Everyday Attacks Against Verisign-Operated DNS Infrastructure

Matt Weinberg and Piet Barber

May 9, 2015

What Does Verisign Do?

Core Verisign Edge Services

- Authoritative Domain Name System (DNS) for
  - .COM and .NET ~130 million domains
  - Country-Code Top-Level Domains (ccTLDs): .cc and .tv
  - Other Top-Level Domains (TLDs) including .jobs, .gov, .edu, .name and more

- One of twelve Root Server Operators
  - A-root and J-root

Core Verisign Edge Services

- Distributed Denial of Service (DDoS) Mitigation Service
- Managed DNS
- Recursive DNS

Uptime, Uptime…. and Uptime

- We must provide uninterrupted service for all DNS products

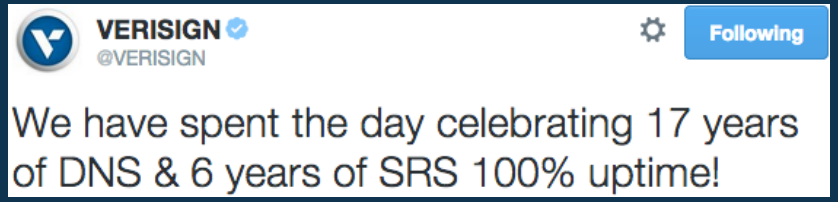

- Typical day is 110 billion DNS queries

Verisign Points of Presence

- 17 large sites at major Internet exchange points
  - Host all Verisign Edge products
  - Access via transit and private peering
- 69 (and growing!) small regional sites
  - Bring .COM/.NET and J-Root DNS closer to the user

On The Map…


Greenland
Iceland
Finland
Sweden
Norway
Russia
Canada
United States
Mexico
North Pacific Ocean
North Atlantic Ocean
Spain
Italy
Ukraine
Kazakhstan
Mongolia
China
South Korea
Turkey
Iraq
Iran
Afghanistan
Pakistan
India
Thailand
Indonesia
Papua New Guinea
Algeria
Libya
Egypt
Saudi Arabia
Mali
Niger
Sudan
Chad
Nigeria
Ethiopia
DR Congo
Tanzania
Angola
Namibia
Botswana
South Africa
Madagascar
Indian Ocean
Australia
New Zealand
Venezuela
Colombia
Peru
Bolivia
Argentina
South Pacific Ocean
South Atlantic Ocean

Mitigation Strategies

Technical Architecture Considerations

- Maximum uptime
- Ability to sustain large-scale traffic events
- Reduced latency

Mitigation Strategies

- Build Big, Build Wide
  - Advantage: Gives us a bit of breathing room
  - Disadvantage:
    - Resource-intensive
    - Risk of reflection attacks

Network Capacity

- 2+ Tbps network capacity and growing
- Dedicated backbone available at most Edge sites
  - Peering relationships with over 700 networks at 1,400 points of interconnection
  - About 80% of all network traffic delivered via peering relationships
    - Improve latency
    - Added network diversity
  - QoS and MPLS
  - BGP FlowSpec for filter deployment

DNS Server Capacity

- Massive compute capacity
- Custom in-memory database of all zone data
- Bare-metal vs. Virtualization

Tools for Mitigation

- Custom, in-house developed software for where it makes sense
  - Load Balancers
  - Name Servers
  - Filter deployment tools for both LB and NS
  - Heads-Up Display for real time monitoring
  - Linux Kernel enhancements and performance tuning

Traffic Filtering Capabilities At Multiple Tiers

- Core routers
  - ACLs, FlowSpec, QoS, MPLS TE
- Custom load balancers can filter based on:
  - Packet size, Query type, Rate limits
  - Anything we can isolate, we can filter
- Kafka/Storm cluster for real-time filter recommendations
  - SNMP shows high interface utilization
  - NetFlow shows high traffic prefixes/attacked services
  - Orchestration tools for routing policy adjustments or filter deployment

Traffic Filtering Capabilities At Multiple Tiers

- Proprietary name server software
  - Highly-tuned for the product it serves
  - Real-time reports stats for our HUD
  - Can filter on:
    - Packet size
    - Query type, RR
    - Can perform rate limits
  - Real-time visibility of filter efficacy

We’re Under Attack!

- What should we do?

DDoS Mitigation Options – DO SOMETHING!

- Filter the offending traffic
- Isolate attack traffic between sites
  - Manipulate BGP announcements
- Isolate traffic within a site
  - Send all traffic to a subset of network and/or server resources

DDoS Mitigation Options – DO SOMETHING!

- Dynamically allocate resources
  - Global bandwidth/Circuits
  - Physical sites
  - Compute resources within a site
  - Reduce or segregate resources to contain impact
- Filter at appropriate layer
  - Priority on fastest deployment
  - Move towards origin

DDoS Mitigation Options – DO NO HARM!

Some mitigation techniques can make it worse

- In corner cases, blocking all traffic = RETRIES

- Too small – no real impact

- Just let Response Rate Limiting (RRL) do its thing!

Real-World Attacks

Attack 1: Random QNAME Reflectors

- 10 Aug 2014
- (Random).www.jd7777.com/ IN / A
  (Random).www.lt8005.com / IN / A
- About 3 million qps
- Lots of source addresses, IPv4 & IPv6
  - 135,000 unique /32s within 96,800 unique /24s
  - Spot-check sources against the Open Resolver Project, 100% correlation
- Conclusion: Real name servers hitting us

Attack 1: Random QNAME Reflectors

- Why are real name servers hitting the .COM/.NET name servers?
  - The jd7777.com went NXDOMAIN moments before we saw the traffic spike
  - One nasty side effect of the random QNAME attack: It hits name servers higher in the DNS hierarchy when NXDOMAIN
  - Root servers also see Random QNAME attacks –
    - attackers made a typo for the attack query
    - e.g.: (Random).www.host.tld+(Literal Period)

Attack 1: Random QNAME Reflectors

- This chart shows before attack, after attack ramp-up (2 steps), after we activate filters
- Attack changes from one domain to a new one at about 13:00
- Once we start filtering, the caching name servers start retry storms, and traffic jumps to 14 million QPS

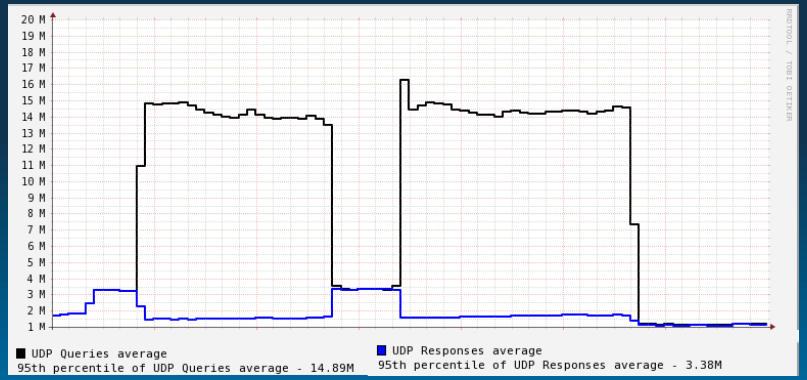

Attack 1: Random QNAME attacks

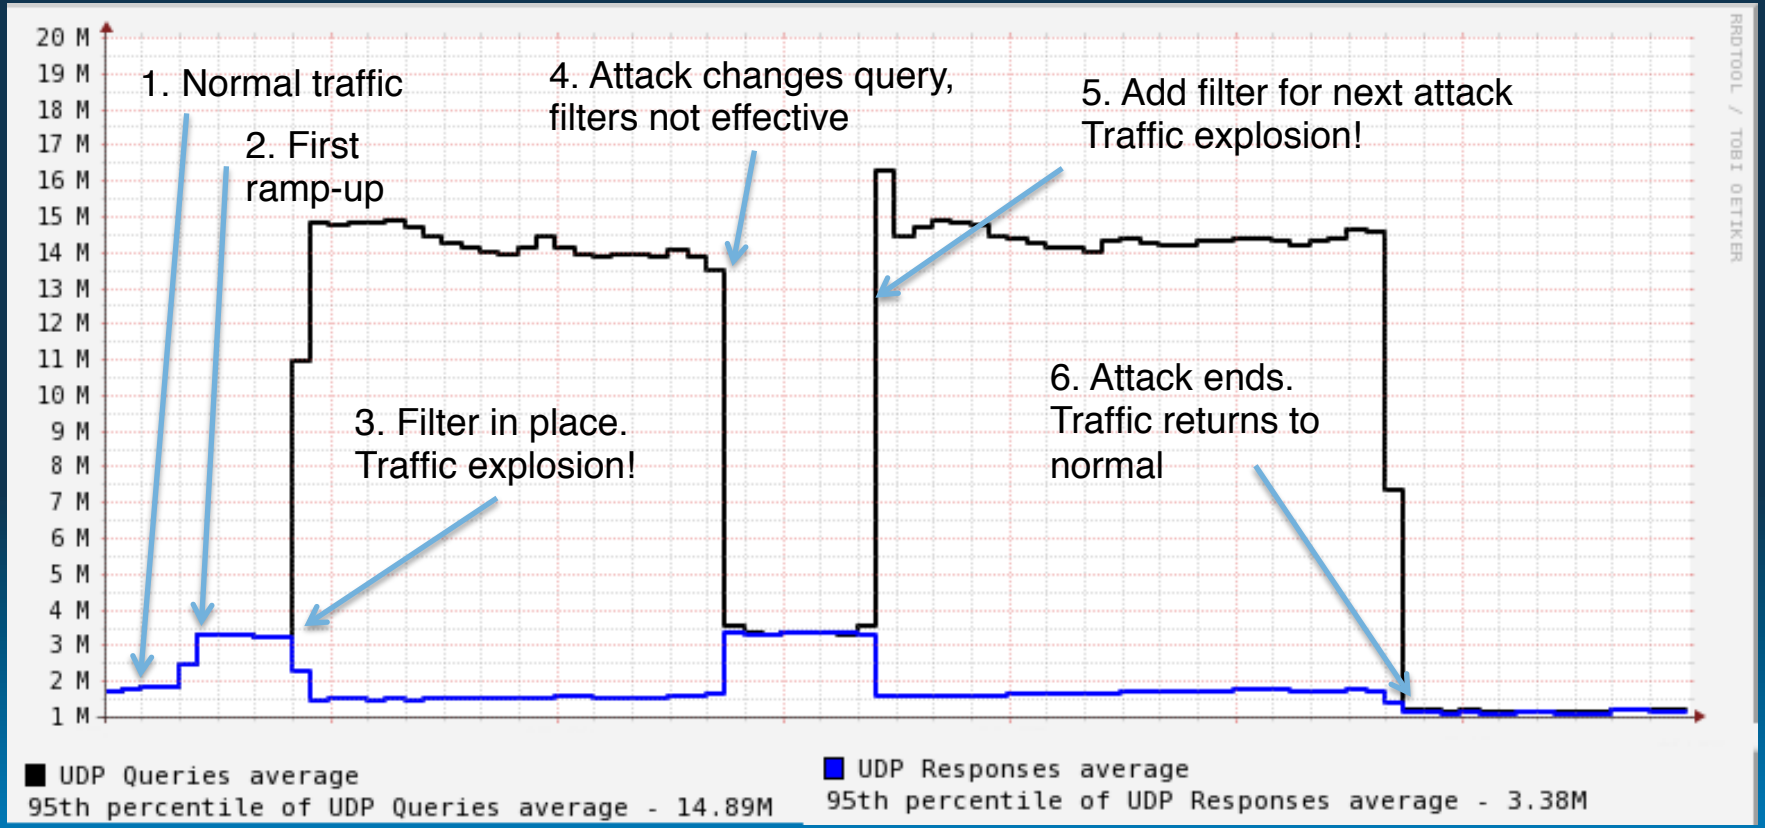

Attack 1: Random QNAME Reflectors

- Our big sites OK
- b.gtld-servers.net had some loss
- Red on RIPE only **after** we put the filter in place
- Valuable Lessons Learned about filtering

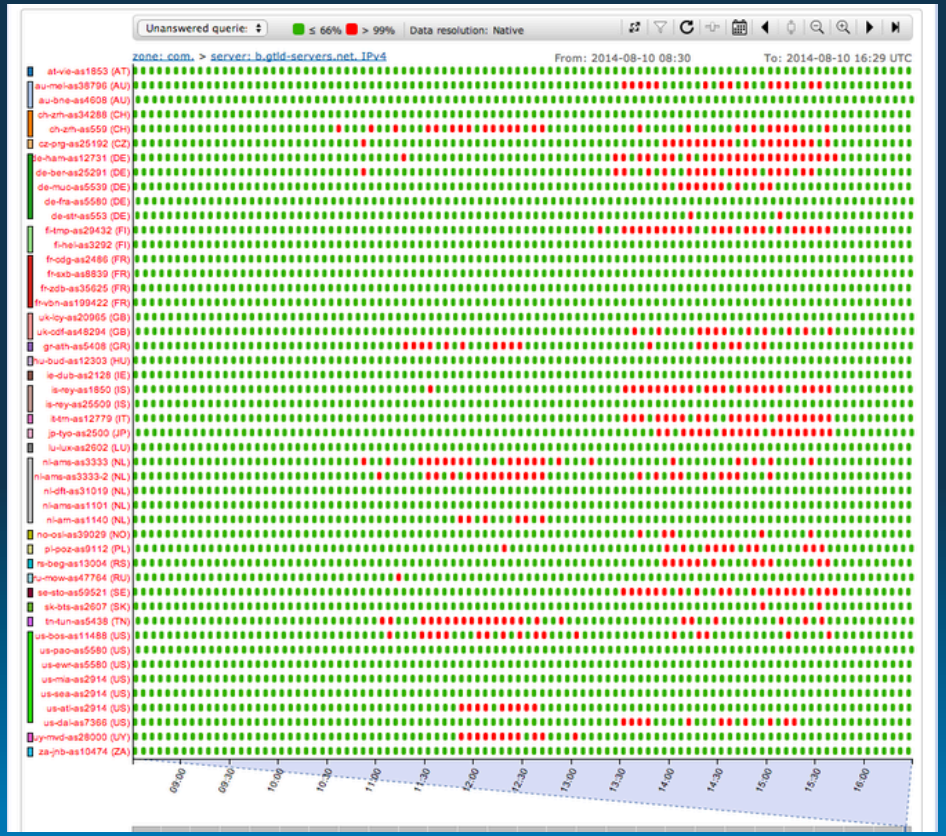

Attack 1: Random QNAME Reflectors

- Real name servers
  - MUCH better to rate-limit
  - 100% drop causes retries
- Caching name servers can retry at 4x (or more)

Attack 1: Random QNAME Reflectors

- 100% filter-drop random QNAME attacks will increase traffic volume

- If you can't filter it, what do you do?
  - Rate Response Limit?
  - Ask caching name servers to "Stop that!" (Good luck tracking down all 135,000 IP addresses!)
  - Anything else?

Alternatives To Dropping Random QNAME Attacks

- TLD operators temporarily take over the offending domain
  - Harder to do with some TLDs
  - It's already NXDOMAIN or you wouldn't be seeing it
  - The queries won't come to you if domain in question is delegated
- Delegate the (random.).domain.tld domain to some sacrificial name servers
  - Offloads traffic
  - Prevents retry storms
  - No fancy filtering software necessary

Attack 2: EDNS0 bufsize=9000

- 2013: frequent reflector attacks
  - Usually apex-name queries
  - Several different attacks to .cc, .com, .jobs
  - Sometimes root, as well (usually from a typo)
  - Verisignlabs.com / IN / ANY
    - (big DNSSEC response)
  - It looked like the attack came from a small range of IP addresses

Attack 2: EDNS0 bufsize=9000

- Impacted many DNS operators, not just Verisign
- Bad guys found big pay-off
  - 32 bytes in, 2000+ bytes out
  - Hard to trace because of forged sources
  - TLDs with big infrastructure handle the load nicely
  - Freely-available source code to perform exploit.
- This attack was in-style around 2013
  - Haven’t seen recently, but still a viable attack strategy

Attack 2: EDNS0 bufsize=9000

- Do you recognize this sort of thing?

```
12:46:53.308200 IP 77.98.44.228.19220 > 10.63.32.81.53: 16468+ [1au] ANY? name. (32)
   0x0000:  4500 003c 035e 0000 ef11 237d 4d62 2ce4  E..<.^....#}Mb,.
   0x0010:  0a3f 2051 4b14 0035 0028 3806 4054 0100  .?.QK..5.(8.@T..
   0x0020:  0001 0000 0000 0104 6e61 6d6e 0000 ff00  ........name....
   0x0030:  0100 0029 2328 0000 0000 0000            ...)#(......
```

- This specific attack was against .name
- Similar seen on .com, .net, .tv, .cc, .jobs
- Usually has ANY or DNSKEY as Resource Record

Attack 2: EDNS0 bufsize=9000

- If you know what DNS looks like at the packet level, you know this is uncommon:

```
12:46:53.308200 IP 77.98.44.228.19220 > 10.63.32.81.53: 16468+ [1au] ANY? name. (32)
   0x0000:  4500 003c 035e 0000 ef11 237d 4d62 2ce4  E..<.^....#}Mb,.
   0x0010:  0a3f 2051 4b14 0035 0028 3806 4054 0100  .?.QK..5.(8.@T..
   0x0020:  0001 0000 0000 0104 6e61 6d6e 0000 ff00  ........name....
   0x0030:  0100 0029 2328 0000 0000 0000            ...)#(......
```

- The EDNS0 section of a DNS query
  - Requesting 9000 (0x2328) bytes worth of DNS Response
- Usually we see 512, 1024, 2048, 4096. Never 9000.

Attack 3: zz.com-Style Attacks

- 15+ separate events in 2012
- Method:
  - High query volume for international gaming sites
  - Verisign used as a reflector
    - Pre-RRL days
  - Possible motive: Censorship/deletion of the domain?
  - Rate limiting is the answer here
  - 100% filter completes the attack

Future Plans

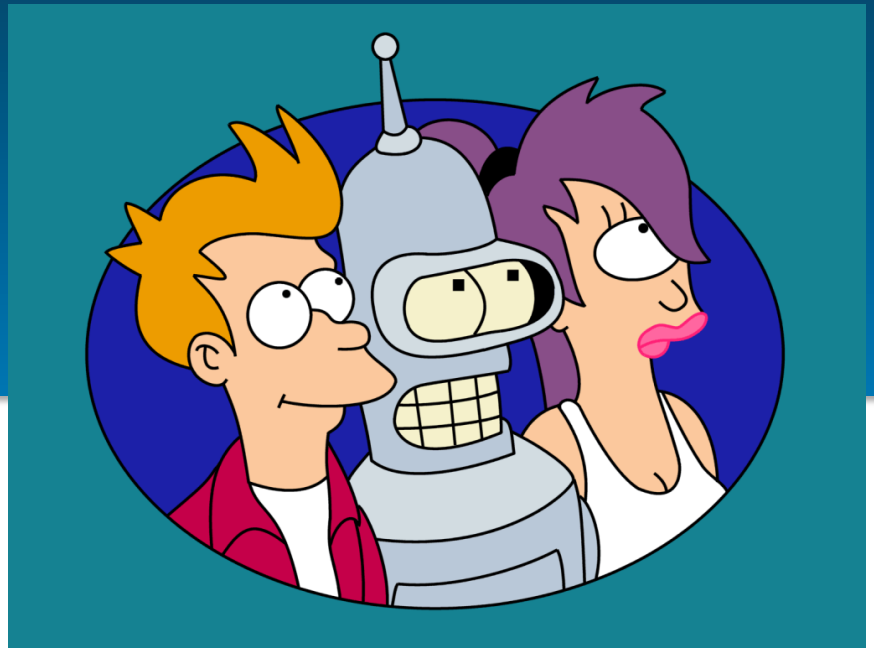

Capacity Enhancements

- Increase NETWORK capacity
- Increase SERVER capacity
- Increase number of deployments worldwide
  - Shameless plug: You too can help!
  - More information: http://rirs.verisigninc.com

Response Rate Limiting

- “RRL helps mitigate DNS denial-of-service attacks by reducing the rate at which authoritative servers respond to high volumes of malicious queries.” [1]

- Continued tuning of RRL capabilities
  - Gradual, measured, and ever-evolving

1. https://kb.isc.org/article/AA-01000/0/A-Quick-Introduction-to-Response-Rate-Limiting.html

Direct Announce from Name Servers

- Leverage Intel DPDK and FreeBSD Netmap
  - OS network stack is a performance bottleneck for us at the server level
  - DPDK and Netmap allow our code to bypass the OS network stack, communicating directly with the NIC from user space.
- ~6 Million queries per second per server
  - Full 10 Gbps of response data with our .COM/.NET custom name server
  - Industry-leading DNS server capacity

Direct Announce from Name Servers

- Name server to announce directly to upstream router
  - Diversity strategy at load balancing layer
  - Improved scale
  - Frees us of ECMP limitations from various router vendors

Questions?

powered by

VERISIGN™