FOSSETCON 2015:
BIND9 – Recursive Client Rate limiting

Chuck Aurora, Technical Support Engineer
Internet Systems Consortium, Inc.

Special Thanks to:
- Cathy Almond, ISC Support Team Leader
- Victoria Risk, ISC BIND9 Product Manager
- ISC Support Customers

Presenter

Chuck Aurora

ISC Technical
Support Engineer

Agenda

1. **Pseudo-random subdomain attack**
2. Recognizing the attack
3. Recommended mitigation
4. Results from live environments
5. Any questions?

The attack - unusual queries

high volume of queries for non-existent sub-domains

<randomstring>.www.example.com
<anotherstring>.www.example.com

does not exist	exists
??	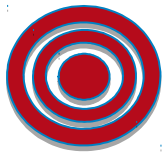

The source

- Open resolvers
  - your servers
  - your clients (CPE devices/proxies and forwarders)

- Compromised clients (botnets)
- Compromised devices

Attack begins

Clients

compromised devices

1. Requests for randomstring.www.example.com

nothing about this in the cache

ISP resolvers

2. Attempt to resolve

example.com

Target of the DDOS Authoritative provider

Initiator of DDoS traffic

Initially, the target responds

Clients

compromised devices

Initiator of DDoS traffic

4. Reply (NXDOMAIN)

ISP resolvers

3. Server replies "no such domain"

example.com

Target of the DDOS Authoritative provider

More requests flood in

Clients

compromised devices

1. Requests for randomstring2.www.example.com

ISP resolvers

Initiator of DDoS traffic

example.com

Target of the DDOS Authoritative provider

Target is overwhelmed

compromised devices

Initiator of DDoS traffic

ISP resolvers

2. Attempt to resolve

3. Server is unresponsive

example.com

Target of the DDOS Authoritative provider

Resolver is degraded

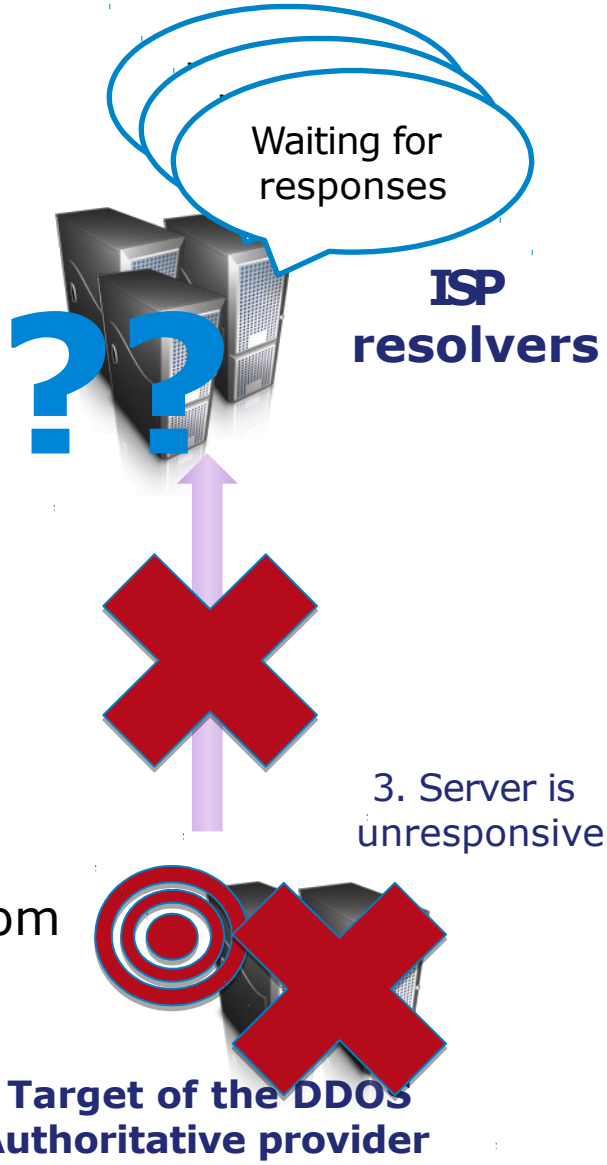

Legitimate queries fail

All Clients

Request for www.othersite.com

Reply SERVFAIL

ISP resolvers

No more resources available to handle new queries!

Target of the DDOS Authoritative provider

2. recognizing the attack

Symptoms

- ✓ Many SERVFAIL responses
- ✓ Increased inbound client queries
- ✓ Resolution delays to clients
- ✓ Dropped responses
- ✓ Increased memory consumption
- ✓ Increased NXDOMAIN responses
- ✓ Firewall connection table overflows

Evidence

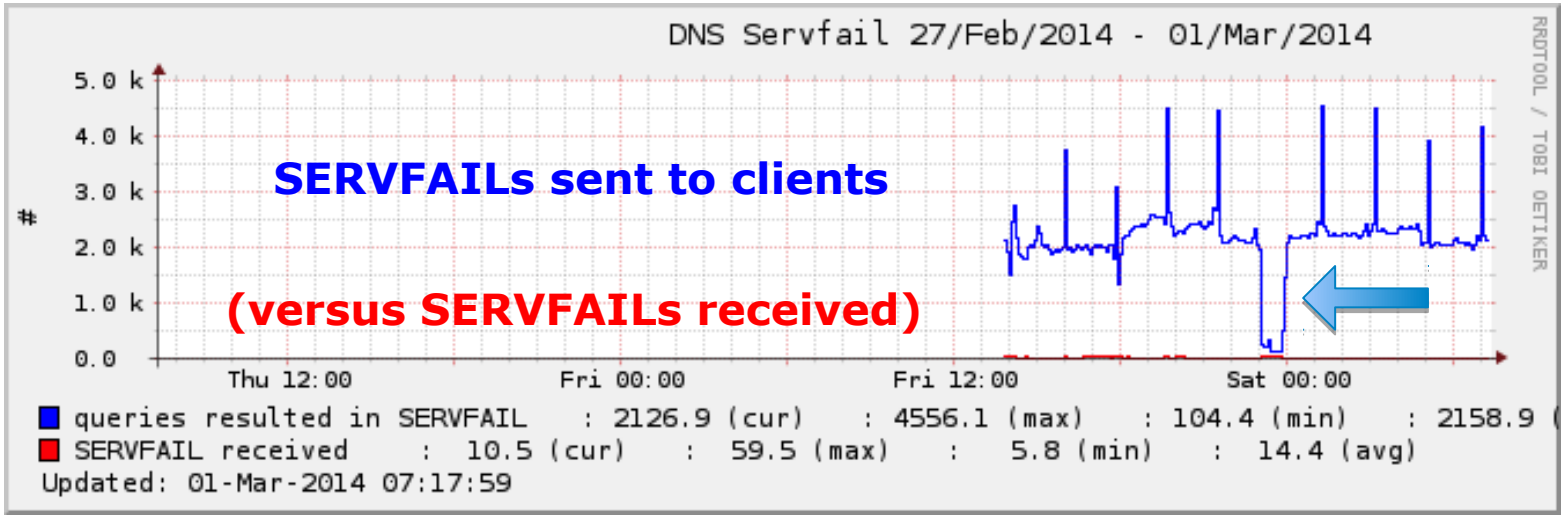

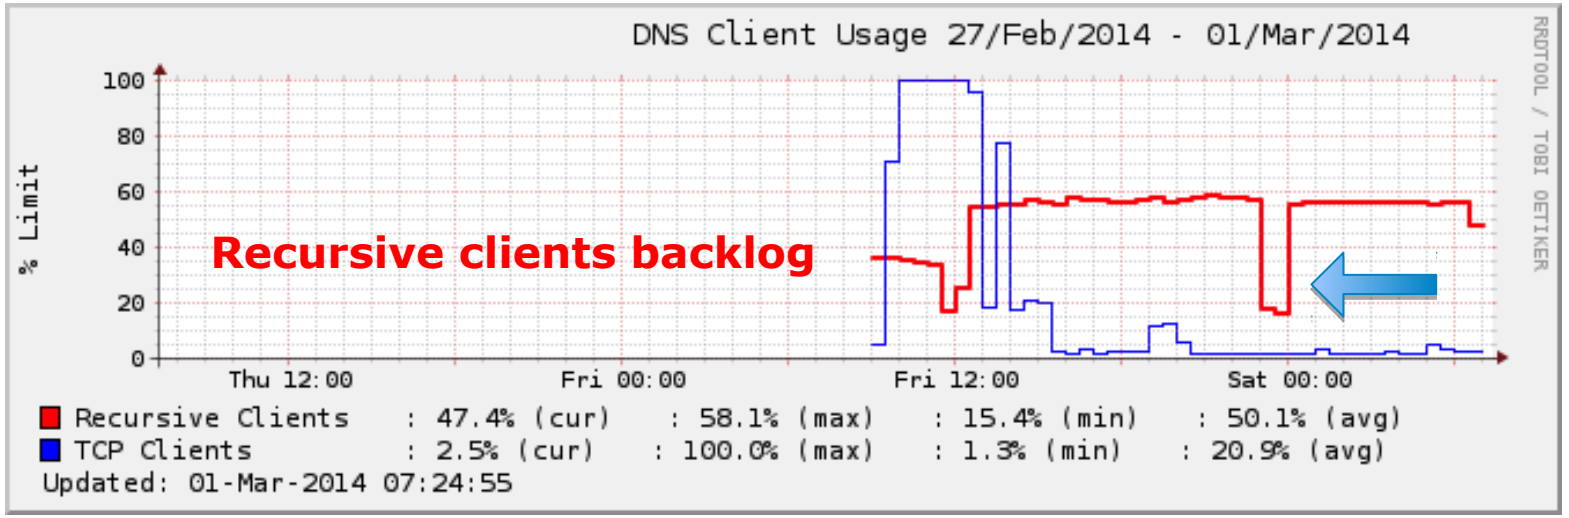

Accurate diagnosis

1. Do you have a significant (and unusual for you) backlog of recursive client contexts?

   rndc status

   recursive clients: 0/1900/2000

   rndc recursing
2. What are those queries for?
3. Why are they in the backlog?
4. Where are they coming from?

3. Mitigation

Mitigation Goals

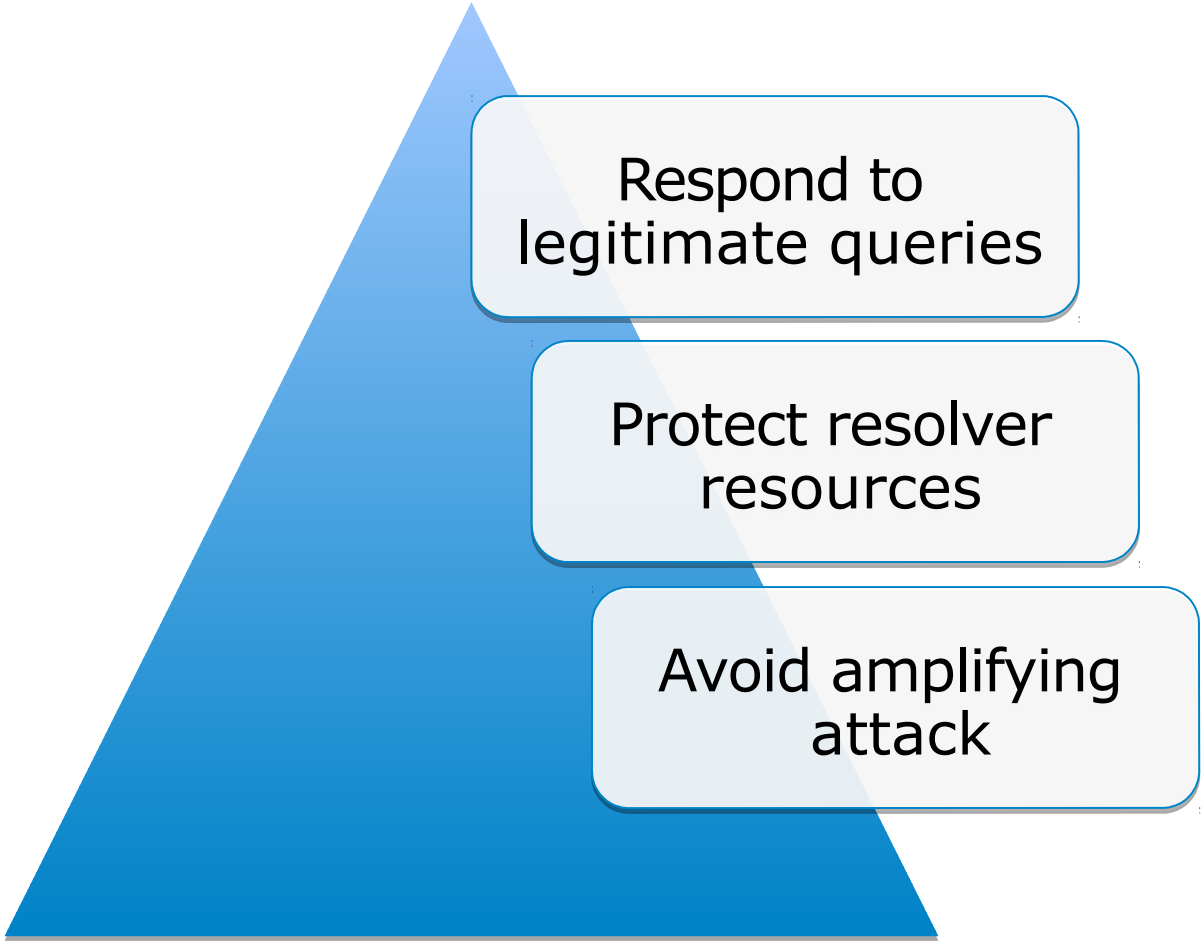

Don’t…

- Panic!!
- Assume that increasing server resources (e.g. recursive client contexts, sockets, network buffers etc..) is going to help *
- Block your clients (although, it depends…)

* For very large/busy resolvers, take a look at BIND 9.10 and new configuration option --with-tuning=large

Step 1: Lie if necessary

- Make recursive server temporarily authoritative for the target domain
  - Local zone
  - DNS-RPZ (*qname-wait-recurse no;)

- *Manual configuration change*
- *Need to undo the mitigation afterwards*

Step 2: Filtering

(Near) Real Time Block Lists

- Detect ‘bad’ domain names or just the problematic queries & filter them
- Local auto-detection scripts that dynamically add local authoritative zones (potential false-positives)
- BIND DNS-RPZ *
- Costs associated with feeds

* Requires ‘qname-wait-recurse no;’

Step 3: Rate-limiting

PER ZONE

PER SERVER

NEW: fetches-per-server

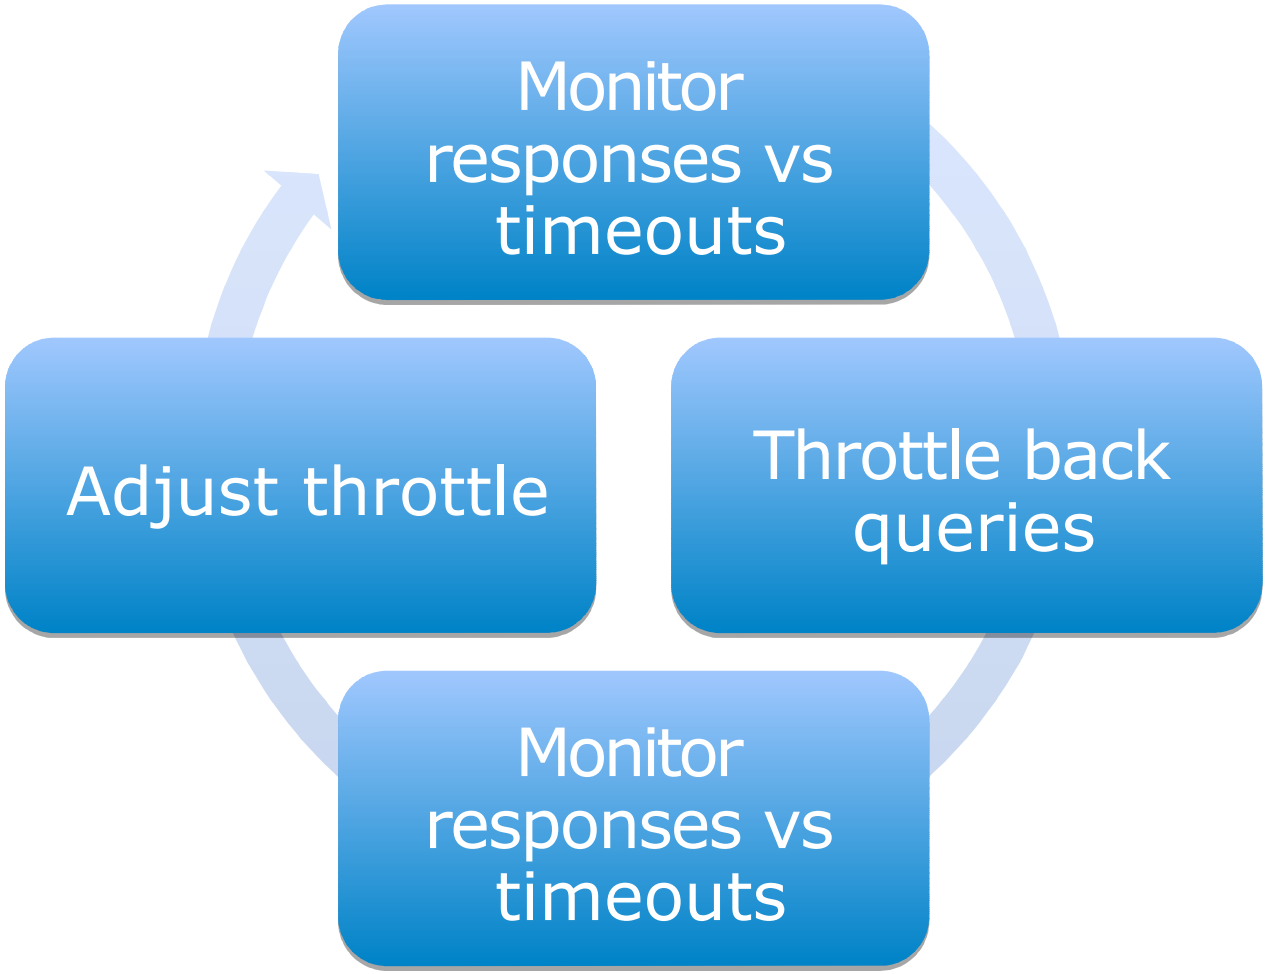

fetches-per-server

- Per-server quota dynamically resizes itself based on the **ratio of timeouts to successful responses**
- Completely non-responsive server eventually scales down to fetches quota of 2% of configured limit.
- Similar (loosely) in principle to what NLnet Labs is doing in Unbound

NEW: fetches-per-zone

- Works with unique clients (as does fetches-per-server)
- Does NOT auto-adjust
- Tune larger/smaller depending on normal QPS
- Use as a ‘backstop’ for fetches-per-server

Mitigation Summary

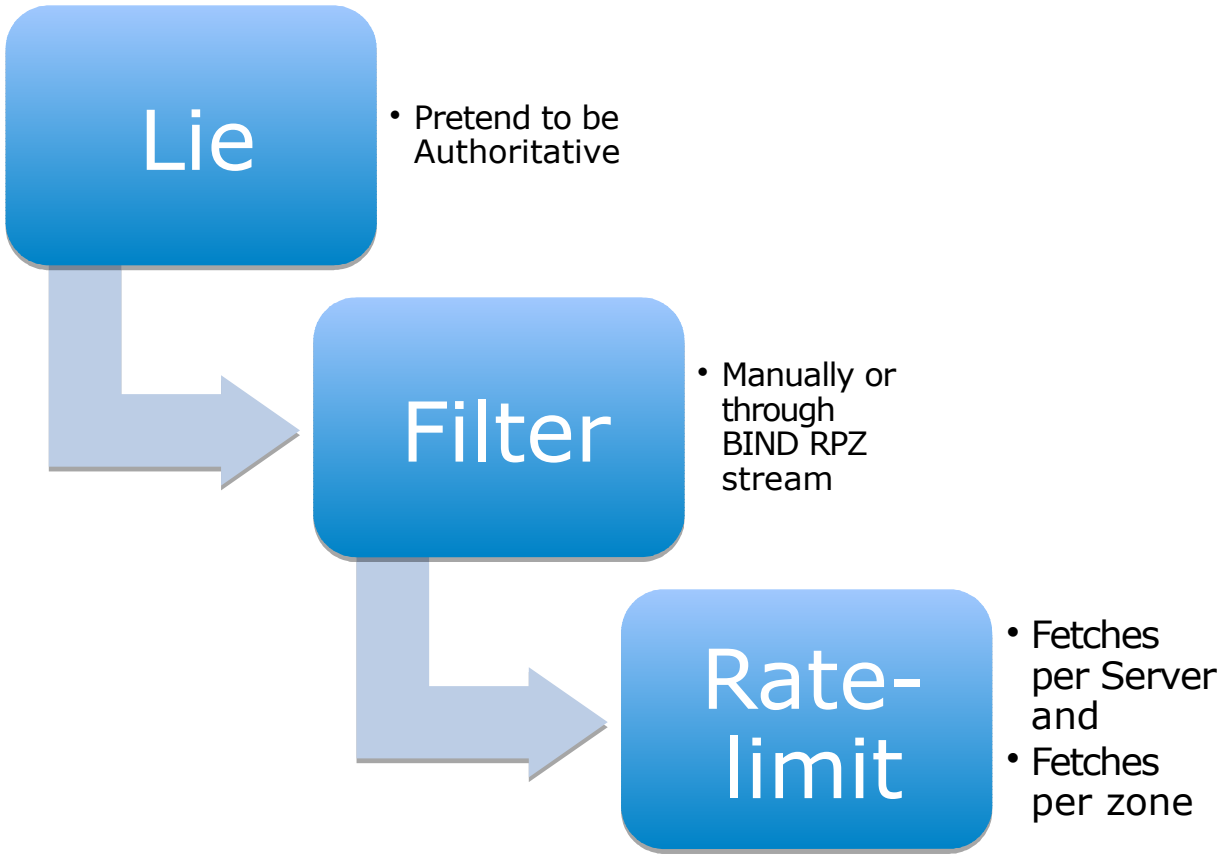

4. Results FROM LIVE PRODUCTION SYSTEMS

fetches-per-zone

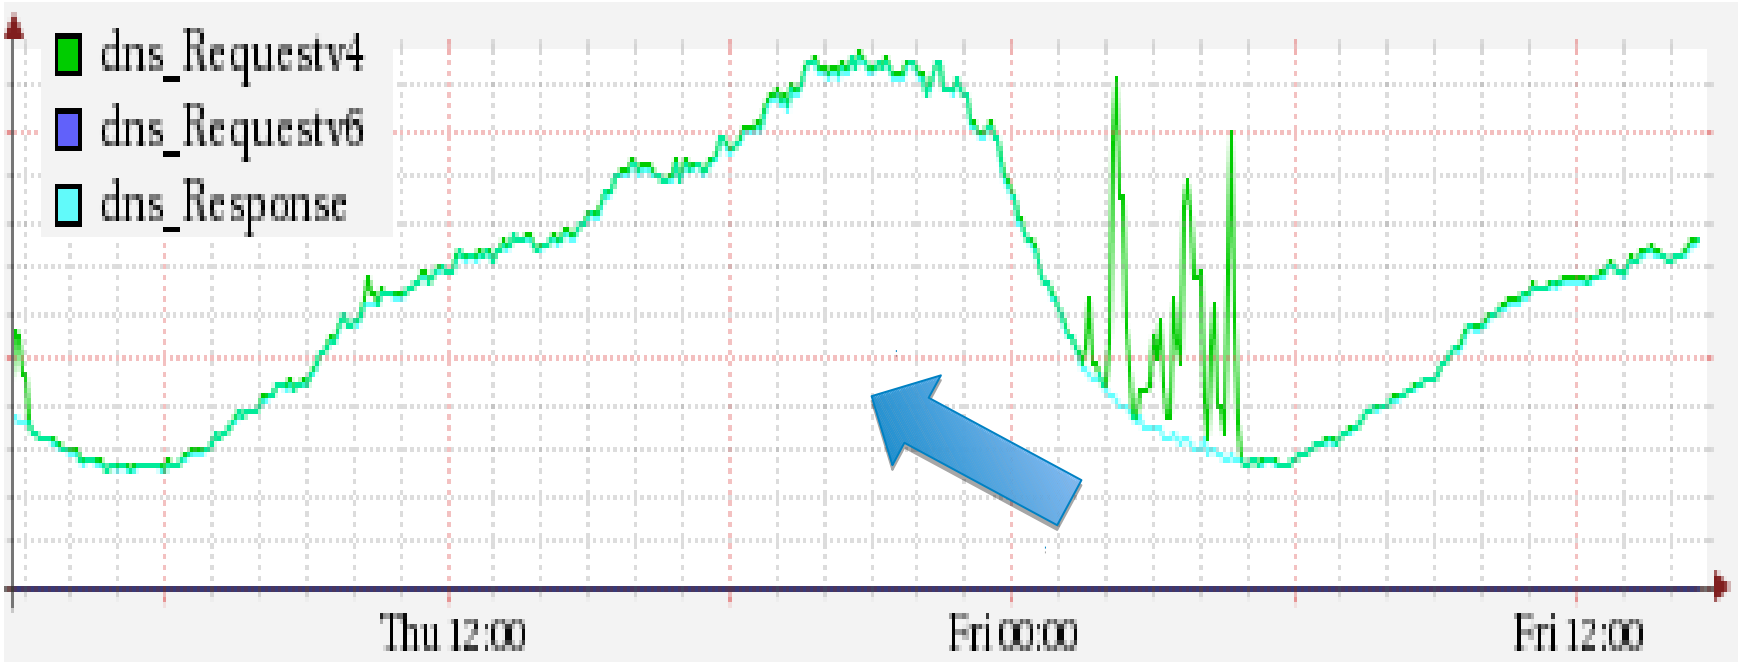

Spanish triple-play ADSL carrier & ISP
Roberto Rodriguez Navio, Jazztel Networking Engineering
used with permission

More on fetches per zone

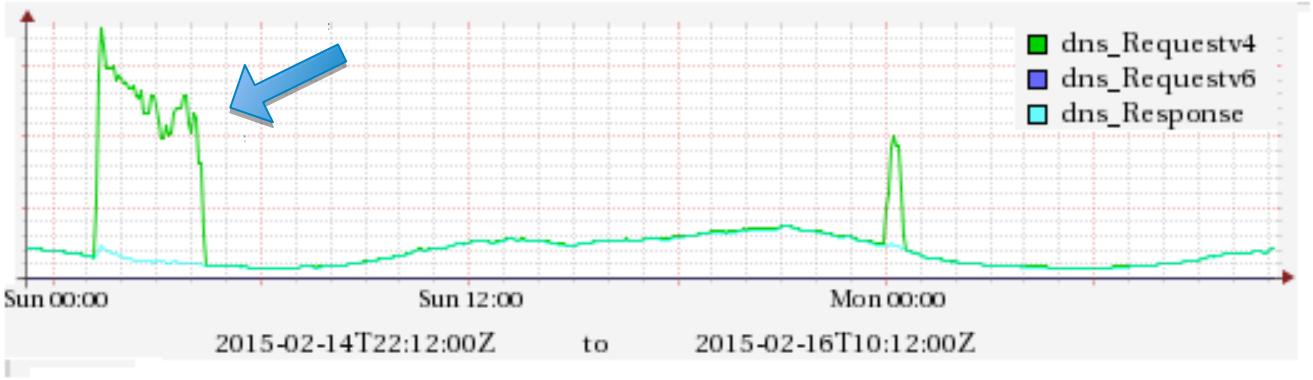

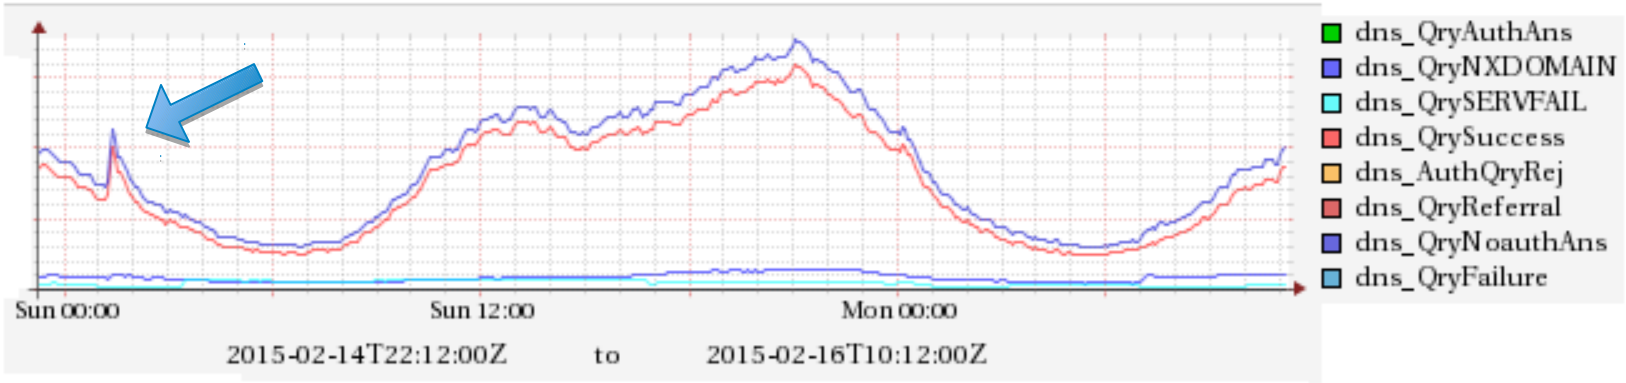

Spanish triple-play ADSL carrier & ISP
Roberto Rodriguez Navio, Jazztel Networking Engineering
used with permission

fetches-per-server

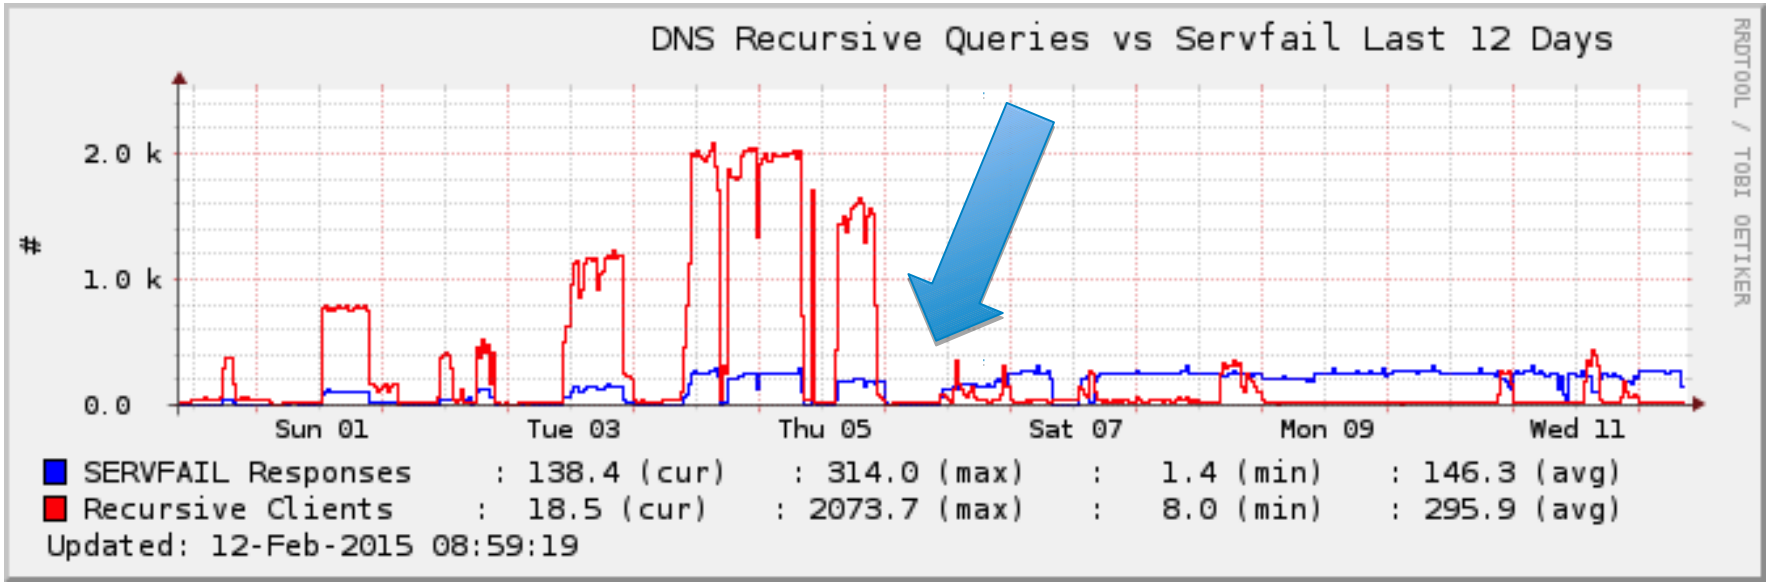

per-zone v. per-server

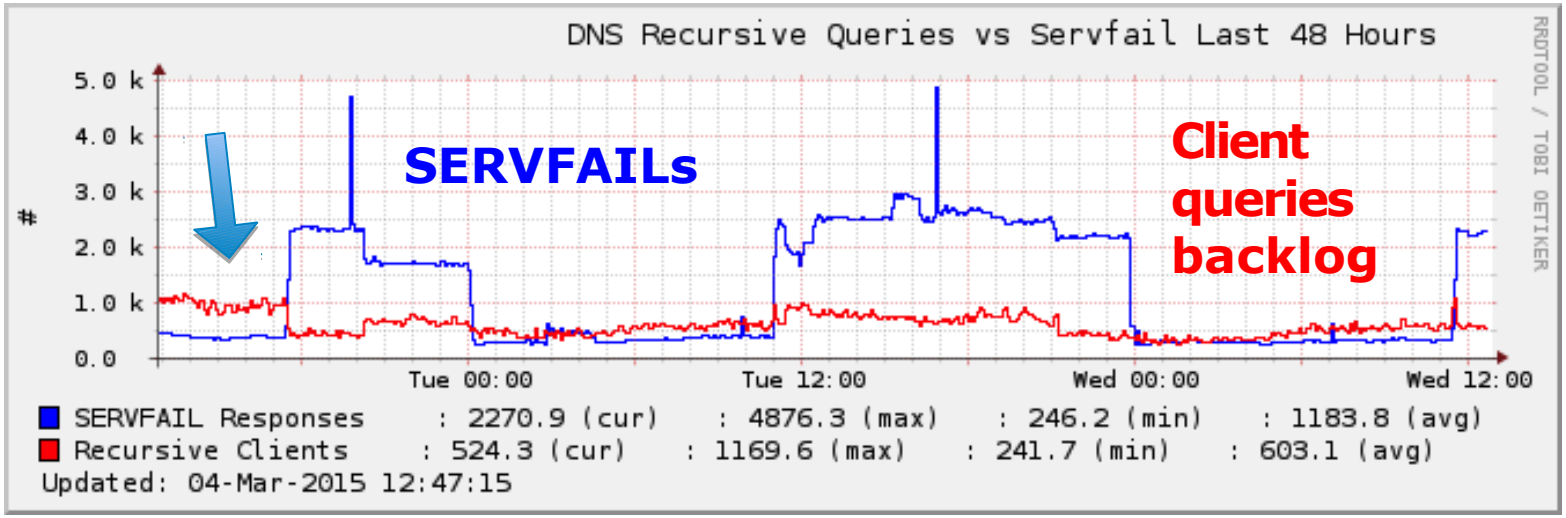

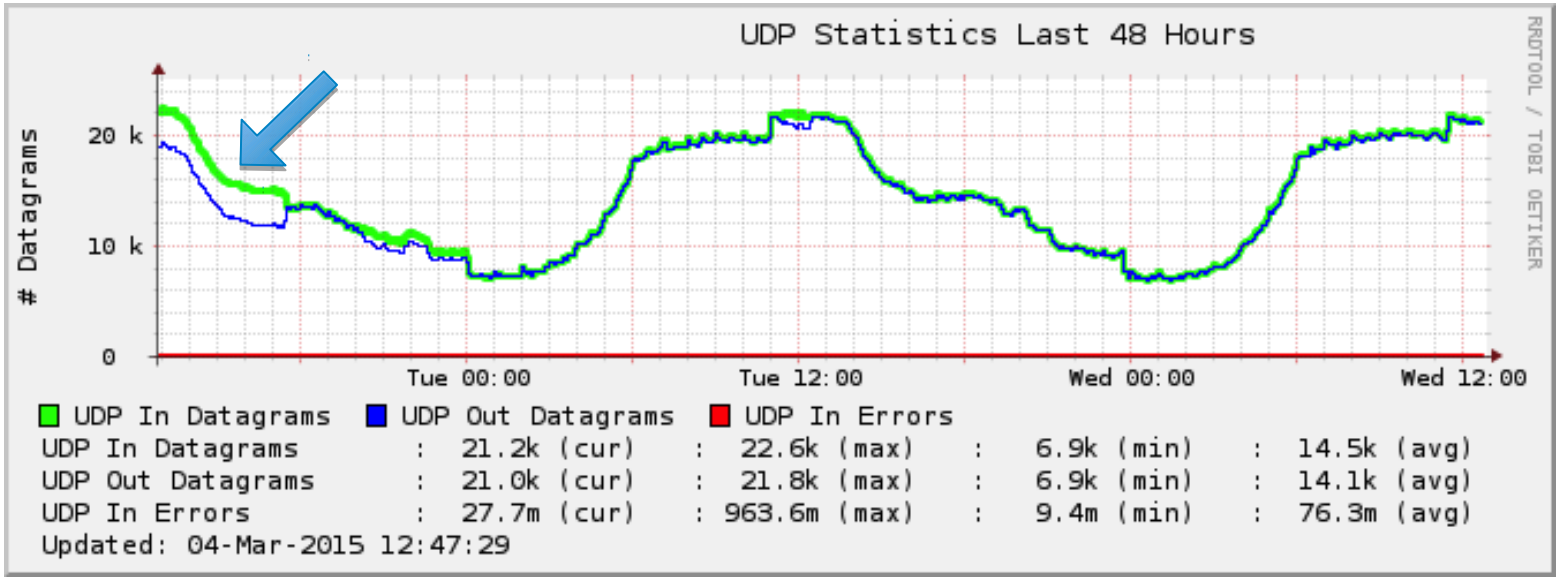

Comparison

Fetches Per Server

- Rate-limits per server
- Impacts queries for all zones served by the same machine
- Dynamically re-sizes based on the ratio of timeouts to successful responses

Fetches Per Zone

- Rate-limits per zone
- Manually tuned
- Set to larger value on higher-performance machines

What will the user see?

- Situation normal – no change to their usual experience (for most)
- (Some) SERVFAIL responses to names in zones that are also served by under-attack authoritative servers (collateral damage)
- NXDOMAIN responses for names in legitimate zones for which we ‘lie’

Client gets ...

No Response

** fetches-per-zone*

Legitimate queries will retry

Could be a problem for forwarding servers when the forwarder 'doesn't respond

SERVFAIL

** fetches-per-server*

- Legitimate queries will retry
- Doesn't protect resolver as much, but is the 'correct' response when authoritative server is overwhelmed

NXDOMAIN

Stops client from retrying

Same response the authority would send for the DDoS queries

(May be) wrong response to genuine clients

** Default behavior (configurable, except for NXDOMAIN)*

Further Resources

- Recursive Client Rate Limiting
  - available now in BIND 9.9.8 and 9.10.3
  - https://kb.isc.org/article/AA-01304
- Feature Webinar Recording available (8 July 2015)
  https://www.isc.org/mission/webinars/
- FAQs:
  https://kb.isc.org/article/AA-01316

Questions

info@isc.org, bind-users@isc.org, cba@isc.org

https://kb.isc.org/article/AA-01304