SSAC Advisory SAC008
DNS Distributed Denial of Service (DDoS) Attacks

A Report from the ICANN
Security and Stability
Advisory Committee
(SSAC)
March 2006

Executive Summary

In early February 2006, name servers hosting Top Level Domain zones were the repeated recipients of extraordinary heavy traffic loads. Analysis of traffic by TLD name server operators and security experts at large confirmed that DNS packets comprising the attack traffic exhibited characteristics associated with previously attempted DDoS attacks collectively known as *amplification attacks*.

This advisory describes representative incidents, identifies the impacts, and recommends countermeasures that TLD name server operators can employ for immediate and long-term relief from the harmful effects of these attacks. Certain countermeasures may adversely affect legitimately operated domain name resolvers whose configurations contribute to the success of DDoS attacks; specifically, by operating in the manner they do, some resolvers facilitate DNS amplification attacks. Countermeasure that name server operator might implement to assist in their timely restoration of normal service could also adversely affect name server operators who rely on the service they provide. TLD operators may need to take specific measures to assure they do not worsen the effects of the attacks.

Respected security organizations and advisory groups worldwide encourage name server operators to adopt measures to disable open recursive service and to protect their infrastructures against DDoS attacks. SSAC joins these organizations and makes the following recommendations:

Recommendation (1): For the long term, SSAC recommends that the most effective means of mitigating the effects of this and numerous DoS attacks is to adopt source IP address verification.

Recommendation (2): SSAC specifically recommends that each ROOT and TLD name server operator should:

i. Document operational policies relating to countermeasures it will implement to protect its name server infrastructure against attacks that threaten its ability to offer service, give notice when such measures are implemented, and identify the actions affected parties must take to have the measures terminated.

ii. Respond faithfully and without undue delay to all questions and complaints about unanswered traffic, and

iii. Act with haste to restore service to any blocked IP address if the owner of that IP address can demonstrate that it has secured its infrastructure against the attack.

Recommendation (3): SSAC recommends that name server operators and Internet Service Providers consider the possible remedies described in Section 3 of this Advisory. In particular, SSAC urges name server operators and ISPs to disable open recursion on name servers from external sources and only accept DNS queries from trusted sources to assist in reducing amplification vectors for DNS DDoS attacks.

Table of Contents

1 Problem Description

On Sunday 5 February 2006, from 20:44 through 20:58 GMT, name servers operated by a key TLD name server operator received an average of 60 Mbps of traffic subsequently classified as attack traffic at each interface of every public name server node it operates. The aggregated attack traffic received was later determined to be approximately 1 Gigabit per second. Traffic analysis during the attack period showed that the operator was receiving abnormally large UDP messages (in excess of 1500 bytes), resulting in IP packet fragmentation.

The TLD name server operator employs a “screen and choke” firewall arrangement to protect its name server infrastructure. In this configuration, the screening firewall devices (routers) block and silently discard malformed traffic and traffic for unsupported services. During the attack period, technical staff observed that the screening devices were discarding an extraordinary number of IP packets. A closer inspection showed that these IP packets contained second and subsequent fragments of multi-packet Extended DNS (EDNS0) messages, conveyed in UDP datagrams. The initial IP fragments of these UDP messages, with destination ports set to 53/UDP, were allowed to pass through the screening routers, but were blocked and discarded by the “choke” firewall, the name server operator's second line of defense.

Block and discard actions partly reduced the traffic load at each screening router by about 30 Mbps; however, in excess of 30 Mbps of DNS traffic satisfied access control rules and were forwarded on to the choke firewall. Quick calculations showed that these were a combination of normal queries of less than 100 bytes and attack packets of 1500 bytes, resulting in an average of 1450 bytes per packet. Thus, even after the first line of defense discarded approximately half of the incoming traffic, about 96.4% of the packets forwarded to the choke firewall were attack packets. Since these packets were much larger than typical DNS query packets, they represented about 99.7% of the traffic.

Under normal conditions, traffic exhibits packet sizes averaging less than 100 bytes. Because there are currently no authoritative TLD zones that would require query sizes of greater than 512 bytes, the name server operator enforces a firewall rule to discard UDP/53 packets in excess of 512 bytes plus header, and the attack traffic was thus dropped.

By 20:58 GMT, traffic had dropped to normal levels.

Two days later (7 February 2006), at 23:54 GMT and continuing through 00:08 on 8 February 2006, the TLD name server operator was again attacked. Analysis by technical staff revealed that this attack followed the same pattern and appeared to be the same exploit attempt by the same attackers. The attack spanned all of the name server operator’s nodes and IP addresses. However, the second attack fully saturated two-thirds of the name server operator’s access circuits and partially saturated the remaining circuits, which provide an aggregate bandwidth of 2.4 Gbps (Gigabits per second). The name server operator’s three global transit providers have since provided data showing that each provider carried in excess of 2.5 Gbps of attack traffic during the 14 minute attack period.

1.1 DNS amplification attack

The attacks against TLD name servers are all forms of a DNS amplification attack [1, 2]. This type of attack uses IP address spoofing, a type of impersonation technique, where the attacker transmits packets with a forged source IP address rather than its own. The attack also exploits DNS servers that allow open recursion. *Recursion* is a method of processing a DNS query in which a name server pursues the query for a client at (typically) the authoritative name server for the name. When recursion is performed for any client as opposed to a trusted set of clients, a name server is said to be an *open recursive server*. Finally, the attack uses *amplification,* where the attacker sends a small request with the expectation of invoking a much larger response.

In the DNS attacks, the amplification component of the attack uses a recent extension to the DNS protocol, EDNS0 [3, 4]. First, the attacker composes a DNS query for a resource record that he knows will evoke a response that is significantly larger than the request. There are many ways for the attacker to know the size of this resource record in advance; for example, the attacker may have previously attacked and compromised a name server and has modified this server's zone file to include the amplification resource record.

Next, the attacker gathers a list of open recursive name servers that will recursively query for, and then return the amplification record he created. Even a list of known name servers may be sufficient. Anecdotal evidence suggests that there is a 75% chance that any known name server is also an open resolver, so a copy of a TLD zone file may suffice.

The attacker also needs a large number of attack sources. DoS attackers commonly use *botnets*, a collection of hosts that have been compromised, typically by email-borne worms [5]. Once a worm infects a host, it can be programmed to install software agents that an attacker can remotely control. Remote software agents can be directed to initiate a DoS attack: such agents are called *zombies*. In the DNS attacks, IP spoofing enables the attacker to redirect DNS response messages resulting from DNS queries made by the attack sources comprising the botnet to the targeted name server.

In the observed cases, very large numbers of very large UDP messages containing DNS response messages were delivered to a targeted name server infrastructure. The DNS servers answered every query they received, but the name server's operators' communications infrastructures were overloaded.

Imagine that an attacker decides to attack a name server at 10.10.1.1. He compromises and installs DNS DDoS agents on a large number of hosts. The attacker creates a large (4000 byte [6]) DNS TXT resource record in a zone file on a compromised authoritative name server: this is the amplification record. The attacker scans for and compiles a list of open recursive DNS servers that will recursively query for, and then return the amplification record he created. The response will be delivered as IP fragments that can be reassembled into the original TXT DNS record.

The attacker now directs his zombie agents to send DNS queries requesting the amplification record to open recursive servers. The zombies write the IP address of the targeted name server (10.10.1.1) in the source address field of the IP header used to transmit the DNS request message and set the port numbers in the UDP encapsulation to 53/DNS. (Other spoofing variants are possible. The spoofed IP address can be any potential target, i.e., other IP addresses in the target's network block, upstream devices along the route to the target, e.g., intermediate ISPs, router interfaces, and name servers. The attacker may spoof IP addresses across several blocks but still within the target's path to make filtering more difficult.

Many networks do not employ appropriate (if any) forms of source IP address validation [7, 8, 9], so large numbers of these spoofed DNS requests will be delivered to the open recursive servers. If the open recursive DNS servers have not received a prior request for this record and do not already hold the amplification record in their cache, they will issue a DNS request message of their own to the compromised authoritative server to retrieve it. The open recursive DNS servers will then compose responses to the spoofed DNS queries. If the open recursive servers have cached the amplification record, the attack is even more efficient: subsequent requests for the record are answered without additional load on the name server that is hosting the corrupted zone file. In either case, the purpose of IP spoofing is fulfilled: the large DNS response messages will be forwarded to the target name server 10.10.1.1, to port 53 UDP rather than being returned to the attack sources that originated the requests.

The targeted name server at 10.10.1.1 never actually issued any DNS record queries but now receives a torrent of DNS query responses to the spoofed requests. Because the approximate 4000 byte DNS response message exceeds the maximum (Ethernet) transmission unit, it will be fragmented into 3 IP datagrams.

A trace of IP packets representative of the traffic received by a targeted name server follows:

```
     IP (tos 0x0, ttl  45, id 13325, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  46, id 13731, offset 1480, flags [+],
     length: 1500) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  45, id 13919, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  45, id 14050, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  45, id 14277, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  45, id 14473, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  45, id 14494, offset 2960, flags
     [none], length: 1119) 10.1.1.1  192.168.1.1: udp
     IP (tos 0x0, ttl  46, id 14527, offset 0, flags [+],
     length: 1500) 10.1.1.1.53  192.168.1.1.31753:  49252 1/2/1
     DOMAIN.WITH.4k.TXT. TXT[|domain]
```

The presence of non-zero values in the offset fields in the IP packet trace illustrates that the DNS response message doesn't fit inside a standard Ethernet frame of 1500 bytes, so it is broken into multiple IP packet fragments. Forcing reassembly at the targeted name server increases the processing at the target and enhances the deception: since the response spans many fragments it may not be immediately apparent that the attack is DNS-based.

This DDoS attack is most effective when launched via a large number of open recursive servers: distribution increases the traffic and decreases the focus on the sources of the attack. The impact on the misused open recursive servers is generally low, but the effect on the target is high. The amplification factor is estimated at 1:73. Attacks based on this method have exceeded seven (7) Gigabits per second.

Figure 1.1 depicts the DNS DDoS attack.

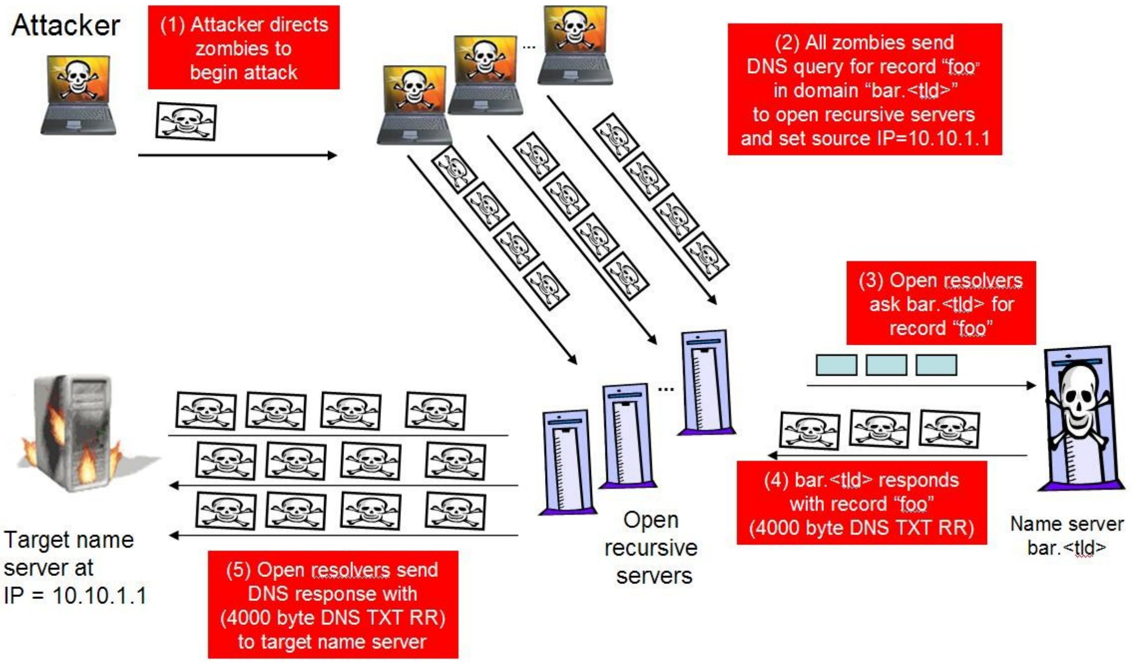
Attacker
(1) Attacker directs zombies to begin attack
(2) All zombies send DNS query for record "foo" in domain "bar.<tld>" to open recursive servers and set source IP=10.10.1.1
(3) Open resolvers ask bar.<tld> for record "foo"
(4) bar.<tld> responds with record "foo" (4000 byte DNS TXT RR)
Name server bar.<tld>
Open recursive servers
(5) Open resolvers send DNS response with (4000 byte DNS TXT RR) to target name server
Target name server at IP = 10.10.1.1

2 Impact

The RIPE NCC *dnsmon* service continuously measures the responsiveness of high level DNS name servers around the world from more than sixty locations across the Internet, predominantly in Europe. The principle of *dnsmon* is that of active measurements. *dnsmon* probes send normal DNS queries to the target name servers and note whether a correct reply is received as well as the time it took for the reply to arrive. In essence dnsmon acts as a DNS client that uses the name server in question.

Because *dnsmon* actively measures reply accuracy and response times from many locations across the Internet at the same time, it is easy to tell whether any problem is local to the probe or close to the name server in question. If most probes observe service degradation at the same time, it is reasonable to conclude that the problem lies close to the name server.

The figures in Sections 2.1 and 2.2 help illustrate the impact of the DDoS attacks described in Section 1. In the figures, individual name servers, generically identified as tld1, tld2 ..., are represented on the y-axis numerically (0-6). The graphs show the number of times the name servers did not answer a query even after a retry. Measurements are taken from more than 60 probes; figure 2.1 is a representation of more than 40,000 individual measurements. The graphs are color-coded from green to red. Green represents normally answered queries; yellow bars represent unanswered queries with a full yellow bar denoting 50% unanswered queries; the full bar is colored orange when more than two thirds of the queries remained unanswered and red when more than 90% of the queries are unanswered. These graphs depict averages of averages and as such are only useful as a general indication; graphs showing all data for a particular server are used to analyze particular events. These are not reproduced here; they available on the *dnsmon* web site [10, 11, 12].

Figure 2.3 in Section 2.3 illustrates traffic load at the thirteen root name servers during a DDoS incident. In the figure, individual root name servers are identified and are represented on the y-axis numerically (0-13). The color-coding and interpretation of the vertical and horizontal patterns are the same for Figure 2.3 as Figures 2.1 and 2.2.

2.1 February 5th Attack

Figure 2.1 shows that four out of six servers that are being probed by *dnsmon* were significantly unresponsive to most probes during this attack. The detailed graphs show that for some servers this has a strong dependency on the location of the probes. This suggests that the attack overloaded some parts of the infrastructure significantly more than others.

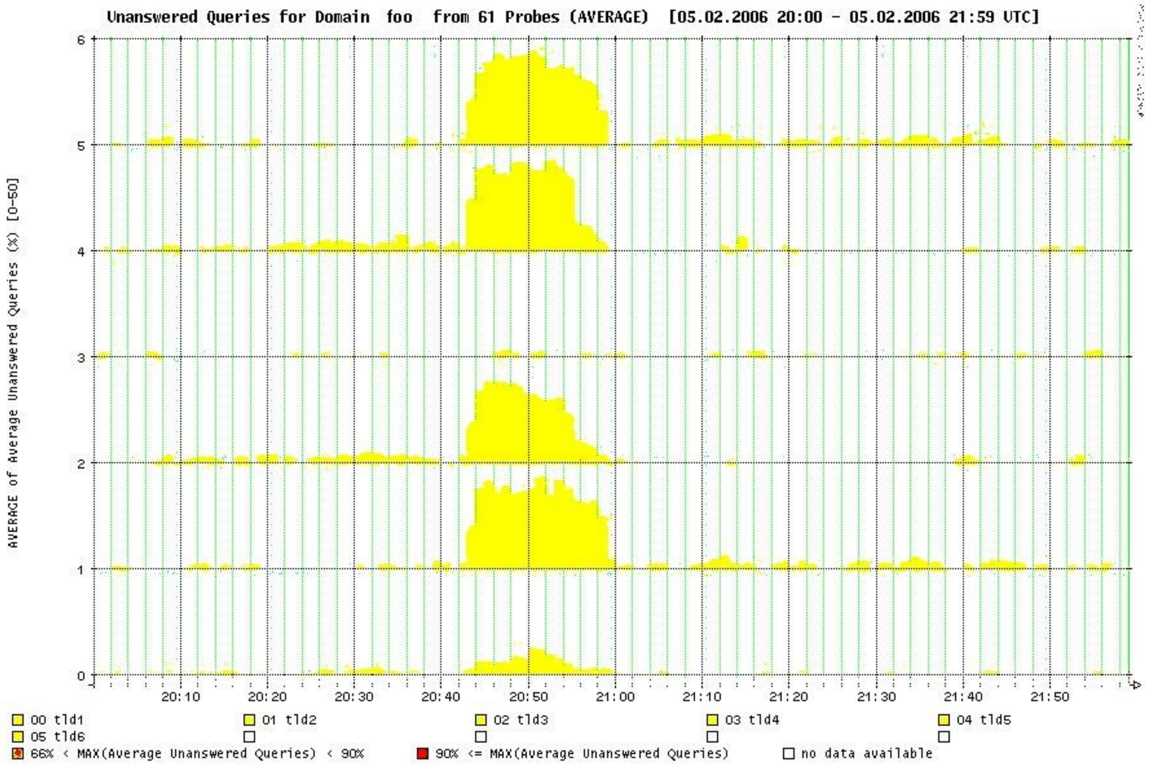

Figure 2.1. Unanswered Queries for gTLD during 5 February 2006 Attack

2.2 February 7th and 8th Attacks

Figure 2.2 shows a more severe degradation of service. This is consistent with reports from the target of a much higher level of attack traffic. The figure shows that all servers for the TLD that were being probed by *dnsmon* were significantly affected. Two servers were practically taken out of service, two others were unavailable from almost all probes and the service level of the remaining two was degraded significantly. This attack was certainly visible to some users and would have been very noticeable if sustained for a longer time.

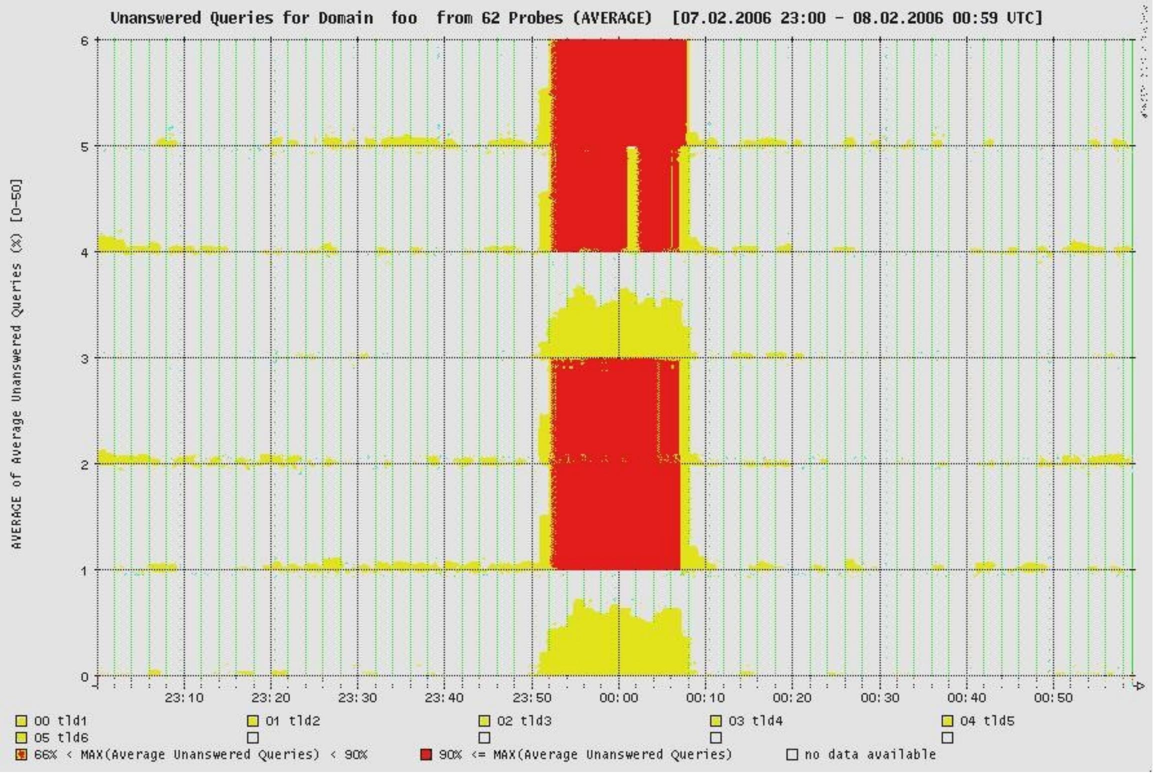

Figure 2.2. Unanswered Queries for gTLD during 7-8 February 2006 Attacks

2.3 February 15th Probes of root name servers

On 15 February 2006, probes against three of the 13 root name servers took out one server, degraded another and was noticeable on a third. Analysis by the operators of the two most affected servers suggests that this was indeed an attack using the same reflective amplification as the one described in section 1.

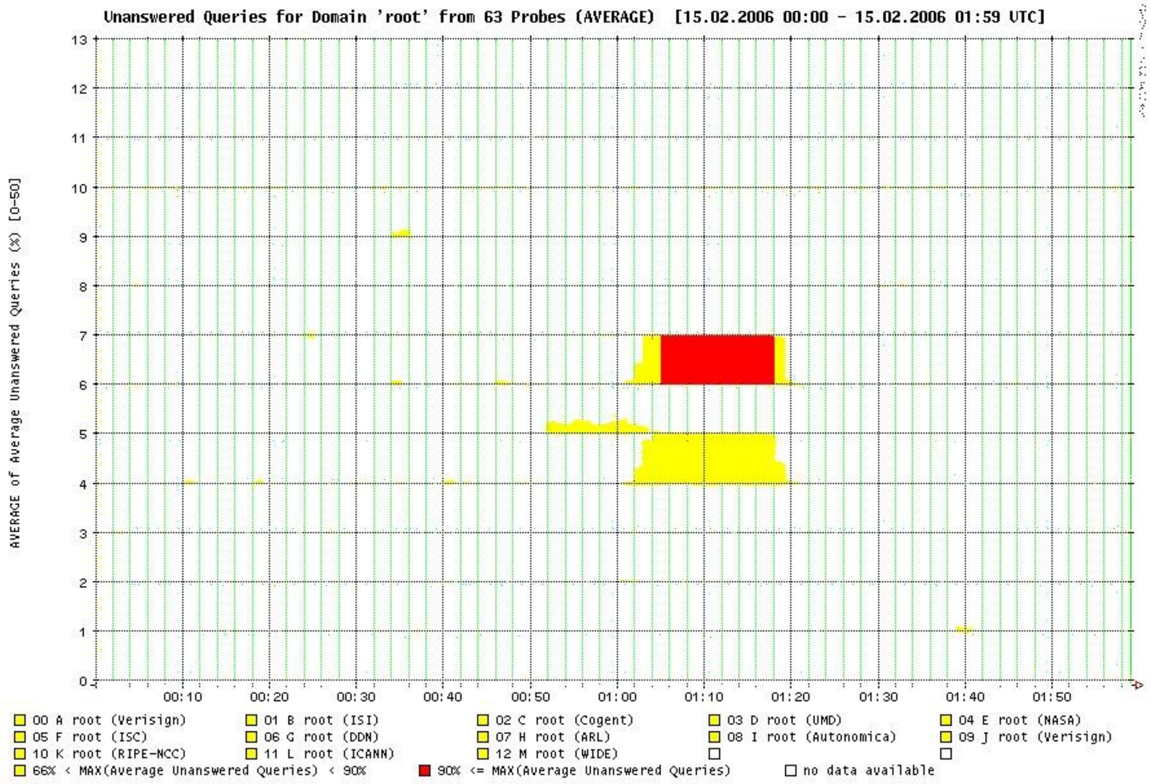

Figure 2.3. Probes on Root Servers

If we assume that the attack on F-root was of the same strength as the others, it is possible to conclude that massive anycasting is indeed a good defense against such attacks. Insufficient data are currently available to more than speculate at this time. It is also impossible to surmise why other root servers were not attacked.

3 Possible Remedies

Some measures name server operators and ISPs can implement to reduce or mitigate DDoS amplification attacks are enumerated below. It is important to note that these measures are useful for all name server operators and ISPs and not solely applicable to victims of DDoS because these and similar attacks leverage easily exploitable third parties all over the Internet.

3.1.1 Source address validation

The single action that would most significantly mitigate the effects of the kind of attack described in this document, as well as other attacks that make use of IP address spoofing, is for all network service providers to perform source IP address verification at the edge (see BCP 38, RFC 2827, *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing* [7] and SSAC004, *Securing The Edge* [13]).

3.1.2 Securely configure DNS Servers

Operating securely configured DNS application services on servers running securely configured operating systems reduces the number of servers attackers can exploit through DNS cache poisoning and privilege escalation and other system compromise attacks. This reduces the number of innocent systems that can host large DNS resource records used as amplification domain records. [14] identifies a number of secure configuration considerations.

3.1.3 Disable Open Recursive DNS

Disabling open recursion on name servers from external sources and only accepting recursive DNS from trusted sources greatly reduces the amplification vector. Available data indicate that an overwhelming percentage of DNS servers operate as open recursive servers. The Measurement Factory [6] reports that over 75% of domain name servers of roughly 1.3 million sampled allow recursive name service to arbitrary querying sources. This opens a name server to both cache poisoning and denial of service attacks. [15] and [16] describe how to disable open recursive DNS on Windows 2003 Server and BIND implementations.

3.1.4 Implement blocking and filtering

SSAC believes that name server operators and ISPs have a duty to protect the security and stability of the domain name infrastructure, and the responsibility and authority to do so, providing they

1. Document and make public the measures they will take in response to attacks,
2. Make reasonable efforts to ensure that parties whose servers are vectors for such attacks have been notified that protective measures are pending, and that

Collateral damage caused when measures are put into effect is minimized.

TLD name server operators and network operators/ISPs may consider the following measures to further defend against DNS amplification attacks.

- Block invalid DNS messages at the network edge. This includes blocking IP packets carrying UDP messages exceeding the standard 512 bytes issued to or from the DNS port (53). TLD name server operators should recognize that as the profile of normal, legitimate traffic changes, the deployment of future protocol extensions and enhancements may require changes to this filtering rule.
- Rate limit traffic sources.
- In extreme cases, prohibit or block queries from open recursive servers that are being utilized by attackers to reflect spoofed query answers to the TLD servers until those open recursive servers are reconfigured to not allow such attacks.
- Apply anti-spoofing filters to prevent the attacks from originating from networks you administer or from your customers' networks.
- Scan to detect open recursive servers that are operating on networks you administer and disable these. Scan your customers' networks (if applicable) as well, and encourage customers to disable open recursive service.
- Use Unicast Reverse Path Forwarding to mitigate problems that are caused by malformed or forged IP source addresses that are passing through a router [17].

4 Recommendations

Respected security organizations and advisory groups worldwide [1, 18] encourage name server operators to adopt measures to disable open recursive DNS and to protect their infrastructures against DDoS attacks. SSAC joins these organizations and makes the following recommendations:

Recommendation (1): For the long term, SSAC recommends that the most effective means of mitigating the effects of this and numerous DoS attacks is to adopt source IP address verification.

Recommendation (2): SSAC specifically recommends that each ROOT and TLD name server operator should:

i. Document operational policies relating to countermeasures it will implement to protect its name server infrastructure against attacks that threaten its ability to offer service, give notice when such measures are implemented, and identify the actions affected parties must take to have the measures terminated.

ii. Respond faithfully and without undue delay to all questions and complaints about unanswered traffic, and

iii. Act with haste to restore service to any blocked IP address if the owner of that IP address can demonstrate that it has secured its infrastructure against the attack.

Recommendation (3): SSAC recommends that name server operators and Internet Service Providers consider the possible remedies described in Section 3 of this Advisory. In particular, SSAC urges name server operators and ISPs to disable open recursion on name servers from external sources and only accept DNS queries from trusted sources to assist in reducing amplification vectors for DNS DDoS attacks.

5 References

[1] The Continuing Denial of Service Threat Posed by DNS Recursion http://www.us-cert.gov/reading_room/DNS-recursion121605.pdf

[2] Distributed Denial of Service attacks and their defenses http://www.lancs.ac.uk/postgrad/pissias/netsec/links/index.html

[3] Extension Mechanisms for DNS (EDNS0), http://www.rfc-editor.org/rfc/rfc2671.txt

[4] RFC 3226, DNSSECand IPv6 A6 aware server/resolver message size requirements, http://www.ietf.org/rfc/rfc3226.txt

[5] Distributed Denial of Service (DDoS) Attacks/tools, http://staff.washington.edu/dittrich/misc/ddos/

[6] DNS Amplification Attacks, http://www.isotf.org/news/DNS-Amplification-Attacks.pdf

[7] BCP 38, RFC 2827, Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, http://www.ietf.org/rfc/rfc2827.txt

[8] Three Practical Ways to Improve Your Network, Kevin Miller, *CMU* http://www.usenix.org/publications/library/proceedings/lisa03/tech/miller.html

[9] Firewall Best Practices - Egress Traffic Filtering, David Piscitello and Nathan Buff http://hhi.corecom.com/egresstrafficfiltering.htm

[10] *dnsmon* at RIPE NCC, http://dnsmon.ripe.net/

[11] *dnsmon* Domain View at RIPE NCC, http://dnsmon.ripe.net/dns-servmon/domain/information

[12] *dnsmon* Server View, http://dnsmon.ripe.net/dns-servmon/server/information

[13] SAC004, Securing The Edge, http://www.icann.org/committees/security/sac004.txt

[14] Fixing Open DNS Servers, http://www.dnsstuff.com/info/opendns.htm

[15] Disable recursion on the Windows 2003 Server DNS service, http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/ServerHelp/e1fe9dff-e87b-44ae-ac82-8e76d19d9c37.mspx

[16] Secure Bind Template http://www.cymru.com/Documents/secure-bind-template.html

[17] Unicast Reverse Path Forwarding enhancements for the Internet Service Provider http://www.cisco.com/warp/public/732/Tech/security/docs/urpf.pdf

[18] SANS Twenty Most Critical Internet Security Vulnerabilities: DNS Vulnerabilities, http://www.sans.org/top20/#c6

Appendix A
Contributors

Alain Aina, Consultant

Jaap Akkerhuis, NLnet Labs

KC Claffy, CAIDA

Steve Crocker, Shinkuro (Chairman)

Daniel Karrenberg, RIPE/NCC

Johan Ihrén, Autonomica

Rodney Joffe, UltraDNS

Mark Kosters, Verisign

Allison Mankin, Consultant

Ram Mohan, Afilias

Russ Mundy, SPARTA, Inc

Frederico Neves, registro.br

Jon Peterson, NeuStar

David Piscitello, ICANN SSAC Fellow

Ray Plzak, ARIN

Mike St. Johns, Nominum

Doron Shikmoni, ForeScout, ISOC-IL

Bruce Tonkin, Melbourne IT; Chairman, Generic Names Supporting Organization

Paul Vixie, ISC

Suzanne Woolf, ISC