Serving DNS during a DDoS

freedns.afraid.org
Joshua Anderson

About the freedns project

- Free public shared DNS provider
- Started in 2001 for sharing of vanity/recreational domain names
- 2 million DNS records (A/MX/CNAME/TXT etc)
- 5.3 billion DNS queries September 2010 (2k/sec)
- 56k accounts active in last 90 days
- 7 Mbps of bandwidth consumed at peak

The thing about DNS is ...

- Infrastructure type of service
- Reliability is extremely important for user trust, outage means massive collateral damage
- Free users can be rowdy, and often disruptive to paid users experience
- 9,815 banned accounts in 2010 so far (34 per day average)

Measuring the damage

- Disruptive multi-gigabit attacks for freedns.afraid.org have occured 1-2 times a year since 2006
- Smaller attacks occur fairly regularly (<1gbps)
- If a DDoS attack succeeds at taking a site down, attackers become very encouraged, often greatly extending the attacks duration

DNS provider lacks any incoming delegation control

- After receiving hours of malicious traffic while receiving protection from provider filtering, I must ask a domain owner via email to remove domain delegation - they don't always respond + propigation times
- Presently no way to remove incoming delegation to nameservers, though effort is being made at http://dnscert.org/ to address this

Long time architecture

- 4 DNS servers, known as ns1 ns2 ns3 ns4
- Everyone (both free and paid members) use the same nameservers

- ...is 4 nameservers really not redundant enough? (It is plenty for securing against hardware failure).

Methods of dealing with attack

- Variety of countermeasures, tuned timeouts, automatic packet filtering scripts, netstat extended state watchers, http timeout watchers, automatic 15 second packet capture during traffic bursts
- None of these tools help with multi-gigabit ”valid DNS query” attacks
- Hard to find places willing to provide a ”deny of all non-port 53 traffic” rule

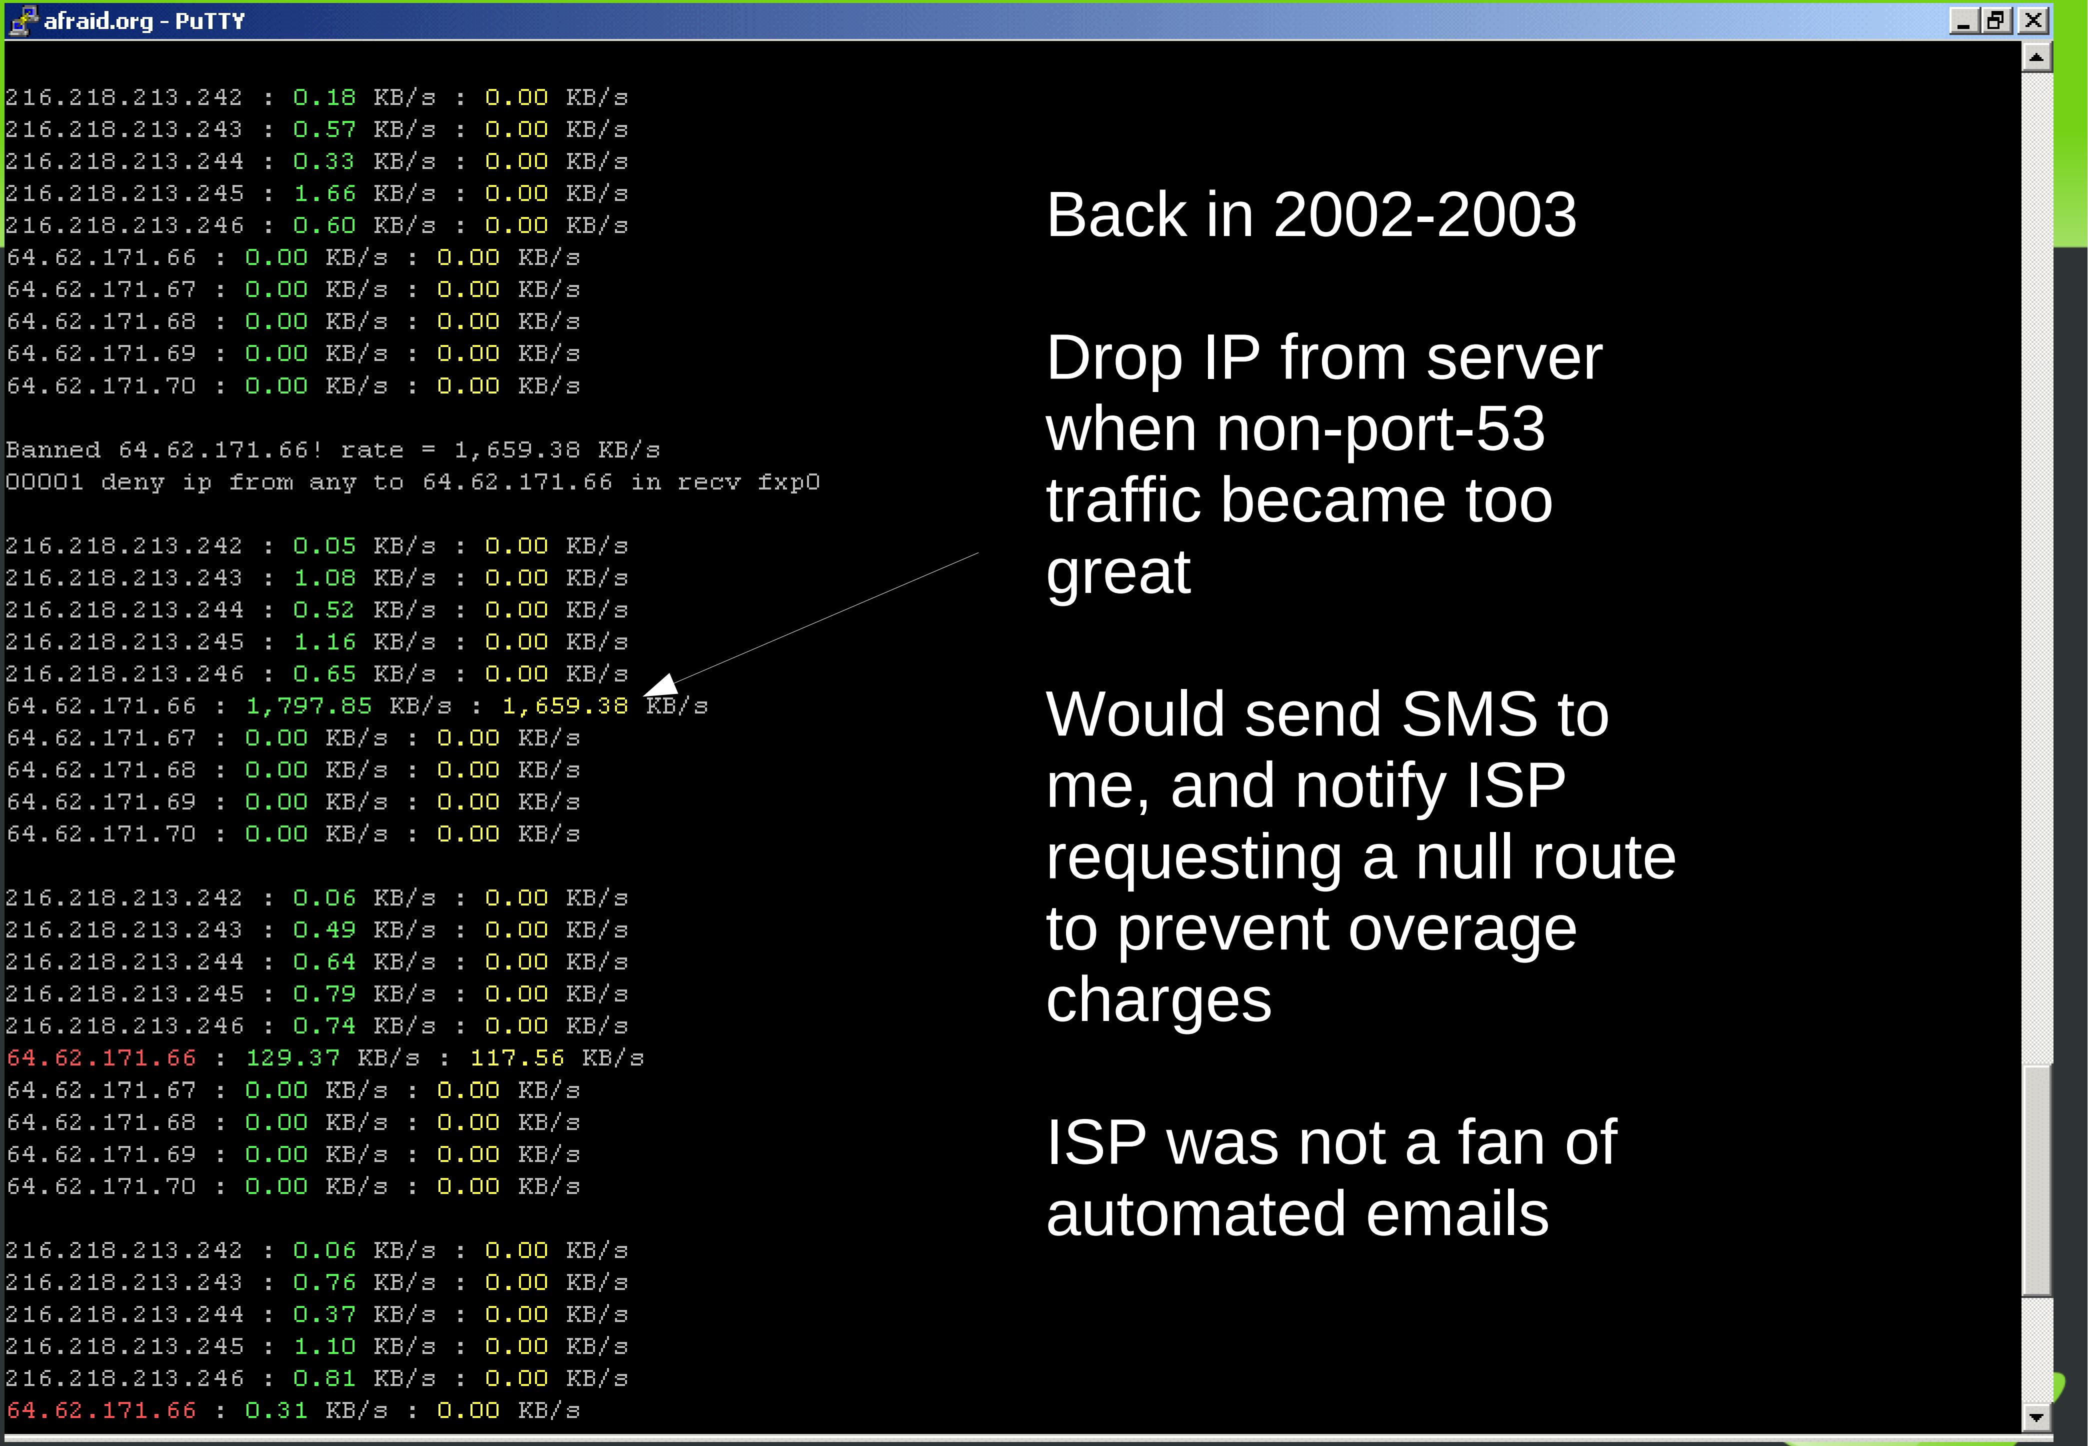

```
216.218.213.242 : 0.18 KB/s : 0.00 KB/s
216.218.213.243 : 0.57 KB/s : 0.00 KB/s
216.218.213.244 : 0.33 KB/s : 0.00 KB/s
216.218.213.245 : 1.66 KB/s : 0.00 KB/s
216.218.213.246 : 0.60 KB/s : 0.00 KB/s
64.62.171.66 : 0.00 KB/s : 0.00 KB/s
64.62.171.67 : 0.00 KB/s : 0.00 KB/s
64.62.171.68 : 0.00 KB/s : 0.00 KB/s
64.62.171.69 : 0.00 KB/s : 0.00 KB/s
64.62.171.70 : 0.00 KB/s : 0.00 KB/s

Banned 64.62.171.66! rate = 1,659.38 KB/s
00001 deny ip from any to 64.62.171.66 in recv fxp0

216.218.213.242 : 0.05 KB/s : 0.00 KB/s
216.218.213.243 : 1.08 KB/s : 0.00 KB/s
216.218.213.244 : 0.52 KB/s : 0.00 KB/s
216.218.213.245 : 1.16 KB/s : 0.00 KB/s
216.218.213.246 : 0.65 KB/s : 0.00 KB/s
64.62.171.66 : 1,797.85 KB/s : 1,659.38 KB/s
64.62.171.67 : 0.00 KB/s : 0.00 KB/s
64.62.171.68 : 0.00 KB/s : 0.00 KB/s
64.62.171.69 : 0.00 KB/s : 0.00 KB/s
64.62.171.70 : 0.00 KB/s : 0.00 KB/s

216.218.213.242 : 0.06 KB/s : 0.00 KB/s
216.218.213.243 : 0.49 KB/s : 0.00 KB/s
216.218.213.244 : 0.64 KB/s : 0.00 KB/s
216.218.213.245 : 0.79 KB/s : 0.00 KB/s
216.218.213.246 : 0.74 KB/s : 0.00 KB/s
64.62.171.66 : 129.37 KB/s : 117.56 KB/s
64.62.171.67 : 0.00 KB/s : 0.00 KB/s
64.62.171.68 : 0.00 KB/s : 0.00 KB/s
64.62.171.69 : 0.00 KB/s : 0.00 KB/s
64.62.171.70 : 0.00 KB/s : 0.00 KB/s

216.218.213.242 : 0.06 KB/s : 0.00 KB/s
216.218.213.243 : 0.76 KB/s : 0.00 KB/s
216.218.213.244 : 0.37 KB/s : 0.00 KB/s
216.218.213.245 : 1.10 KB/s : 0.00 KB/s
216.218.213.246 : 0.81 KB/s : 0.00 KB/s
64.62.171.66 : 0.31 KB/s : 0.00 KB/s
```

Back in 2002-2003

Drop IP from server when non-port-53 traffic became too great

Would send SMS to me, and notify ISP requesting a null route to prevent overage charges

ISP was not a fan of automated emails

Methods of dealing with attack (cont)

- Seek providers who offer DDoS mitigation plans to scrub out bad packets before they reach the server
- Such providers can become prohibitively expensive, unable to protect past a certain rate, or require expensive monthly commits, whether needed or not. Also not 100% reliable, still need to open a ticket and human supervision.

Methods of dealing with attack (cont)

- Adding more nameservers has diminishing returns
- 1 DNS server to 2, very helpful
- 2 DNS servers to 3, sort of helpful
- 3 DNS servers to 4, not as helpful
- 4 servers successfully covers most small attacks, only the major attacks left to solve which will take a topology change

Idea 1 of 3, move paid users to own pool

- Free users often the targets of attack
- Seperate ”free” domains from ”paid” user domains
- Problem: All members that are source of revenue become exposed in 1 target location

Idea 2 of 3, anycast!

- Anycast could thinly distribute traffic to a greater number of servers
- Question: As a consumer of 7 mbps of traffic normally, how many mostly idle anycast servers with quad port gigabit cards would it take to smoothly survive a surprise 15+ gigabit attack?
- A wonderful solution, but how expensive?

Other ideas?

- Does there exist any realtime fast deep packet inspection solution to conditionally deny certain types of valid DNS patterns before they reach the application layer of a DNS server?
- If so, then that should be combined with the following...

Idea 3 of 3, segment pools of nameservers

- Setup mini nameserver instances such as:
- Pool 1: ns1-ns32
- Pool 2: ns33-ns64
- Pool 3: ns65-ns96
- Pool 4: ns97-ns128
- Leasing IPs not more than $1/mo each
- Naming would be using named adjectives, rather then numerical digits

Nameserver segmentation continued

- Allocate authorative nameservers in a way that does not overlap with other allocations, whether free or paid.
- Leverage exponential combinations, 32x32x32x32 becomes 1048576 possible unique combinations without incoming delegation overlap
- Ordered example: 1111 1112 1113 1114 1121 1122 1123 1124 1131 1132.. (precalculated ordering to be more random)

Principle differences vs Anycast

- Makes no attempt to withstand a 15 gigabit attack to a single domain, instead it attempts to isolate the damage from other users
- Scales pretty well, paid and free members are both isolated from one another, not just paid
- Though could potentially be more confusing for users to add new domains

Closing thoughts

- Not an end all solution, should be combined with as many other mitigation solutions as possible
- Seemingly wasteful of IPs, looming IPv4 exhaustion.
- Reachable at : josh@afraid.org