ausregistry
INTERNATIONAL

The TLD
Specialists

DNSSEC with Parent and Child Zones on the same Name Server

Chris Wright

CTO – AusRegistry International

Structure of .au

- Multi-level registration model all managed by the one Registry System
  - .au
    - Not open for registration, but does contain some records
  - 2LDs e.g. com.au, gov.au
    - Open for registration
    - Main area of registration in .au
    - Different policies and eligibility criteria for 2LDs
  - 3LDs e.g. vic.gov.au, nsw.edu.au
    - Open for registration
    - State level government and education departments

DNS Setup

- DNS Servers
  - Most servers are authoritative for more than one level of the hierarchy (some not yet, legacy issues)
  - Most serve au, gov.au, <state>.gov.au i.e. all 3 levels of the hierarchy
  - Name servers are in zone and glued at parent
- Why?
  - Name Servers will deliver the most specific response they can
  - In zone servers are glued, thus addresses will be included in referrals
  - Reduces the overall query load on the system as a whole
  - Decreases response time
  - When root refers to .au servers for www.health.vic.gov.au because the au servers are authoritative for vic.gov.au they will return the 'health' referral straight away

How does DNSSEC effect this?

- It doesn’t break it, nor does DNSSEC break our ability to do this
- It does potentially reduce the benefit
- It does mean that we need to plan for a significant increase in the number of queries
  - Beyond the normal expected increases when DNSSEC signing a zone i.e.
    - Query size increase
    - Query count increase with queries for the DNSKEY record
    - Query count increases when resolvers switch to TCP (demonstrated most in .org)
- Lets show why…

The Test System

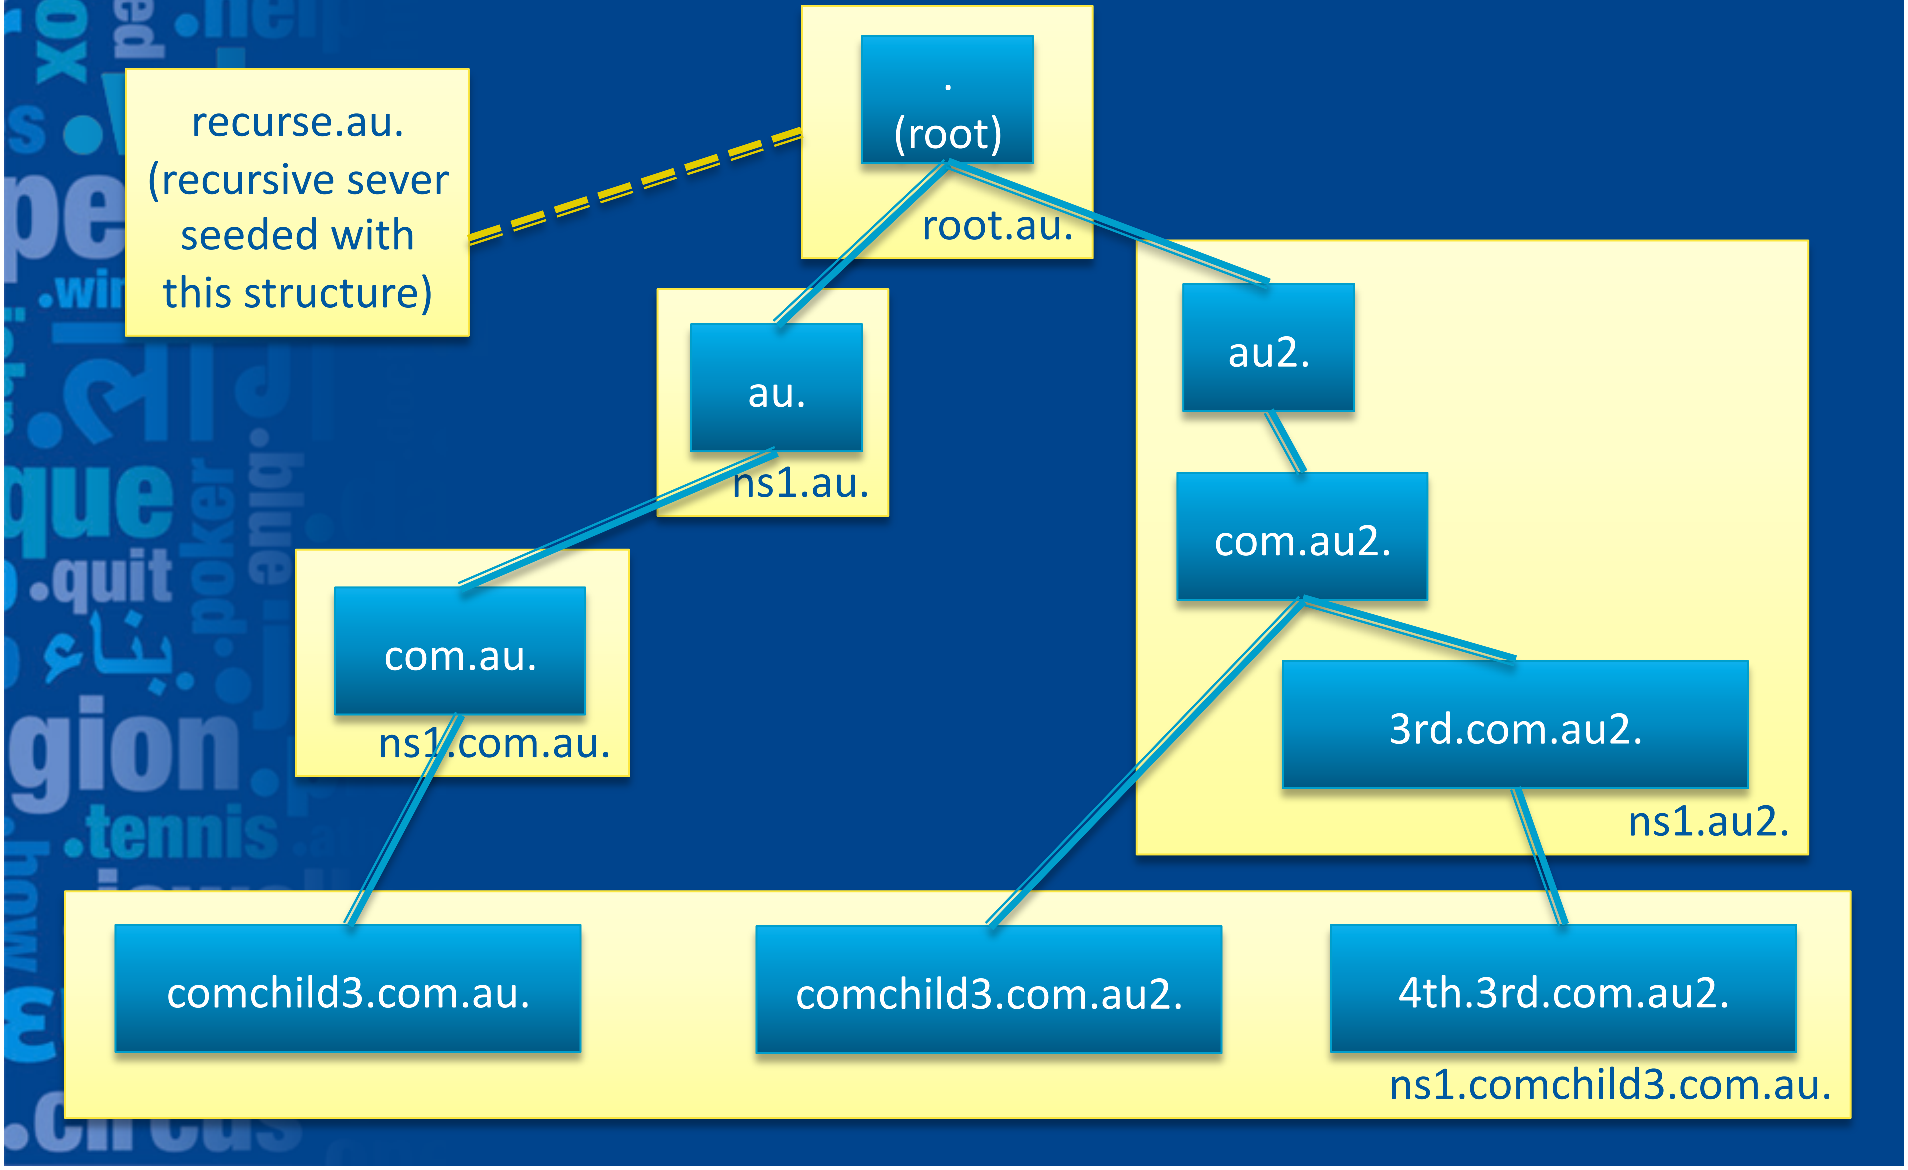

Non-DNSSEC 3rd level Query – No Common Name Servers

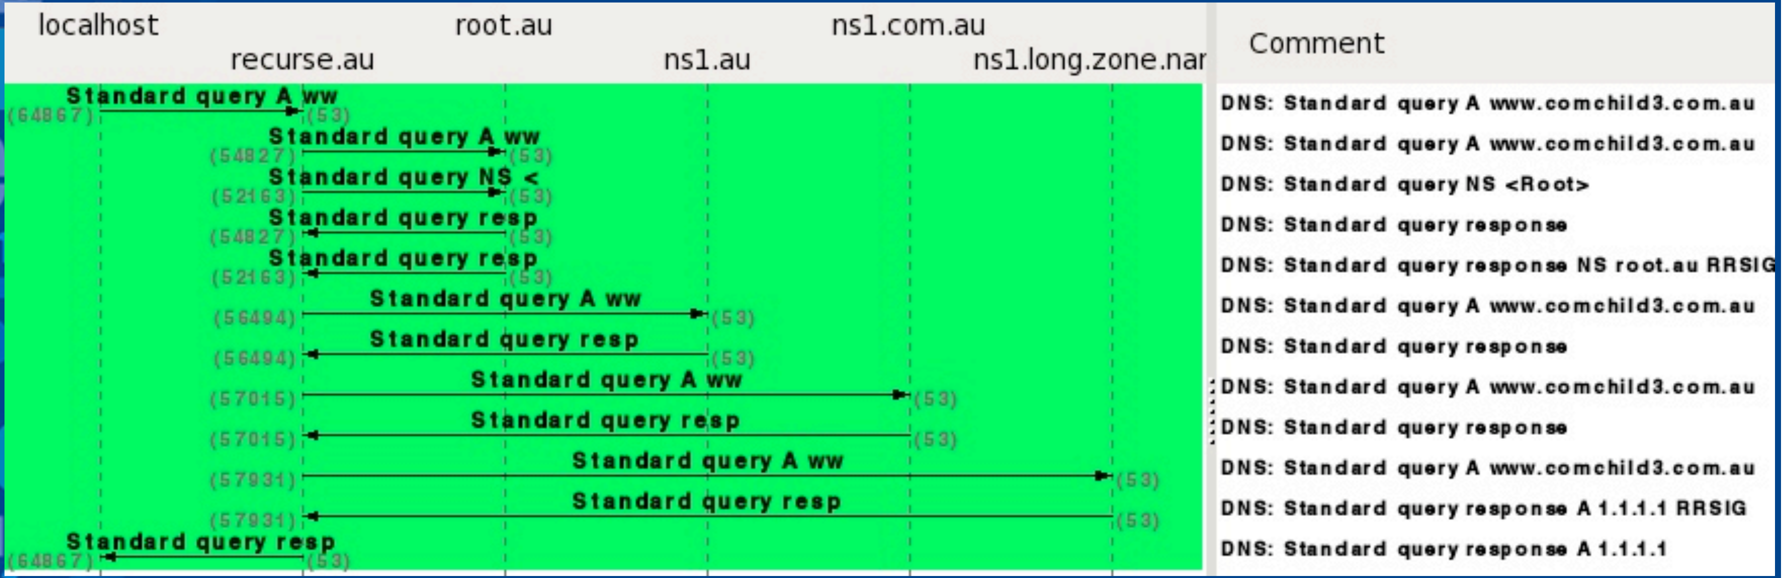
localhost
recurse.au
root.au
ns1.au
ns1.com.au
ns1.long.zone.nar
Comment
(64867) Standard query A ww (53)
(54827) Standard query A ww (53)
(52163) Standard query NS < (53)
(54827) Standard query resp (53)
(52163) Standard query resp (53)
(56494) Standard query A ww (53)
(56494) Standard query resp (53)
(57015) Standard query A ww (53)
(57015) Standard query resp (53)
(57931) Standard query A ww (53)
(57931) Standard query resp (53)
(64867) Standard query resp (53)

```
DNS: Standard query A www.comchild3.com.au
DNS: Standard query A www.comchild3.com.au
DNS: Standard query NS <Root>
DNS: Standard query response
DNS: Standard query response NS root.au RRSIG
DNS: Standard query A www.comchild3.com.au
DNS: Standard query response
DNS: Standard query A www.comchild3.com.au
DNS: Standard query response
DNS: Standard query A www.comchild3.com.au
DNS: Standard query response A 1.1.1.1 RRSIG
DNS: Standard query response A 1.1.1.1
```

- Query for www.comchild3.com.au
- 4 queries required (1 on root, 2 on .au infrastructure, 1 on registrant)
  - root
  - au
  - com.au
  - comchild3.com.au

Non-DNSSEC 3rd level Query – Common Name Server for TLD and 2LD

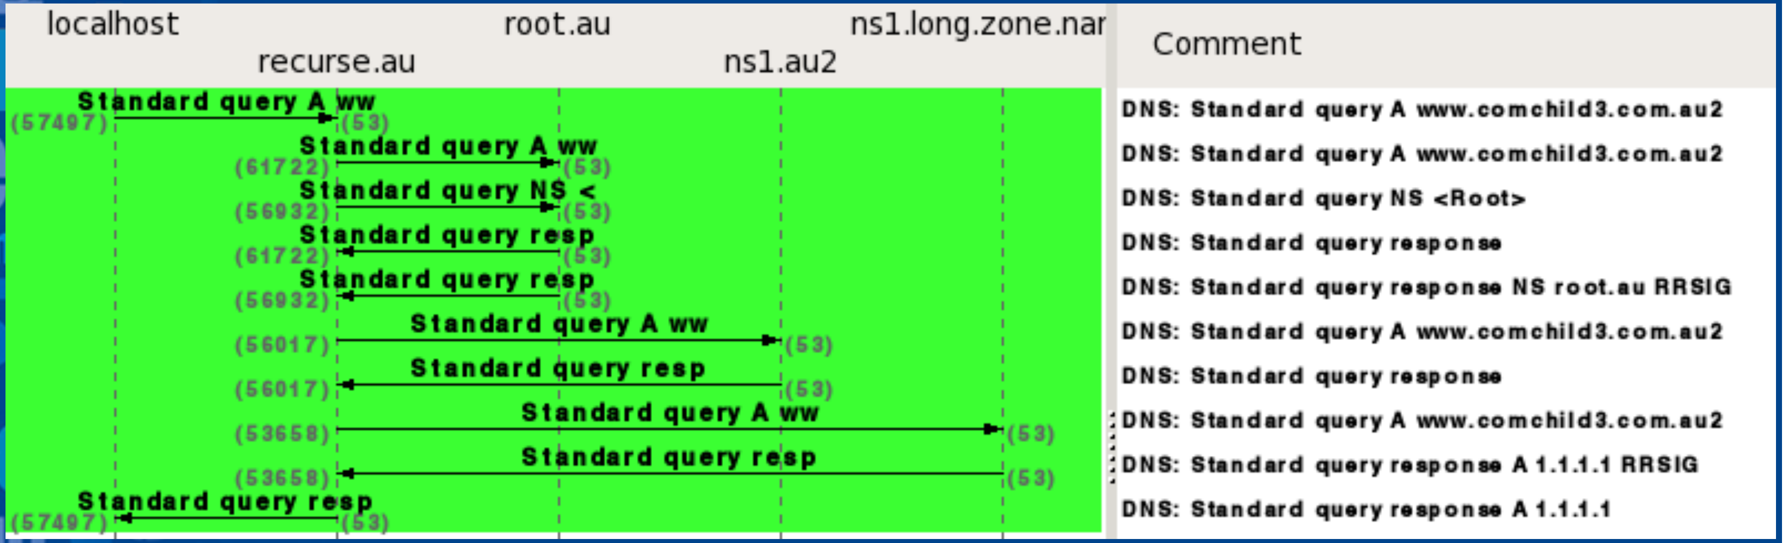

```
Comment
DNS: Standard query A www.comchild3.com.au2
DNS: Standard query A www.comchild3.com.au2
DNS: Standard query NS <Root>
DNS: Standard query response
DNS: Standard query response NS root.au RRSIG
DNS: Standard query A www.comchild3.com.au2
DNS: Standard query response
DNS: Standard query A www.comchild3.com.au2
DNS: Standard query response A 1.1.1.1 RRSIG
DNS: Standard query response A 1.1.1.1
```

- Query for www.comchild3.com.au2
- Only 3 Queries required (1 on root, 1 on .au infrastructure, 1 on registrant)
  - root
  - au2/com.au2
  - comchild3.com.au2

Non-DNSSEC 4th level Query – Common Name Server for TLD, 2LD & 3LD

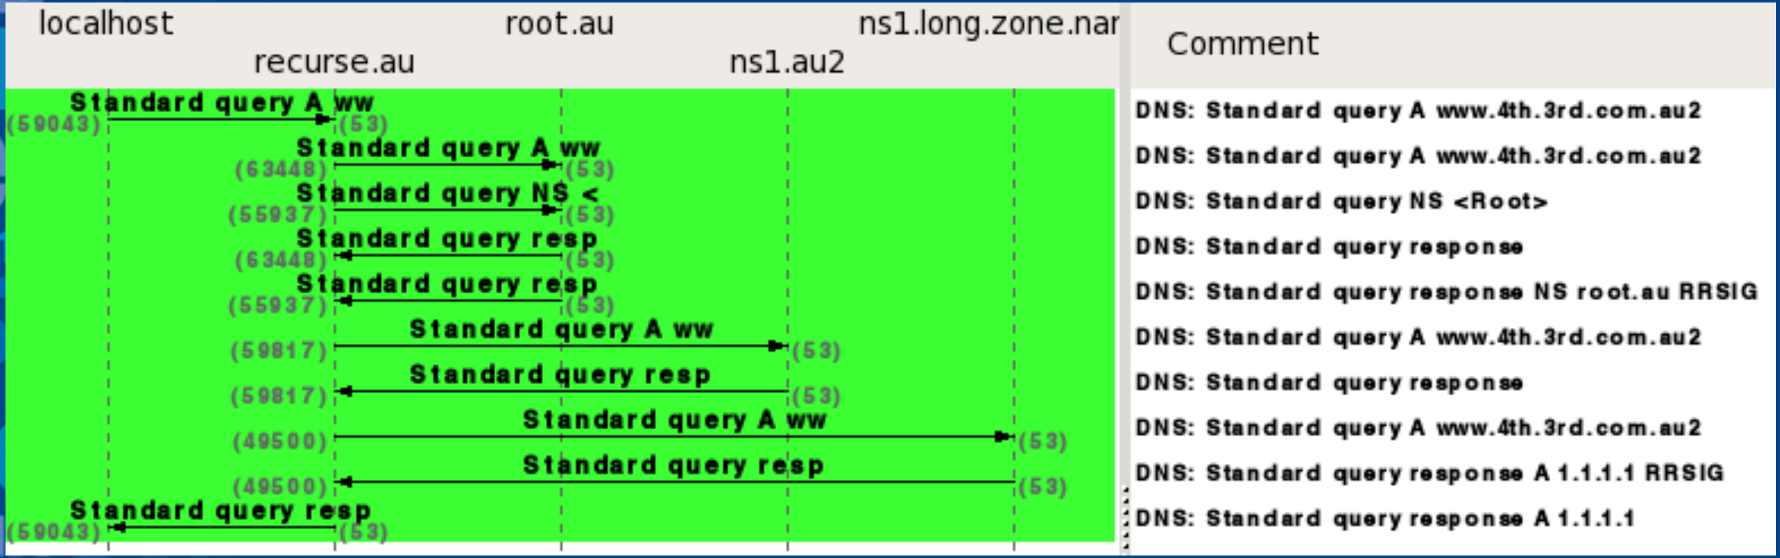

Comment

```
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query NS <Root>
DNS: Standard query response
DNS: Standard query response NS root.au RRSIG
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query response
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query response A 1.1.1.1 RRSIG
DNS: Standard query response A 1.1.1.1
```

- Query for www.4th.3rd.com.au2
- Still only 3 queries required (1 on root, 1 on .au infrastructure, 1 on registrant)
  - root
  - au2/com.au2/3rd.com.au2
  - 4th.3rd.com.au2
- As apposed to 5 if individual name servers were used (1 on root, 3 on .au infrastructure, 1 on registrant)

DNSSEC 3rd level Query – No Common Name Servers

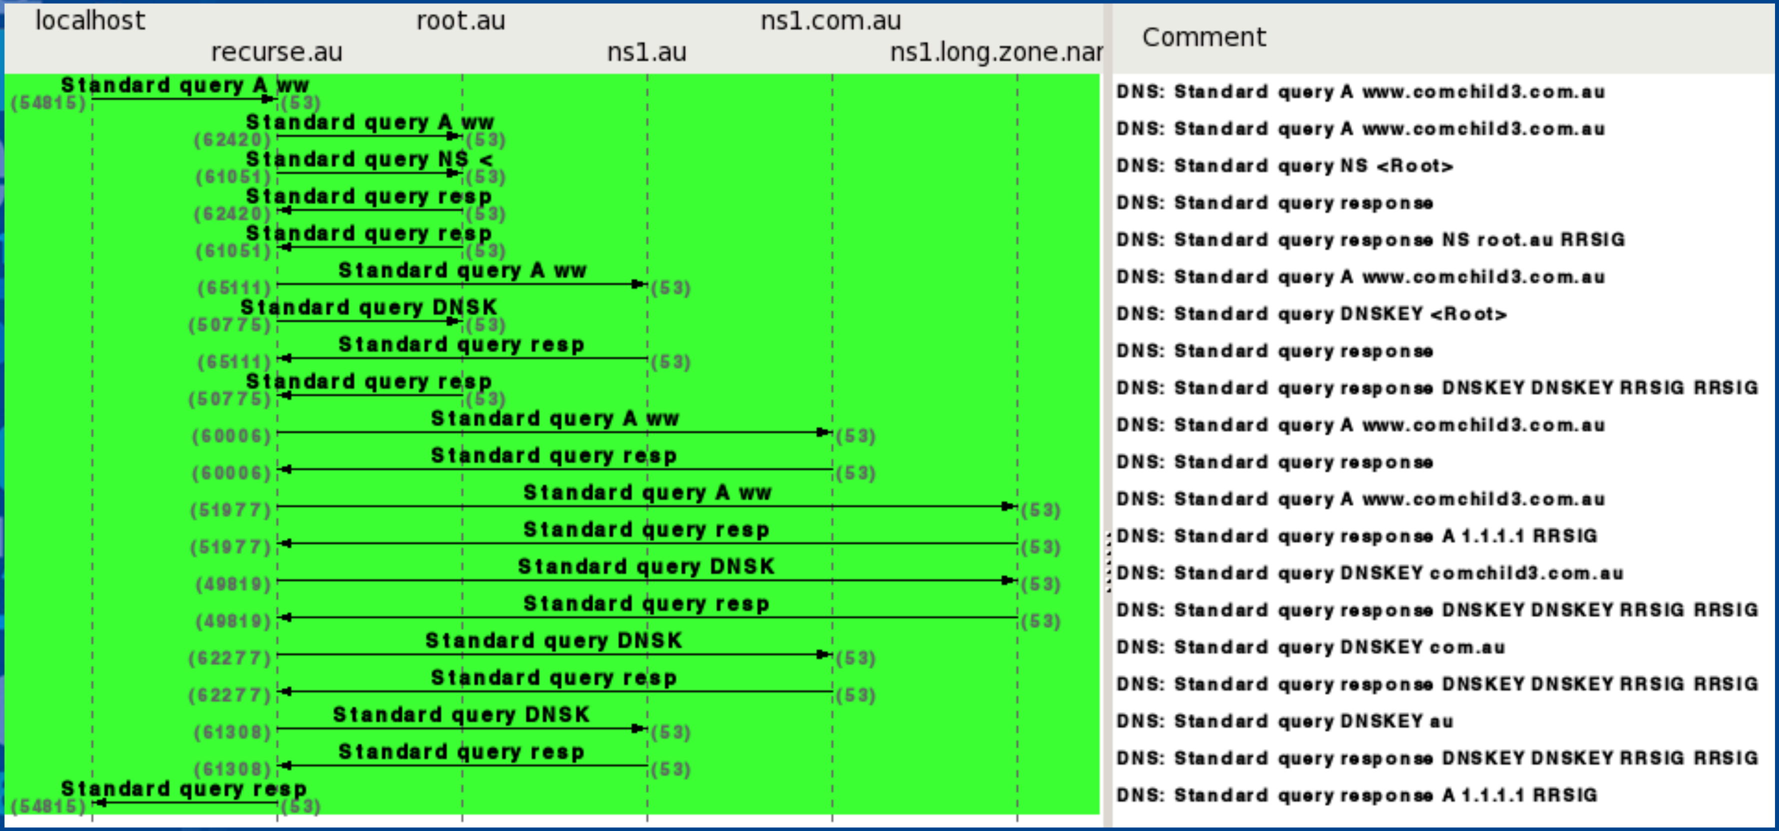

			Comment
Standard query A ww	(54815)	(53)	DNS: Standard query A www.comchild3.com.au
Standard query A ww	(62420)	(53)	DNS: Standard query A www.comchild3.com.au
Standard query NS <	(61051)	(53)	DNS: Standard query NS <Root>
Standard query resp	(62420)	(53)	DNS: Standard query response
Standard query resp	(61051)	(53)	DNS: Standard query response NS root.au RRSIG
Standard query A ww	(65111)	(53)	DNS: Standard query A www.comchild3.com.au
Standard query DNSK	(50775)	(53)	DNS: Standard query DNSKEY <Root>
Standard query resp	(65111)	(53)	DNS: Standard query response
Standard query resp	(50775)	(53)	DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query A ww	(60006)	(53)	DNS: Standard query A www.comchild3.com.au
Standard query resp	(60006)	(53)	DNS: Standard query response
Standard query A ww	(51977)	(53)	DNS: Standard query A www.comchild3.com.au
Standard query resp	(51977)	(53)	DNS: Standard query response A 1.1.1.1 RRSIG
Standard query DNSK	(49819)	(53)	DNS: Standard query DNSKEY comchild3.com.au
Standard query resp	(49819)	(53)	DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query DNSK	(62277)	(53)	DNS: Standard query DNSKEY com.au
Standard query resp	(62277)	(53)	DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query DNSK	(61308)	(53)	DNS: Standard query DNSKEY au
Standard query resp	(61308)	(53)	DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query resp	(54815)	(53)	DNS: Standard query response A 1.1.1.1 RRSIG

DNSSEC 3rd level Query – No Common Name Servers

- DNSSEC query for www.comchild3.com.au
- 8 queries in total as expected (2 on root, 4 on .au, 2 on registrant)
  - root referral
  - root DNSSKEY
  - au referral
  - au DNSSKEY
  - com.au referral
  - com.au DNSSKEY
  - comchild3.com.au answer
  - comchild3.com.au DNSKEY
- DS records ‘piggy back’ with referral responses
- Chain of trust at zone cuts (.,au.,com.au.,comchild3.com.au.) established as tree is walked
- Increase of 4 over non-signed (100% increase)
- Increase of 2 on .au infrastructure (100% increase)

DNSSEC 3rd level Query – Common Name Server for TLD and 2LD

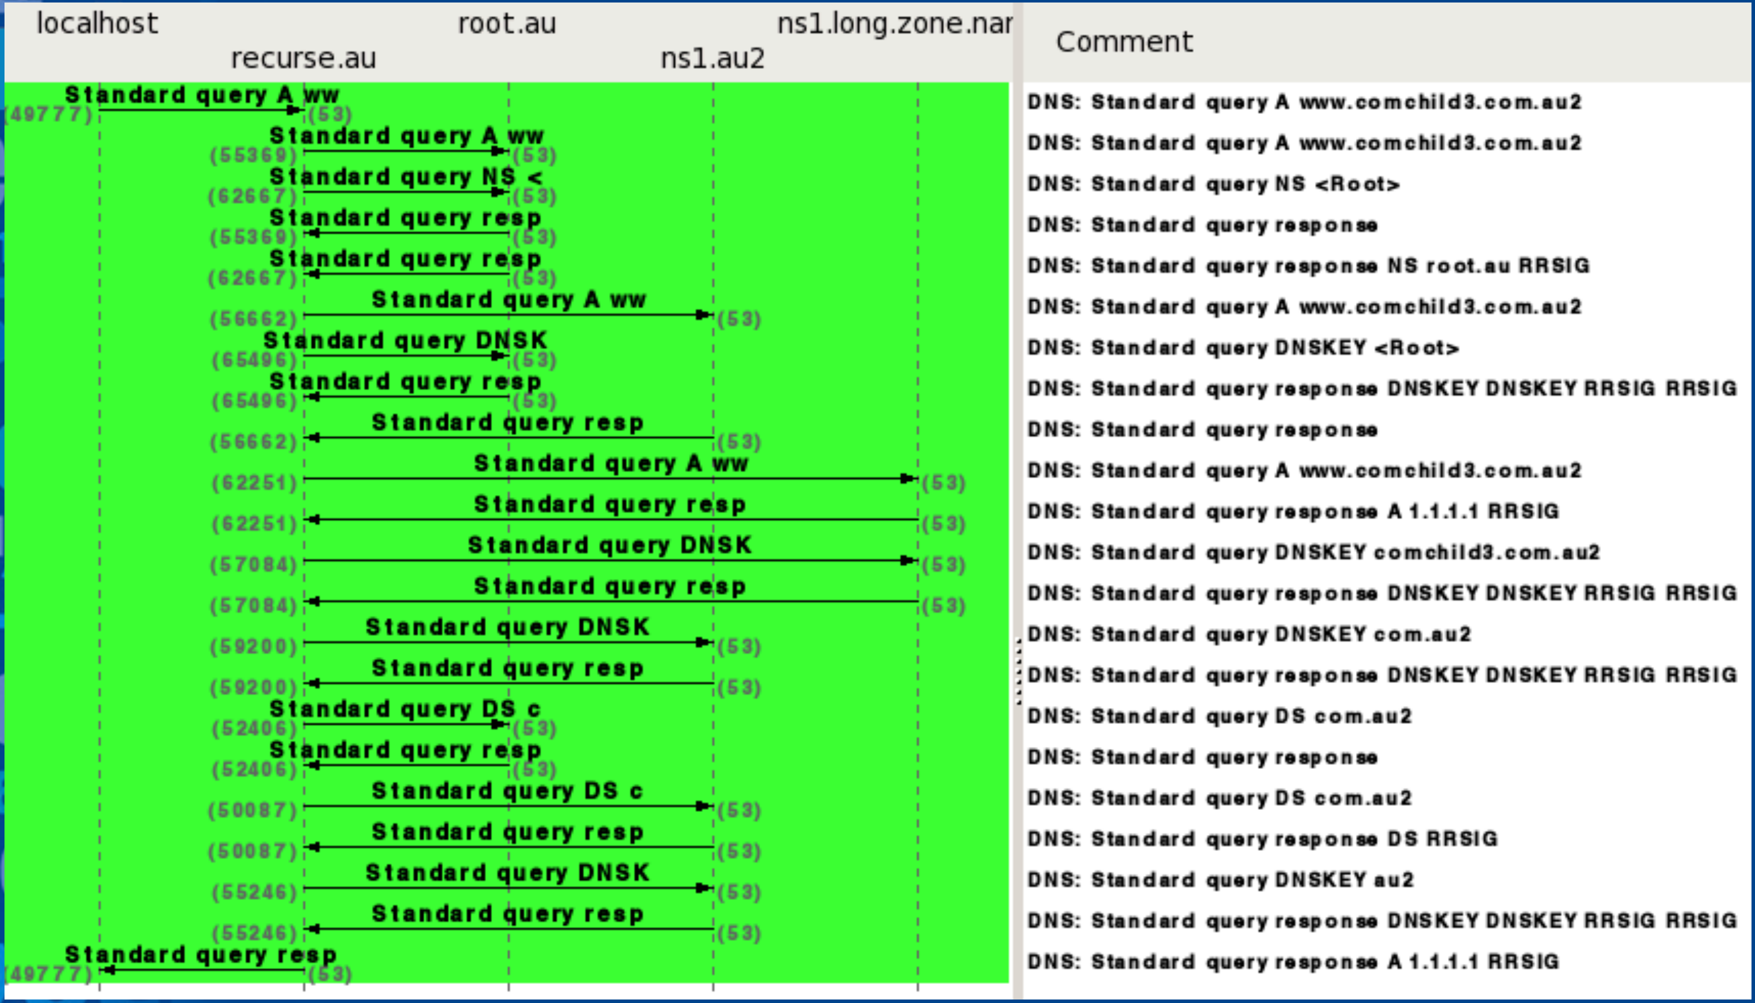

```
localhost        root.au        ns1.long.zone.nar        Comment
      recurse.au        ns1.au2

Standard query A ww       (49777) (53)    DNS: Standard query A www.comchild3.com.au2
Standard query A ww       (55369) (53)    DNS: Standard query A www.comchild3.com.au2
Standard query NS <       (62667) (53)    DNS: Standard query NS <Root>
Standard query resp       (55369) (53)    DNS: Standard query response
Standard query resp       (62667) (53)    DNS: Standard query response NS root.au RRSIG
Standard query A ww       (56662) (53)    DNS: Standard query A www.comchild3.com.au2
Standard query DNSK       (65496) (53)    DNS: Standard query DNSKEY <Root>
Standard query resp       (65496) (53)    DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query resp       (56662) (53)    DNS: Standard query response
Standard query A ww       (62251) (53)    DNS: Standard query A www.comchild3.com.au2
Standard query resp       (62251) (53)    DNS: Standard query response A 1.1.1.1 RRSIG
Standard query DNSK       (57084) (53)    DNS: Standard query DNSKEY comchild3.com.au2
Standard query resp       (57084) (53)    DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query DNSK       (59200) (53)    DNS: Standard query DNSKEY com.au2
Standard query resp       (59200) (53)    DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query DS c       (52406) (53)    DNS: Standard query DS com.au2
Standard query resp       (52406) (53)    DNS: Standard query response
Standard query DS c       (50087) (53)    DNS: Standard query DS com.au2
Standard query resp       (50087) (53)    DNS: Standard query response DS RRSIG
Standard query DNSK       (55246) (53)    DNS: Standard query DNSKEY au2
Standard query resp       (55246) (53)    DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
Standard query resp       (49777) (53)    DNS: Standard query response A 1.1.1.1 RRSIG
```

DNSSEC 3rd level Query – Common Name Server for TLD and 2LD

- DNSSEC query for www.comchild3.com.au2
- 9 queries in total (3 on root, 4 on .au, 2 on registrant)
  - root referral
  - root DNSSKEY
  - com.au2 referral
  - com.au2 DNSSKEY
  - comchild3.com.au2 answer
  - comchild3.com.au2 DNSKEY
  - root referral for com.au2 DS (shouldn’t this be cached from above?)
  - au2 for com.au2 DS
  - au2 DNSKEY
- DS records ‘piggy back’ with referral responses but no referral from au2 to com.au2 was ever received (wasn’t required as server had a more specific answer) thus it had to be specifically queried for, this involved walking the tree again!
- Chain of trust at zone cuts (.,au.,com.au.,comchild3.com.au.) required extra queries to be established as referral was ‘missed’

DNSSEC 3rd level Query – Common Name Server for TLD and 2LD

- Worse case
  - 9 queries in total (3 on root, 4 on .au, 2 on registrant)
  - Increase of 6 over non signed (200% increase)
  - Increase of 3 on .au infrastructure (300% increase)
  - Additional 1 query (12.5% above) the no common server scenario
  - Same number of queries as the no common server scenario on .au infrastructure
- Best Case (assuming cache is used)
  - 8 queries in total (2 on root, 4 on .au, 2 on registrant)
  - Increase of 5 over non signed (166.6% increase)
  - Increase of 3 on .au infrastructure (300% increase)
  - Same number of queries as the no common server scenario overall and on .au
- Will come back to why there is a best and worse case soon

DNSSEC 3rd level Query – Common Name Server for TLD, 2LD & 3LD

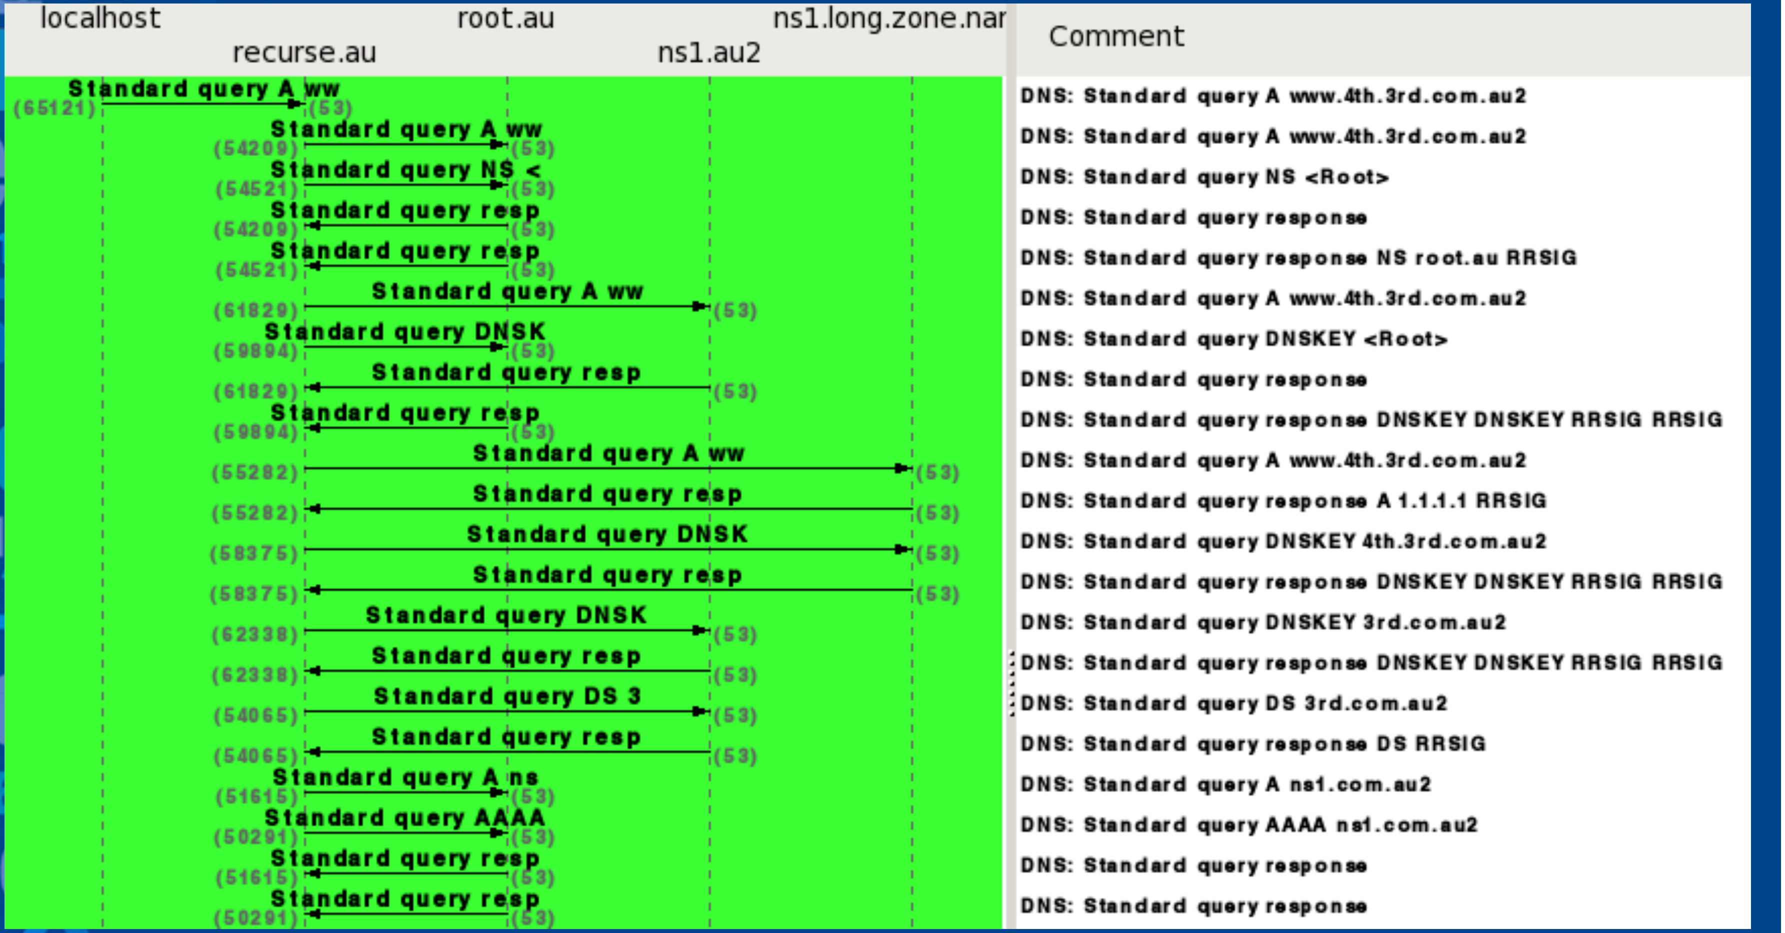
localhost
recurse.au
root.au
ns1.au2
ns1.long.zone.nar
Standard query A ww (65121) (53)
Standard query A ww (54209) (53)
Standard query NS < (54521) (53)
Standard query resp (54209) (53)
Standard query resp (54521) (53)
Standard query A ww (61829) (53)
Standard query DNSK (59894) (53)
Standard query resp (61829) (53)
Standard query resp (59894) (53)
Standard query A ww (55282) (53)
Standard query resp (55282) (53)
Standard query DNSK (58375) (53)
Standard query resp (58375) (53)
Standard query DNSK (62338) (53)
Standard query resp (62338) (53)
Standard query DS 3 (54065) (53)
Standard query resp (54065) (53)
Standard query A ns (51615) (53)
Standard query AAAA (50291) (53)
Standard query resp (51615) (53)
Standard query resp (50291) (53)

```
Comment
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query NS <Root>
DNS: Standard query response
DNS: Standard query response NS root.au RRSIG
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query DNSKEY <Root>
DNS: Standard query response
DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
DNS: Standard query A www.4th.3rd.com.au2
DNS: Standard query response A 1.1.1.1 RRSIG
DNS: Standard query DNSKEY 4th.3rd.com.au2
DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
DNS: Standard query DNSKEY 3rd.com.au2
DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
DNS: Standard query DS 3rd.com.au2
DNS: Standard query response DS RRSIG
DNS: Standard query A ns1.com.au2
DNS: Standard query AAAA ns1.com.au2
DNS: Standard query response
DNS: Standard query response
```

DNSSEC 3rd level Query – Common Name Server for TLD, 2LD & 3LD

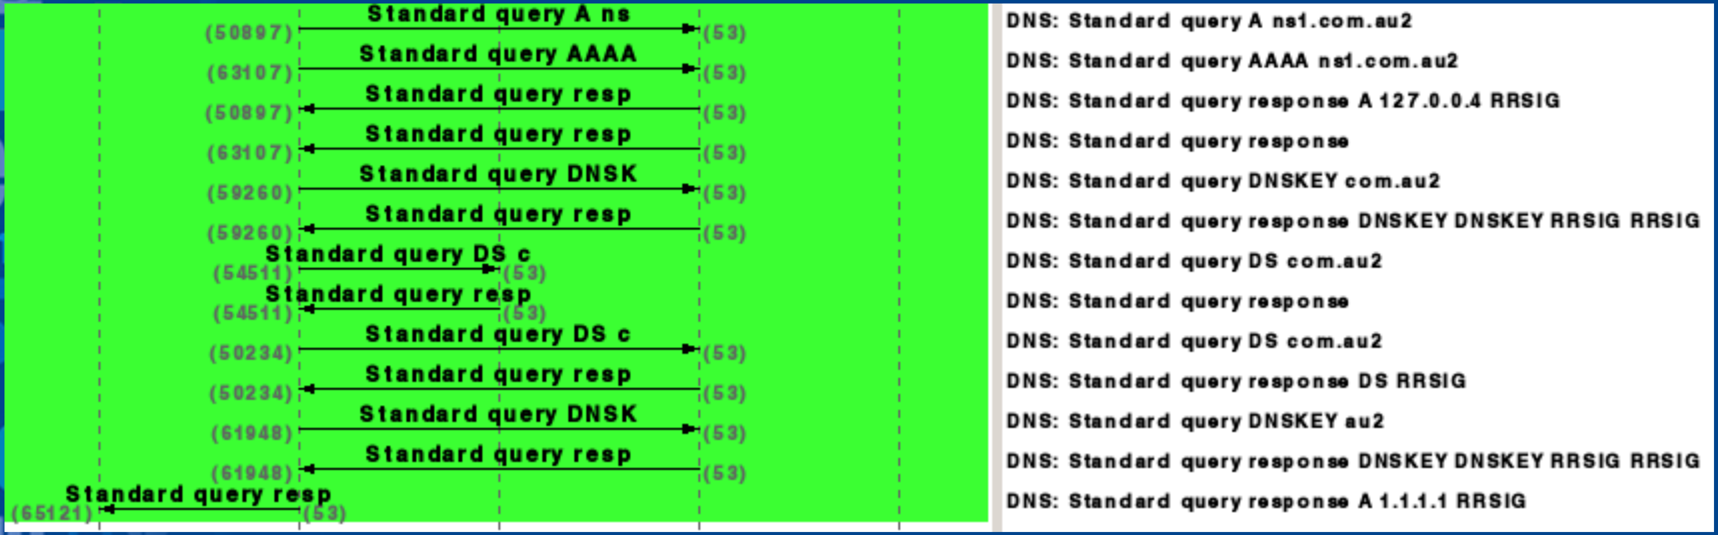

```
DNS: Standard query A ns1.com.au2
DNS: Standard query AAAA ns1.com.au2
DNS: Standard query response A 127.0.0.4 RRSIG
DNS: Standard query response
DNS: Standard query DNSKEY com.au2
DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
DNS: Standard query DS com.au2
DNS: Standard query response
DNS: Standard query DS com.au2
DNS: Standard query response DS RRSIG
DNS: Standard query DNSKEY au2
DNS: Standard query response DNSKEY DNSKEY RRSIG RRSIG
DNS: Standard query response A 1.1.1.1 RRSIG
```

- The situation is so much worse that the packet trace wont even fit on one slide!

DNSSEC 3rd level Query –
Common Name Server for TLD, 2LD & 3LD

- DNSSEC query for www.4th.3rd.com.au2
- 20 queries in total (8 on root, 10 on .au, 2 on registrant)
  - root referral
  - root DNSSKEY
  - 3rd.com.au2 referral
  - 3rd.com.au2 DNSSKEY
  - 4th.3rd.com.au2 answer
  - 4th.3rd.com.au2 DNSKEY
  - root referral for 3rd.com.au2 DS (cache?)
  - com.au2 for 3rd.com.au2 DS
  - com.au2 DNSKEY
  - root referral for com.au2 DS (cache?)
  - au2 for com.au2 DS
  - au2 for DNSKEY
  - root referral for 3rd.com.au2 name server A and AAAA records (2 queries, cache?)
  - 3rd.com.au2 answer for name server A and AAAA records (2 queries)
  - root referral for com.au2 name server A and AAAA records (2 queries, cache?)
  - com.au2 answer for name server A and AAAA records (2 queries)

DNSSEC 3rd level Query – Common Name Server for TLD, 2LD & 3LD

- DS records ‘piggy back’ with referral responses but no referral from au2 to com.au2 or from com.au2 to 3rd.com.au2 was ever received (wasn’t required as server had a more specific answer) thus it had to be specifically queried for, this involved walking the tree again!
- Chain of trust at zone cuts (.,au.,com.au., 3rd.com.au.,4th.3rd.com.au) required extra queries to be established as referral was ‘missed’
- If this same query was performed to non-common servers none of the ‘extra’ queries would be required

DNSSEC 3rd level Query – Common Name Server for TLD, 2LD & 3LD

- Worse case
  - 20 queries in total (8 on root, 10 on .au, 2 on registrant)
  - Increase of 17 over non signed (566.6% increase)
  - Increase of 9 on .au infrastructure (900% increase)
  - Additional 10 queries (100% above) the no common server scenario
  - Additional 4 queries (66.6% above) the no common server scenario on .au infrastructure
- Best Case (assuming cache is used)
  - 14 queries in total (4 on root, 8 on .au, 2 on registrant)
  - Increase of 11 over non signed (366.6% increase)
  - Increase of 7 on .au infrastructure (700% increase)
  - Additional 4 queries (40%) above the no common server scenario
  - Additional 2 queries (33.3% above) the no common server scenario on .au infrastructure
- So we are actually no better, and in some cases WORSE OFF using common servers in a DNSSEC scenario

Why is the cache not getting used during the validation phase of this lookup?

- Or ‘Why is there a best case and worst case result?’
- Two theories
  - BIND wont use the data because it hasn’t been able to validate it yet (because that is what it is actually in the process of doing)

  Unlikely because…
  - the results are in-consistent sometimes it appears as if the cache is actually being used thus the worse case / best case scenarios… most likely it is because…
  - BIND is trying to speed things up and doing lots of its queries in parallel thus it is a timing issue as to whether the answers from previous queries have been received yet, yet alone have been put in the cache
- Further investigation is required here

Why is it not all bad?

- The query counts referred to are for the first time a recursive sever attempts to resolve the name
- The DS, DNSKEY etc records will be cached for subsequent children lookups
- The real impact will depend on the TTL of these records and how many unique recursive servers attempt queries
- This makes the ‘real’ impact of these side effects hard to quantify in a testing environment
- Initially the amount of DNSSEC validating clients will be small

Conclusion

- Searching for missing DS is required to form the chain of trust so these queries are unavoidable, if the recursive sever didn’t get these in the referral it has to hunt them down, this can result in queries all the way back to the root again.
- When deploying DNSSEC, how you design your DNS infrastructure is more important than ever, and can have great impacts on the query volume you will need to process
- With DNSSEC you can be WORSE off having multiple levels of the hierarchy on the same server
- Recursive server implementations are still relatively young with respect to DNSSEC and need time to improve

- Note: tests were conducted using BIND 9.7.0a1

ausregistry
INTERNATIONAL

The TLD Specialists