RIPE NCC
RIPE NETWORK COORDINATION CENTRE

RIPE NCC DNS Update

Anand Buddhdev | Oct 2016 | RIPE 73

The DNS team

Anand

Colin

Iñigo

Paul

Florian

Romeo

K-root

AS 25152

Status

- Active at 44 sites
  - Five “core” sites - multi-server, high capacity
  - 39 “hosted” sites - single server
- Bird, Cisco, ExaBGP and Juniper for routing
  - BGP anycast for high availability and low latency
- BIND 9.10, Knot 1.6 and NSD 4
  - Update to Knot 2 soon
- More applications from hosts in the queue

Authoritative DNS

AS 197000

Features

- BGP anycast from three sites
- RIPE NCC's forward and reverse zones
- Secondary for other RIRs' forward and reverse zones
- Secondary for ccTLDs
- Secondary for large reverse zones of RIPE NCC members
- Diversity with BIND, Knot and NSD, as well as Cisco and Juniper routers

ccTLD Status

- RIPE 663 published in December 2015
- We are evaluating all ccTLDs
  - 28 ccTLDs do not qualify based on zone size
  - Some ccTLDs do not qualify based on name server set
  - Process will continue until mid-2017
- ccTLDs that qualify will have to sign an agreement with RIPE NCC to receive secondary DNS service
  - Agreement to be renewed periodically

Other Secondary Zones

- Forward and reverse DNS zones of other RIRs will remain
- Internet infrastructure zones will remain, eg.
  - root-servers.org
  - as112.net
- Zones of other operators, especially commercial ones, will stop receiving service

Resiliency for ripe.net

- Improve resiliency for ripe.net in the face of bigger and more frequent DDoS attacks
- Open request-for-proposal process from July to September 2016
- We received three proposals, and selected the one that best satisfied our requirements
  - Verisign is now providing secondary DNS for ripe.net and related zones
  - Contract to be reviewed annually

DNSMON

Features

- Distributed DNS monitoring system
- Based on the RIPE Atlas infrastructure
  - Measurements from more than 50 RIPE Atlas anchors
- SOA, hostname.bind and version.bind queries, with NSID enabled
  - Both over UDP and TCP
- https://atlas.ripe.net/dnsmon/

DNSMON Visualisation

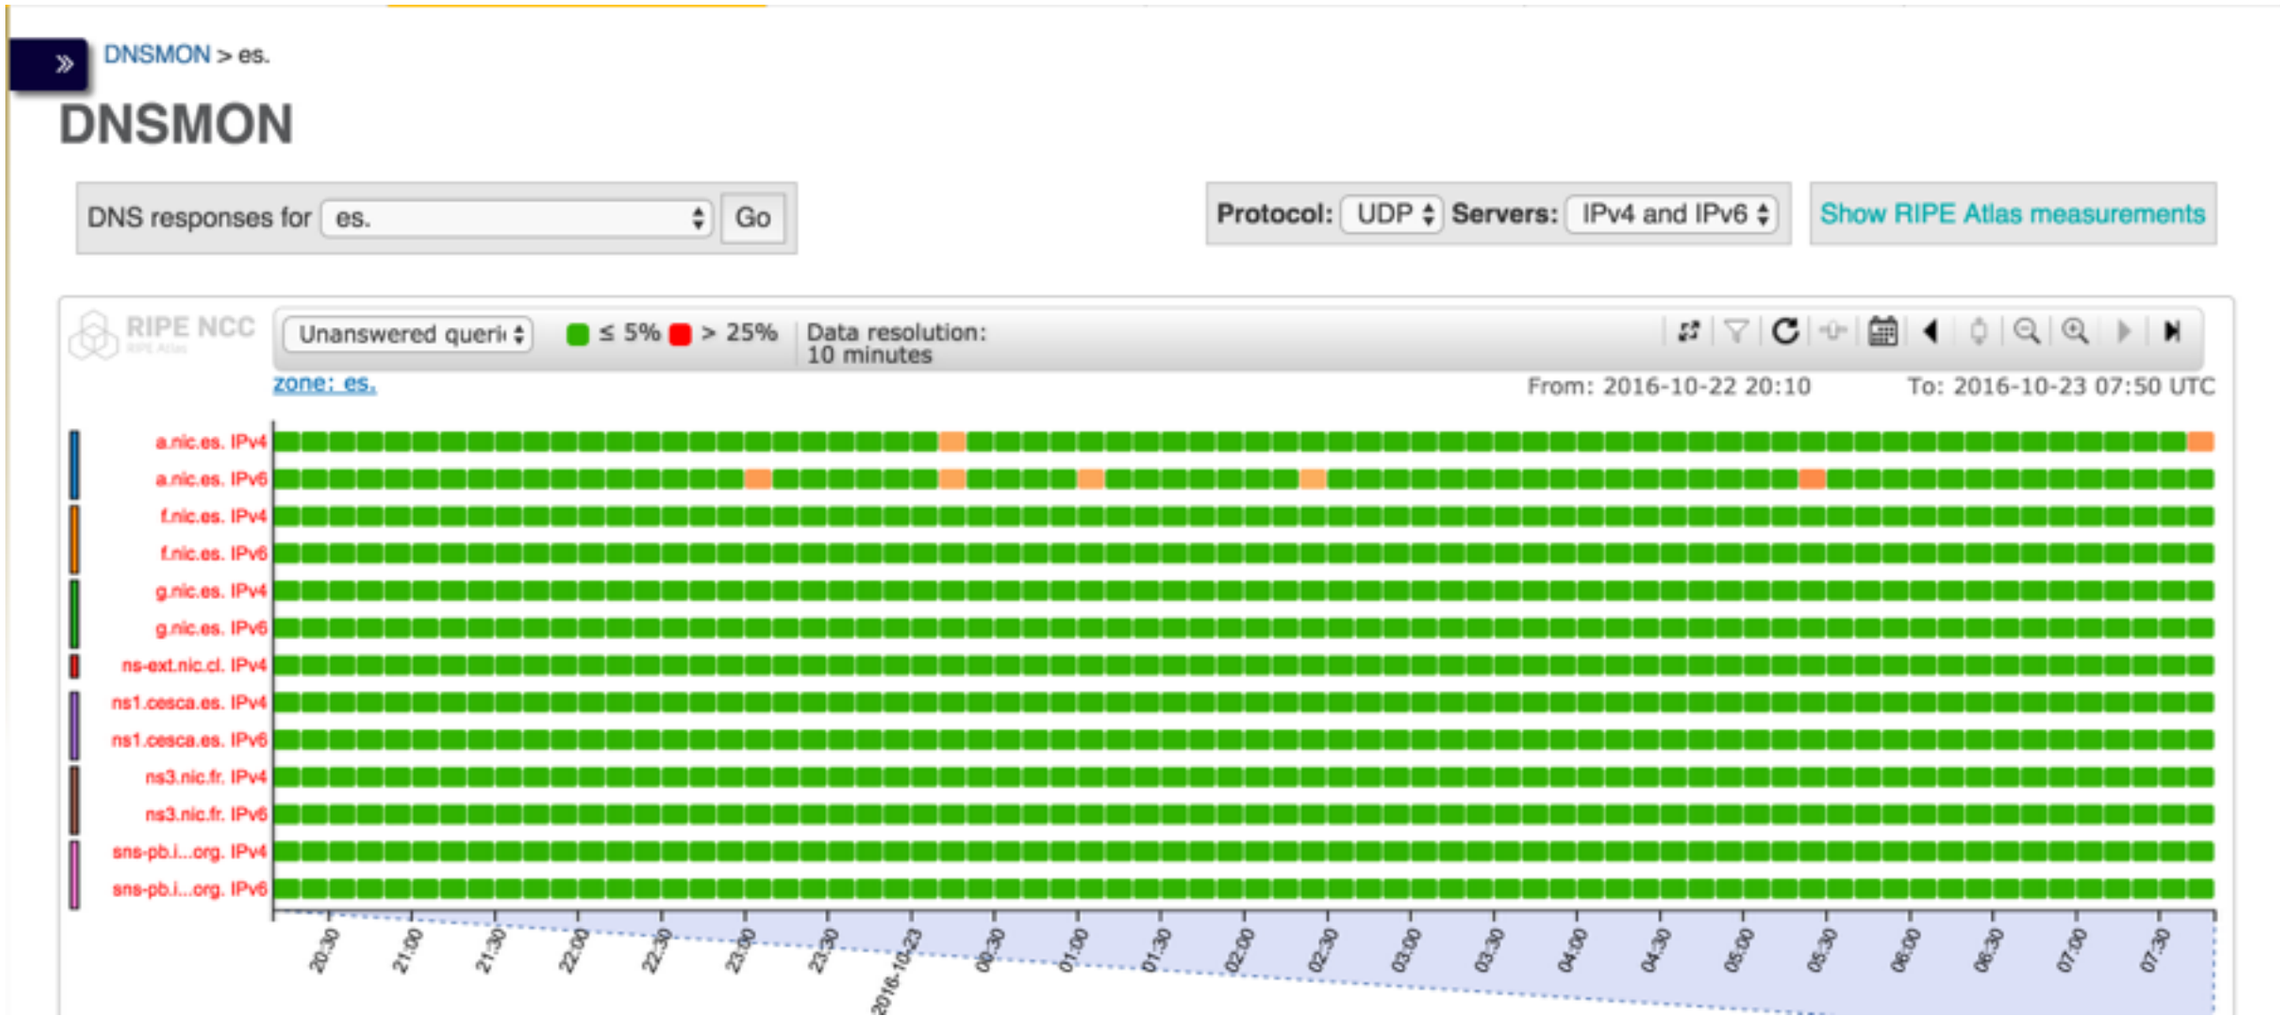

Domains in DNSMON

- RIPE 661 has criteria about which ccTLDs qualify
  - ccTLDs (including IDN variants) in the RIPE NCC service region
  - ccTLDs not in the RIPE NCC service region, whose administrative or technical contacts in the IANA database are RIPE NCC members

DomainMON

- DNSMON's little brother
- Monitor any domain - as long as you have RIPE Atlas credits
- Monitoring from probes, rather than anchors
- Visualisation similar to DNSMON

Server Benchmarking

Motivation

- DDoS attacks are becoming bigger
- We want to have a more flexible upgrade path, including (multiple) 10G connections
- We need to understand how our OS and DNS software performs at such speeds

Test Setup

- One switch with 10 Gbit/s ports
- Three Dell servers with 10 Gbit/s interfaces
  - Intel Xeon E5-2470 2.4 Ghz, 10-core processor
  - First server is the source of queries
  - Second server runs the DNS software
  - Third server is the sink
- BIND 9.10, Knot DNS 1.6 and 2, NSD 4, Yadifa 2.2 and PowerDNS 4

Query Source Server

- tcpreplay, compiled with Quick_TX
  - Writes packets directly to an interface
- 5-minute pcap file with queries from a K-root server
- 6,882,162 packets sent each time
  - Should generate 6,882,142 responses (20 dud query packets)

DNS Server

- All DNS software configured with root, arpa and root-servers.net zones
- Routing set up to send all responses to sink server
- iptables with PREROUTING and OUTPUT chains of “raw” table to count queries and responses
- Target NOTRACK avoids keeping state

Response Sink Server

- iptables with PREROUTING chain of “raw” table to count and drop responses

Test Runs

- Three runs of tcpreplay, recording the best result
  - Started at 100,000 q/s
  - Ramped up by 100,000 each time until name server shows loss
  - Finished with maximum rate (tcpreplay's -t option)

Summary of Results

- CentOS 6 doesn't work - its old kernel/drivers lost 85% of the packets
  - Need to upgrade to CentOS 7
- NSD 4 is the best performer, as long as:
  - “server-count” is increased from 1 to number of CPUs
  - “reuseport” is set to “yes”
- Faster CPUs and more cores are required to saturate a 10 Gbit/s network interface

Reverse DNS

Pre-delegation testing

DNScheck -> Zonemaster

- DNScheck has been abandoned
- Zonemaster:
  - Has better tests
  - Handles newer DNSSEC algorithms
  - Is being actively developed by IIS and AFNIC
- Migration from DNScheck to Zonemaster will happen in the coming weeks
  - Users may see slightly different diagnostic output

Inter-RIR Transfers

Reverse DNS

Reverse DNS provisioning

- Reverse DNS delegation accepted through RIPE Database
- Automatic provisioning in parent zones
  - Software adjusts itself for transferred space
  - Publishes zonelets for other RIRs to pick up
  - Downloads zonelets from other RIRs to merge in delegation records
- Some delegation can take up to 24 hours to be published

Questions

anandb@ripe.net
@aabdnn