

# Past, Present, and Future of Spoofed-source IP Packets

Paul Vixie  
Chairman and Founder  
Internet Systems Consortium

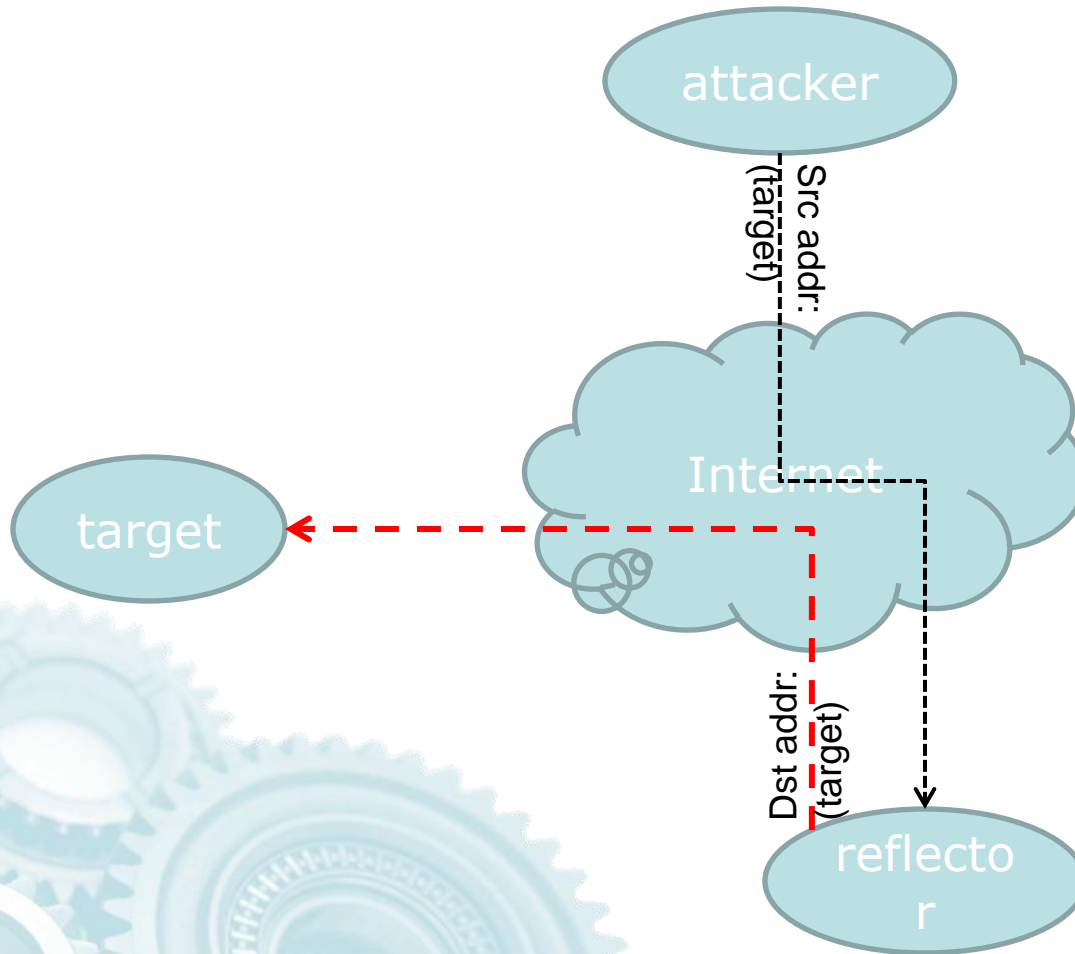


# Overview

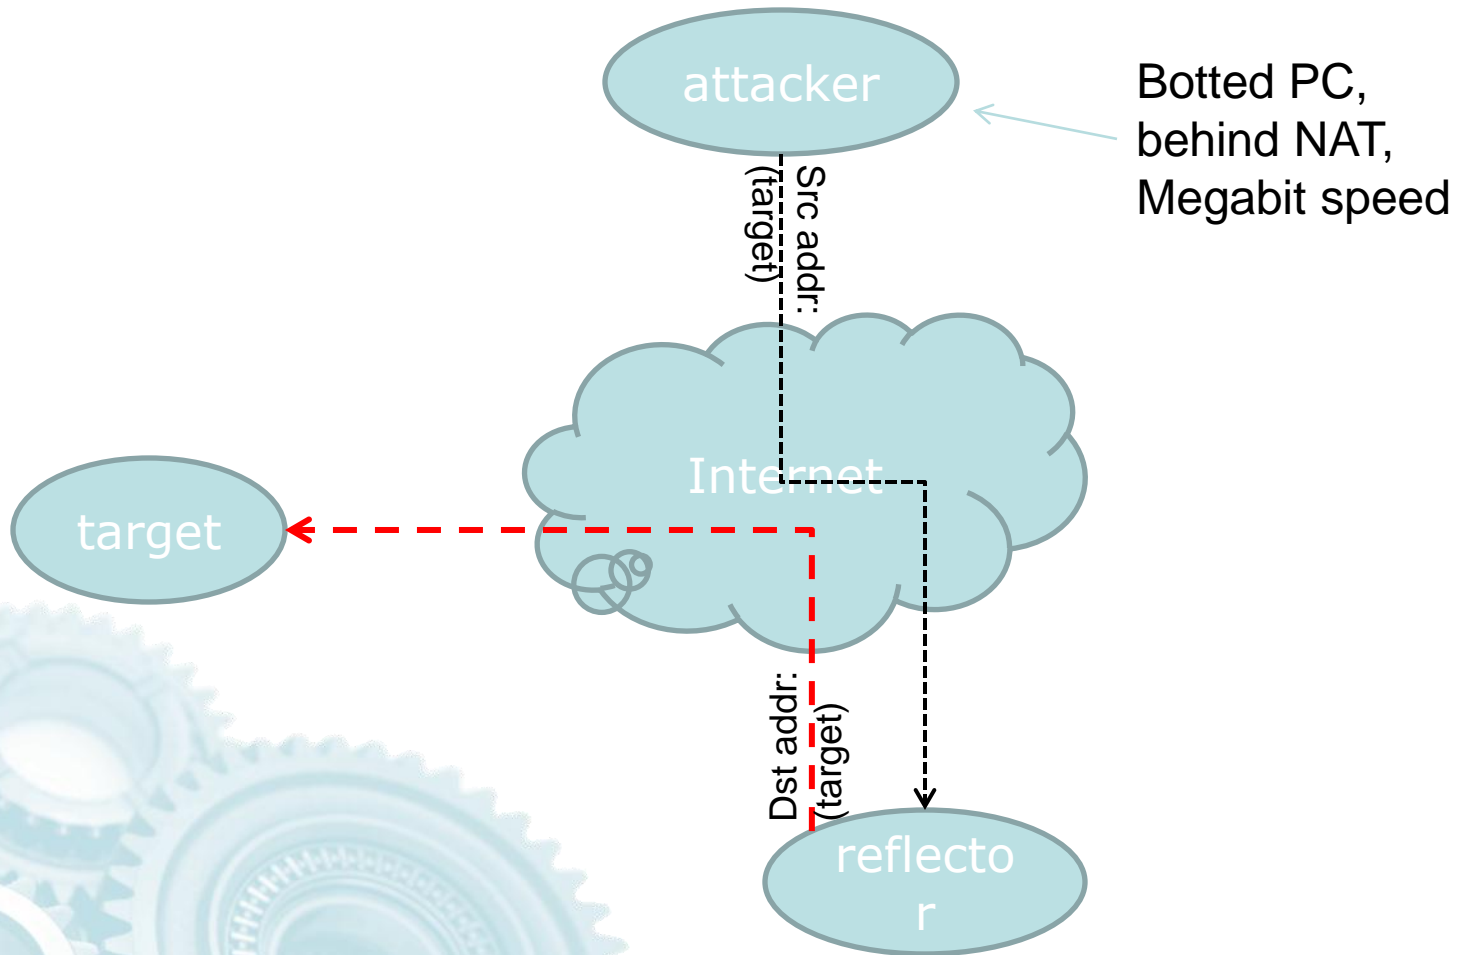
- IP Spoofing: the root of most evil
- DNS RRL: radical DDoS opt-out




# Spoofered Source Attacks: Essence




# Spoofered Source Attacks: Past



# Spoofed Source Attacks: Present



# Spoofer Source Attacks: Future



# Crazy Lessons of History

- Wide area UDP services must never amplify
  - In this light, DNS was crazy
  - And: DNSSEC is even crazier
  - But: NTP is (strangely) OK
- Promoting data to executable code is crazy
  - Like: Java, Flash, ActiveX, Autorun, JavaScript, or the conficker worm's "click to permit" hack
- Expecting users to be sysadmins is crazy
  - Like: PC, Mac, cloud servers, smart phones

# Action Items for Industry

- All recursive name servers need access control
  - They should *only* answer for their customers
- All authority name servers need rate limiting
  - Quickly repeated responses are *never* necessary
- Edge networks should validate their src addrs
  - This *can't* be done closer to the Internet "core"
- Cloud/VM providers should offer sys admin
  - Webmasters *can't* be expected to update Joomla
- References
  - BCP38, "Network Ingress Filtering", 2000
  - SAC004, "Securing the Edge", 2002



# RRL On The Wire

```
[nsa:amd64] repeat 25 \  
    dig +novc +ignore +retries=0 +time=1 vix.com aaaa \  
        @ns.sql1.vix.com \  
    | grep tc  
;; flags: qr aa tc rd ad; QUERY: 1, ANS: 0, AUTH: 0, ADD: 1  
;; flags: qr aa tc rd ad; QUERY: 1, ANS: 0, AUTH: 0, ADD: 1  
;; flags: qr aa tc rd ad; QUERY: 1, ANS: 0, AUTH: 0, ADD: 1
```




# RRL Configuration

```
options {  
    directory "/var/local/named";  
    pid-file "/var/run/named-nsa.pid";  
    query-source address 149.20.48.227 port *;  
    listen-on-v6 { ::1; 2001:4f8:3:30::3; };  
    listen-on { 127.0.0.1; 149.20.48.227; };  
    recursion yes;  
    notify yes;  
    dnssec-enable yes;  
    dnssec-lookaside . trust-anchor dlv.isc.org.;  
    dnssec-validation yes;  
    rate-limit {  
        responses-per-second 5;  
        window 5;  
    };  
};
```

# Using RRL In Your Servers

- In authority servers
  - RRL has no negative impact on real flows, because real clients have caches, will retry with UDP, will try TCP if given a truncated response
- In recursive servers
  - RRL would have a negative impact on real flows, because real clients do not have caches
  - It should not be necessary, just use ACLs

# RRL In Action: Afilias



# Final Thoughts: DNS RRL

- RRL was first implemented in BIND but is intended for use in *all* name servers
- NSD added it in 3.2.15, February 2012
- Please study the DNS RRL specification carefully, it's intended to be implemented literally
- Specification, patches, pointers, and specification are available online
  - <http://www.redbarn.org/dns/ratelimits>

# Final Thoughts: IP Spoofing

- Economics at the edge aren't just misaligned, they're pessimal
- There will always be spoofing, although regulation isn't impossible
- Meanwhile we have to get rid of all DDoS amplifiers
- Fortunately, the economics are better aligned for this