall that allows public access at least one service on that host, regardless of whether or not the host runs that service. This effectively removes from consideration hosts that cannot be detected.

3.4 Firewall Responses

Some firewalls respond to packets when they drop them. Naïvely, such responses would be counted the same as responses from a host at that IP address, even though that host, if it exists, did nothing. Firewall responses can account for more than 30% of all responses and 78% of measured listeners. Unless firewall responses are detected and removed, they skew the results.

Firewalls may be configured to respond to packets for a variety of reasons. One example is to respond to Ident [16] TCP SYN packets with TCP RST. Some FTP servers still attempt to connect back using Ident to clients to determine the user connecting to them. If the user's firewall simply drops such packets, users must wait for the Ident connection to timeout. However, if the user's firewall responds with TCP RST packets, the connection immediately fails, avoiding the response delay. When a firewall responds, it does not indicate a host is at that IP address, much less whether or not that host is listening for the particular service being tested.

To detect firewalls responses, pairs of IP addresses close to responding IP address were tested for being aliases of each other using IPid verification [2]. This tests if the IP identification (IPid) field in the IP header [26] of two IP addresses are correlated. Most machines use a common counter to set the IPid of response packets. If pairs of close IP addresses were tested and the IPid field values were correlated, then the responses are likely to be from the same host. This could be due to IP aliasing or a firewall intercepting the packets and sending responses itself.

Routers and web servers are often given many IP aliases. Web servers most commonly have multiple IP addresses to support virtual hosting under HTTP 1.0. Generally, a system administer will give such a web server adjacent IP addresses. Thus, the further two IP addresses are away, the less likely they are to be the same due to IP aliasing. Routers have an IP address for each interface. Routers are commonly given IP addresses at the beginning (.1) of the subnet in which they belong, although the end (.254 for a 24-bit network mask) is also common. To avoid falsely concluding that web server responses come from firewalls, the selected IP addresses should have large separation and multiple pairs of IP addresses should be tested. To avoid falsely concluding that responses from routers come from firewalls, the selected IP addresses should have different ending bit patterns and multiple pairs of IP address should be tested.

Five pairs of IP addresses were tested for each IP address. Each pair includes the original IP address and an IP address with the same three leading octets. Keeping the same three leading octets makes it likely that a firewall protecting the original IP address is also protecting the other.

This does not detect firewalls protecting smaller networks, but these are believed to be less common and such firewalls skew the results much less than, say, a firewall protecting a /8 network. Each of the five pairs includes the original IP address. The other IP address in each of the five pairs differs from the original IP address by 7, 13, 27, 44, and 69 respectively.

4 Results

	Raw		Firewall		Host	
Service	Responders	Listeners	Responders	Listeners	Responders	Listeners
Chargen	26388	667	4632	124	21756	543
DNS	22080	3049	3025	264	19055	2785
Echo	26577	725	4833	139	21744	586
Finger	26444	880	4927	217	21517	663
FTP	33361	8513	6836	2746	26525	5767
Gnutella	28172	385	5555	118	22617	267
HTTP	27087	9024	4679	1212	22408	7812
Ident	40655	5423	12749	4254	27906	1169
Kazaa	28018	823	5229	214	22789	609
NetBIOS	16233	176	3646	38	12587	138
NNTP	31959	1904	8890	332	23069	1572
Ping	27399	27327	1080	1079	26319	26248
POP3	32792	4060	8953	691	23839	3369
SMTP	31730	8143	7408	3103	24322	5040
SSH	34465	4287	9059	972	25406	3315
Telnet	28697	8180	6717	2816	21980	5364
UDP	22306	0	1751	0	20555	0

Figure 2: Summary of per-service results. Host counts exclude firewall responds.

Figure 2 shows the raw results. The raw columns give the number of responders and listeners found for each server using the testing method described in Section 3.3. The firewall responders count is the number of responding IP addresses that were detected as responses coming from a firewall, according to the firewall detection algorithm of Section 3.4. The firewall listening count is the number of those responses that were considered listening. The host responders and listeners columns show the corrected numbers. For example, the host responder count is the raw responder count minus the firewall responder count. After this correction, firewall responses and listeners are treated the same as receiving no response at all.

4.1 Effect of Large Domains

Some of the counts of Figure 2 are affected by large domains. The practice of a large domain can represent a large fraction of a particular behavior, especially when that behavior is otherwise unusual. This happens with firewall listeners counts. One surprising result is the large number of firewalls which responded with listening replies to Ident, SMTP, FTP, and Telnet. Many of these

Domain	Service	Type	Cnt	Pct
army.mil	POP3	FW responder	4182	47%
army.mil	NNTP	FW responder	4103	46%
army.mil	SSH	FW responder	4176	46%
army.mil	Ident	FW responder	4022	32%
army.mil	FTP	FW responder	2055	30%
army.mil	Telnet	FW responder	1993	30%
army.mil	SMTP	FW responder	1960	27%
army.mil	Ident	Host listener	274	23%
direcpc.com	NetBIOS	Host listener	32	23%
bbtec.net	DNS	Host listener	449	16%
usmc.mil	NNTP	Host listener	241	15%
nrao.edu	NetBIOS	Host listener	20	15%
ripe.net	NetBIOS	Host listener	18	13%
direcpc.com	Gnutella	Host listener	33	12%
bbtec.net	HTTP	Host listener	852	11%
aol.com	Kazaa	Host listener	61	10%

Figure 3: Domains contributing more than 10% of the IP addresses exhibiting some behavior for a service (raw numbers and firewall listening numbers omitted for brevity and clarity).

IP addresses (1836 of 3103 for SMTP, 1312 of 2746 for FTP, 1974 of 2816 for Telnet, and 3893 of 4254 for Ident) are USA military networks in the .mil top level domain, which includes all of 55.0.0.0/8. Hand testing five IP addresses reported as firewall listeners for SMTP revealed that nearby IP addresses returned the same SMTP header, indicating that the requests are processed by the same machine. For FTP, ten IP addresses reported as firewall listeners that were manually tested all closed the connections without sending even a header (as if running TCP wrappers). Based on IPid behavior, nearby IP addresses yield responses from the same machine. Thus, a firewall is sending listening responses, either directly or by redirecting the requests to another host. The USA military also represents a large fraction of the NNTP firewall listeners: 59 of the 332.

Fortunately, large domains primarily affect firewall responses. Figure 3 shows domains that represented more than 10% of firewall response or host listening counts for some service. This paper focuses primarily on host responses, where no single domain represents more than 10% of the responses. The largest effect by a single domain for host responses is bbtec.net, which represents 7.9% of the host responses to UDP. Thus, large domain effects do not greatly affect the results.

4.2 How Prevalent are Firewalls?

Of the 56,297 IP addresses that responded to some service, only 4,023 responded to all services. Thus, at least 93% of hosts attached to the Internet are behind a filtering device of some type. Because this excludes hosts behind firewalls that block all incoming connection attempts, the true percentage is even higher than 93%. Clearly, firewalls are an important consideration when trying to understand the Internet.

4.3 What do Firewalls Block?

Service	% of Hosts Responding	% of Hosts Listening	Total Public Servers (95% Confidence Interval)
Chargen	38.6%	1.0%	1,070,000 – 1,268,000
DNS	33.8%	4.9%	5,761,000 – 6,207,000
Echo	38.6%	1.0%	1,159,000 – 1,365,000
Finger	38.2%	1.2%	1,317,000 – 1,536,000
FTP	47.1%	10.2%	12,067,000 – 12,709,000
Gnutella	40.2%	0.5%	507,000 – 646,000
HTTP	39.8%	13.9%	16,406,000 – 17,152,000
Ident	49.6%	2.1%	2,369,000 – 2,659,000
Kazaa	40.5%	1.1%	1,206,000 – 1,416,000
NetBIOS	22.4%	0.2%	249,000 – 350,000
NNTP	41.0%	2.8%	3,211,000 – 3,547,000
Ping	46.8%	46.6%	55,687,000 – 57,053,000
POP3	42.3%	6.0%	6,993,000 – 7,483,000
SMTP	43.2%	9.0%	10,527,000 – 11,126,000
SSH	45.1%	5.9%	6,879,000 – 7,365,000
Telnet	39.0%	9.5%	11,213,000 – 11,832,000
UDP	36.5%	0.0%	0 – 8,000

Figure 4: Percentage of hosts responsive to at least one service that respond and listen to each service, as well as the estimated total number of public servers for each service.

Figure 4 shows what percentage of the hosts that respond to at least one service respond to each of the services. The first two columns correspond to the last two columns of Figure 2, except that they are expressed as percentages of hosts that respond to at least one service. The total public servers column gives the 95% confidence interval for the number of public servers running each service, based on the number of host listeners found.

Unexpectedly, Ident is the most responsive service[1]. That is, Ident is the service most rarely blocked by firewalls. This is most likely due to FTP servers attempting to use Ident to determine the user attaching to them. As mentioned above, if a firewall simply drops these packets, the user must wait for the FTP server to timeout, which can be thirty seconds or longer. To avoid this, firewalls must be configured to either allow such requests or respond to the packets themselves[2].

The least responsive service is NetBIOS. That is, NetBIOS is the service most commonly blocked by firewalls. This is expected, as NetBIOS is a network file service usually meant to be accessed only locally. As such, remote access is rarely desirable. The estimated number of public NetBIOS servers is also small.

[1] Technically, Ident is only the most responsive of the services tested. For brevity, when comparing behavior of services, this paper will assume that the seventeen services represent all services

[2] Another option would be to allow Ident queries only from machines with which FTP connections have been established. To the authors' knowledge, no commercial firewall provides a simple way to do this.

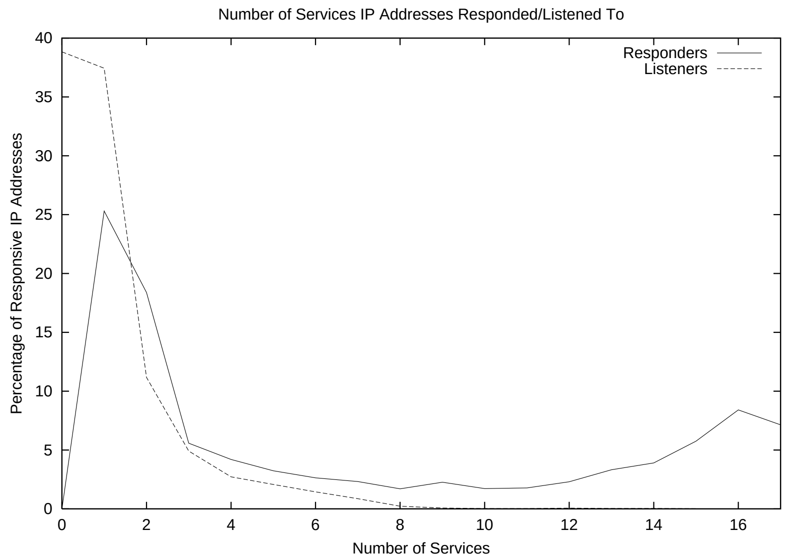

Figure 5: Distribution of the number of services IP addresses responded to and listened to.

4.4 How Much do Firewalls Block?

There are two basic firewall configuration methods: block by default and block by exception. One way to estimate how common these two configuration methods are is to look at how many services to which hosts respond. IP addresses beyond firewalls that block by default are responsive to few services. IP addresses beyond firewalls that block by exception are responsive to many services. Figure 5 shows the distribution of the number of services to which hosts respond. 31,927 of the tested IP addresses are responsive to five or fewer services. Presuming that a firewall configured to block by default would allow no more than five services, 56.7% of IP addresses are beyond a firewall configured to block by default.

An alternative method to estimate the percentage of IP addresses behind a firewall configured to block by default is to look at a random port. The best examples of this are Gnutella and Kazaa, because they are new, more often run on different ports, and run on ports above 1024[3]. 31,707 of the tested IP addresses respond to neither Kazaa nor Gnutella. Presuming that a firewall configured to block by default would block both, 56.3% of the IP addresses are behind firewalls that block by default. These two methods agree that about 56% of IP addresses are behind a firewall configured to block by default.

The average number of services that an IP address responds to is 6.83. The median is four, lower than the average. The distribution is bimodal, with maximums at one and sixteen. The maximum at one represents hosts responsive to a single service. This is a statement that only one service produces enough value for the security risk it represents, perhaps because the host only runs that service. The maximum at sixteen represents hosts that are accessible using all services except one. This represents a conclusion that only one service is worthy of blocking, a damning statement about that service. This could be either because the value of that service is low or the security risk for that service is high. Figure 6 shows how often these two behaviors are observed for each service.

[3]Some systems require special access to listen to ports below 1024.

Service	Responded all but this		Responded only to this	
	Count	Percent	Count	Percent
Chargen	59	0.1%	224	0.4%
DNS	566	1.0%	1017	1.8%
Echo	24	0.0%	172	0.3%
Finger	42	0.1%	137	0.2%
FTP	88	0.2%	476	0.8%
Gnutella	33	0.1%	287	0.5%
HTTP	271	0.5%	1023	1.8%
Ident	42	0.1%	3488	6.2%
Kazaa	57	0.1%	436	0.8%
NetBIOS	2578	4.6%	175	0.3%
NNTP	28	0.0%	115	0.2%
Ping	364	0.6%	3906	6.9%
POP3	17	0.0%	236	0.4%
SMTP	116	0.2%	1258	2.2%
SSH	40	0.1%	287	0.5%
Telnet	227	0.4%	590	1.0%
UDP	183	0.3%	417	0.7%

Figure 6: Number of IP addresses with unusual response behavior for each service. This table only refers to host responses, not firewall responses. Percentages are the fraction of hosts that respond to at least one service.

NetBIOS is, by far, the most common service to block if blocking only one service. Recall from Figure 2 that NetBIOS is also the least common service to elicit responses. Part of the reason so many IP addresses are responsive to every service except NetBIOS is that some large ISPs block only it. Such ISPs contribute a large part of the numbers: bbtec (921), Comcast (294), Road Runner (243), and AOL (153). Excluding these ISPs given leaves 968 IP addresses for which only NetBIOS is blocked, still 70% more than DNS, the second-most-common service to block if blocking only one service.

Ident and Ping are the two most common services to allow if allowing only one service. Neither of these have single domains representing a large fraction of the IP addresses responding only to it. They are allowed for different reasons. Ping is allowed because it is useful for measurement and debugging and it has a low security risk. Ident is allowed because some FTP servers try to use it. Ident is unsurprising, because, as noted above, it is also the most common service allowed through a firewall. Ping, in contrast, is the sixth-most-common service allowed through a firewall.

Since most IP addresses are responsive to only a few services, it is difficult to determine what hosts are accessible from the Internet. On a large set of IP addresses, testing a large set of services is intrusive and time consuming. Thus, to determine which hosts are accessible from the Internet, as few services as possible should be tested. One idea is to use only one service. 27,906 of the tested IP addresses respond to Ident, more than any other service. This represents 49.6% of the IP addresses that responded to at least one service. That is, each of the services tested are blocked

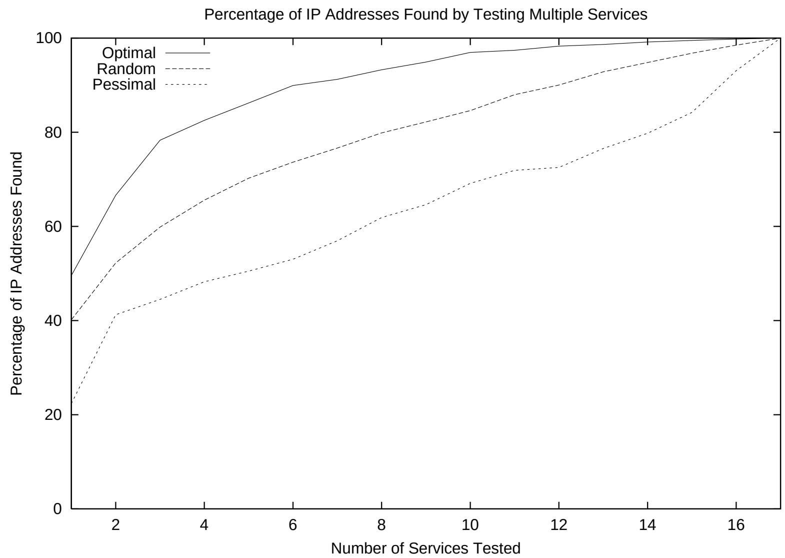

Figure 7: IP addresses found by using multiple services with optimal, random, and pessimal service selection.

for more than half the responsive IP addresses. Thus, testing only one service will, at best, find approximately half the IP addresses that would be found if many services were tested. Obtaining a more reasonable estimation, such as 90% of the accessible hosts, requires testing multiple services.

Figure 7 shows what fraction of responsive IP addresses would be found if a subset of services were tested. If services are selected optimally, only six services are necessary to obtain more than 90% of the IP addresses: DNS, FTP, Gnutella, HTTP, Ident, and Ping. This set of services is interesting for two reasons. Firstly, DNS elicits the second-smallest number of responses, yet it is included. Secondly, NNTP, POP3, and SSH are excluded, despite the fact that they elicit more responses than DNS, Gnutella, and HTTP. Obviously, behaviors between services are correlated. This will be discussed in more detail below. If services are selected randomly, the number of services required to obtain at least 90% of the IP addresses jumps to 12, declining to all but one of the services being required if the services are selected poorly. The issue here is that some services have many IP addresses that respond only to them.

4.5 How is Firewall Behavior Correlated?

As mentioned above, behavior of different services are correlated. Because the measures here are boolean, a non-standard definition of correlation is used. Define the correlation between two services to be the probability that an IP address responds to the second service given that it responds to the first. By formula, let a and b be two services, and R_x the set of IP addresses responding to the service x. Then, the correlation between a and b, C_a^b, is given by:

$$C_a^b = P[x \in R_b | x \in R_a] = \frac{\|R_a \cap R_b\|}{\|R_a\|}$$

Figure 8 shows the correlation between services. The bottom quartile is between 34% and 57%. The top quartile is between 80% and 94%. Presuming IP addresses respond to two services

Service		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
Chargen	A	x	56	93	86	84	86	64	80	84	48	87	68	81	77	83	83	67
DNS	B	64	x	65	67	66	66	67	76	67	42	67	68	73	69	71	61	77
Echo	C	93	57	x	87	85	89	65	81	86	48	89	69	82	78	84	81	68
Finger	D	87	59	88	x	87	85	67	84	85	50	94	74	85	81	89	77	74
FTP	E	69	47	70	71	x	69	60	67	70	38	77	58	71	71	81	69	58
Gnutella	F	83	56	85	81	81	x	63	76	92	46	83	68	76	73	79	76	67
HTTP	G	63	57	63	65	71	64	x	75	64	43	67	65	76	73	70	63	57
Ident	H	63	52	63	65	64	62	60	x	62	42	66	60	72	67	68	57	54
Kazaa	I	80	56	82	80	81	91	63	76	x	45	82	67	77	73	78	74	67
NetBIOS	J	82	64	84	86	80	82	77	92	81	x	85	72	92	87	86	72	67
NNTP	K	82	56	84	87	89	81	65	80	81	47	x	68	83	81	87	77	70
Ping	L	56	49	57	60	58	58	55	64	58	34	60	x	59	57	62	53	60
POP3	M	74	59	75	77	79	72	71	84	73	49	80	65	x	87	83	68	63
SMTP	N	69	54	70	72	77	68	67	77	68	45	77	62	86	x	80	67	60
SSH	O	71	53	72	75	85	70	62	75	70	43	79	65	78	77	x	69	64
Telnet	P	82	53	80	75	83	78	65	72	77	41	81	63	74	74	80	x	65
UDP	Q	71	71	72	77	75	74	62	73	75	41	78	77	73	71	79	69	x

Figure 8: Correlation of IP addresses responding to two services, expressed as percentages. The upper-right value of 67 corresponds to $C_A^Q = 67\%$.

independently each with probability 40%, the correlation would be 40%. Thus, services are more strongly correlated than randomness would imply. From above, 37% of IP address are behind firewalls that block by exception. Thus, the correlation should be higher than randomness would imply.

The largest correlation is 94%. Finger is 94% correlated with NNTP, so most IP addresses responsive to Finger are responsive to NNTP. Chargen is 93% correlated with Echo and Echo and Chargen are 93% correlated. These services are often lumped with Discard, Daytime, and Time under the term "tiny services," as they are simplistic services. We expected less correlation, with Chargen blocked more often than Echo. However, this is not the behavior: most IP addresses that respond to one service respond to the other.

The lowest correlations are with NetBIOS. For example, Ping is 34% correlated with NetBIOS, the lowest correlation. This reflects the fact that only 22.4% of IP addresses are responsive to NetBIOS. Ignoring NetBIOS, the lowest two correlations are 47% and 49%. FTP is 47% correlated with DNS and Ping is 49% correlated with DNS. This derives from the fact that DNS was the second-least responsive service, with 33.8% of IP addresses responding to DNS.

One theory of why DNS so rarely induces responses is that it uses UDP for transport (at least, as commonly used and as tested), while most of the services tested employ TCP. Since TCP is connection-based, it is far simpler for firewalls to understand TCP communications than UDP communications, and as a result, UDP may be considered less secure. Ping, which uses ICMP, is the sixth-most responsive service with 46.8% of IP addresses responding to it. However, ICMP is used to communicate error messages and Ping itself is generally considered low-risk. UDP, like DNS, uses UDP (unsurprisingly). In fact, slightly more IP addresses respond to UDP packets on

Service		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
Chargen	A	x	20	4	8	24	11	24	30	13	6	12	33	18	22	21	12	17
DNS	B	26	x	25	24	37	27	26	36	27	12	27	36	27	30	32	28	16
Echo	C	4	19	x	7	23	10	24	30	12	6	11	33	17	21	21	12	17
Finger	D	9	18	8	x	22	12	23	28	13	5	8	30	16	20	18	15	13
FTP	E	12	22	11	9	x	15	22	34	14	8	9	37	16	19	13	12	17
Gnutella	F	9	19	7	10	25	x	24	32	6	7	13	33	20	23	23	14	16
HTTP	G	23	19	22	21	31	25	x	33	25	9	24	35	20	23	29	23	23
Ident	H	15	16	14	12	31	19	20	x	20	4	16	33	13	20	23	21	19
Kazaa	I	10	19	9	10	24	5	24	32	x	7	13	33	19	23	23	15	16
NetBIOS	J	26	25	26	25	38	28	29	37	29	x	28	39	28	31	33	30	28
NNTP	K	8	19	7	4	18	12	22	29	12	6	x	32	14	17	16	12	13
Ping	L	23	20	22	19	37	24	26	37	25	12	25	x	27	31	30	27	16
POP3	M	13	16	12	10	23	17	17	24	16	3	12	33	x	11	17	17	17
SMTP	N	16	19	15	13	24	19	19	29	19	5	13	35	10	x	19	18	19
SSH	O	12	18	11	8	16	16	22	29	16	6	10	32	13	16	x	14	14
Telnet	P	11	22	12	14	24	16	24	35	17	10	15	36	22	24	23	x	18
UDP	Q	20	12	19	16	31	21	27	36	21	12	19	30	24	27	26	22	x

Figure 9: Anticorrelation of IP addresses responding to two services. The upper-right value of 17 corresponds to $A_A^Q = 17\%$.

ports used by traceroute than DNS queries: 36.5% of IP addresses respond to UDP.

UDP is 71% correlated with DNS and DNS is 77% correlated with UDP. Although no other service is more correlated with either DNS or UDP, these correlations are not remarkably high when compared to other correlations with DNS or UDP: NetBIOS is 64% correlated with DNS and Finger is 74% correlated with UDP. The high correlation between DNS and UDP implies that, indeed, part of the reason DNS gains fewer responses is that it uses UDP. However, these correlations are not so high that it fully explains the small number of responses. Firewalls do not block DNS more often simply because it uses UDP.

The correlation C_a^b considers only IP addresses behind firewalls that allow access to service a. Consider instead IP addresses behind firewalls that block access to service a. It is useful to define "anticorrelation" to better understand their behavior. Let $R = \cup_x R_x$, the set of all hosts that respond to at least one service. Define the anticorrelation of a to b, A_a^b, as:

$$A_a^b = P[x \in R_b | x \in R \text{ and } x \notin R_a] = \frac{\|R_b \cap \bar{R}_a\|}{\|R \cap \bar{R}_a\|}$$

A low anticorrelation A_a^b means that most IP addresses that do not respond to service a also do not respond to service b. A high anticorrelation A_a^b means that IP addresses that respond to service b do not respond to service a.

Anticorrelation, as defined, takes into account only IP addresses that respond to at least one of the services tested. This overstates the fraction of responsive IP addresses behind a firewall that allows service b but blocks service a, as it ignores firewalls that block all inbound traffic. Administrators of such firewalls are not making a decision on the relative merits of allowing two services, they are making a blanket decision to block all traffic from the general Internet for the services tested.

Figure 9 shows the anticorrelations. The anticorrelation of every service with NetBIOS is low, showing that most firewalls that block any traffic block NetBIOS. Ignoring NetBIOS, the lowest two anticorrelations are Finger and NNTP, Chargen and Echo, and Echo and Chargen, all 4%. The low anticorrelations between Chargen and Echo and between Echo and Chargen show that IP addresses behavior to these two services are similarly. This is not surprising, given their high correlations. Finger and NNTP are interesting because the anticorrelation between NNTP and Finger is 8%, twice as high as the anticorrelation between Finger and NNTP. In particular, 20,160 IP addresses respond to both Finger and NNTP. While 2,911 IP addresses respond to NNTP but not Finger, only 1,365 IP addresses respond to Finger but not NNTP. Both of these services are traditional UNIX services. Finger is considered a large security problem, as it reveals usernames and, historically, had a buffer-overflow flaw that was exploited by the Morris Worm [34]. NNTP provides a service, USENET, that requires that most NNTP servers be accessible from the Internet.

The large anticorrelations are between rare services and Ident. This derives from the fact that Ident has a relatively low number of IP addresses that do not respond to it but do respond to another service. Some of the other large anticorrelations are more surprising. For example, Ping and SMTP have approximately the same number of responses, yet they are anticorrelated in both directions. 14,998 of the IP addresses tested respond to both services, 9,326 respond only to SMTP, and 11,323 respond only to Ping. In fact, many services are more anticorrelated with Ping than expected, given the number of IP addresses responsive to Ping. Ping is, as mentioned, generally considered low risk and useful for debugging. Thus, firewall administrators are likely to allow Ping even if blocking other services.

These correlations imply a dependency, where an IP address in unlikely to respond to service x if it does not respond to service y. Let there be an arc from x to y if $A_x^y \leq 8\%$. Thus, an arc implies that at least 92% of the IP addresses that have firewalls that block service x also block service y. Note that this relationship is neither symmetric nor transitive. The resulting graph, shown in Figure 10, has 24 arcs on twelve of the services.

There are four completely connected cliques in this graph. The first one is Echo and Chargen, which is expected given their high correlation. The second is NNTP and Finger and the third is NNTP and Chargen. Echo, Chargen, and Finger are traditional services that, historically, ran on most UNIX boxes, and NNTP generally runs on UNIX or UNIX-like systems, so this is somewhat expected. The third clique is Kazaa and Gnutella. This clique results not only because they provide similar services, but also because they are the only TCP services tested that used ports above 1024. Interestingly, 936 of hosts responded only to Gnutella and Kazaa, about 0.2% of hosts. For these three cliques, firewall administrators generally configure firewalls equally for each service in the clique.

Five services are missing from the graph: DNS, HTTP, Ping, Telnet, and UDP. Four additional services have no inbound arcs and only one outbound arc, going to NetBIOS: FTP, Ident, POP3, and SMTP. These represent services that are blocked relatively independently from other services. Except for UDP, which has no listeners, these are eight of the nine services with the largest number of hosts listening. Thus, part of their high anticorrelation comes from public servers running the service, which requires that the firewall allows packets to reach the host.

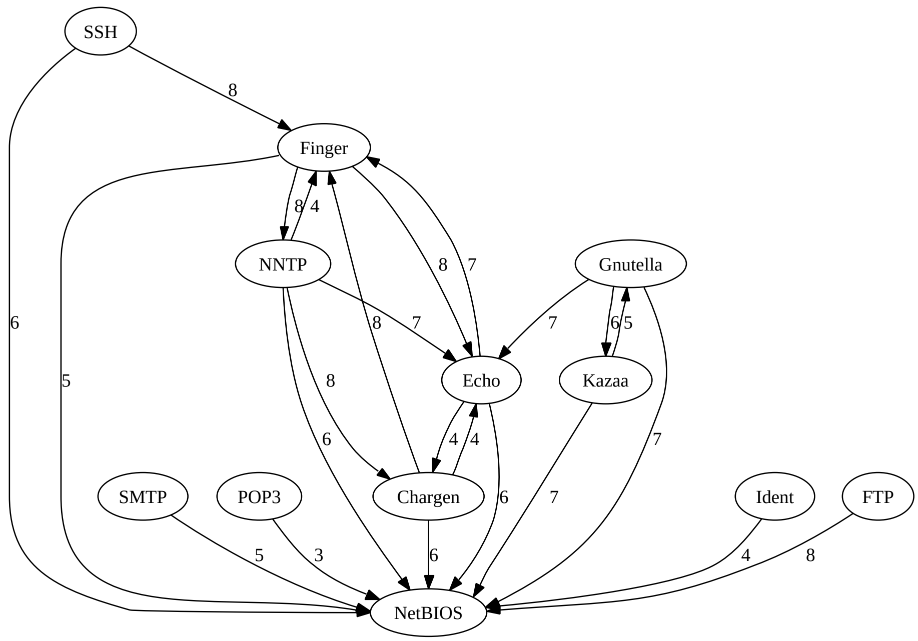

Figure 10: Dependency digraph of firewall blocking. An arc from service x to service y implies that $A_x^y \leq 8\%$. That is, of the tested IP address that respond to some service but do not respond to service x, no more than 8% respond to service y.

4.6 What Accessible Services do Hosts Run?

Figure 4 shows the percentage of IP addresses that run each services as well as the estimated total number of public servers. Not surprisingly, ICMP ping was the most common listening service. Excluding that, HTTP, FTP, Telnet, and SMTP were the next four most common. HTTP, FTP, and SMTP are the three basic services of the Internet: file transfer, e-mail transfer, and web servers. DNS was only the ninth most-common service, despite the fact that DNS is required for almost every other service to properly work. The fourth most common service, Telnet, is insecure as usually implemented, without encryption. It is surprising that Telnet is still in large use, given the availability of a secure alternative with more functionality, SSH. The number of HTTP servers varies greatly from the 50 million number reported by Netcraft [21]. Netcraft counts domains, not servers. A single server can host millions of domains.

4.7 How are Running Servers Correlated?

Although not directly related to the question of firewall behavior, another interesting question is how often two services run on the same machine. This requires looking at the correlation of IP addresses listening to two service (the listening correlation). Figure 11 shows the listening correlations of the services. Note that this is based on listening and available to the public Internet. A host may be running services that are not publicly available. Moreover, the host may be running local filters such as TCP wrappers that make the service not usable by the public, even though machines on the public Internet can connect to the service.

82% of the IP addresses running POP3 servers run SMTP servers. A host running POP3 is a

Service		A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
Chargen	A	x	11	82	42	45	27	45	23	23	16	31	56	38	38	36	34	0
DNS	B	2	x	4	1	21	0	33	2	0	0	2	73	15	19	13	12	0
Echo	C	76	19	x	31	39	22	45	23	22	12	25	64	31	35	26	35	0
Finger	D	35	4	27	x	43	22	35	24	18	13	37	52	46	36	39	52	0
FTP	E	4	10	4	5	x	2	54	8	2	1	23	43	33	51	37	45	0
Gnutella	F	0	1	49	54	42	x	56	49	49	33	54	38	53	48	60	38	0
HTTP	G	3	12	3	3	40	2	x	4	3	1	8	59	27	37	21	22	0
Ident	H	11	6	11	14	37	11	29	x	11	7	20	35	32	36	38	32	0
Kazaa	I	20	1	21	20	20	22	43	21	x	11	25	59	29	23	24	18	0
NetBIOS	J	64	1	51	62	56	63	76	63	46	x	62	29	71	77	75	37	0
NNTP	K	11	3	9	15	83	9	42	15	10	5	x	9	49	82	67	70	0
Ping	L	1	8	1	1	9	0	18	2	1	0	1	x	5	7	6	11	0
POP3	M	6	13	5	9	57	4	63	11	5	3	23	39	x	82	39	24	0
SMTP	N	4	11	4	5	59	3	58	8	3	2	26	37	55	x	38	30	0
SSH	O	6	11	5	8	65	5	50	13	4	3	32	48	39	58	x	46	0
Telnet	P	3	6	4	6	48	2	32	7	2	1	20	53	15	29	28	x	0
UDP	Q	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	x

Figure 11: Correlation of IP addresses listening to two services, expressed as percentages. The second value on the first line, 11, corresponds to $LC_A^B = 11\%$.

mail server, so it is unsurprising that it also runs SMTP. However, in the case of a corporate mail gateway, it must run SMTP to allow incoming electronic mail, but, if it runs POP3, the firewall should block access from the Internet, as users accessing their electronic mail should be on the corporate network (physically or virtually). Indeed, only 55% of IP addresses running SMTP also run POP3.

Listening correlations show which services are run by servers and which ones are run by end-systems. For example, Ping, run by every host, although perhaps blocked by a firewall, has low correlation with every other service. 61.1% of IP addresses responding to at least one service were listening to at least one service. 30.6% of IP addresses were listening only to Ping. Thus, more than half of the IP addresses listening to at least one service listen only to Ping.

All services have low listening correlations with DNS, and DNS has low listening correlation with all services except Ping. Thus, while it is common to combine mail servers (SMTP) and web servers (HTTP) (58% of mail servers are also web servers), it is rare to combine nameservers (DNS) with other services.

In responsiveness, Chargen and Echo behaved similarly. This similarity is observed again in listening correlation. In contrast, Kazaa and Gnutella behaved similarly in terms of responsiveness, but are strikingly different in terms of listening. For example, listening to Gnutella is 54% correlated with Finger, but listening to Kazaa is only 20% correlated with Finger.

Listening correlations are somewhat biased by the default configuration of operating systems. Because default installations and near-default installations are common, decisions by operating system distributors such as RedHat affect numerous machines.

Figure 12 shows the dependency digraph of listening to the services tested. An arc is included

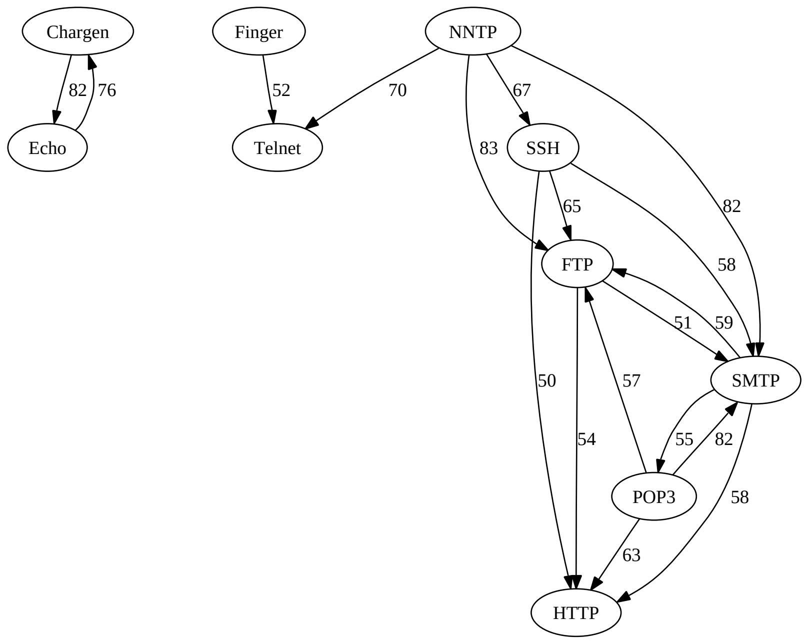

Figure 12: Dependency digraph of listening to services. An arc from service x to service y implies that $LC_x^y \geq 50\%$. That is, at least half the IP addresses listening for service x also listen for service y.

in the graph from x to y if $LC_x^y \geq 50\%$. The graph contains 18 edges on ten nodes. Ping is excluded from consideration, as the definition of listening for this service is unusual. Also excluded are NetBIOS and Gnutella. Both of these services would have many out-arcs. However, few IP addresses listen to these services, so the correlations have high error, making it unclear if these out-arcs are valid.

The dependency digraph of listening has three cliques: FTP and SMTP, Chargen and Echo, and POP3 and SMTP. POP3 and SMTP were previously discussed, as were Chargen and Echo. The listening correlation between FTP and SMTP is 51% or 59%, depending on direction. Both of these services commonly run on hosts running UNIX-like operating systems. It is more surprising that SMTP servers run FTP servers than the reverse. An FTP server may run SMTP to allow clients to send electronic mail to operators of the server. However, there is no particular reason why the reverse must be true. In fact, because SMTP servers deal with incoming and outgoing electronic mail for a set of hosts, SMTP is generally considered to require more security than an FTP server. Thus, it is not clear why this forms a clique.

4.8 How Prevalent are Anonymous Hostnames?

Reverse queries were sent to determine the hostnames of the 968,482 IP addresses tested for responsiveness. 133,496 of these lookups were successful in finding a hostname. There are two common ways to implement anonymous DNS: give the same dummy hostname, such as "unknown" or "nothing" for all IP address or derive the hostname from the IP address. Hostnames can be derived from IP addresses in many ways. Ways to encode the IP address 10.45.127.11 in the hostname include:

Hostname Type	# IP Addresses	# Responding	Pct Responding
Based on IP address	94635	21189	22.4%
Dummy hostnames	3108	73	2.3%
Other hostname	35753	10532	29.5%
Successful lookups	133496	31794	23.8%
Non-existent domain responses	735225	21107	2.8%
Lookup error	99761	2933	2.9%
All IPs	968482	55834	5.8%

Figure 13: Prevalence of anonymous hostnames, non-existent hostnames, and DNS errors. The percentage is the percentage of IP addresses with that type of hostname that responded to at least one service.

1. The hostname contains the entire IP address as a substring, either forwards or backwards. Examples include "10-45-127-11", "Yahoo010045127011", "dialup-10.45.127.11", and "h10s45a127n11".

2. The hostname contains only the last two or three octets of the IP address, in forward or reverse order. Examples include "ws127011", "PA011127", "011-127-045", and "rdu11-127-45".

3. The hostname contains the entire IP address or all but one octet in an unusual order. Examples include "11.Red-10-45-127", "ca-avignon-3-247.w10-45", "h11-10-45-127", and "host10451270b".

4. The hostname contains the IP address in hexadecimal, in either forward or reverse order. Examples include "p0a2d7f0b", "0A2D7F0B", and "0B.7F2D0A".

5. The hostname contains the entire IP address in hexadecimal or all but one octet in an unusual order. Examples include "p2d7f0b", "0b.7f.0a2d", and "0a2d7f-11".

Figure 13 shows how often anonymous DNS occurs in both the entire set of IP addresses tested and those IP addresses responding to at least one of the services tested. Hostnames not based on the IP address will be called "meaningful". Note that there are other ways to encode part of the IP address in the hostname that are not included herein, so this overestimates the number of meaningful hostnames.

A majority of hostnames are anonymous. Of 133,496 successful lookups, 97,743 (73.2%) of the hostnames returned are based, in part, on the IP address or are dummy hostnames. For IP addresses accessible from the Internet, a bare majority have hostnames. Only 31,794 (56.9%) of the 55,834 IP addresses that respond to at least one service have hostnames. Excluding errors (most of which were server failure messages from a name server), the percentage rises to 60.1%. Clearly, many active IP addresses do not have hostnames. Even those IP addresses that do have hostnames often have anonymous hostnames. 21,262 (66.9%) of the IP addresses that have hostnames and respond to at least one service have hostnames that are not meaningful.

It is difficult to come to a definitive conclusion regarding whether or not an anonymous or meaningful hostname belongs to an active IP address. The data show that anonymous hostnames are less likely to represent IP addresses that are responsive from the Internet. However, this does not necessarily mean the machines are not active. In fact, one might expect that a network which employs split DNS is more likely to be protected from the Internet by a restrictive firewall.

141,969 of the DNS lookups fail with a non-existent domain response from ARIN, indicating that no DNS server is responsible for reverse queries on that IP address. Of these, 419 IP addresses were responsive. These IP addresses represent 58 different assigned network blocks; none of them are within unassigned space. It is not clear why these IP addresses do not have any DNS server associated with their reverse lookup. Aggressive filtering based on the top level domain database for in-addr.arpa would reduce the number of IP addresses tested by 14.7%, but miss 0.7% of the responsive IP addresses.

5 Conclusions

In this paper, we performed the first study on firewall behavior and anonymous DNS on the Internet. We proposed a methodology for such a study and our analysis led to surprising findings.

We discovered that more than 93% of the hosts on the Internet are behind a firewall or filter device of some sort. Thus, almost every packet delivered on the Internet passes through at least one such device, greatly muddying the notion of "connected to the Internet". These filtering devices differ in the level and type of filtering they employ.

Approximately 56% of IP addresses are protected by a firewalls that block by default. This means that 44% of IP addresses are behind a firewall that blocks by exception. In terms of security, blocking by exception is a much worse policy, as it depends on knowing which services are bad. It is much better, for security reasons, to assume by default that services are bad and only allow access to services on specific machines where the service is useful enough to accept the security risk. The high percentage (44%) of IP addresses behind a block-by-exception firewall raises a security concern.

Firewall behavior is correlated between services. For some pairs of services, such as Finger and NNTP, firewall behavior is heavily correlated. In such cases, the perceptions of the two services are similar. In other cases, blocking a certain service makes it likely to block another service but not vice-versa. For example, only 5% of responsive IP addresses that are not responsive to POP3 are responsive to NetBIOS, but 28% of responsive IP addresses that are not responsive to NetBIOS are responsive to POP3. In such cases, one of the services is perceived either as more risky or less valuable. In this case, NetBIOS is both riskier and less valuable than POP3, in terms of remote access.

It is also surprising to find out that certain insecure services are still prevalent. For example, Telnet is much more popular than SSH. There are approximately eleven million Telnet servers and approximately seven million SSH servers.

The accessible services running on hosts are also correlated. Machines listening to some services are likely to run others. For example, most machines than run either SMTP or POP3 run both. In other cases, the relationship is only one way. For example, 83% of NNTP servers run FTP servers, but only 23% of FTP servers run NNTP servers. When most servers running service A run some other service B, any security problem in service B will affect most A servers. In this case, if FTP servers could be compromised, then most NNTP servers are compromised as well. This leads to some unexpected security implications, where a security flaw in service B causes problems with A servers despite the fact that the services are not obviously related.

Anonymous DNS is common. 70.9% of hostnames for IP addresses are based, at least in part, on the IP address. Even hosts which are responsive on the Internet rarely have meaningful names. Only 56.9% of responsive hosts have a hostname at all. For responsive hosts with hostnames, 66.9%

do not have meaningful hostnames. Because hostnames are rarely meaningful, our study shows that DNS, contrary to prior belief, is a poor indication of whether or not a host is active.

Our study leads to a new area of Internet measurement – a security study of the Internet, including how security defense mechanisms are deployed and configured on the Internet. Such a study provides important information for vulnerability assessment and assists in designing new defense mechanisms. We hope that our study will encourage new research and efforts in security measurement. We plan to continue our work to monitor how security-relavant behaviors on the Internet change over time.

References

[1] Lee Bain and Max Engelhardt. *Introduction to Probability and Mathematical Statistics*. PWS-Kent Publishing Company, 1991.

[2] Hal Burch. Scalable IP alias resolution. Unpublished draft.

[3] J. D. Case, M. Fedor, M. L. Shoffstall, and C. Davin. Simple network management protocol. RFC 1157, May. 1990.

[4] Bill Cheswick. The design of a secure Internet gateway. In *USENIX Summer Conference*, pages 233–237, Anaheim, California, 1990.

[5] R. Clements. Who talks ICMP, too? – survey of 18 february 83. RFC 844, Feb. 1983.

[6] D. W. Dodds. August, 1974, survey of new-protocol telnet servers. RFC 701, Aug. 1974.

[7] D. W. Dodds. November, 1974, survey of new-protocol telnet servers. RFC 669, Dec. 1974.

[8] D. W. Dodds. September, 1974, survey of new-protocol telnet servers. RFC 702, Sep. 1974.

[9] D. W. Dodds. February, 1975, survey of new-protocol telnet servers. RFC 669, Feb. 1975.

[10] D. W. Dodds. July, 1975, survey of new-protocol telnet servers. RFC 703, Jul. 1975.

[11] R. Fielding, J. Gettys, J. Mogul, H. Frystk, L. Masinter, P. Leach, and T. Berners-Lee. Hypetertext transfer protocol – HTTP/1.1. RFC 2616, Jun. 1999.

[12] Gnutella. http://gnutella.wega.com.

[13] NetBIOS Working Group in DARPA, Internet Activities Board, End-to End Services Task Force. Protocol standard for a NetBIOS service on a TCP/UDP transport: Concepts and methods. RFC 1001, Mar. 1987.

[14] ISC. ISC internet domain survey. http://www.isc.org/index.pl?/ops/ds/, 1998 - 2004.

[15] V. Jacobson. traceroute. ftp://ftp.ee.lbl.gov/traceroute.tar.Z, 1989.

[16] M. St Johns. Identification protocol. RFC 1413, Feb. 1993.

[17] B. Kantor and P. Lapsley. Network news transfer protocol. RFC 977, Feb. 1986.

[18] Kazaa. http://www.kazaa.com.

[19] J. Klensin. Simple mail transfer protocol. RFC 2821, Apr. 2001.

[20] P. D. Lebling. Survey of FTP mail and MLFL. RFC 751, Dec. 1978.

[21] Netcraft LTD. Netcraft web server survey. http://www.netcraft.com.

[22] P. Mockapetris. Domain names – implementation and specification. RFC 1035, Nov. 1987.

[23] J. Myers and M. Rose. Post office protocol - version 3. RFC 1939, May. 1996.

[24] Netsizer. http://www.netsizer.com, 2000.

[25] J. Postel. Internet control message protocol. RFC 792, Sept. 1981.

[26] J. Postel. Internet protocol. RFC 791, Sept. 1981.

[27] J. Postel. Character generator protocol. RFC 864, May. 1983.

[28] J. Postel. Echo protocol. RFC 862, May. 1983.

[29] J. Postel. Telnet protocol specification. RFC 854, May. 1983.

[30] J. Postel. File transfer protocol. RFC 959, Oct. 1985.

[31] Niels Provos and Peter Honeyman. Scanssh - scanning the internet for ssh servers. UMich Technical Report CITI-TR-01-13, Oct. 2001.

[32] RIPE. Ripe region hostcount. http://www.ripe.net/ripencc/pub-services/stats/hostcount/-index.html.

[33] D. Smallberg. Survey of smtp implementations. RFC 876, Sep. 1983.

[34] Eugene H. Spafford. The internet worm program: An analysis. Purdue Technical Report CSD-TR-823, Nov. 1988.

[35] A. Westine, D. Smallberg, and J. Postel. Summary of smallberg surveys. RFC 846, Feb. 1983.

[36] T. Ylönen. Secure login connections over the internet. In *Proceedings of 6th USENIX Security Symposium*, pages 37–42, Jul. 1996.

[37] D. Zimmerman. The finger user information protocol. RFC 1288, Dec. 1991.