

# DNS: the Root, DNSSEC and Operations ... living in a DNSSEC enabled world

Jim Martin, ISC Operations Manager

[jrmii@isc.org](mailto:jrmii@isc.org)



# The Plan

- The DNS Root
- A bit of History
- Several DNSSEC operational incidents
- Tools and Recommendations
- **NOTE:** This is the impact of DNSSEC for the network and systems types. DNS Admins will need to go much deeper.

# DNS Root Myths

- Where all DNS Queries begin!

*Well, if you don't have anything cached*

- Holds all the “IMPORTANT” zones

*Only needs to hold “.”*

- Consists of 13 servers, A-Root through M-Root

*13 Instances, but MANY more servers*



# DNS Root Servers January 2011

# And then came DNSSEC

- DNSSEC signing the root was a BIG deal
- Months of planning
- Repeated DITLs (Day in the Life)
- Then added the DURZ (Deliberately Unvalidatable Root Zones)
- Finally signed legitimately on July 15, 2010

# And of course the TLDs

- Many TLDs and ccTLDs signed already ... .org, .net, .edu, .gov, .se, .au, .jp, .cz, .fr, and many more (though some may still be in testing)
- .nz is close ... scheduled end of Jan 2011
- .com is scheduled for end of Q1 2011

# Network Impact of DNSSEC

- Signed DNS responses are BIG
  - Have DS, NSEC, DNSKEY, & RRSig data
  - Dramatically increases query response sizes
  - 512 byte UDP packets just don't cut it
  - EDNS0 is no longer “nice to have”

# Just how much bigger?

- Without DNSSEC

```
Jim@131-203-50-204:~/>dig +nodnssec www.isc.org

;<<>> DiG 9.6.0-APPLE-P2 <<>> +nodnssec www.isc.org
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 2203
;; flags: qr rd ra; QUERY: 1,ANSWER: 1,AUTHORITY: 4,
ADDITIONAL: 8

;; QUESTION SECTION:
;www.isc.org.          IN      A

;; ANSWER SECTION:
www.isc.org.          547     IN      A      149.20.64.42

{ LOTS REMOVED }

;; Query time: 40 msec
;; SERVER: 131.203.1.5#53(131.203.1.5)
;; WHEN: Wed Jan 26 17:53:27 2011
;; MSG SIZE rcvd: 320
```

- With DNSSEC

```
Jim@Bikeshed:/data/users/jrmii>dig www.isc.org. a +dnssec @204.152.187.13

;<<>> DiG 9.6.2-P2 <<>> www.isc.org. a +dnssec @204.152.187.13
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 51546
;; flags: qr rd ra ad; QUERY: 1,ANSWER: 2,AUTHORITY: 5,ADDITIONAL: 13

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;www.isc.org.          IN      A

;; ANSWER SECTION:
www.isc.org.          442     IN      A      149.20.64.42
www.isc.org.          442     IN      RRSIG  A 5 3 600 20110221233210
20110122233210 26982 isc.org. ZPrxCONvy/
c2FEKmcEgKD7rS3YCI f4RL9Du3hIw6/XcuIYOAzhFA33ZG j/
Q2d9GqG5oTWkflkTyVDg68fOrpNhvc0nKIOTUoT7GWu4Q6odMx0iAh
II l/dlchktSd2amBap3MOLcMcPAlY4AKfaceDss8DIHrrQTQOWyhn4RI IWw=

{ LOTS MORE REMOVED }

;; Query time: 1 msec
;; SERVER: 204.152.187.13#53(204.152.187.13)
;; WHEN: Wed Jan 26 05:45:12 2011
;; MSG SIZE rcvd: 1623
```

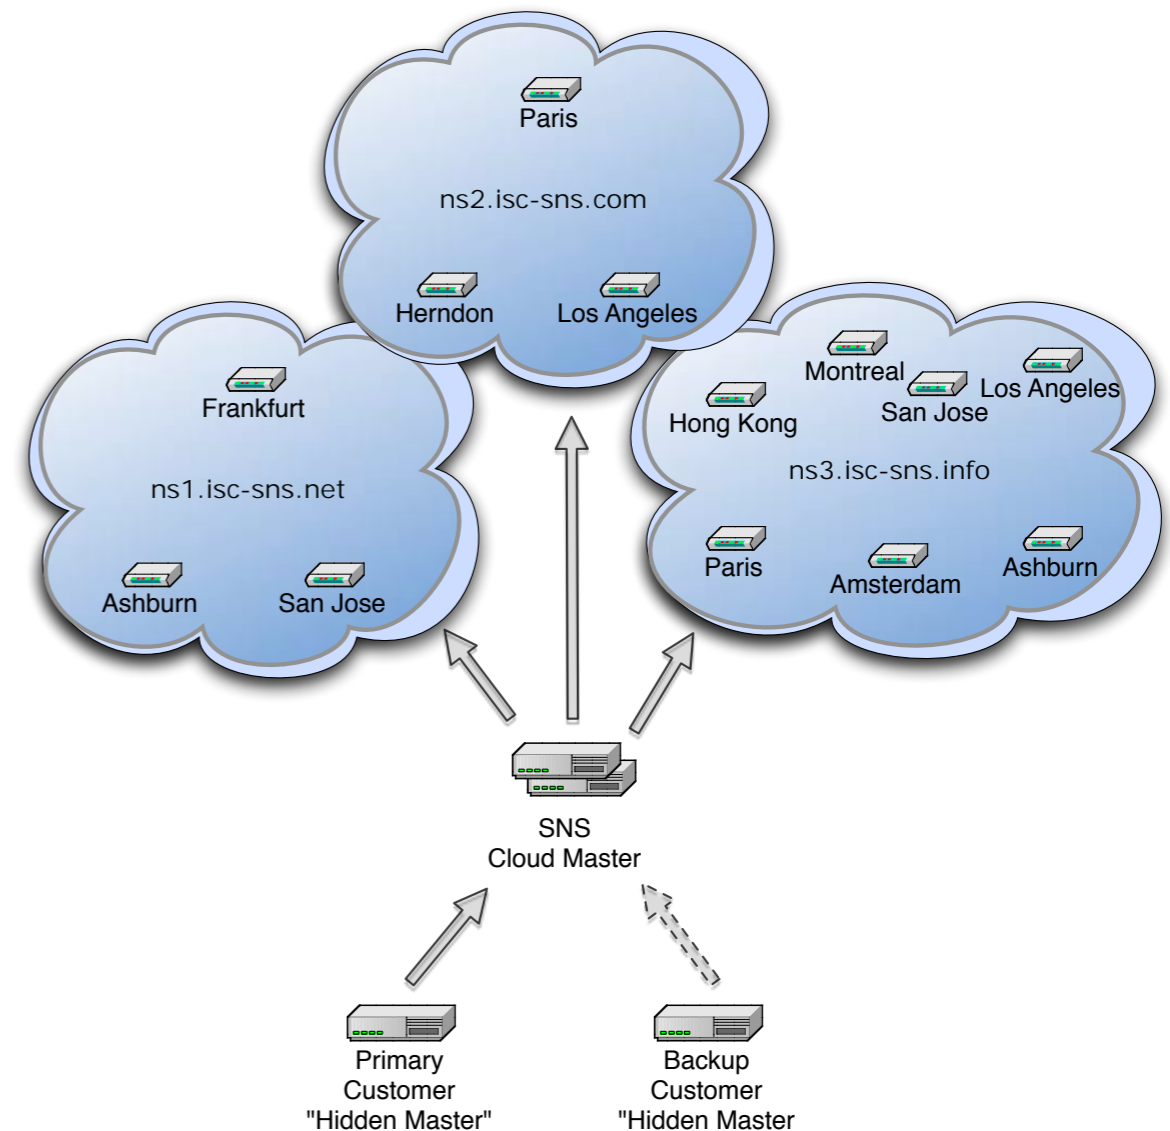


# EDNS0

- Extension Mechanisms for DNS-RFC 2671
  - Allows for bigger DNS messages
  - Uses IP Fragments
  - Recommended maximum of 4K
  - NOT on by default in all firewalls
  - NOT a given in all home “routers”

# ISC Secondary Name Service (SNS)

- Provide both free (SNS-PB) and SLA-Backed (SNS-Com) DNS Secondary Service
- 3 separate AnyCast Clouds with multiple providers
- Largely IPv6 Enabled
- Fully DNSSEC capable



# Simple DNSSEC Failure Behavior

- Content provider DNSSEC signs their zones and gets proper DS records installed
- End user tries to look up FQDN in that zone (eg, www.foo.org) and it fails
- End user believes that www.foo.org is “down”

# Simple DNSSEC Failure Cause

- Client unwittingly sets the bits to allow validation
- {Firewall, middlebox, CPE} not EDNS0-compliant
- {Firewall, middlebox, CPE, host firewall} explicitly blocks IPv4 Fragments
- Recursive resolver not DNSSEC capable or doesn't have the root anchors installed

# More Obscure Failure Behavior

- We host a large multi-national Internet property in SNS
- Their zones were NOT signed
- Some users couldn't successfully resolve records in that domain.

# The cause of that more obscure failure

- The zone of the content provider were NOT signed
- NS records for the zone referenced records that WERE in a signed zone
- The resultant responses popped above the 512 byte limit, and we're back at the same behavior as the simple case

# Key Rollover

- Periodically the Signing Keys (KSK/ZSK) should be changed
- During the period while ANY server could be passing out the old key, both the old and new keys are included in responses
- Yupp, that response just got bigger!
- This mostly impacts places where EDNS0 is enabled, but a “conservative” (often 1K) limit is chosen

# But I don't have a problem... really!

- You sure? Use the OARC Reply Size test!
- When it works:

```
Jim@131-203-50-204: />dig +short rs.dns-oarc.net txt
rst.x4091.rs.dns-oarc.net.
rst.x3837.x4091.rs.dns-oarc.net.
rst.x3843.x3837.x4091.rs.dns-oarc.net.
"Tested at 2011-01-26 03:52:33 UTC"
"202.53.189.253 sent EDNS buffer size 4096"
"202.53.189.253 DNS reply size limit is at least 4091"
```

- When it doesn't:

```
rivendel:~ jim$ dig +short rs.dns-oarc.net txt
rst.x476.rs.dns-oarc.net.
rst.x485.x476.rs.dns-oarc.net.
rst.x490.x485.x476.rs.dns-oarc.net.
"68.87.76.181 DNS reply size limit is at least 490"
"68.87.76.181 lacks EDNS, defaults to 512"
"Tested at 2011-01-26 05:56:41 UTC"
```



# BIND

- Has been doing EDNS0 since 8.3.0
- Got DNSSEC (bis) in 9.3.0
- BUT has known flaws for anything before 9.4-ESV
- It's highly recommended that you run 9.7.2-P3 or 9.6-ESV-R3

# Key Take Aways

- Make sure your BIND install is 9.7.2 or 9.6-ESV
- Make sure all your network elements that touch DNS can do EDNS0 and allow 4K responses
- Make sure any network security elements allow IP Fragments
- Use the reply size tester to validate your systems and to identify customer problems

# References

- **DNS Root Servers**

<http://www.root-servers.org/>

<http://www.root-dnssec.org/>

- **DNSSEC**

<http://dnssec.net/>

- **Test Tools**

<https://www.dns-oarc.net/oarc/services/replysizetest>

- **Name Server (BIND)**

<http://www.isc.org>

<http://www.isc.org/software/bind/versions>

# Questions?

- While you're thinking of questions:
  - If you want to peer with F-Root, send mail to [peering@isc.org](mailto:peering@isc.org)
  - We host public-benefit organizations through our Hosted@ and SNS-PB programs. Contact [{hosted,sns}@isc.org](mailto:{hosted,sns}@isc.org)
  - Remember ISC is a public-benefit and survives through donations, forum memberships, SNS-Com and support contracts.
  - We appreciate any help, and need it to keep doing good work!