# IPv6 How-To for a Registry
## 17th CENTR Technical Workshop

Amsterdam, October 2007

Alvaro Vives (alvaro.vives@consulintel.es)
Jordi Palet (jordi.palet@consulintel.es)


# Introduction

- Main steps to be undertaken and issues to be taken into account by an organization in charge of a Top Level Domain (TLD) which wants to offer IPv6 [1] support to its DNS [2][3] service.
- In addition the TLD receives the service of root servers, so TLDs also have to ask for this and use it.

# Status (I)

- Since July 2004 with the ICANN's announcement [4] of IPv6 addresses support within the root servers, several TLDs have added them [5].
- Root Servers actually (root zone 2004012900) have no official IPv6 address [6].
- There is an unofficial list [7] where some Root server have IPv6 addresses: B, F, H, K and M.
- It is possible to manually configure the DNS resolvers with those IPv6 addresses.

# Status (II)

- ICANN is working on the IPv6 on Root Servers issue by means of the RSSAC [27] and SSAC [26]. Output:
  - Draft report about glue AAAA on root servers [25]: root zones and root hints.
  - Study of real traffic on some Root Servers being done.
  - Transport status: B, F, H, I, J, L and M IPv6 capable but not yet with production quality connectivity.

# Data vs. Transport

- We consider the introduction of IPv6 from two points of view:
  - The data stored in the DNS servers (Resource Records, RR).
  - The transport of that data.
- As stated in RFC 3596 [8]:
  - The IP protocol version used for querying resource records is independent of the protocol version of the resource records; e.g., IPv4 transport can be used to query IPv6 records and vice versa.

# Data

- New definitions [8]:
  - **AAAA** RR: Resource record used to map a domain name to an IPv6 address.
  - **IP6.ARPA** domain name to support lookups of domain names based on address [9].
- IPv6 notation, address length, etc. must be taken into account within all the software used to configure and update the RRs stored into the DNS servers and the Data Bases used to store user's information.

# Response Size (I)

- To be taken into account in the configuration of authoritative NSs for a TLD zone and the A and AAAA glue records for them.
- The introduction of IPv6 and its coexistence with IPv4 means that the amount of data stored and sent into a response will increase [10]. Also the size of IPv6 addresses is four times the IPv4 one.

# Response Size (II)

- ***Glue* records** must be configured within the ***root* servers**. In case of several servers authoritative for a TLD, the number of AAAA glue records is limited and should be calculated [11][12][25]. These are the **delegation responses** that contains the NSs authoritative for a domain, a TLD in our case, and additional data with the IP address(es) of those NSs.
- Positive responses could also have problems in case of several A and AAAA records.

# Response Size (III)

| HEADER SECTION | QUESTION SECTION | ANSWER SECTION | AUTHORITY SECTION | ADDITIONAL SECTION |
|---|---|---|---|---|

- **DELEGATION RESPONSE PARTS:**
- Header Section: fixed length (12 octets)
- Question Section: original query (name, class, type)
- Answer Section: (empty)
- Authority Section: NS RR set (name server names)
- Additional Section: A and AAAA RR sets (name server addresses)

# Response Size (IV)

- If the packet is bigger than 512 bytes it is truncated eliminating one or more RRsets and the appropriated flag is set to inform the client, who could establish a TCP connection to query again. This last behavior is much less effective and not desired.
- The reason of the limit in the number of glue records (A and AAAA) that can be sent within a DNS reply comes from the DNS standard ([3] section 4.2.1), that limits the message size to 512 bytes. In fact because part of the request is included in the response, the limit is even lower. The idea is that the more AAAA RR within additional section the bigger will be the answer.
- DNS label compression ([3] section 4.1.4) allows to optimize the use of the 512 bytes. So, giving all name server names a common parent (such as GTLD-SERVERS.NET or ROOT-SERVERS.NET) saves space.

# Response Size (V)

- Although EDNS0 [13] permits larger responses by mutual agreement of the requestor and responder, in practical this limit stays because in the short or medium term it is not foreseen to have it in all the resolvers.
- As a conclusion from [10], adding two to five IPv6 name server address records (AAAA RRs) to a prototypical delegation would not have a significant negative operational impact on the domain name system.
- Both the number of AAAA RRs used and the DNS label compression are the mechanisms that can be used to optimize the response size.

# Transport

- Differentiate between IPv4 and IPv6 transport used for DNS requests and replies.
- Four things to be provided:
  - IPv6 capable servers-> not a problem: BIND, nsd, newbie, maradns and djbdns [14][15].
  - IPv6-capable network-> not a problem nowadays: several production products.
  - IPv6 addresses for server(s) and nework(s) -> not a problem with IPv6 :-)
  - IPv6 Connectivity: Native or tunneled.

# Recommended Guidelines (I)

- IPv4 and IPv6 will coexist, 3 types of name servers:
  - IPv4 only -> reachable only by IPv4
  - IPv6 only -> reachable only by IPv6
  - Dual-stack -> -> reachable by both
- Name space fragmentation should be avoided, this happens when the recursive resolution process is broken (e.g. when only IPv6 NS are authoritative for a domain, resulting that IPv4 only DNS server won't be able to follow the resolution chain)
- IDEA: **backward compatibility**

# Recommended Guidelines (II)

- Administrative policies [16]:
  - Every recursive name server should be either IPv4-only or dual stack.
  - Every DNS zone should be served by at least one IPv4-reachable authoritative name server.

# ICANN/SSAC/RSSAC Report

- Considers the issues related to the inclusion of the IPv6 addresses for the root level of the DNS.
- IPv6 addresses are already included for Top Level Domain Name Servers in the root zone file.
- Operators of a number of root name servers have assigned IPv6 addresses to their servers.
- Root Servers not yet reachable by IPv6.
- Recommends that type AAAA records for all root name servers so addressed should be included in the root hints and root zone files and that they be returned in priming responses from root name servers as soon as practically possible.

# Other Considerations

- One of the fist steps is the education. In the DNS case, things won't change too much.
- IPv6 connectivity.
- Front-office and back-office applications.
- Introduction of security techniques defined for DNS [17]-[22] are not affected by the IPv6 support.

# References (I)


- **[1]** Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- **[2]** Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- **[3]** Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- **[4]** Next-generation IPv6 Address Added to the Internet's Root DNS Zone: http://www.icann.org/announcements/announcement-20jul04.htm
- **[5]** IANA Administrative Procedure for Root Zone Name Server Delegation and Glue Data: http://www.iana.org/procedures/delegation-data.html
- **[6]** Root Zone Hints File in IANA Popular Links: http://www.iana.org/popular.htm
- **[7]** http://www.root-servers.org/
- **[8]** Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.

# References (II)


- **[9]** Bush, R., "Delegation of IP6.ARPA", BCP 49, RFC 3152, August 2001.
- **[10]** Paul Vixie and Akira Kato, “DNS Referral Response Size Issues”, Internet-Draft draft-ietf-dnsop-respsize-07, February 2007.
- **[11]** Mohsen Souissi, “DNS Response Size and Name Compression”, http://w6.nic.fr/dnsv6/resp-size.html
- **[12]** Ronald van der Pol and Daniel Karrenberg, “Adding IPv6 glue to the root zone”, October 2003: http://www.nlnetlabs.nl/ipv6/publications/v6rootglue.pdf
- **[13]** P. Vixie, “Extension Mechanisms for DNS (EDNS0)”, RFC 2671, August 1999.
- **[14]** Internet Systems Consortium http://www.isc.org

# References (III)


- **[15]** DeepSpace6 - Current Status of IPv6 Support for Networking Applications http://www.deepspace6.net/docs/ipv6_status_page_apps.html
- **[16]** A. Durand and J. Ihren, “DNS IPv6 Transport Operational Guidelines”, RFC 3901, September 2004.
- **[17]** R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, “DNS Security Introduction and Requirements”, RFC 4033, March 2005.
- **[18]** R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, “Resource Records for the DNS Security Extensions”, RFC 4034, March 2005.
- **[19]** R. Arends, R. Austein, M. Larson, D. Massey, S. Rose, “Protocol Modifications for the DNS Security Extensions”, RFC 4035, March 2005.
- **[20]** Vixie, P., Gudmundsson, O., Eastlake, D. and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.

# References (IV)


- **[21]** S. Kwan, P. Garg, J. Gilroy, L. Esibov, J. Westhead, R. Hall, “Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)”, RFC 3645, October 2003.
- **[22]** S. Weiler, J. Ihren, “Minimally Covering NSEC Records and DNSSEC On-line Signing”, RFC 4470, April 2006.
- **[23]** R. Bush, A. Durand, B. Fink, O. Gudmundsson, T. Hain, “Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)”, RFC 3363, August 2002.
- **[24]** R. Austein, “Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6)” ,RFC 3364, August 2002.
- **[25]** “Accommodating IP Version 6 Address Resource Records for the Root of the Domain Name System”, ICANN, March 2007. http://www.icann.org/committees/security/sac018.pdf
- **[26]** ICANN Security and Stability Advisory Committee (SSAC). http://www.icann.org/committees/security/
- **[27]** ICANN DNS Root Server System Advisory Committee (RSSAC). http://www.icann.org/committees/dns-root/