IPv6 and DNS

Name resolution in a network with IPv6

Carsten Strotmann

CREATED: 2025-01-30 THU 09:40

Agenda

- IPv6 support in DNS Server
- The IPv6 AAAA Address Record
- Reverse Lookup
- DNS and IPv6 pitfalls
- Resolver configuration
- IPv6 in URLs

IPv6 and DNS

IPv6 Transport and IPv6 Payload

- There are two things to adhere with DNS and IPv6
  - IPv6 can be the transport protocol to send IPv4 and IPv6 DNS payload
  - IPv6 Resource Records served by a name server

- Both can be used independently
  - Answer queries for IPv6 Address(AAAA) Requests over IPv4
  - Answer queries to IPv4 Address (A) Requests over IPv6

IPv6 Support in DNS server software

DNS Resolver

- BIND 9 from ISC
- Windows DNS (since 2008) from Microsoft
- Unbound from NLnetLabs
- PowerDNS Recursor from PowerDNS B.V.
- Knot-Resolver from cz.nic

Authoritative DNS Server

- BIND 9 from ISC
- Windows DNS (since 2008) from Microsoft
- NSD from NLnetLabs
- PowerDNS from PowerDNS B.V.
- Knot from cz.nic

The IPv6 Address Record

AAAA - the IPv6 Address Record

- IPv6 Addresses are stored using AAAA (Quad-A) Records
- The AAAA Record is similar to the IPv4 Address Record, just use an IPv6 Address notation (↪RFC 3596 *DNS Extensions to Support IP Version 6*)

AAAA - the IPv6 Address Record

```
$TTL 3600
example.org.     IN SOA  dns1.example.com. hostmaster.example.org. (
                                 20210816   ; serial
                                     1d   ; refresh
                                     2h   ; retry
                                    40d   ; expire
                                     2h ) ; neg TTL
example.org.       IN NS   dns1.example.com.
example.org.       IN NS   dns2.example.com.

example.org.       IN MX   10 mail.example.org.
www.example.org.   IN A    192.0.2.44
www.example.org.   IN AAAA 2001:db8:100::66

mail.example.org. IN AAAA 2001:db8:100:0:AF:999:A800:1
mail.example.org. IN A 192.0.2.42
```

What about the A6 Record?

- The A6 Resource Record (RFC 2874, aimed to ease renumbering of IPv6 Networks) is not used anymore
  - However if you sniff the network, you might see A6 queries from older operating systems
  - See ↪RFC 6563 (March 2012): “Moving A6 to Historic Status”

The HTTPS Record

- The `HTTPS` record is a relatively new record. It's type number is 65 (aka TYPE65)
  - It has been first observed "in the wild" in Summer 2020
  - It is being used in the new Apple operating systems (iOS, iPadOS, macOS) since fall 2020
  - It is now the 3rd most queried DNS record in the Internet (after `A` and `AAAA`)
  - ↪RFC 9460 - Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records) (Nov 2023)

The HTTPS Record

- The `HTTPS` delivers connection information for an HTTPS service
  - IPv4-Addresses of the service
  - **IPv6-Addresses** of the server
  - The public key of the server to be used to initiate an encrypted TLS "client hello"
  - Protocol selection: HTTP/2 (TCP) or HTTP/3 (QUIC)
  - One or more DoH-Resolver for the service

The HTTPS Record

- Example of a `HTTPS` Record for a service offering HTTP/2 and HTTP/3 (preferred)

```
example.com 3600 IN HTTPS 1 . alpn=”h3,h2”
```

The HTTPS Record

- Example of a HTTPS Record for a service with both IPv4 and IPv6 Addresses

```
example.com 3600 IN HTTPS 1 . alpn=”h3,h2” (
            ipv4hint=”192.0.2.1”
            ipv6hint=”2001:db8::1”
            )
```

The HTTPS Record

- Benefits of the HTTPS records
  - Browser don't need to request explicit IPv6 and IPv4 Addresses (faster connection)
  - The `HTTPS` Records is a signal to the web-browser to only allow TLS secured/encrypted connections for this domain name. This prevents downgrade attacks (DNSSEC recommended!).
  - With the `HTTPS` Record it is possible to create Domain-Alias definitions for whole zones (not possible with the `CNAME` Record)

The HTTPS Record

- The BIND 9 DNS server and tools (`dig`, `host`) support the `HTTPS` record since version 9.16.21 (September 2021)

Happy Eyeballs v3 the HTTPS DNS-Record

- The HTTPS-Record (↪RFC 9460) in DNS can give applications guidance on which IP protocol to choose
- The IETF is working on a 3rd iteration on Happy Eyeballs to take HTTP/3 and the HTTPS record into account: ↪"Happy Eyeballs Version 3: Better Connectivity Using Concurrency"

Reverse Lookup

IPv6 PTR Records

- Domain-Names in Reverse Zones for IPv6 Addresses are stored using normal PTR (Pointer) Records
- In the Owner-name, the IPv6 IP Address is written in reverse order and in hexadecimal Nibble Notation (Nibble = 4bit)
- The second-level domain for IPv6 Reverse zones is `ip6.arpa.`
- Pre-population of IPv6 reverse zones is not feasible
    - there are too many IPv6 addresses in one /64 subnet

- Reverse DNS name resolution is optional for most applications and protocols

Example reverse DNS zone

```
$TTL 3600
$ORIGIN 0.0.0.0.0.0.1.0.8.b.d.0.1.0.0.2.ip6.arpa.
@            IN SOA  dns1.example.org. hostmaster.example.org. (
                                 20210116   ; serial
                                     1d   ; refresh
                                     2h   ; retry
                                    40d   ; expire
                                     2h ) ; neg TTL
             IN NS  dns1.example.org.
             IN NS  dns2.example.org.

6.6.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 IN PTR     www.example.com.
1.0.0.0.0.0.8.a.9.9.9.0.f.a.0.0 IN PTR     mail.example.com.
```

Strategies for IPv6 reverse zones

- Several real-world solutions for populating an IPv6 reverse zone exist:
  - Manually manage PTR records (for important machines)
  - Update PTR records from DHCPv6
  - Update PTR records from an IP Address Management System
  - Use DNS Wildcards to provide generic PTR records

Example DNS wildcard use

```
$ORIGIN .
$TTL 7200       ; 2 hours
0.0.0.0.3.a.5.e.7.e.2.4.3.d.f.ip6.arpa IN SOA ns1.example.com. hostmaster.example.com. (
                               2014110509 ; serial
                               28800      ; refresh (8 hours)
                               7200       ; retry (2 hours)
                               604800     ; expire (1 week)
                               7200       ; minimum (2 hours)
                               )
$TTL 86400      ; 1 day
                        NS     ns1.example.com.
$ORIGIN 0.0.0.0.0.3.a.5.e.7.e.2.4.3.d.f.ip6.arpa.

$TTL 7200       ; 2 hours

;; manual managed PTR records
25.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 PTR mail.example.com.
53.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 PTR ns1.example.com.
;; Wildcard PTR records
*                              PTR some-node.example.com.
```

wildcard DNS record matches all PTR requests that match the prefix and have no dedicated PTR entry

DNAME records

- DNAME records (↪RFC 6672) redirects a non-delegated subtree on an existing namespace
  - Like CNAME, but for all names below a domain-name
  - Can be used to map one IPv6 reverse zone over another (given that the host-ids are identical for both IPv6 network prefixes)

- Example: the ULA Prefix `fd55:10:20:30::/64` and the GUA Prefix `2001:db8:20:30::/64` contain the same machines
  - With DNAME, only one reverse zone need to be maintained

Example DNAME records

```
$TTL 3600
$ORIGIN ipv6hosts.example.com.
@                           IN SOA ns1.example.com. hostmaster.example.com. (
                                                      20141105 1d 2h 41d 1h )
                            IN NS  ns1.example.com.

f.2.c.a.1.e.e.f.f.f.f.2.b.0.2 IN PTR host1.example.com.
d.c.1.2.a.c.e.f.f.f.2.f.a.1.0 IN PTR ns1.example.com.
```

Zone: ipv6hosts.example.com. containing PTR records

```
$TTL 3600
$ORIGIN 0.3.0.0.0.2.0.0.0.1.0.0.5.5.d.f.ip6.arpa.
@                           IN SOA ns1.example.com. hostmaster.example.com. (
                                                      20141105 1d 2h 41d 1h )
                            IN NS  ns1.example.com.
0                           IN DNAME ipv6hosts.example.com.
```

Zone: 0.3.0.0.0.2.0.0.0.1.0.0.5.5.d.f.ip6.arpa

```
$TTL 3600
$ORIGIN 0.3.0.0.0.2.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
@                           IN SOA ns1.example.com. hostmaster.example.com. (
                                                      20141105 1d 2h 41d 1h )
                            IN NS  ns1.example.com.
0                           IN DNAME ipv6hosts.example.com.
```

Zone: 0.3.0.0.0.2.0.0.8.b.d.0.1.0.0.2.ip6.arpa.

DNS Data and IPv6

IPv6 Addresses to be published in DNS

- on public authoritative DNS server:
  - Global Unicast Addresses
- on private internal authoritative DNS server (not visible to the public Internet):
  - Global Unicast Addresses
  - Multicast addresses
  - Unique-local addresses

Service Names in DNS

- It is recommended that you publish service names in DNS in addition to host-names
  - SRV records are not widely supported in applications
  - Separates the service function from the physical machine
  - Use `HTTPS` records for web-services

- Example Hostname: `m68112.datacenter.la.example.com`
- Example Service-Name: `www.example.com`

IPv6 Addresses **not** to be published in DNS

- Link-local addresses
- Temporary addresses (privacy extensions)
- Teredo/Tunnel Addresses

TTL for IPv6 and IPv4 service names in DNS

- If IPv4 and IPv6 Address records are present for the same service, the TTL of the records should be identical
  - Especially important for *critical* additional data, such as addresses of DNS Servers (NS-Records)
  - Glue-Records (AAAA/A-Records for delegation NS-Records) will not be queried **on-demand** by DNS resolver. Having different TTL between A and AAAA records can lead to nameserver only reachable over one IP protocol.

Example: Bad use of TTL

```
example.org. 2h  IN SOA  dns1.example.org. hostmaster.example.org. (
                                20110116   ; serial
                                    1d   ; refresh
                                    2h   ; retry
                                   40d   ; expire
                                    2h ) ; neg TTL
example.org. 1d  IN NS  dns1.example.org.
example.org. 1d  IN NS  dns2.example.org.

www.example.org. 2h IN A  192.0.2.44
www.example.org. 1h IN AAAA 2001:db8:100::66
```

```
dns1.example.org. 2h IN AAAA 2001:db8:100:0:AF:999:A800:1
dns1.example.org. 1h IN A 192.0.2.42
```

IPv6 for authoritative DNS Server

- Most DNS server software can run on a dual stack machines (IPv4 and IPv6)
- Name server should be able to communicate over IPv4 as long as many Addresses and Internet Name server are still on IPv4
- Most ccTLD and gTLD DNS server are available on IPv6 networks

IPv6 and DNS round robin

- DNS round robin is used in pure IPv4 networks to create a "poor mans load distribution"
  - It makes use of the fact that many applications only choose the topmost address record from a record set

```
; <<>> DiG 9.9.2-vjs287.12 <<>> yahoo.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 31107
;; flags: qr rd ra; QUERY: 1, ANSWER: 3, AUTHORITY: 7, ADDITIONAL: 8

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;yahoo.com.         IN  A

;; ANSWER SECTION:
yahoo.com.      1800    IN  A   98.138.253.109
yahoo.com.      1800    IN  A   206.190.36.45
yahoo.com.      1800    IN  A   98.139.183.24
```

resource record set

IPv6 and DNS round robin

- IPv6 nodes sort all destinations received from DNS to select a destination to use
  - This applies to IPv4 and IPv6 addresses and it disables DNS round robin selection for IPv4.
  - ↪DNS Round Robin and Destination IP address selection (Archived version from MS Technet)
  - ↪RFC 6724 "Default Address Selection for Internet Protocol Version 6 (IPv6)"

DNS Client support

Provisioning the DNS resolver addresses

- With IPv6, the DNS resolver addresses for a client can be configured ...
  - Manual
  - Using DHCP (v4/v6)
  - Through Router Advertisements (RDNSS, RFC 6106)
- see ↪RFC 4339 - "IPv6 Host Configuration of DNS Server Information Approaches"

Dynamic Updates

- Windows OS IPv6 clients can get their IPv6 Address by Auto-configuration (SLAAC) and register themselves in DNS
  - There is no equivalent function in an Unix stub-resolver today to enter this Information into DNS (forward and reverse zones)
  - RFC 9686 style IP address registration towards a DHCPv6 server opens the option to use the DHCPv6 server as a dynamic DNS proxy
    - The SLAAC client will register it's addresses towards the DHCPv6 server, the DHCPv6 server will update the name to address binding in DNS

Applications & IPv6

Applications & IPv6

- Not all Applications can use IPv6 Addresses or can communicate over IPv6
- Services on hosts with AAAA Addresses should be able to use IPv6
- Recommendation: During initial test and IPv6 deployment, configure a separate subdomain for IPv6 enabled hosts and services.
- Example:
  - Web-Service on IPv4: `www.example.com`
  - Web-Service on IPv6: `www.ip6.example.com`

Literal IPv6 addresses in URLs

- Literal IPv6 addresses in URLs must be enclosed in brackets. To reach a web-server on the IPv6 address `2001:db8::2:3:4` use the URL `https://[2001:db8::2:3:4]:443/index.html`
  - This is not intended for “End-Users”
  - But is helpful for Administrators and for Troubleshooting

Literal IPv6 addresses in URLs

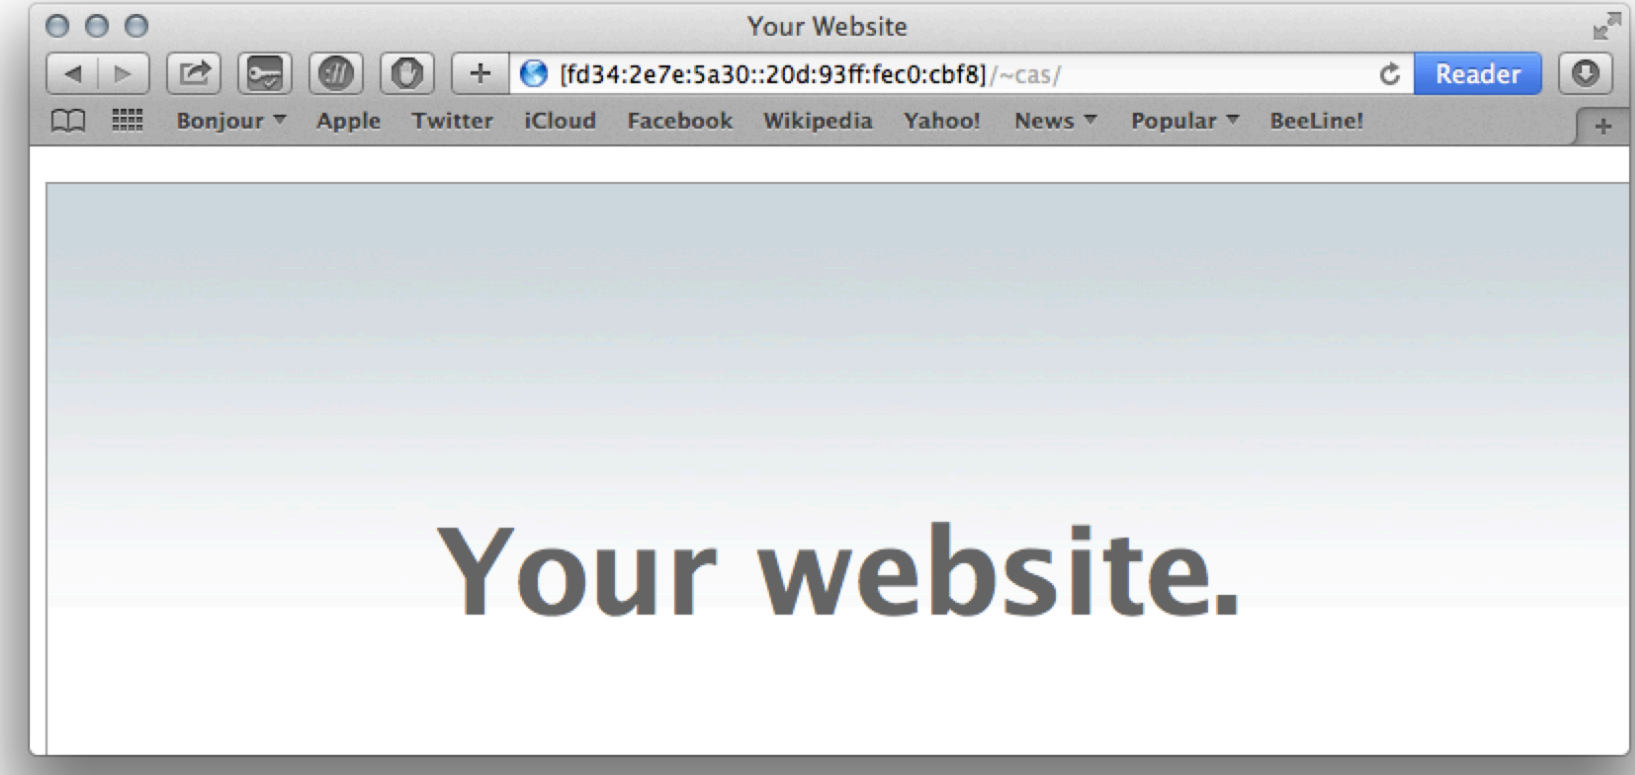

Literal IPv6 addresses in URLs

- ↪RFC 6874 specifies how to encode a zone identifier in URLs
  - Zone identifiers are needed on link-local scoped IPv6 addresses
  - Zone identifiers are separated from the IPv6 address by a encoded "%" sign: %25

- Example:

```
https://[fe80::226:b0ff:fed6:a4e0%25en0]:8088/index.html
```

ipv6-literal.net

- Many „legacy“ applications cannot handle IPv6 addresses, but they can work with DNS Names
- For legacy applications on modern Windows operating systems (Vista, 7, 8, 10, 11): IPv6 addresses can be entered as DNS names
  - To turn an IPv6 address into a DNS Name, all colons `:` in the address must be replaced by dashes – …
  - … and the 2nd level domain `ipv6-literal.net` must be appended

ipv6-literal.net

- A Windows TCP/IP Stack will not send these DNS Names ending in `ipv6-literal.net` to a resolver, instead it will connect directly to the derived IPv6 address.
- Example: IPv6 Address `2001:db8:0:0:0:0:fea:1` is converted into the pseudo name `2001-db8-0-0-0-0-fea-1.ipv6-literal.net.`

IPv6, e-mail and DNS

MX Records

- For receiving E-Mail over IPv6
  - The mail server must listen on an IPv6 address
  - The mail-host part of the MX record must resolve into an IPv6 address = AAAA DNS record

- For the MX record we have a choice
  - Generic MX record, mail-host resolves into both A and AAAA record(s) (IPv4 and IPv6)
  - Separate MX records for IPv4 and for IPv6

MX-Records and IPv6

DNS zonefile

```
$TTL 3600
example.org.  IN SOA dns1.example.org. hostmaster.example.org. (
                                20110116   ; serial
                                    1d   ; refresh
                                    2h   ; retry
                                   40d   ; expire
                                    2h ) ; neg TTL
example.org.  IN NS  dns1.example.org.
example.org.  IN NS  dns2.example.org.

example.org.  IN MX  10 mailhost-dual-stack.example.org.

mailhost-dual-stack.example.org.  IN A  192.0.2.44
mailhost-dual-stack.example.org.  IN AAAA 2001:db8:100::66
```

one mail-host, dual-stack with an IPv4 and IPv6 address

MX-Records and IPv6

DNS zonefile

```
$TTL 3600
example.org.  IN SOA dns1.example.org. hostmaster.example.org. (
                                20110116   ; serial
                                    1d   ; refresh
                                    2h   ; retry
                                   40d   ; expire
                                    2h ) ; neg TTL
example.org.  IN NS  dns1.example.org.
example.org.  IN NS  dns2.example.org.

example.org.  IN MX  10 mailhost-ipv6.example.org.
example.org.  IN MX  10 mailhost-ipv4.example.org.

mailhost-ipv4.example.org.  IN A  192.0.2.44
mailhost-ipv6.example.org.  IN AAAA 2001:db8:100::66
```

two MX records, one for each IP protocol

DNS IPv6 issues

Old DNS-Server and Load-Balancer

- Some (very old) DNS server and some (not so old) DNS load balancer products do not respond correctly when seeing a AAAA record query
  - They ignore the query
  - They sending a NAME ERROR (NXDOMAIN)
  - Return the AAAA record with an IPv4 address
  - Other wrong answers

- Wrong or no answer on AAAA in some cases can cause a connection failure when IPv4 transit (and DNS) would work correctly
- For example, a web browser that fails to connect to a web server it could otherwise reach (via IPv4)

IPv6 and Fragmentation (1)

- DNS, as designed in 1983 (RFCs 1034, 1035 & 1036), had a limitation of 512-Byte DNS payload over UDP.
  - The 512B limitation was lifted with the EDNS0 extension, RFC 2671 (Aug 1999) and RFC 6891 (April 2013).

- UDP DNS answers > 1280 byte can cause fragmentation
- IPv6 fragmentation is broken in the Internet: ↪RFC 7872 - Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World

IPv6 DNS troubleshooting

- Modern DNS troubleshooting tools support IPv6
  - `dig`, `delv`, `host` (BIND 9)
  - `drill` (ldns)
  - `unbound-host` (unbound)

- Avoid the old, unmaintained and obsolete `nslookup`

DNS Query examples

Example DNS Query

- Asking for IPv6 addresses:

```
# dig aaaa www.isc.org

; <<>> DiG 9.9.5 <<>> aaaa www.isc.org
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 61098
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 7

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

;; ANSWER SECTION:
www.isc.org.            60      IN      AAAA    2001:4f8:0:2::69

;; AUTHORITY SECTION:
isc.org.                7200    IN      NS      ns.isc.afilias-nst.info.
isc.org.                7200    IN      NS      ord.sns-pb.isc.org.
isc.org.                7200    IN      NS      sfba.sns-pb.isc.org.
isc.org.                7200    IN      NS      ams.sns-pb.isc.org.

;; ADDITIONAL SECTION:
ams.sns-pb.isc.org.     7200    IN      A       199.6.1.30
ams.sns-pb.isc.org.     7200    IN      AAAA    2001:500:60::30
ord.sns-pb.isc.org.     7200    IN      A       199.6.0.30
ord.sns-pb.isc.org.     7200    IN      AAAA    2001:500:71::30
sfba.sns-pb.isc.org.    7200    IN      A       149.20.64.3
sfba.sns-pb.isc.org.    7200    IN      AAAA    2001:4f8:0:2::19

;; Query time: 622 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Mon Oct 27 14:20:12 CET 2014
;; MSG SIZE  rcvd: 299
```

Example DNS Query

- Asking for names of an IPv6 address (Reverse DNS lookup)

```
# dig -x 2001:4f8:0:2::19

; <<>> DiG 9.9.5 <<>> -x 2001:4f8:0:2::19
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 37284
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;9.1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.0.0.0.0.0.0.0.8.f.4.0.1.0.0.2.ip6.arpa. IN PTR

;; ANSWER SECTION:
9.1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.2.0.0.0.0.0.0.0.8.f.4.0.1.0.0.2.ip6.arpa. 7200 IN PTR sfba.sns-pb.isc.org.

;; AUTHORITY SECTION:
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      sec2.authdns.ripe.net.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      ord.sns-pb.isc.org.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      sfba.sns-pb.isc.org.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      ams.sns-pb.isc.org.

;; Query time: 932 msec
;; SERVER: 172.22.1.22#53(172.22.1.22)
;; WHEN: Mon Oct 27 14:22:32 CET 2014
;; MSG SIZE  rcvd: 219
```

Example DNS Query

- Sending a query towards a DNS server via IPv6:

```
# dig -6 @sec2.authdns.ripe.net. 8.f.4.0.1.0.0.2.ip6.arpa. soa

; <<>> DiG 9.9.5 <<>> -6 @sec2.authdns.ripe.net. 8.f.4.0.1.0.0.2.ip6.arpa. soa
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 44834
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 1
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;8.f.4.0.1.0.0.2.ip6.arpa.      IN      SOA

;; ANSWER SECTION:
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      SOA     ns-int.isc.org. hostmaster.isc.org. 2014102100 28800 1800 2592000 3600

;; AUTHORITY SECTION:
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      ams.sns-pb.isc.org.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      sfba.sns-pb.isc.org.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      ord.sns-pb.isc.org.
8.f.4.0.1.0.0.2.ip6.arpa. 7200  IN      NS      sec2.authdns.ripe.net.

;; Query time: 63 msec
;; SERVER: 2001:67c:e0::4#53(2001:67c:e0::4)
;; WHEN: Mon Oct 27 14:25:05 CET 2014
;; MSG SIZE  rcvd: 211
```

Quiz

- An IPv6 enabled application that runs on a dual stack operating system wants to connect to `somehost.example.com`.
  - It will send a DNS query for
  - A record for `somehost.example.com`
  - `ANY` record for `somehost.example.com`
  - `AAAA` record for `somehost.example.com`
  - both `AAAA` and `A` records for `somehost.example.com`
  - `HTTPS` record for `somehost.example.com`

Questions?