Geographic Split-horizon DNS

RIPE 71, Bucharest

Jan Včelák • jan.vcelak@nic.cz • 2015 November 19

CZ DOMAIN REGISTRY

Implementation challenges

- Mapping of resources to locations
- Scope of the split-horizon
- Multiple servers deployment
- Compliance with global resolvers
- DNSSEC

EDNS Client Subnet

- draft-ietf-dnsop-edns-client-subnet

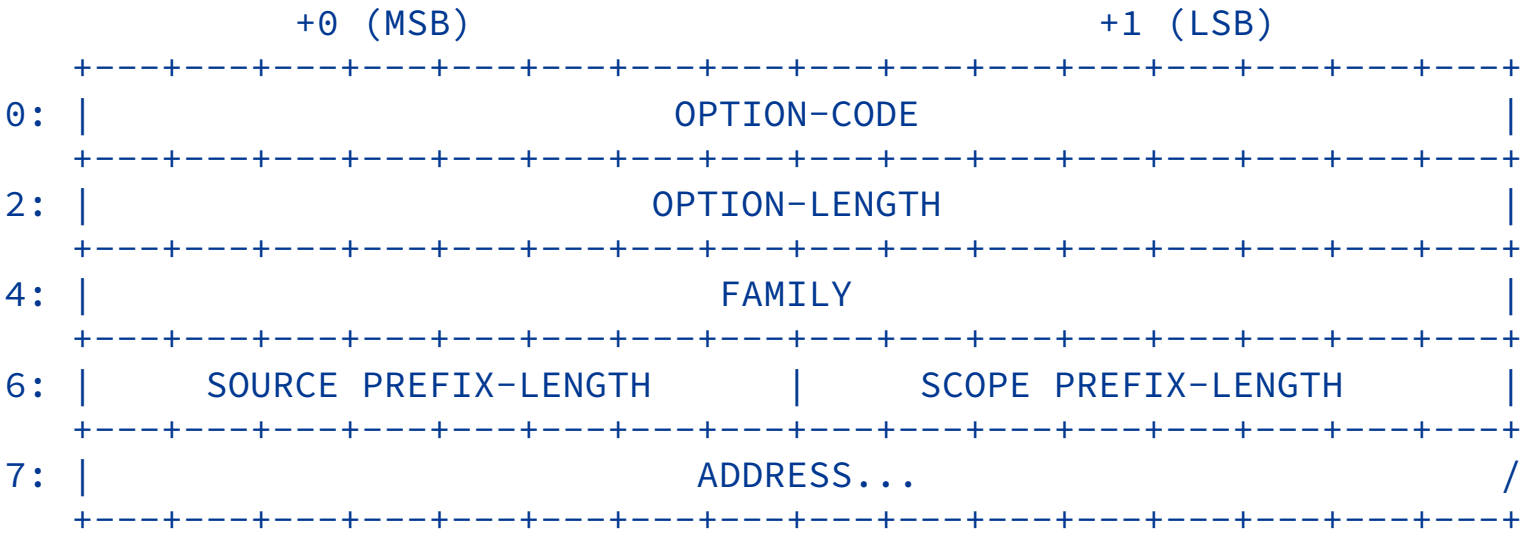

- Supported by a lot of (non-anycast) CDN providers
- Essential for global resolver providers

Existing implementations

- geoipdns
- GeoDNS
- PowerDNS
- gdnsd

- Knot DNS

Existing implementations – geoipdns

- based on djbdns
- no EDNS Client Subnet, no DNSSEC
- works at the resource record level

```
%prague:192.0.2.0:24
%amsterdam:198.51.100.0:24

+www.example.com:192.0.2.42:1200::prague
+www.example.com:198.51.100.42:1200::amsterdam
+www.example.com:203.0.113.42:1200::nomatch
```

Existing implementations – GeoDNS

- patch for BIND (IPv4 only)
- no EDNS Client Subnet, limitted DNSSEC support
- works at the zone level

```
view "america" {
  match-clients { country_US; country_CA; country_MX; ... };
  zone "example.com" {
    type master;
    file "example.com/america.db";
  };
};
view "other" {
  match-clients { any; };
  zone "example.com" {
    type master;
    file "example.com/fallback.db";
  };
};
```

Existing implementations – PowerDNS

- geoipbackend
- no EDNS Client Subnet, DNSSEC supported (Front-Signing)
- works at the resource record level (synthesizes CNAMEs)

```
domains:
- domain: example.com
  ttl: 30
  records:
    example.com:
      - soa: ns.example.com root.example.com 1234 600 300 3600 30
      - ns: ns.example.com
      - mx: 10 mx.example.com
    cz.eu.example.com:
      - a: 192.0.2.42
      - aaaa: 2001:db8::42
  services:
    www.example.com: '%co.%cn.example.com'
```

Existing implementations – gdnsd

- most advanced implementation
- great support for EDNS client subnet, no DNSSEC
- works at the resource record level

```
maps => {
  manual_map => {
    datacenters => [dc-eu, dc-us, dc-fail],
    map => { EU => [ dc-eu, dc-fail ], NA => [ dc-us, dc-fail ] }
  },
},
resources => {
  prod_www => {
    map => manual_map
    dcmap => {
      dc-eu => 192.0.2.42, dc-us => 203.0.113.42, dc-fail => fail.example.com.
    }
  }
}

www    600 DYNA geoip!prod_www
```

Existing implementations – Knot DNS

- prototype, *module-geoip* branch in the repository
- EDNS client subnet supported, DNSSEC supported (online-signing)
- works at the resource record level

```
mod-geoip:
 - id: default
   database: /usr/share/GeoLite2/GeoLite2-Country.mmdb
mod-online-sign:
 - id: default
zone:
  - domain: example.com
    module: [ mod-geoip/default, mod-online-sign/default ]
```

```
www.cz      60 A      192.0.2.42
            60 AAAA   2001:db8::42
www.us      60 A      203.0.113.42

www         60 TYPE65280  \# 9 7777772E256320656E   ; "www.%c en"
```

IP address lookup

- Libraries:
  - libGeoIP – GPL 2.1+, legacy library
  - libmaxminddb – APL 2.0
  - ...
- Databases:
  - GeoLite – CC BY-SA 3.0
  - IpToCountry – donationware
  - IP2Location Lite – free, requires attribution
  - ...

IP address lookup – libmaxminddb

1. Open the database:

```
MMDB_open("GeoLite2-City.mmdb", 0, &db);
```

2. Lookup the entry by the IP address:

```
struct sockaddr ip = /* ... */;
match = MMDB_lookup_sockaddr(&db, &ip, &error);
```

3. Retrieve the entry attribute:

```
MMDB_get_value(&match.entry, &value,
                "country", "iso_code", NULL);
```

4. Profit!

```
printf("%.*s\n", value.data_size, value.utf8_string);
```

Thank you!

Jan Včelák • jan.vcelak@nic.cz