Falling Trees

or If a DNS Server is Lame but Nobody Queries It, Should You Send an E-mail?

Shane Kerr <shane@isc.org>

DNS Working Group, RIPE 59

Lisbon, 2009-10-08

Background

The RIPE NCC implemented a system which checks for lame DNS servers in the part of the reverse DNS tree they maintain.

I thought maybe we should check to see if lameness is a real problem.

The RIPE NCC asked me if I would be willing to investigate this.

Co-operation

- The RIPE NCC DNS department provided data & explanation
- OARC provided the system to hold the data & perform analysis

Methodology

- RIPE NCC runs checks
  1. Resolve NS to A/AAAA record
  2. SOA lookup of zone at A/AAAA
- RIPE NCC captures traffic
  - At RIPE NCC's master, `ns-pri.ripe.net`
  - For a set of 1-hour periods
- Compare actual traffic with lame zones

On the Checking of Traffic

- Parse each reply to get NS RRSET
- A *very* small mismatch with data sets
  - Lameness check run at different time than packet captures
  - 3 NS in an hour of packet captures
- If no NS, stop
  - NXDOMAIN, SOA query, errors, …
- If NS, check for errors

Skinning Cats

or, There's More Than One Way to Make a Lame DNS Server

1. The NS entry may not resolve to an A or AAAA record.

```
example.org NS ns1.example.cmo
            NS ns2.exmaple.info
```

2. Nothing at the IPv4 or IPv6 address answers.

```
ns.example.net A 0.0.1.16
```

Diagrammatically Explained

or Explained Diagrammatically

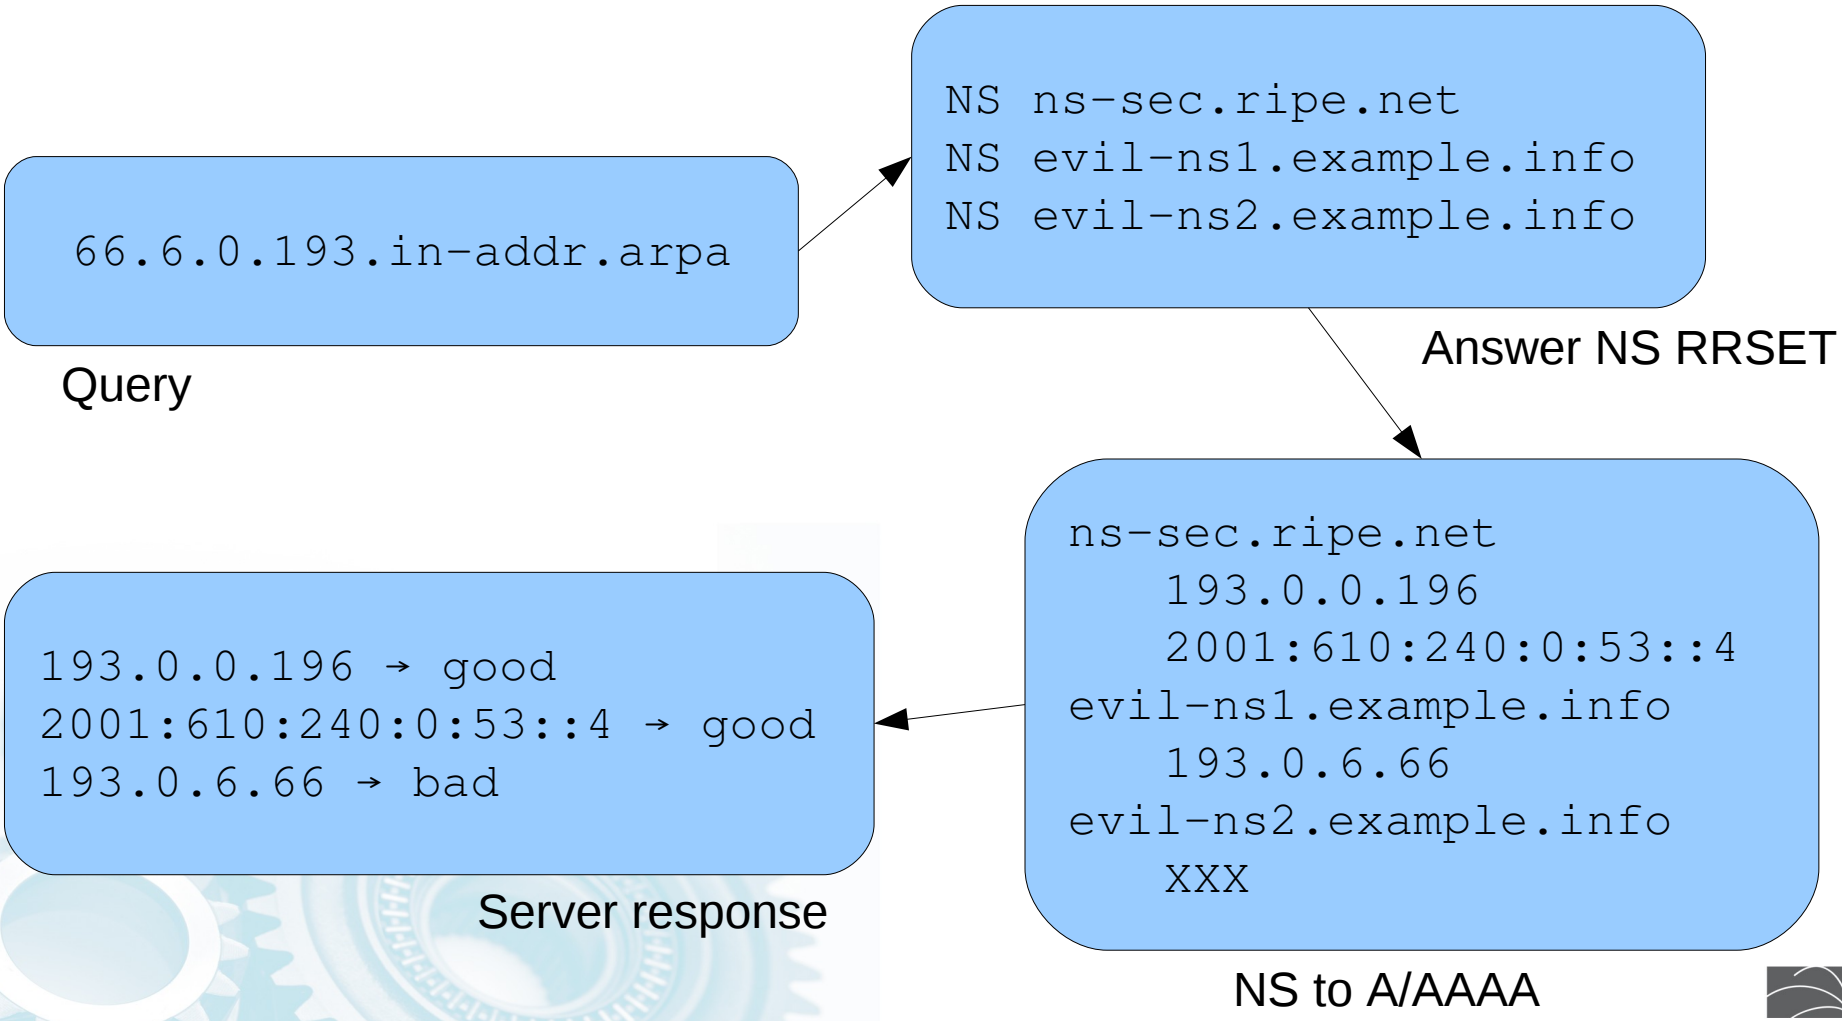

Impact of Bad NS in RRSET

- Might be an immediate failure
  - `ns.exxample.org` gives NXDOMAIN
- Might be a set of timeouts
  - All NS in the RRSET may fail
- Current data does not distinguish

Chance of Using a Bad NS

- Two NS, one bad
  - 50% chance bad NS is chosen 1st
- Three NS, one bad
  - 33% chance bad NS is chosen 1st
- Three NS, two bad
  - 66% chance bad NS is chosen 1st
  - 50% chance bad NS is chosen 2nd
- And so on...

Impact of Bad Servers

- Server at IP address does not answer
- Set of timeouts
- Delay for the user!

- Logic about chance of using a bad server same as for using a bad NS...

Applicability Statement
An Aside Before the Results

If anyone running a delegation-only server wants to perform similar analysis on their domains, you can have the code or work with me to perform similar analysis.

The Numbers

Total packets scanned:	16618986
Replies sent from server:	8309072
Unparsable replies:	9350
Packets with non-0 RCODE:	3165031
Packets with stale NS:	17839
Packets with NS in RRSET:	5096414

Estimate about 25030.984 lookups have NS lookup failures
Estimate about 39303.618 NS lookup failures
Estimate about 64100.080 lookups use non-answering IP
Estimate about 77258.599 queries to non-answering IP

A Carta

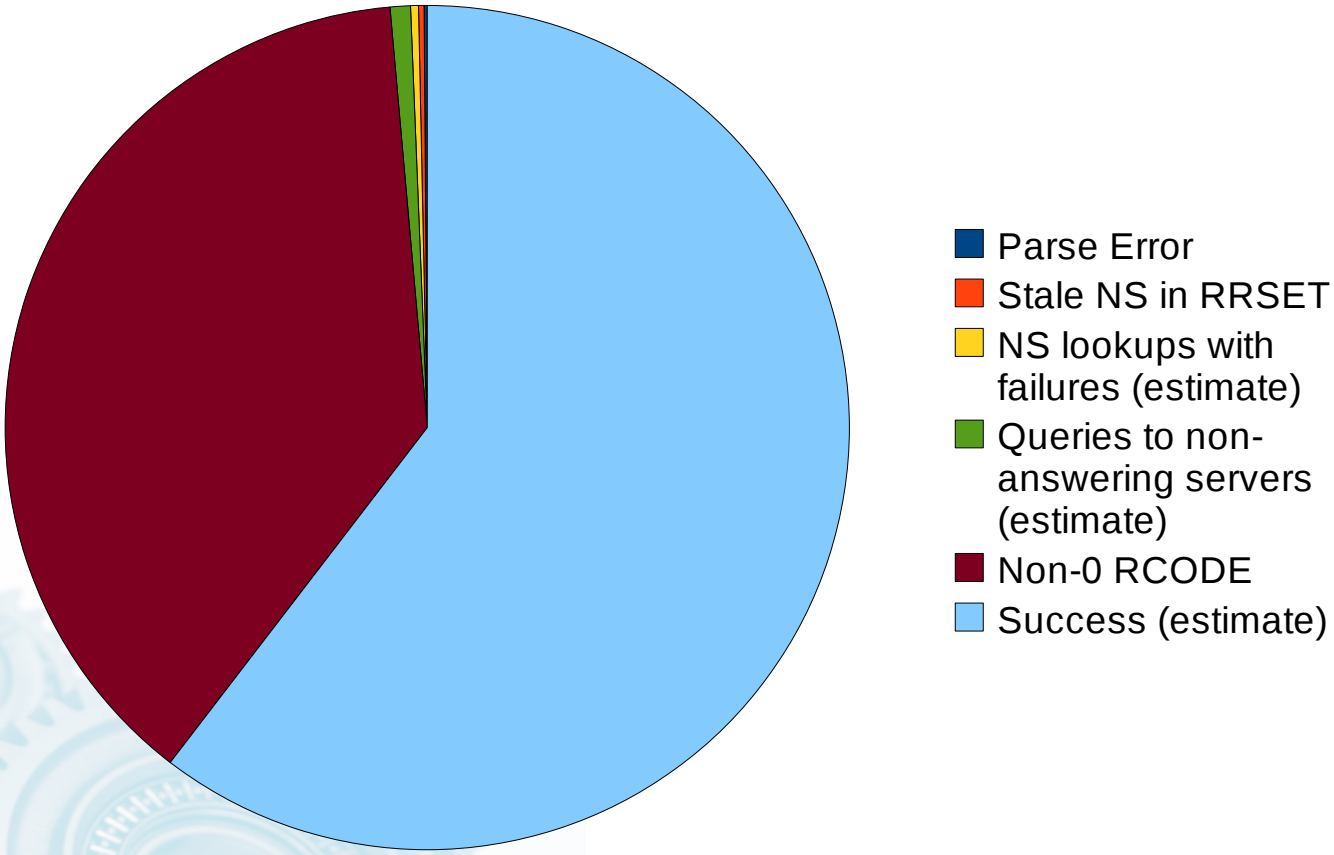
Parse Error
Stale NS in RRSET
NS lookups with failures (estimate)
Queries to non-answering servers (estimate)
Non-0 RCODE
Success (estimate)

Results

- 0.3% of queries had bad NS
- 0.8% of queries had an A/AAAA resulting in a DNS lookup failure
- About 1% of queries affected by lameness

Caveats

- Results do not account for caching
- Code based on IP, rather than IP+domain (oops)
- Effects of bad NS very uncertain
- Effects of bad servers slightly uncertain

Discussion

- But 5% of servers lame!
  - Unused lame servers don't get fixed
  - Used lame servers *DO* get fixed
- 1% is still a bit much...
  - DNSMON error rate for ns-pri was 0.3%
  - Average DNSMON error rate is 0.8%

Proposals

1. No more blanket e-mails
2. Annual report to LIRs?
3. Targeted e-mails to most-impacted users?