IPv6 and DNS

- 'A' record of DNS(IPv4)
  - `www.kame.net  A  203.178.141.212`
- 'Qaud A' record of DNS(IPv6)
  - `www.kame.net  AAAA`
    `3ffe:501:4819:2000:5254:ff:fedc:50d2`
- A6 records - experimental
- Reverse: IP6.int, IP6.arpa - nibble boundary

IPv6 & DNS

- IPv6 specific RR's & labels in the DNS
- Address Aggregation & Prefix Delegation
- Zone file examples
- Configuring Bind v9 for IPv6
- Issues

IPv6 RR/1

- AAAA resource record (rfc 1886)
  – NAME [TTL] TYPE ADDRESS Example:
  `vh  AAAA  2001:0780:0:1::1`
- Usage is similar to A record.

IPv6 RRs/2

- Reverse mapping & Nibble labels
- example of an IPv6 address:
  `3ffe:1900:4545:2:2d0:9ff:fef7:6d2c`
- This would map in the IPv6 .ip6.int. tree as follows:
- `c.2.d.6.7.f.e.f.f.f.9.0.0.d.2.0.2.0.0.0.5.4.5.4.0.0.9.1.e.f.f.3.ip6.int. PTR ns.example.com.`
- IPv6 Reverse delegation with ip6.int. is described in rfc 1886
  - Nibble labels are also described in rfc 1886
  - Nibble = 4 bits presented by a hexadecimal number.

IPv6 RRs/3

- A6 resource record (rfc 2874)
- NAME [TTL] TYPE BITS ADDRESS [REFERRAL]
- Example
  - At a host site:
    - `eg.example.com.  A6 96 ::fef7:6d2c ipv6.isp.com.`
  - And at the provider site:
    - `ipv6.isp.com. A6  0 3ffe:1900:4545:2:2d0:9ff::`
- Defined in RFC 2874.

IPv6 RRs/4

- Reverse mapping & Binary labels
  - example of an IPv6 address
    - `3ffe:1900:4545:2:2d0:9ff:fef7:6d2c`
  - This would map in the IPv6 .ip6.arpa. tree as follows.
    - `\[x3ffe19004545000202d009fffef76d2c/128].ip6.arpa. PTR ns.example.com.`
- Binary labels start with \[and end with]
- x defines the base (this case hexadecimal number).
- /128 defines count of significant bits.
- Binary labels are described in RFC 2637
- delegation at bit boundaries instead of nibble boundaries (as with ip6.int.)

IPv6 RRs/5

- DNAME resource record
  - NAME [TTL] DNAME REFERRAL
  - Example
    - At a provider site:
      - \[x3ffe1900454500020 2d009ff/96].ip6.arpa. DNAME example.com.
    - And at a host site:
      - \[xfef76d2c/32].example.com. PTR eg.example.com.
    - Resolver will aggregate it (eventually) as:
      - \[x3ffe1900454500020 2d009fffef76d2c/128].ip6.arpa. PTR eg.example.com.
- DNAME might be used only in ip6.arpa. reverse zones.
- Defined in RFC 2672.

RRs summary

- IPv4
  - A - name to address translation
  - PTR - address to name trasnlation
  - CNAME - alias, 1 level recursion
- IPv6
  - AAAA - as in IPv4
  - A6 - partial name to address translation (0-128 bits)
  - PTR - as in IPv4
  - DNAME -similar to CNAME, partial replace, multiple level of recursion

Forward Zone file example

At a host site:

```
$TTL 1D
$ORIGIN example.com.
@    IN SOA   ns1 hostmaster 2002101802 8H 2H 1W 1D
        NS    ns1
        NS    ns1.example2.com.

roy     AAAA    3ffe:1900:4545:2:2d0:9ff:fef7:6d2c
roy     A6    96  ::fef7:6d2c  ip6-referral.myprovider.com.
```

Entries at a provider site:

```
$ORIGIN myprovider.com.
.....
ip6-referral A6    0   3ffe:1900:4545:2:2d0:9ff::
.....
```

Reverse Zone Example file

```
For the ip6.int. tree. (nibble label example)

$TTL 1D

$ORIGIN f.f.9.0.0.d.2.0.2.0.0.0.5.4.5.4.0.0.9.1.e.f.f.3.ip6.int.

@    IN SOA   ns1.example.com. hostmaster.example.com. 2000101802 8H 2H 1W 1D

        NS    ns1.example.com.

        NS    ns1.example2.com.

c.2.d.6.7.f.e.f PTR roy.example.com.

d.2.d.6.7.f.e.f PTR ns1.example.com.


For the ip6.arpa. tree. (binary label example)

Entries at a provider site:

$ORIGIN \[x3ffe19004545/48].ip6.arpa.

\[x000202d009ff/48]   DNAME  rev.example.com.

At a host site.

$TTL 1D

$ORIGIN example.com.

@    IN SOA   ns1 hostmaster 2000101802 8H 2H 1W 1D

        NS    ns1

        NS    ns1.example2.com.

\[xfef76d2c/32].rev PTR roy

\[xfef76d2d/32].rev PTR ns1
```

IPv6 DNS support

- BIND8
  - IPv6 RRs only AAAA)
  - IPv4 transport (IPv6 transport with patch)
- BIND9
  - All IPv6 RRs
  - IPv4/IPv6 transport
- djbns
  - IPv6 RRs only AAAA)
  - IPv4 transport (IPv6 transport with patch)

Bind 9 configuration/1

- named.conf entries
  - Multiple listen-on-v6 options can be used to listen on multiple ports:

```
options {
        listen-on-v6 port   53 { any; };
        listen-on-v6 port 1234 { any; }
};
```

  - To make the server not listen on any IPv6 address (which is the default state), use

```
options {
        listen-on-v6 { none; };
};
```

Bind9 configuration/2

- Zone transfer:

  ```
  transfer-source-v6 1:2:3:4:5:6:7:8;
  ```

- Convert RFC1886-style recursive lookup requests into RFC2874-style lookups with

  ```
  allow-v6-synthesis;
  ```

BIND9 lightweight resolver

- New RRs for IPv6 introduces new complexity into the resolution process. A6/DNAME chains & simultaneous IPv4/IPv6 lookups, DNSSEC.
  - in concept similar to ncsd
  - Resolver library (get*by*() routines) calls resolver daemon.
  - Uses lightweight resolver protocol.
  - listens to IPv4 loopback address, port 921
  - Acts like caching server

Issues

- A6 DNAME/PTR referrals can cause overwhelming dns-traffic.
  - In theory, 127 referrals are possible and every refered NS can have an aggregatable address again.
- A6 DNAME/PTR referrals can cause lookup loop.
  - (A refers to B refers to C refers to A) DoS attack.
- >A6 and DNAME moved to experimental!
- Signing of A6 records means aggregated addresses can be partially signed.
- No IPv6 transport capable root nameserver.- can be resolved with hidden IPv6 roots. - alternative roots are operational - Bill Manning