

DNS Amplification Attacks

Vaughn & Evron

2005-11-01 22:56:05.097681 IP (tos 0x0, ttl 52, id 16855, offset 0, flags [+], length: 1500)53 >81.198.48839: 42684 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.097786 IP (tos 0x0, ttl 50, id 43395, offset 0, flags [+], length: 1500)53 >81.2.63170: 26885 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.097800 IP (tos 0x0, ttl 56, id 55748, offset 0, flags [+], length: 1500)53 >81.17.20879: 44337 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.097854 IP (tos 0x0, ttl 56, id 29724, offset 0, flags [+], length: 1500)53 >81.53.27800: 63317 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.097976 IP (tos 0x0, ttl 51, id 51511, offset 0, flags [+], length: 1500)53 >81.66.28903: 57713 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098119 IP (tos 0x0, ttl 50, id 62401, offset 0, flags [+], DF], length: 1500)53 >81.191.53122: 33503 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098155 IP (tos 0x0, ttl 52, id 64104, offset 0, flags [+], length: 1500)53 >81.41.24793: 47330 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098330 IP (tos 0x0, ttl 46, id 3664, offset 0, flags [+], length: 1500)53 >81.211.8611: 47954 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098372 IP (tos 0x0, ttl 56, id 55745, offset 0, flags [+], length: 1500)53 >81.13.42959: 46611 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098386 IP (tos 0x0, ttl 50, id 30744, offset 0, flags [+], length: 1500)53 >81.63.55986: 31023 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098390 IP (tos 0x0, ttl 57, id 9044, offset 0, flags [+], length: 1500)53 >81.141.3515: 49050 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098435 IP (tos 0x0, ttl 52, id 18930, offset 0, flags [+], length: 1500)53 >81.154.17100: 32920 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098575 IP (tos 0x0, ttl 56, id 10261, offset 0, flags [+], length: 1500)53 >81.129.14305: 59559 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098596 IP (tos 0x0, ttl 46, id 40892, offset 0, flags [+], length: 1500)53 >81.242.49588: 12522 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098600 IP (tos 0x0, ttl 38, id 58544, offset 0, flags [+], length: 1500)53 >81.15.34671: 47017 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098665 IP (tos 0x0, ttl 53, id 29641, offset 0, flags [+], length: 1500)53 >81.235.37975: 17249 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098688 IP (tos 0x0, ttl 48, id 1884, offset 0, flags [+], length: 1500)53 >81.75.33947: 19706 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098767 IP (tos 0x0, ttl 49, id 34677, offset 0, flags [+], length: 1500)53 >81.56.63871: 55960 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098775 IP (tos 0x0, ttl 46, id 12769, offset 0, flags [+], length: 1500)53 >81.152.2405: 53129 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098800 IP (tos 0x0, ttl 51, id 46446, offset 0, flags [+], length: 1500)53 >81.87.40717: 61969 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098854 IP (tos 0x0, ttl 48, id 51236, offset 0, flags [+], length: 1500)53 >81.224.64893: 64302 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098888 IP (tos 0x0, ttl 54, id 59295, offset 0, flags [+], DF], length: 1500)53 >81.95.62365: 62957 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098928 IP (tos 0x0, ttl 47, id 16978, offset 0, flags [+], length: 1500)53 >81.67.7364: 63079 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.098956 IP (tos 0x0, ttl 48, id 61759, offset 0, flags [+], length: 1500)53 >81.57.51718: 23794 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099068 IP (tos 0x0, ttl 47, id 42163, offset 0, flags [+], length: 1500)53 >81.120.40134: 25812 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099096 IP (tos 0x0, ttl 46, id 20111, offset 0, flags [+], length: 1500)53 >81.77.14744: 28858 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099100 IP (tos 0x0, ttl 49, id 56186, offset 0, flags [+], length: 1500)53 >81.75.27962: 17931 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099135 IP (tos 0x0, ttl 46, id 20887, offset 0, flags [+], length: 1500)53 >81.125.50722: 54041 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099217 IP (tos 0x0, ttl 47, id 47286, offset 0, flags [+], length: 1500)53 >81.222.24260: 8686 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099423 IP (tos 0x0, ttl 49, id 32788, offset 0, flags [+], length: 1500)53 >81.13.13267: 5990 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099509 IP (tos 0x0, ttl 55, id 59656, offset 0, flags [+], length: 1500)53 >81.12.61916: 27493 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099564 IP (tos 0x0, ttl 52, id 6780, offset 0, flags [+], length: 1500)53 >81.105.60774: 15646 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099584 IP (tos 0x0, ttl 53, id 18124, offset 0, flags [+], length: 1500)53 >81.56.51438: 18284 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099624 IP (tos 0x0, ttl 48, id 13089, offset 0, flags [+], length: 1500)53 >81.176.48381: 26259 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099657 IP (tos 0x0, ttl 55, id 20940, offset 0, flags [+], length: 1500)53 >81.67.850: 39904 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099663 IP (tos 0x0, ttl 39, id 35434, offset 0, flags [+], length: 1500)53 >81.16.10535: 44987 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099687 IP (tos 0x0, ttl 51, id 9536, offset 0, flags [+], length: 1500)53 >81.206.62136: 42472 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099732 IP (tos 0x0, ttl 56, id 23195, offset 0, flags [+], length: 1500)53 >81.53.50843: 21918 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099767 IP (tos 0x0, ttl 47, id 51216, offset 0, flags [+], length: 1500)53 >81.123.25447: 44605 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099867 IP (tos 0x0, ttl 52, id 11502, offset 0, flags [+], length: 1500)53 >81.60.29083: 18080 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099918 IP (tos 0x0, ttl 45, id 17913, offset 0, flags [+], length: 1500)53 >81.234.15614: 7372 1/2/1 e.....com TXT[[domain]
2005-11-01 22:56:05.099979 IP (tos 0x0, ttl 47, id 41606, offset 0, flags [+], length: 1500)53 >81.121.46332: 46093 1/2/1 e.....com TXT[[domain]

Seq	time	Src	srcport	Dst	DstPort	Length	Proto	Info
1	0.000000	____.____.____.37	53	____.____.____.46	31145	1514	DNS	Standard query response TXT[Packet size limited during capture]

Frame 1 (1514 bytes on wire, 96 bytes captured)

Ethernet II, Src: _____00:5c:db (00:00:00:00:5c:db), Dst: _____00:7d:00 (00:00:00:00:7d:00)

Internet Protocol, Src: _____.____.____.37 (_____.____.____.37), Dst: _____.____.____.46 (_____.____.____.46)

User Datagram Protocol, Src Port: domain (53), Dst Port: 31145 (31145)

Domain Name System (response)

[Packet size limited during capture: DNS truncated]

0000	00 00 00 00 7d 00 00 00 00 00 5c db 08 00 45 00}\....\...E.
0010	05 dc c7 21 20 00 37 11 79 b3 00 00 00 25 00 00	...! .7.y....%..
0020	00 2e 00 35 79 a9 0f b2 f1 66 57 3a 81 80 00 01	...5y....fW:....
0030	00 01 00 01 00 01 01 78 01 70 04 63 74 72 63 02x.p.ctr.
0040	63 63 00 00 ff 00 ff c0 0c 00 10 00 01 00 03 ee	cc.....
0050	ef 0f 5a ff 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e	..Z.....

Seq	time	Src	srcport	Dst	DstPort	Length	Proto	Info
2	0.001249	____.____.____.37		____.____.____.46		1514	IP	Fragmented IP

protocol (proto=UDP 0x11, off=1480)

Frame 2 (1514 bytes on wire, 96 bytes captured)

Ethernet II, Src: _____00:5c:db (00:00:00:00:5c:db), Dst: _____00:7d:00 (00:00:00:00:7d:00)

Internet Protocol, Src: _____.____.____.37 (_____.____.____.37), Dst: _____.____.____.46 (_____.____.____.46)

Data (62 bytes)

```

0000 00 00 00 00 7d 00 00 00 00 5c db 08 00 45 00  ....}\...E.
0010 05 dc c7 21 20 b9 37 11 78 fa 00 00 00 25 00 00  ...!.7.x....%..
0020 00 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e  .....
0030 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e  .....
0040 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e  .....
0050 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e 2e  .....

```


- The attacker constructs a query for all records of the domain and specifies a large UDP payload buffer using the RFC 2671 optional resource record (OPT RR) format.
- Captured ICMP data suggests this UDP payload buffer was as high as 10000 bytes although the actual queries may have used random UDP payload values.
- Observed packet fragment sizes suggest the effective UDP payload value was 1514 bytes which is consistent with a 1514 octet MTU and with the assumption the UDP payload size specification is geared towards producing fragmentation.
- The attacker adds the IP address of the target as the request initiator's IP address in this packet, and then repeatedly sends these queries to their established list of name servers, most likely by the use of a Botnet.
- It is very likely the attackers harvested a large number of DNS IP addresses. They do not appear to care if the name servers offer EDNS nor if they allow recursive lookups. The same effect could be accomplished with IP addresses chosen at random or with malware selecting their ISP's name servers.

- The recursive name servers respond to the request for all records with SOA RRs, Standard Query, and fragmented TXT RRs responses.
- The TXT RR response, in fragments no larger than 1514 bytes, dominates the captured packets.
- SOA RRs were 222 byte records.
- Standard Query responses were one of:
 - root referrals with packet lengths ranging from 501 bytes to 222 bytes;
 - Responses with reply code 3 (NXDOMAIN - no such name) with 132 to 141 byte records;
 - Address A-type responses with 98 byte packets predominating, or;
 - RCODE FORMERR, or SERVFAIL responses of 60, 65 or 82 byte packets.

Using a 8707 incoming packet capture from 7207 unique IP addresses in 2.44 seconds we conclude:

- Packet arrival rate of 3567 packets per second.
- The total packet size of 8.5Mbytes with an average 28Mbps.
- Effective amplification factor between 16 and 19.
- If the 7207 unique IP addresses are all of the exploited name servers, the query rate for the triggering event would be 2953 queries per second.
- This query rate is single host achievable on a connection with an upload rate between 1Mbps and 2Mbps.
- Could either be single machine or from a Botnet using the infected hosts own ISP's DNS name server.
- NANOG post (NANOG #2 Feb/24/2006) supports the malware assumption.

So how is this Amplification?

- Small ALL – type request
- Large (4k) reply
- Fragments
- Other responses

How much smarts are required to launch an attack?

1. Register a domain, host it in a DNS
2. Create a large TXT record, give it a long TTL (24h)
3. Find a router with no egress filtering. Use it to spoof the victim's IP.
4. Send a TXT query to the domain host in order to cache the TXT record. Spoof an IP, of course.
5. Spoof the victim's IP and send thousands of requests to thousands of Name Servers.

Amplification via duplicate replies

- “The amount of duplicates seems to range between 2 and 24(in steps of 2, 4, 8, 12, 16, 20 and 24), where each reply packet is roughly 2.5x(including IP header) larger than the original request(because of the SOA). So, for example one request to a DNS server that sends 24 dups back would roughly equal $60x(24*2.5)$ amplification of data.”

v9@fakehalo.us - bugtraq@securityfocus.com

- About 580 thousand "Open Resolvers" on the Internet
 - Dan Kaminsky and Mike Schiffman
 - ShmooCon (2006).
- "... an estimated 7.5 million external DNS servers on the public Internet. Over 75% of domain name servers (of roughly 1.3 million sampled) allowed recursive name service to arbitrary queriers..."

"Over 40% allow zone transfers from arbitrary queriers.."

--- The Measurement Factory

Solving this problem

- Spoofing countermeasures :
BCP 38 SAC 004
- Limit DNS recursion to local net