

DNS Lame Delegations: A Case-Study of Public Reverse DNS Records in the African Region

Amreesh Phokeer¹(✉), Alain Aina², and David Johnson³

¹ AFRINIC, Ebene, Mauritius
amreesh@afrinic.net

² WACREN, Accra, Ghana
alain.aina@wacren.net

³ CSIR, Pretoria, South Africa
djohnson@csir.co.za

Abstract. The DNS, as one of the oldest components of the modern Internet, has been studied multiple times. It is a known fact that operational issues such as mis-configured name servers affect the responsiveness of the DNS service which could lead to delayed responses or failed queries. One of such misconfigurations is lame delegation and this article explains how it can be detected and also provides guidance to the African Internet community as to whether a policy lame reverse DNS should be enforced. It also gives an overview of the degree of lameness of the AFRINIC reverse domains where it was found that 45% of all reverse domains are lame.

Keywords: Reverse DNS · Misconfigurations · Lame delegation · Non-authoritative nameservers

1 Introduction

The Domain Name System (DNS) is a core functionality of the Internet which allows the translation of domain names into IP addresses i.e. from human-readable host names to machine-interpretable addresses. The DNS has become popular thanks to its distributed architecture which provides a very convenient way for users to publish and propagate their DNS information to the world.

On the Internet today, besides web browsing which involves lots of DNS queries, many applications such as content distribution through CDNs, email, spam filtering, Voice Over IP (VOIP) and telephone number mapping (ENUM), rely heavily on the availability of the DNS service [12]. However, when DNS was designed in the 1980's, engineers focused mainly on making the system scalable rather than secure, a requirement which only came much later.

As the DNS became an indispensable function of the Internet, questions pertaining to security and high availability became very relevant. The critical nature of the DNS makes it prone to multiple types of attack such as cache poisoning [8] and DDoS on DNS servers [10]. Besides the inherent security vulnerabilities,

the reliability of DNS services is also affected by different configuration errors as explained by Pappas et al. in their study on the impact of misconfiguration on the robustness of the DNS [6].

In this paper, we will look at one particular type of error called lame delegations on a subset of publicly DNS records, more specifically, the public DNS records of the AFRINIC¹ reverse tree.

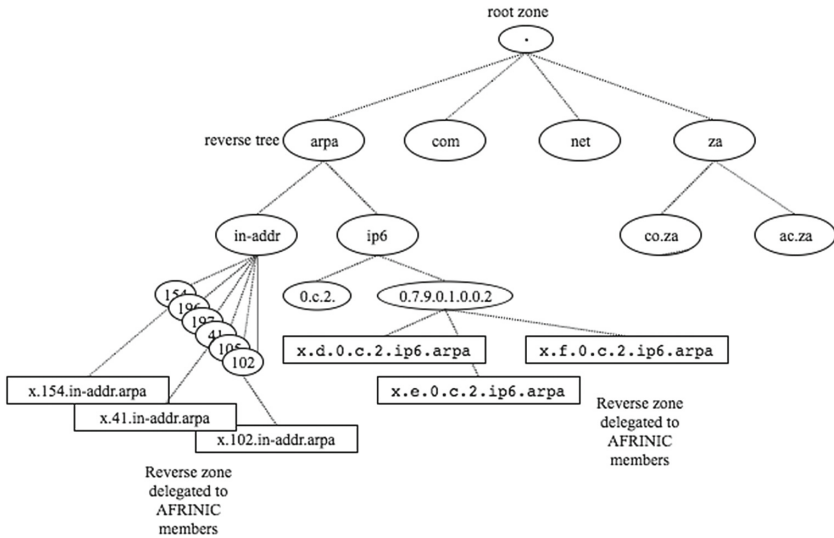


Fig. 1. DNS tree showing reverse delegations to AFRINIC members.

2 Background

AFRINIC manages reverse delegations for the IPv4 and IPv6 address space delegated by IANA. The resources currently managed by AFRINIC are listed on the IANA website^{2,3}. The aim of Reverse DNS entries is to allow applications on the Internet to map an IP address to its host, as opposed to forward DNS entries that map a domain to an IP. An example of a reverse DNS entry is a pointer record (PTR) that maps an IP address to a hostname. PTR records are very important for the many applications on the Internet. For example, some mail servers would enforce the check on reverse entries to make sure the originating IP of a incoming email transfer request is legitimate [7].

¹ The African Network Information Centre (AFRINIC) is the Regional Internet Registry (RIR) for Africa and the Indian Ocean. AFRINIC allocates Internet number resources i.e. IPv4, IPv6 and Autonomous System (AS) numbers to network operators in its constituency.

² <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>.

³ <http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml>.

6.2.216.196.in-addr.arpa domain name pointer www.afrinic.net.

Similarly as any other Top Level Domain (TLD), the DNS reverse tree is managed under the .ARPA zone as shown in Fig. 1. The subdomain for the IPv4 number space is the *in-addr.arpa* and *ip6.arpa* for IPv6. As registry of the IANA allocated space, AFRINIC needs to host an authoritative⁴ DNS server to serve the reverse zones of the space AFRINIC is currently managing. For example, AFRINIC allocates resources from its 41/8 address block and therefore authoritatively serves the *41.in-addr.arpa* zone.

When AFRINIC now allocates an address block to a member for e.g. a 41.10/16, it also delegates the management of the *10.41.in-addr.arpa* zone to the member. As the DNS works in a hierarchy, each child needs to link back to its parent by publishing their name servers in form of NS records. For instance, the NS records for the of the servers managing the *10.41.in-addr.arpa* zone must be published in the *41.in-addr.arpa* zone. Figure 2 shows how a child zone is linked to a parent zone.

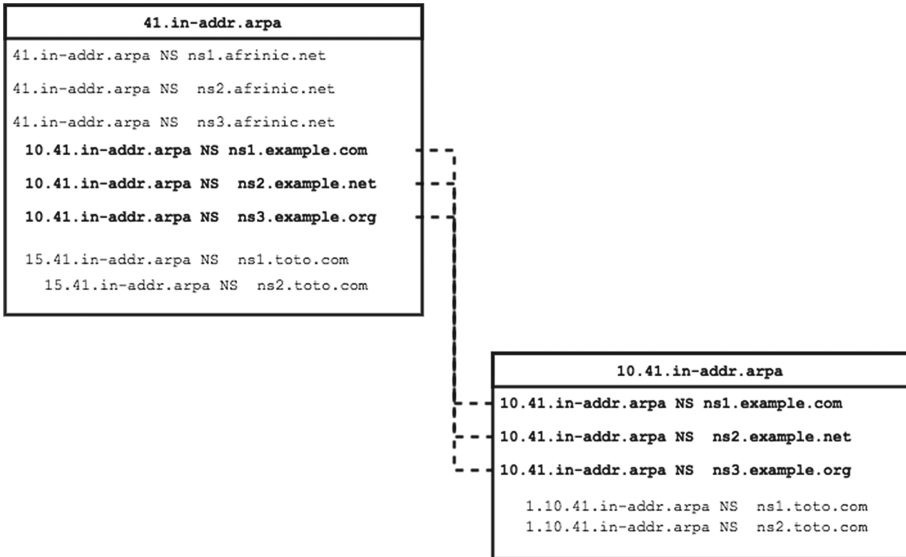


Fig. 2. NS records linking child and parent zones.

All the zones managed by AFRINIC are publicly available data and published on the AFRINIC public repository⁵. By analysing this data set, it gives us an idea of how well reverse delegations are configured in the African region.

⁴ An authoritative name server holds the actual records (A, AAAA, CNAME, PTR, etc.) of the zones, as opposed to a recursive server or resolver that needs to query an authoritative name server to resolve a domain/address.

⁵ <ftp://ftp.afrinic.net/pub/zones>.

3 Definitions and Related Work

In this section, we will provide a definition for lame delegations and give an insight on how other RIRs have dealt with this issue. We shall also provide some insight on the findings of two scientific studies on DNS availability.

3.1 What Is a Lame Delegation?

RFC1912 defines a delegation to be lame when a name server is delegated the responsibility for providing a name service for a zone (via NS records) but it is not actually doing it i.e. the name server is neither set up as a primary nor as a secondary server [2]. This is a classic example of a lame delegation, however there are some more granular cases as described in Sect. 4. A very common example of lame delegation is when a network administrator recently added a new resource record for e.g. *newdomain.example.org* with an NS record pointing to *new-name-server.example.org* in the parent zone, but no name service has yet been deployed on the host.

Basically, if the server does not respond to DNS queries, it is considered lame. Lame delegation is considered as a bad practice as it increases the load on the parent name servers and consequently increases the delay in DNS responses. Many commercial DNS servers now have in-built mechanism to check for lame delegations such as BIND [4]. CISCO Prime Network Registrar, which includes a DNS server, can detect lame delegation by reporting non-matching or missing NS records in the parent zone [3].

3.2 Lame Delegation Policies at Other RIRs

All RIRs run authoritative name servers to serve the reverse zones of the IANA delegated space they manage. LACNIC⁶ and ARIN⁷ have implemented a “Lame delegation policy” which enforces the DNS best practices against lame entries. The APNIC⁸ and the RIPE⁹ community made proposals but have not adopted a lame delegation policy. However, they have implemented checks on their reverse DNS system precluding lame entries [11]. AFRINIC has no lame delegation policy on reverse delegation.

LACNIC periodically revises their *in-addr.arpa* and *ip6.arpa* zones and checks for lame delegation. Their methodology is to check whether a query of a SOA record on a selected server is returned as an authoritative response by the server. If not, the reverse DNS entry is considered as lame and the zone operator is contacted. LACNIC has implemented a “Lame delegation policy” which has helped to curb the number of lame delegations and now has a DNS success rate of 96.80% [9]. Figure 3 shows how the percentage of lame delegation which consistently dropped after implementation of the lame delegation policy in around 2003.

⁶ <http://www.lacnic.net/en/web/lacnic/manual-6>.

⁷ https://www.arin.net/policy/proposals/2014_5.html.

⁸ <https://www.apnic.net/policy/proposals/prop-004>.

⁹ <https://www.ripe.net/ripe/mail/archives/dns-wg/2005-May/001493.html>.

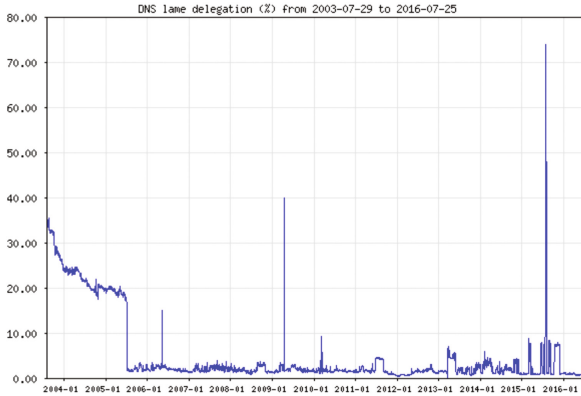


Fig. 3. Percentage of lame delegations in LACNIC database since 2004 [9].

APNIC proposed a slightly different taxonomy of lame delegations in four categories, where a delegation is considered lame if any of the following is true (1) a listed DNS server is unreachable, (2) a listed DNS server is reachable but not responsive on port 53, (3) a listed DNS server is reachable and responds on port 53, but it is not able to answer for the domain, (4) a listed DNS server is reachable and responds on port 53 but serves incorrect data for the domain [1].

3.3 Previous Studies

Pappas et al. conducted passive and active measurements on a university network [6]. They first analyzed DNS traffic exchanges from the university network to external websites and also implemented a specialized resolver to perform DNS queries to a randomly selected list of destinations. They found out that DNS configuration errors are widespread, with more than 15% of delegation being lame, 22% of zones with inconsistency and 2% affected by cyclic dependency [5]. They classified lame delegations in three different categories, depending on the type of error found:

- Type 1: Non responding server
- Type 2: DNS error indication
- Type 3: Non-authoritative answer

Redundancy is another important aspect of availability. A zone can authoritatively be served by multiple redundant name servers. DNS best practices stipulate that it is preferable to have name servers, serving the same zone, spread geographically (both in terms of location and network) [2]. Although Deccio et al., were not specifically targeting lame delegations, they discovered that 14% of DNS entries experience “false redundancy”, meaning that either there is no redundant server (different NS records pointing to the same name server) or the supposedly redundant servers reside on the same network.

4 Methodology

In this section, we will explain how the DNS data was collected and how lame delegations were detected and classified.

4.1 Data Collection

The AFRINIC database contains around 30000 domain objects. Each domain object is associated with at least two name servers. For the purpose of this experiment, we took the whole set of reverse domains and run the experiment against each domain and name server (NS) tuple. A domain can have multiple NS records and each record is considered as an entry in DNS for which we have verified its validity. All the reverse zones were obtained from <ftp://ftp.afrinic.net/pub/zones>. Table 1 shows the breakdown between IPv4 and IPv6 reverse zones and gives the total number of NS records.

Table 1. Total registered domains and corresponding number of NS entries

Type	Domains	NS records
IPv4	29894	72341
IPv6	196	550
Total	29986	72891

As DNS query tool, we used “dig” (Domain information groper) which is commonly found on all Unix machines. It basically performs DNS lookups and returns the answers from the server that has been queried. In our case we used the option *+norec* which instructs the server queried not to retrieve DNS data from “recursive servers” but instead to get the answer from the name server that have been specified or from an authoritative source.

We paid attention to three main elements in the query response: **STATUS**, **FLAGS** and **ANSWER**. The different type of statuses are described in RFC6895¹⁰. In our case, a query is considered “successful” if the STATUS is *NOERROR*, the FLAGS section contains *AA*¹¹ and the ANSWER section is not null (a query is actually also technically “successful” for example if it returns an **NXDOMAIN** with **ANSWER: 0**, however for us here, it would be considered as a “failure”, the domain not being found). Table 2 gives a breakdown of some of the main statuses. Below is example of a dig query and response asking for NS records of the afrinic.net domain from ns1.afrinic.net without recursion:

¹⁰ <https://tools.ietf.org/html/rfc6895>.

¹¹ AA means Authoritative Answer.

```

$ dig NS @ns1.afrinic.net afrinic.net +nored

; <<>> DiG 9.8.3-P1 <<>> NS afrinic.net @ns1.afrinic.net +nored
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 12713
;; flags: qr aa ra; QUERY: 1, ANSWER: 7, AUTHORITY: 0, ADDITIONAL: 12

[...]

;; Query time: 155 msec
;; SERVER: 196.216.2.1#53(196.216.2.1)
;; WHEN: Mon May 2 21:07:13 2016
;; MSG SIZE rcvd: 447

```

Table 2. Meaning of STATUS response (extract)

Status	Description
NOERROR	Domain exists
NXDOMAIN	Domain does not exists
REFUSED	Server refuses to perform query
SERVFAIL	Something went wrong

4.2 Error Classification

We run the experiment on each and every delegation from two different locations (Mauritius and Johannesburg). It is important to use two geographically spread locations to cater for failed queries that may occur due to congestion on the network or due to firewall restrictions. A delegation is considered lame if it fails from both sites. If a query positively responds from one site or the other, we would consider the NS to be successful and if fails on both sites, NS is tagged as erroneous. To simplify the representation of the results, we decided to classify them in four categories as shown in Table 3. The first category is CASE_0 is the null case which means that the NS record is OK.

We developed a simple algorithm to classify the dig results of each delegation found on the public reverse zones of AFRINIC, as per the criteria in Table 3. The algorithm below makes provision for more granular types of lame delegation but in the results section, only the four main categories have been considered.

```

Query: dig domain @nameserver for NS record (without recursion)
begin
  if output = any_of_CASE_1_errors
    then z = CASE_1;
  fi
  if status = (REFUSED||SERVFAIL||NXDOMAIN)
    then z = CASE_2;
  fi
  if status = NOERROR
    if answer = 0
      then z = CASE_2(NO_ANSWER);
    fi
    if flag1 = AA
      if flag2 = RA
        then z = CASE_2(RECURSIVE);
      else
        then z = CASE_3(AUTHORITATIVE)
      fi
    fi
  fi
  if status = NOERROR
    if RAflagispresent
      then z = CASE_0(RECURSIVE);
    else
      then z = CASE_0(AUTHORITATIVE);
    fi
  fi
print(z)

```

5 Results and Observations

5.1 Valid Versus Lame

We found that approximately 55% of IPv4 domain registered in the AFRINIC database do not have any issue and can be considered as valid. For the other 45% considered as lame, it means that at least one of the NS records for the domain is actually lame. For IPv6 domains, 32% is found to be lame. Table 4 gives the number of IPv4 and IPv6 domains that passed the test i.e. tagged as CASE_0.

Table 3. Classification of delegation into different categories

Category	Error response
CASE_0	All of: <ul style="list-style-type: none"> – NS is responsive – NS serves the domain – NS is authoritative – Response status is NOERROR <i>Flag: AUTHORITY(AA)</i>
CASE_1	Either of: <ul style="list-style-type: none"> – Connection timed out – Name or service not known – Connection refused – Network unreachable – Host unreachable – End of file – Communications error – Couldn't get address
CASE_2	Response status is either: <ul style="list-style-type: none"> – REFUSED – SERVFAIL – NOERROR without response from server i.e. ANSWER: 0 <i>Flag: AUTHORITY(AA) or RECURSIVE(RA)</i>
CASE_3	All of: <ul style="list-style-type: none"> – NS is responsive – NS serves the domain – NS is authoritative – Response status is NOERROR <i>Flag: RECURSIVE(RA)</i>

Table 4. Percentage of lame versus non-lame domains

Type	VALID	%	LAME	%	Total
IPv4	39439	54.5	32970	45.5	72409
IPv6	369	68	174	32	543

5.2 Breakdown by Error Type

We classified the 45% of lame delegations found into the three error categories which are CASE_1, CASE_2 and CASE_3. From the results in Table 5, we observed that 75.5% are actually CASE_2 (responsive servers but not serving the zone). Most probably, the name servers that were recorded have been decommissioned by the operators. 23.5% of errors are CASE_1, meaning that the servers are not even reachable, and finally, only 1% of faulty domains have been tagged as CASE_3, meaning that more 99% of all servers queried are authoritative.

Table 5. Percentage of error type vs. address type

Type	CASE#1	CASE#2	CASE#3	Total
IPv4	7803	24941	314	32970
IPv6	19	155	0	174
Total	7822	25096	314	33144
%	23.5	75.5	1	

6 Conclusion

We found that a big chunk (45%) of reverse domains registered at AFRINIC is lame and the predominant cause of lame delegation (more than 75%) is the CASE.2 which means that servers are proper DNS servers and are responsive but they are not serving the zone as indicated by the DNS operator. One reason which could explain this situation is that in our region where resources are constrained, operators do not have the facility of hosting additional servers to ensure redundancy of their name servers. They would therefore register an ad-hoc or sometimes non-operational name server as secondary entry for their zones. Another reason could be a lack of continuity of service for a domain where the server has been recycled to serve another domain. This contributes to pollute the African reverse DNS ecosystem and must definitely have a negative impact on query time, affecting latency of services in general. It is therefore important for AFRINIC to fix those issues and provide a clean and reliable DNS service to the African operators and users on the Internet. It has become clear that to curb the number of lame delegation, AFRINIC needs to come up with a policy or implement stringent operational checks to (1) clear all existing lame delegations and (2) prevent any new lame delegation to be inserted in AFRINIC's database.

7 Future Work

Lame delegation is only a subset of DNS misconfiguration. To ensure full availability, name servers should be truly redundant. By truly redundant, we mean that primary and secondary name servers should be geographically spread and not found on the same host and as far as possible, not on the same network (i.e. on different ASes). In the event of a routing outage and one network is unavailable, the other network would still be reachable. This ensures full redundancy. Furthermore, it would be interesting to see where African network operators are hosting their DNS servers. Mapping the servers by location would give us an indication whether African operators are using local or offshore services, usually reachable on expensive international links. Cyclic zone dependency [6] is another issue that is less known but yet important to tackle as they create dependency loops between DNS servers. The impact is the addition of unnecessary load on those servers ultimately affecting availability on the overall.

References

1. APNIC's operational response to lame delegations. <https://www.apnic.net/manage-ip/manage-resources/reverse-dns/lame-dns-reverse-delegation/apnics-operational-response>
2. Barr, D.: Common DNS Operational and Configuration Errors (1996)
3. CISCO. http://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/network_registrar/8-2/user/guide/CPNR_8.2_User_Guide/UG16_Zon.html
4. Albitz, P., et al.: DNS and BIND, 5th edn. O'Reilly, Sebastopol (2006)
5. Pappas, V., et al.: Distributed DNS troubleshooting. In: Proceedings of the ACM SIGCOMM Workshop on Network Troubleshooting Research, Theory and Operations Practice Meet Malfunctioning Reality - NetT 2004, p. 265 (2004)
6. Pappas, V., et al.: Impact of configuration errors on DNS robustness. *IEEE J. Sel. Areas Commun.* **27**(3), 275–290 (2009)
7. Phokeer, A., et al.: A survey of anti-spam mechanisms from an RIR's perspective. In: Proceedings - IST-Africa (2016)
8. Trostle, J.B., et al.: Protecting against DNS cache poisoning attacks. In: 2010 6th IEEE Workshop on Secure Network Protocols, NPSec 2010, pp. 25–30 (2010)
9. LACNIC. http://lacnic.net/en/lame_delegation.html
10. Wei, M.L., Lu, Y.C., Zhen, M.L.: Alleviating the impact of DNS DDoS attacks. In: NSWCTC 2010 - The 2nd International Conference on Networks Security, Wireless Communications and Trusted Computing, vol. 1, pp. 240–243 (2010)
11. RIPE-NCC. <https://www.ripe.net/manage-ips-and-asns/resource-management/assisted-registry-check>
12. van Adrichem, N.L.M., et al.: A measurement study of DNSSEC misconfigurations. *Secur. Inf.* **4**(1), 8 (2015)