DNS Concepts

FUNDAMENTALS TRAINING

Overview

- What is DNS?
- DNS Features
- Domains and Namespaces
- Zones and Delegation
- Nameservers
- DNS Resource Records
- DNS Query

DNS Overview

DNS is a distributed, hierarchical system for translating objects

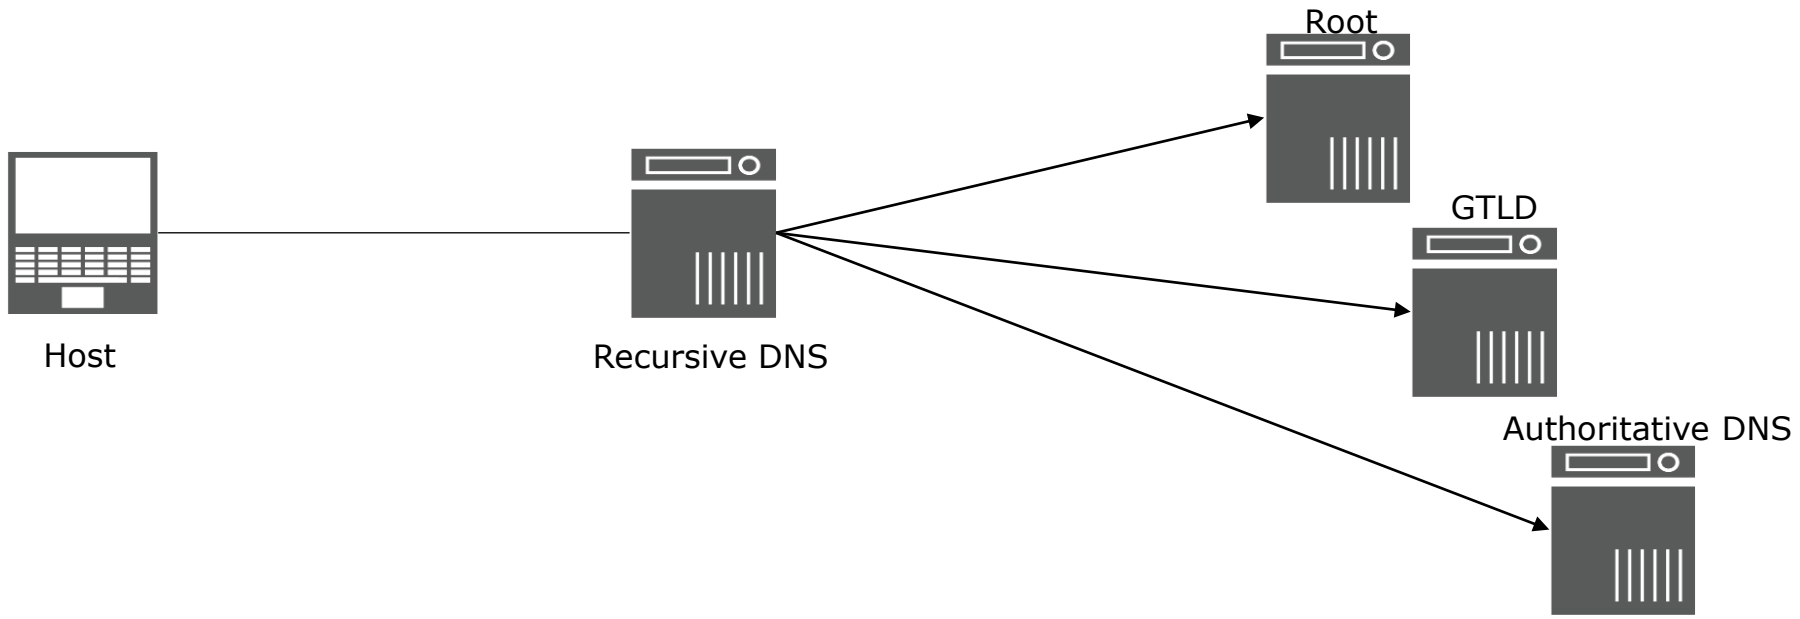

DNS is a critical piece of the Internet infrastructure

IP Addresses vs Domain Names

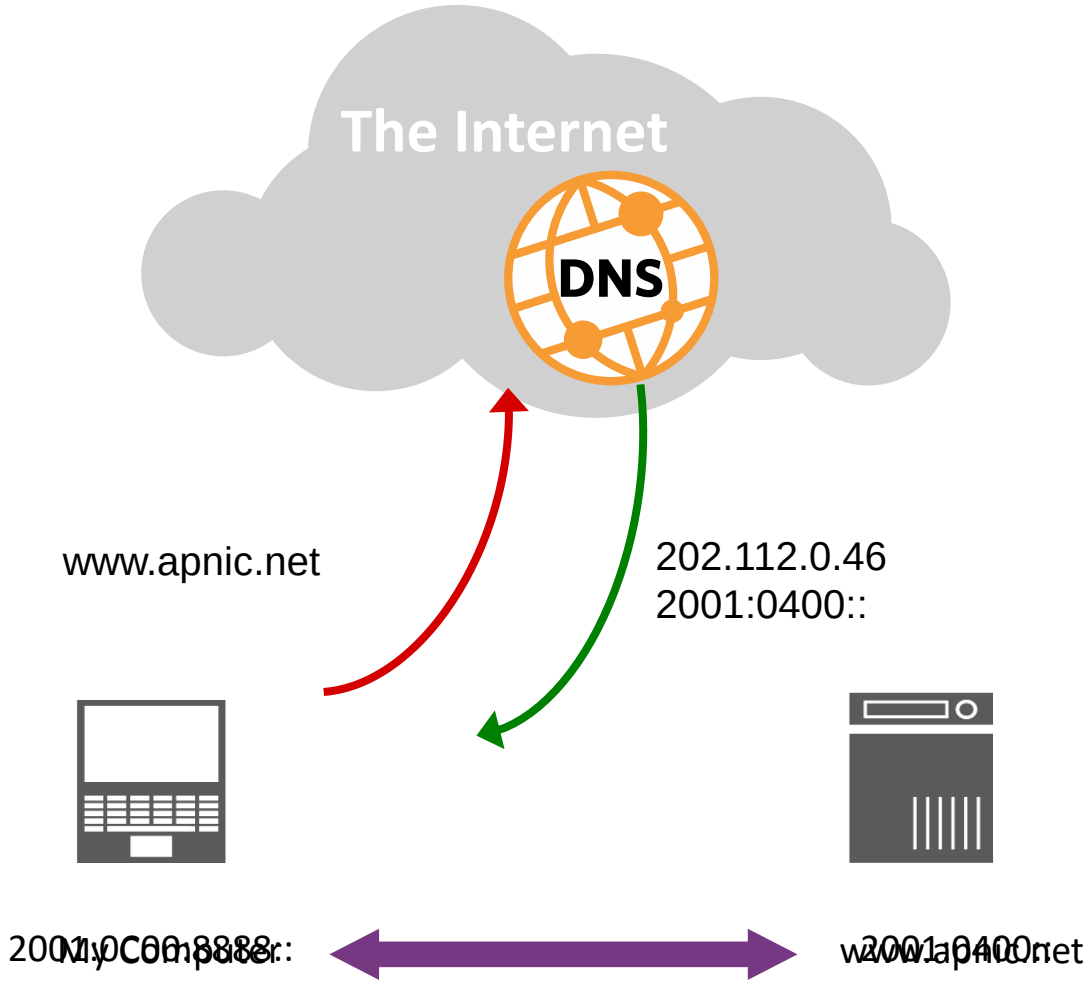

Is DNS important? (Akamai Outage)

A DNS outage just took down a large chunk of the internet

Zack Whittaker @zackwhittaker / 2:55 AM GMT+10 • July 23, 2021

Comment

A large chunk of the internet dropped offline on Thursday. Some of the most popular sites, apps and services on the internet were down, including UPS and FedEx (which have since come back online), Airbnb, Fidelity, and others are reporting Steam, LastPass, and the PlayStation Network are all experiencing downtime.

Many other websites around the world are also affected, including media outlets in Europe.

What appears to be the cause is an outage at Akamai, an internet security giant that provides networking and content delivery services to companies. At around 11 a.m. ET, Akamai reported an issue with its Edge DNS, a service that's designed to keep websites, apps and services running smoothly and securely.

Is DNS important? (Facebook outage)

POSTED ON OCTOBER 5, 2021 TO NETWORKING & TRAFFIC

More details about the October 4 outage

To ensure reliable operation, our DNS servers disable those BGP advertisements if they themselves can not speak to our data centers, since this is an indication of an unhealthy network connection. In the recent outage the entire backbone was removed from operation, making these locations declare themselves unhealthy and withdraw those BGP advertisements. The end result was that our DNS servers became unreachable even though they were still operational. This made it impossible for the rest of the internet to find our servers.

All of this happened very fast. And as our engineers worked to figure out what was happening and why, they faced two large obstacles: first, it was not possible to access our data centers through our normal means because their networks were down, and second, the total loss of DNS broke many of the internal tools we'd normally use to investigate and resolve outages like this.

https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/

Is DNS important?

Cloudflare DNS goes down, taking a large piece of the internet with it

Devin Coldewey @techcrunch / 7:50 AM GMT+10 • July 18, 2020

Comment

Many major websites and services were unreachable for a period Friday afternoon due to issues at Cloudflare's 1.1.1.1 DNS service. The outage seems to have started at about 2:15 Pacific time and lasted for about 25 minutes before connections began to be restored.

https://techcrunch.com/2020/07/17/cloudflare-dns-goes-down-taking-a-large-piece-of-the-internet-with-it/

DNS Features

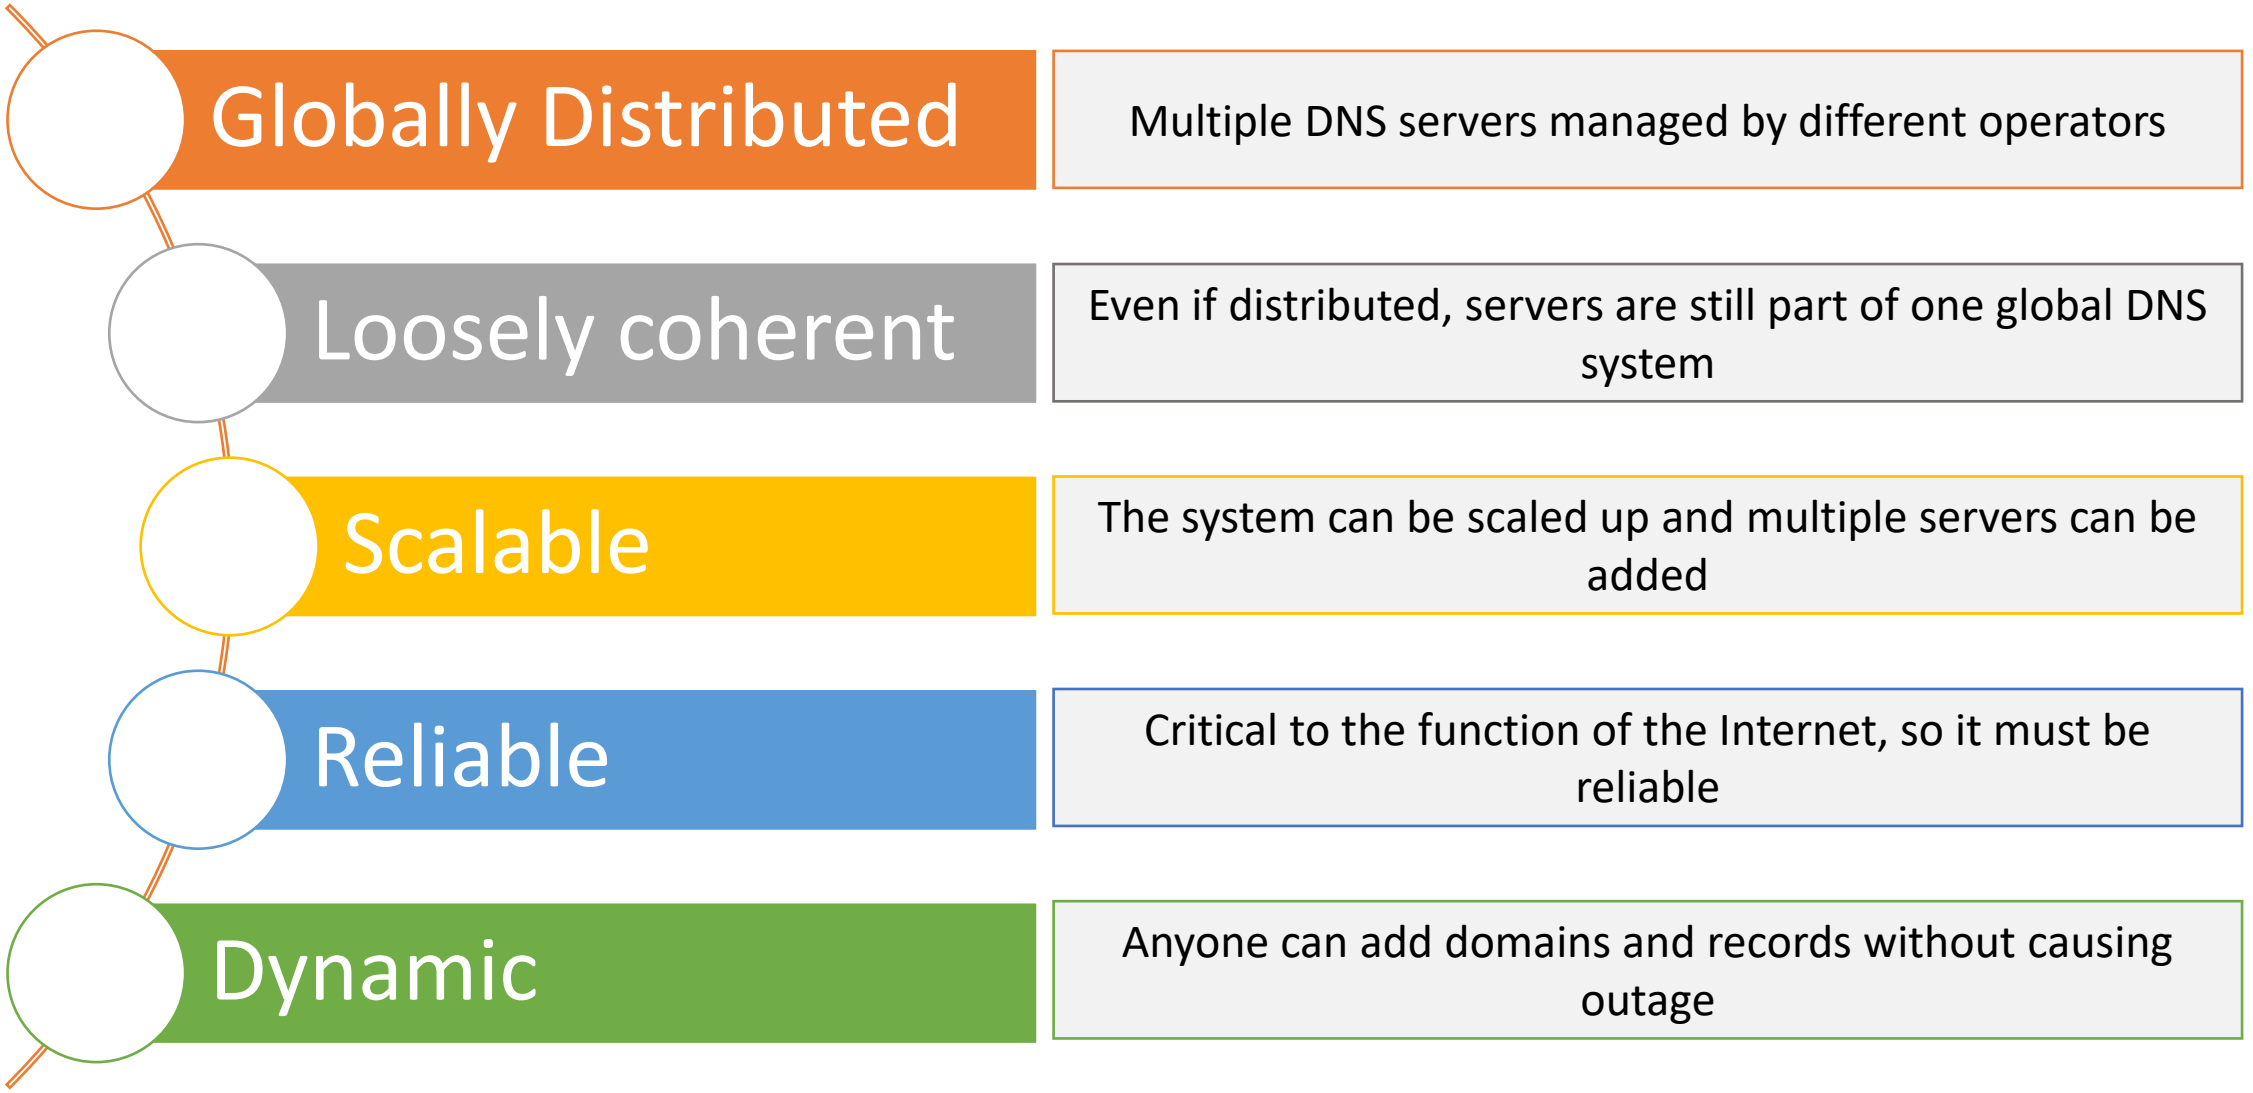

More about DNS

- DNS is a client-server application
  - Client (resolvers) must request, and DNS server responds with information about the record

- Requests and responses are normally sent via UDP port 53

- Occasionally uses TCP port 53 for large requests
  - Ex: Zone transfers

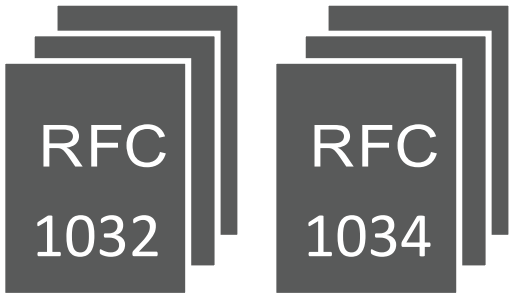

DNS Transport

Name	Description
DNS-over-UDP (Do53)	• Traditional means of DNS transport using port 53 • Limited by message length (<512 bytes) and lacks security
DNS-over-TCP (Do53/TCP)	• Optional transport to allow larger messages (ex: for zone transfers) • Provides TCP advantages like reliable delivery
DNS-over-TLS (DoT)	• Tunnels DNS messages using TLS • Requires a DoT server listening usually on port 853
DNS-over-HTTPS (DoH)	• Tunnels DNS messages over HTTPS using port 443 • Traffic is indistinguishable from web traffic

What is DNS?

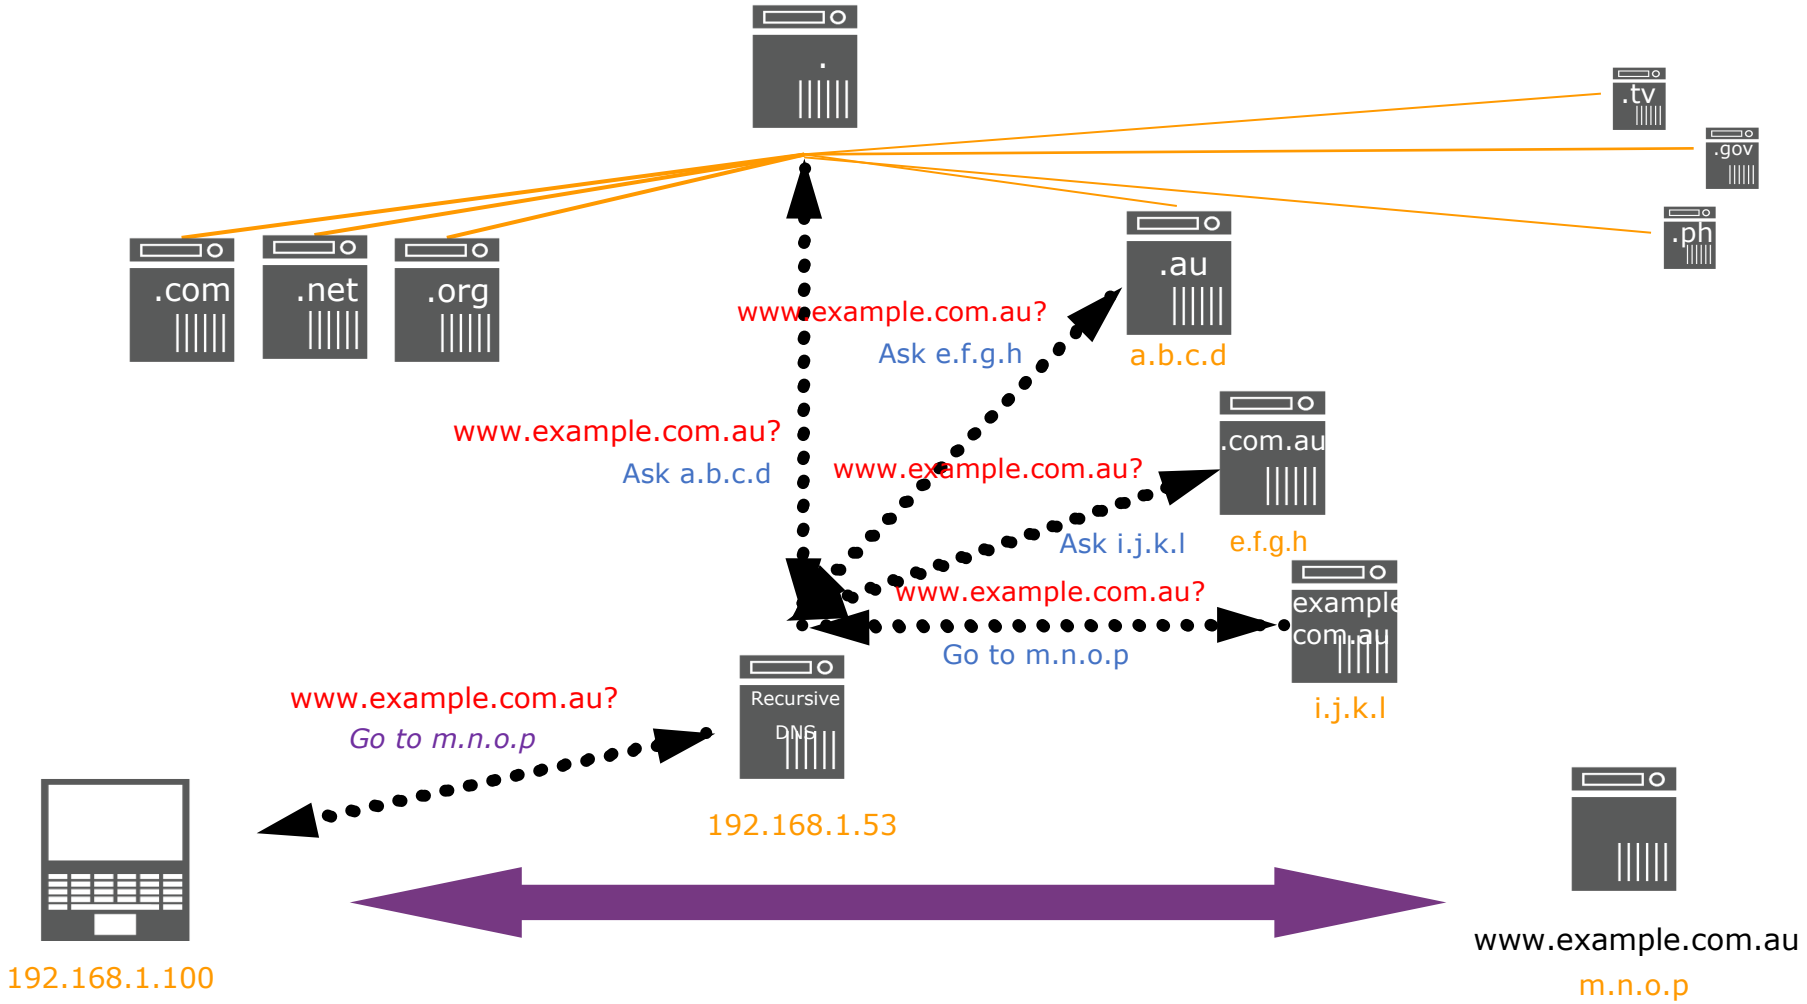
.
.com
.net
.org
.tv
.gov
.ph
.au
a.b.c.d
www.example.com.au?
Ask e.f.g.h
www.example.com.au?
Ask a.b.c.d
.com.au
e.f.g.h
www.example.com.au?
Ask i.j.k.l
www.example.com.au?
Go to m.n.o.p
example
com.au
i.j.k.l
www.example.com.au?
Go to m.n.o.p
Recursive
DNS
192.168.1.53
192.168.1.100
www.example.com.au
m.n.o.p

DNS Hierarchy Tree

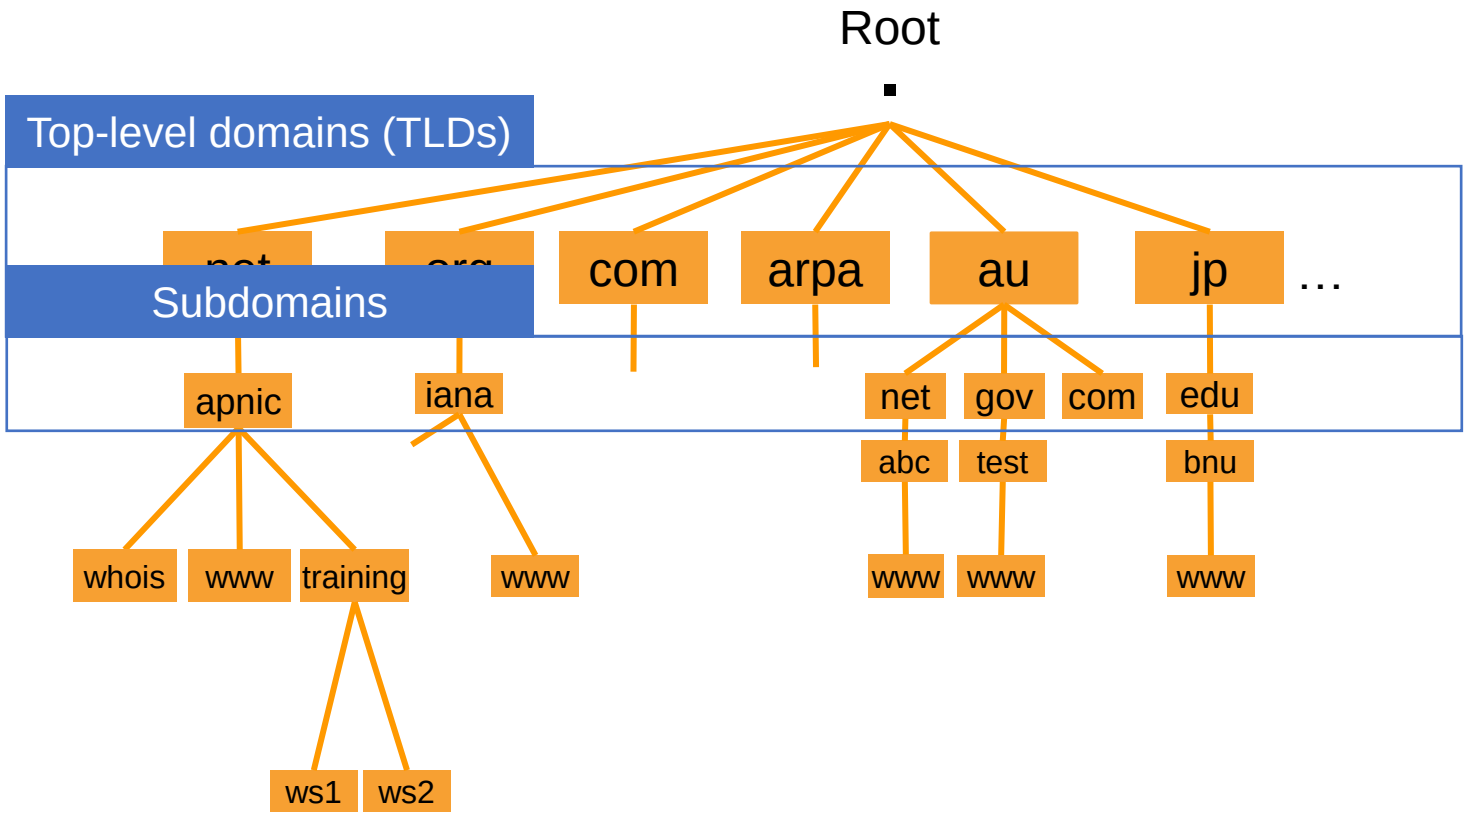

Country-code TLDs (ccTLDs)

Generic TLDs (gTLDs)

Infrastructure TLD

Internationalized TLDs (IDN)

FQDN: ws1.training.apnic.net.

FQDN = **F**ully **Q**ualified **D**omain **N**ame

DNS Components

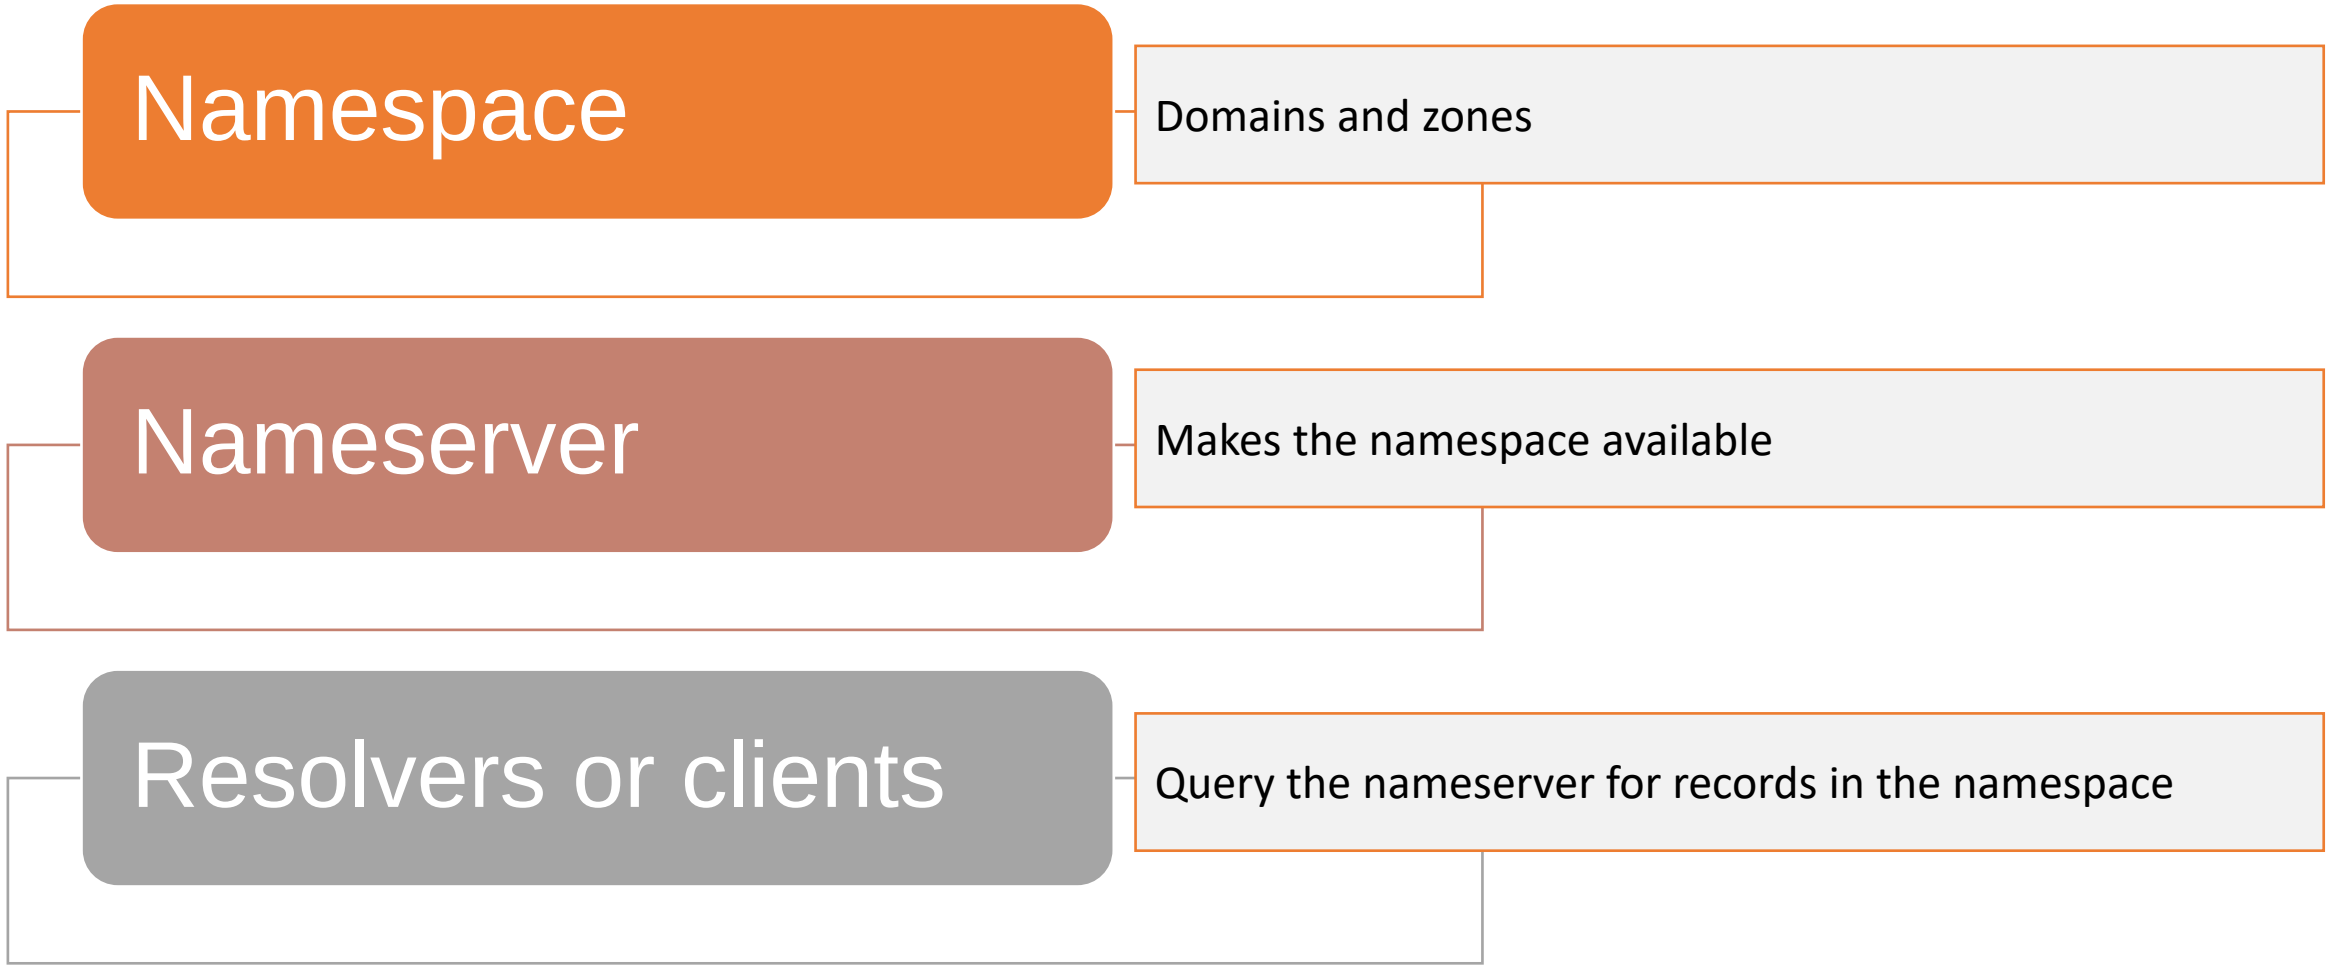

Domains

Domains are “namespaces”

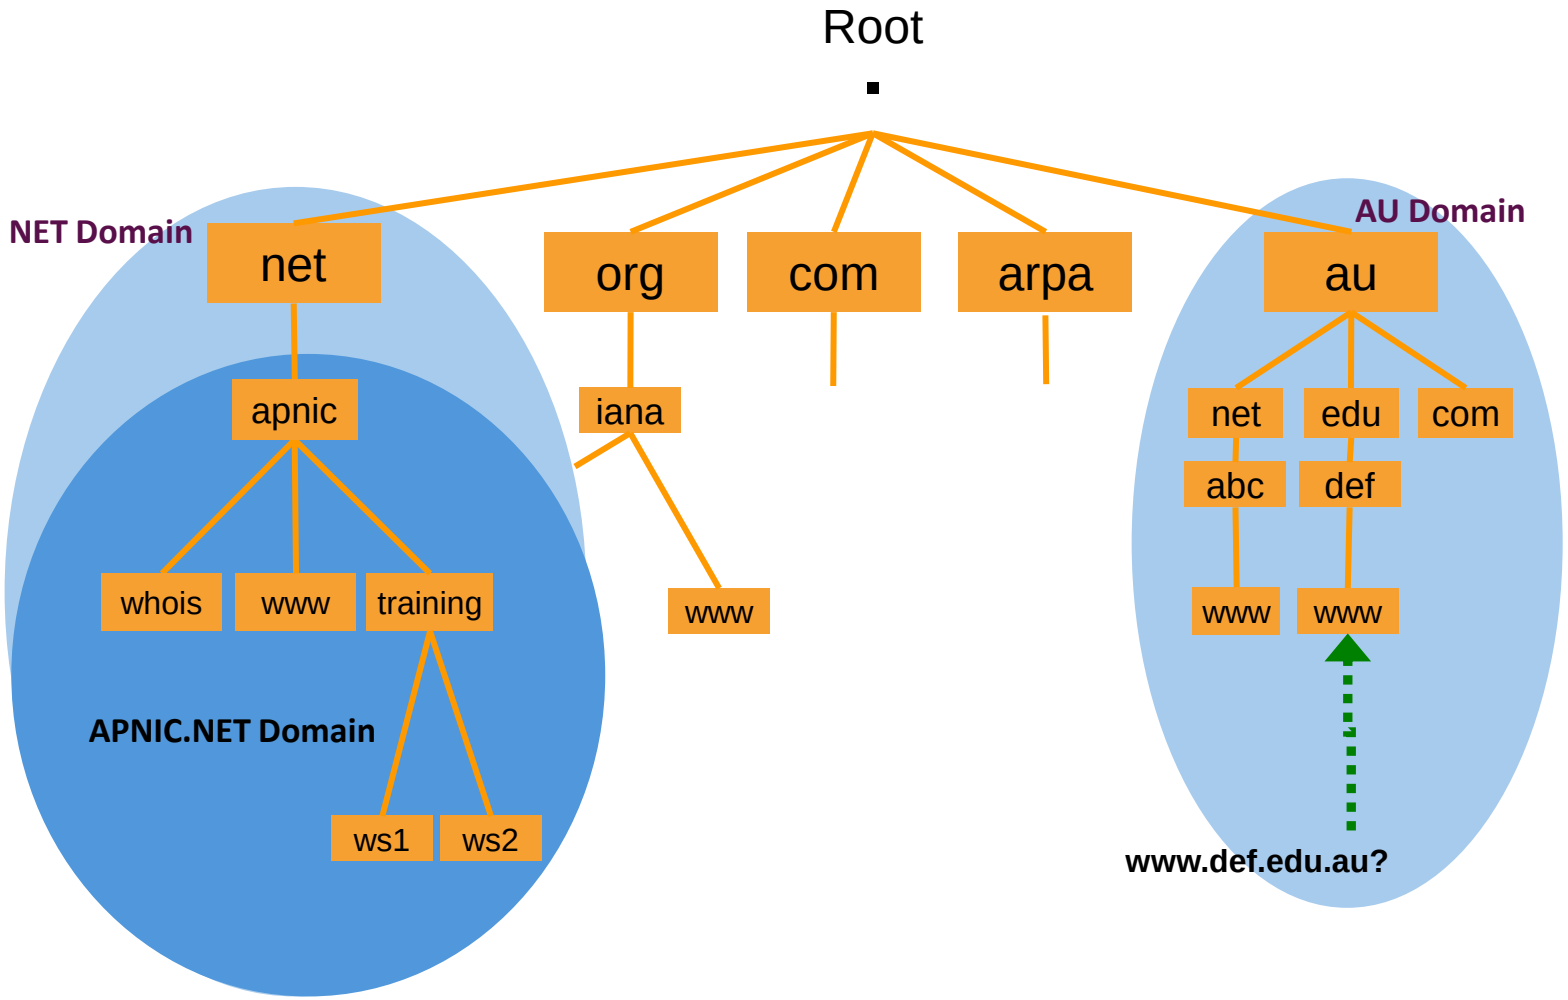

Delegation

Administrators can create subdomains to group hosts

Administrators can delegate responsibility for managing a subdomain to someone else

The parent domain retains links to the delegated subdomain

Zones

Zones are “administrative spaces”

Zone administrators are responsible for a portion of a domain’s name space

Authority is delegated from parent to child

Zones

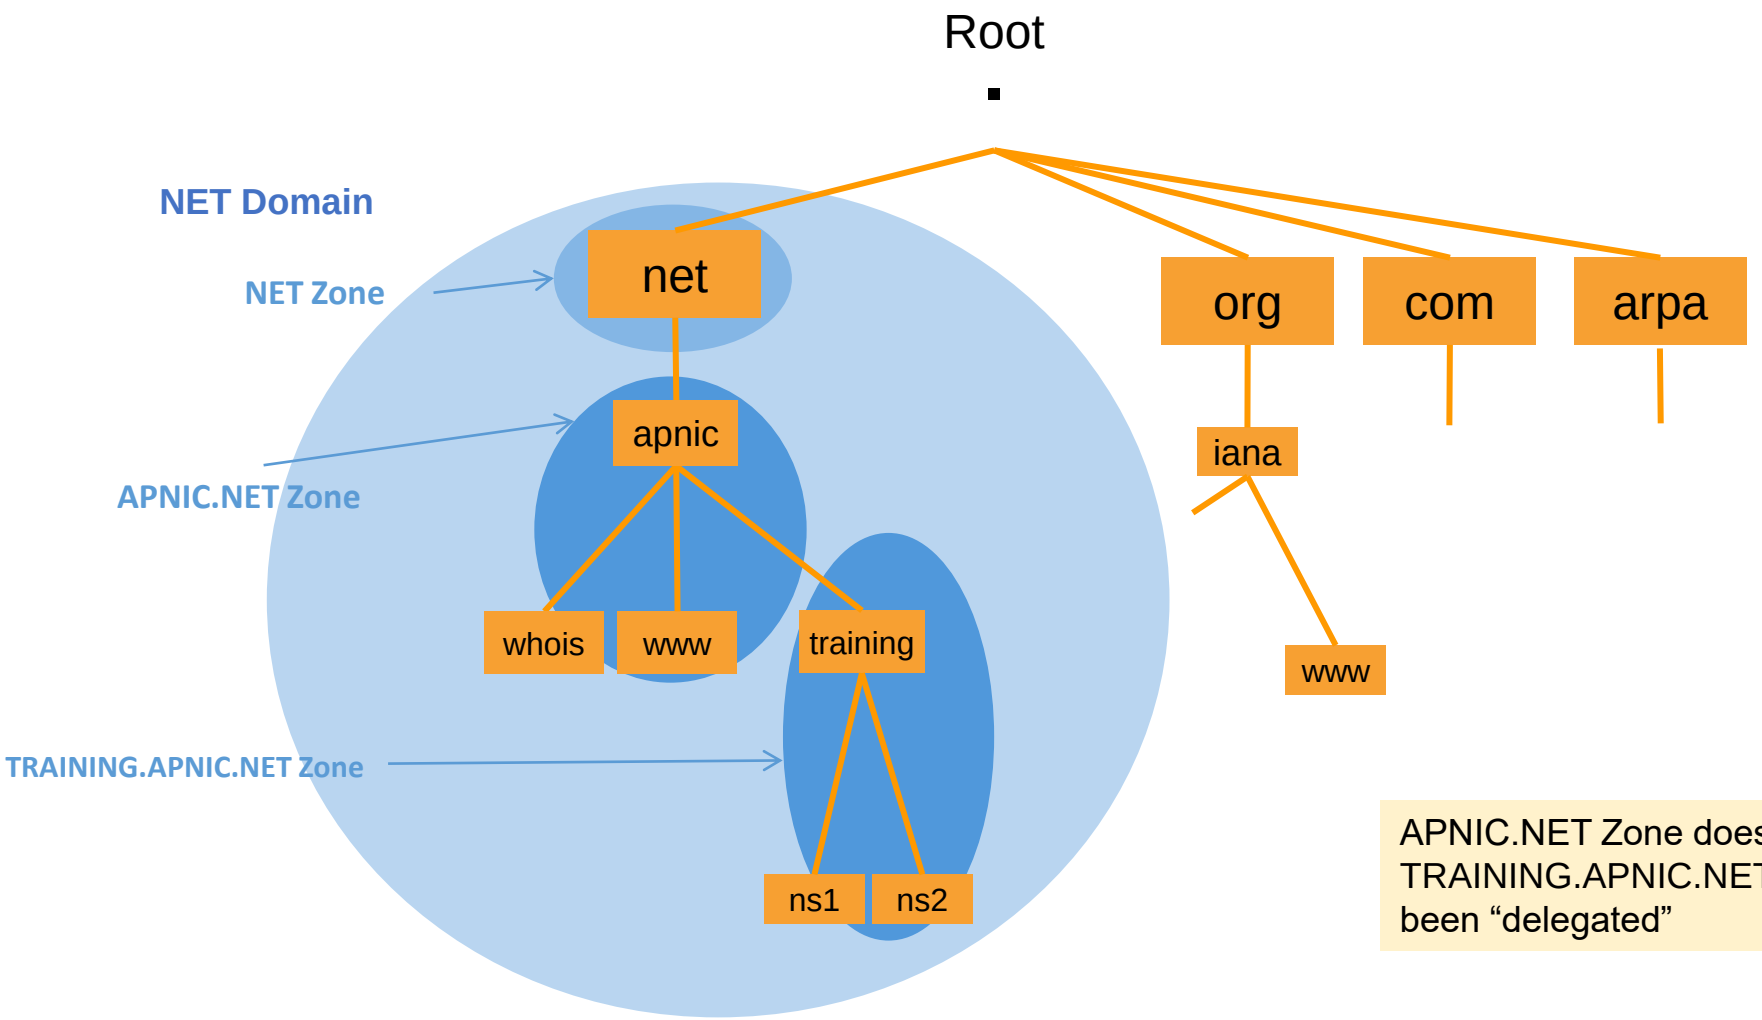

APNIC.NET Zone doesn't include TRAINING.APNIC.NET since it has been "delegated"

Name Servers

- Name servers answer DNS questions

- Several types of name servers
  - Authoritative servers
    - Primary
    - Secondary
  - Recursive servers
    - also caching forwarders

- Mixture of functions

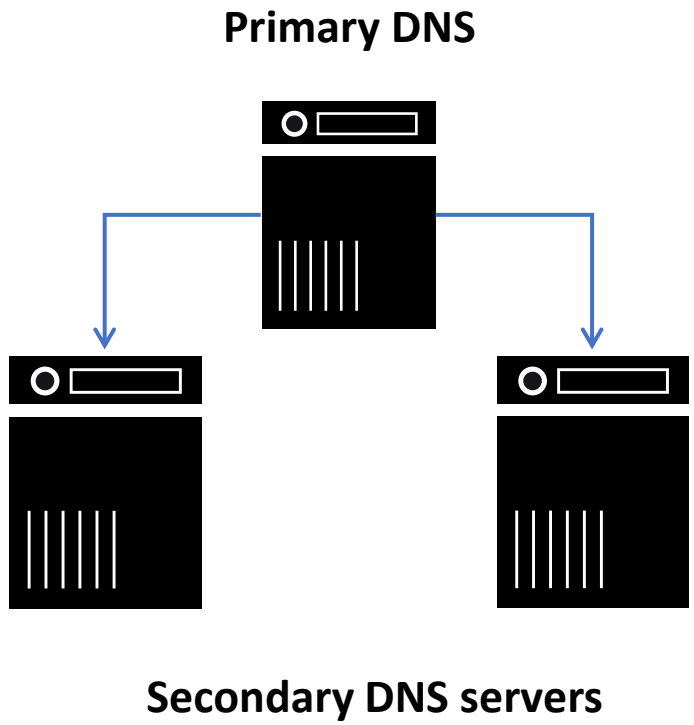

Authoritative Nameserver

- A nameserver that is authorised to provide an answer for a particular domain
  - Can be more than one auth nameserver
- Two types based on management method:
  - Primary (Master) and Secondary (Slave)
- Only one primary nameserver
  - All changes to the zone are done in the primary
- Secondary nameserver/s will retrieve the zonefile from the primary server
  - Secondary polls the primary periodically
- Primary server can “notify” the secondary servers

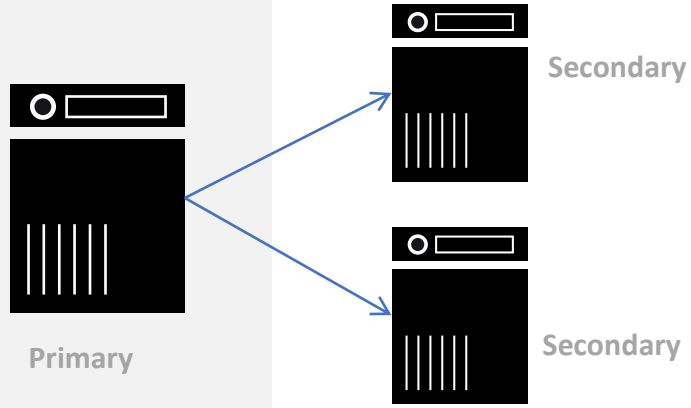

Recursive Nameserver

- The job of the recursive nameserver is to locate the authoritative nameserver and get back the answer
- This process is iterative – starts at the root
- Recursive servers are also usually caching servers
- Prefer a nearby cache
  - Minimizes latency issues
  - Also reduces traffic on your external links

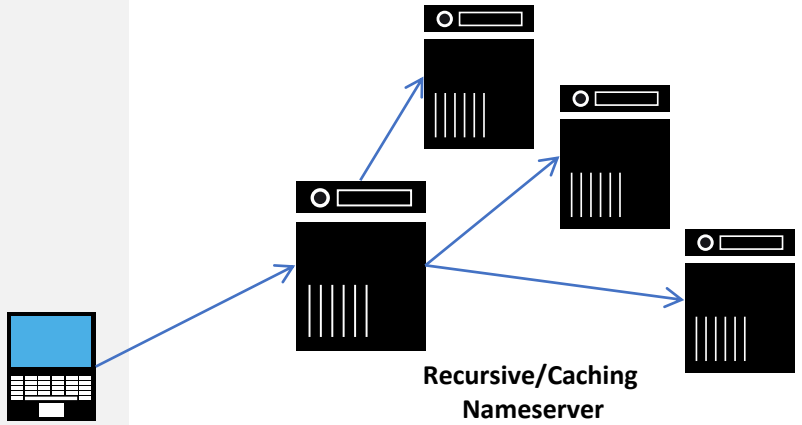

Root Servers

- The top of the DNS hierarchy
- There are 13 root name servers operated around the world

  `[a-m].root-servers.net`
- There are more than 13 physical root name servers
  - Each rootserver has an instance deployed via anycast

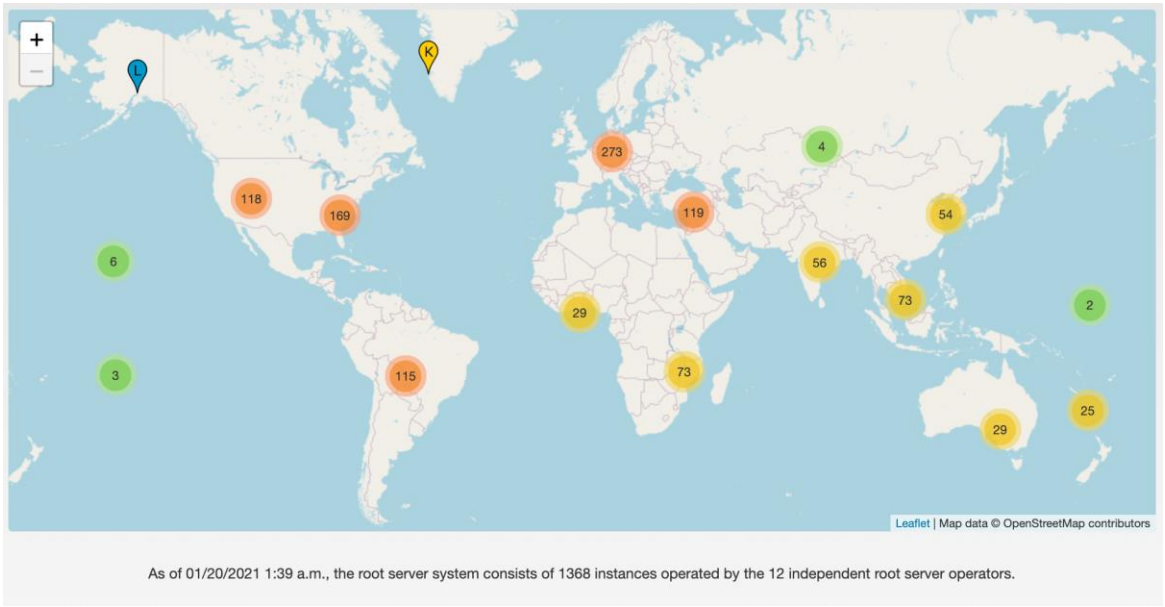

As of 01/20/2021 1:39 a.m., the root server system consists of 1368 instances operated by the 12 independent root server operators.

Src: https://root-servers.org/

Root Server Deployment at APNIC

- Started in 2002, APNIC is committed to establish new root server sites in the AP region
- The aim is to strengthen DNS by deploying additional resources to handle growing Internet traffic.

Timeline of root server deployment

2020	**December** M-Root nameserver installed in Brisbane.
2019	**January** K-Root nameserver installed in Thimphu. **December** K-Root nameserver installed in Yangon.
2018	**July** F-Root nameserver installed in Port Moresby. **December** K-Root nameserver installed in Taipei.
2017	**January** J-Root nameserver installed in Kathmandu.

Ref: https://www.apnic.net/community/support/root-servers/

Resource Records

Entries in the DNS zone file

Resource Record	Function
Label	Name substitution for FQDN
TTL	Timing parameter, an expiration limit
Class	IN for Internet, CH for Chaos
Type	RR Type (A, AAAA, MX, PTR) for different purposes
RDATA	Anything after the Type identifier; Additional data

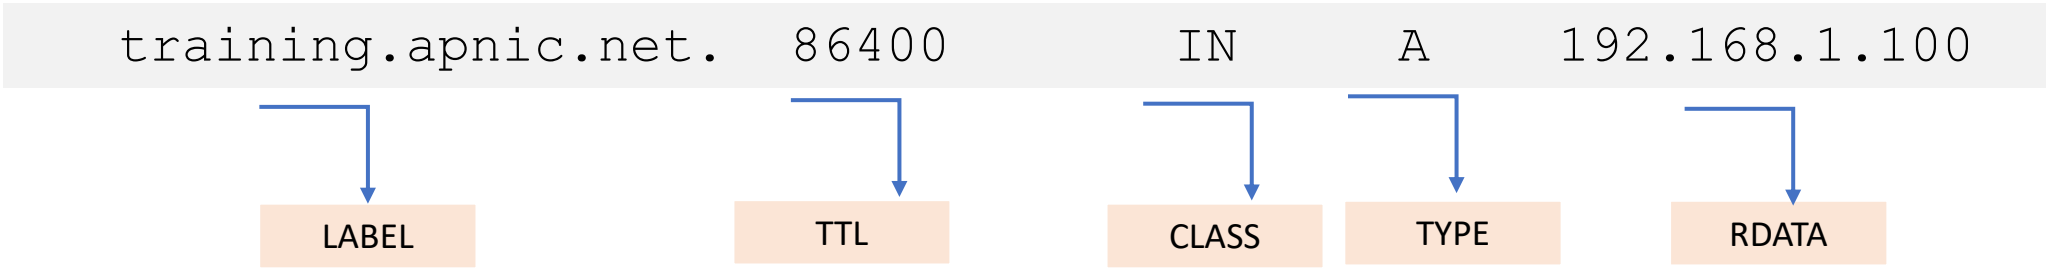

Common Resource Record Types

RR Type	Name	Functions
A	Address record	Maps the domain name to IP address www.example.com. IN A 192.168.1.1
AAAA	IPv6 address record	Maps the domain name to an IPv6 address www.example.com. IN AAAA 2001:db8::1
NS	Name server record	Used for delegating zone to a nameserver example.com. IN NS ns1.example.com.
PTR	Pointer record	Maps an IP address to a domain name 1.1.168.192.in-addr.arpa. IN PTR www.example.com.
CNAME	Canonical name	Maps an alias to a hostname web IN CNAME www.example.com.
MX	Mail Exchanger	Defines where to deliver mail for user @ domain example.com. IN MX 10 mail01.example.com. IN MX 20 mail02.example.com.

Example: RRs in a Zone File

```
apnic.net.      7200        IN   SOA     ns.apnic.net. admin.apnic.net. (
                            2020072001   ; Serial
                            12h          ; Refresh 12 hours
                            4h           ; Retry 4 hours
                            4d           ; Expire 4 days
                            2h )         ; Negative cache 2 hours

apnic.net.      7200        IN    NS            ns.apnic.net.
apnic.net.      7200        IN    NS            ns.ripe.net.

www.apnic.net.  3600        IN    A             192.168.0.2
www.apnic.net   3600        IN    AAAA    2001:DB8::2
```

DNS Data Flow

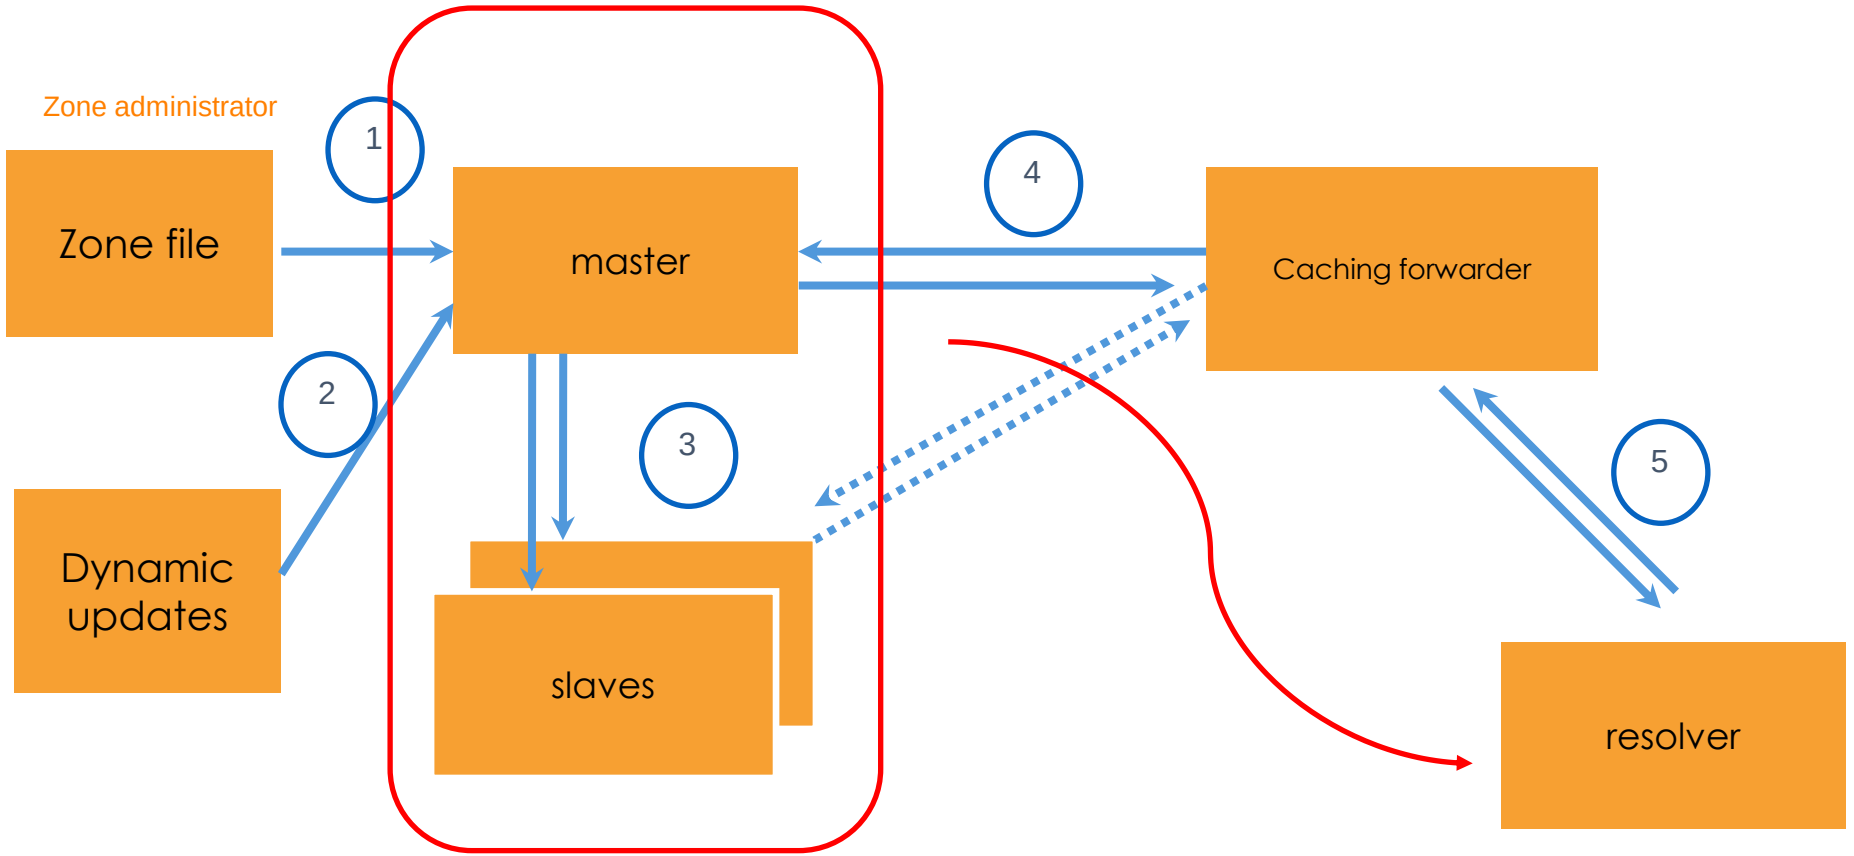

Delegating a Zone

Delegation is done by adding NS records.

In this example, **apnic.net** zone is delegating the subdomain **academy.apnic.net** to these 2 nameservers.

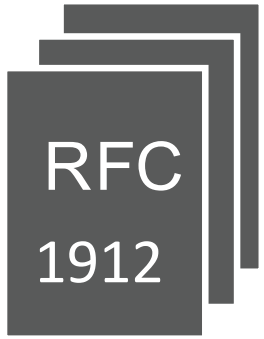


```
;From apnic.net zone, add these records:

academy.apnic.net.    NS    ns1.academy.apnic.net.
academy.apnic.net.    NS    ns2.academy.apnic.net.
```

A client must then go to ns1.academy.apnic.net (or ns2) to query for any of its subdomain.

Now how can we reach ns1 and ns2? We must add a **Glue Record.**

Glue Record

A **glue record** is a non-authoritative data. It is an A record that maps the address of the sub-domain's nameserver.

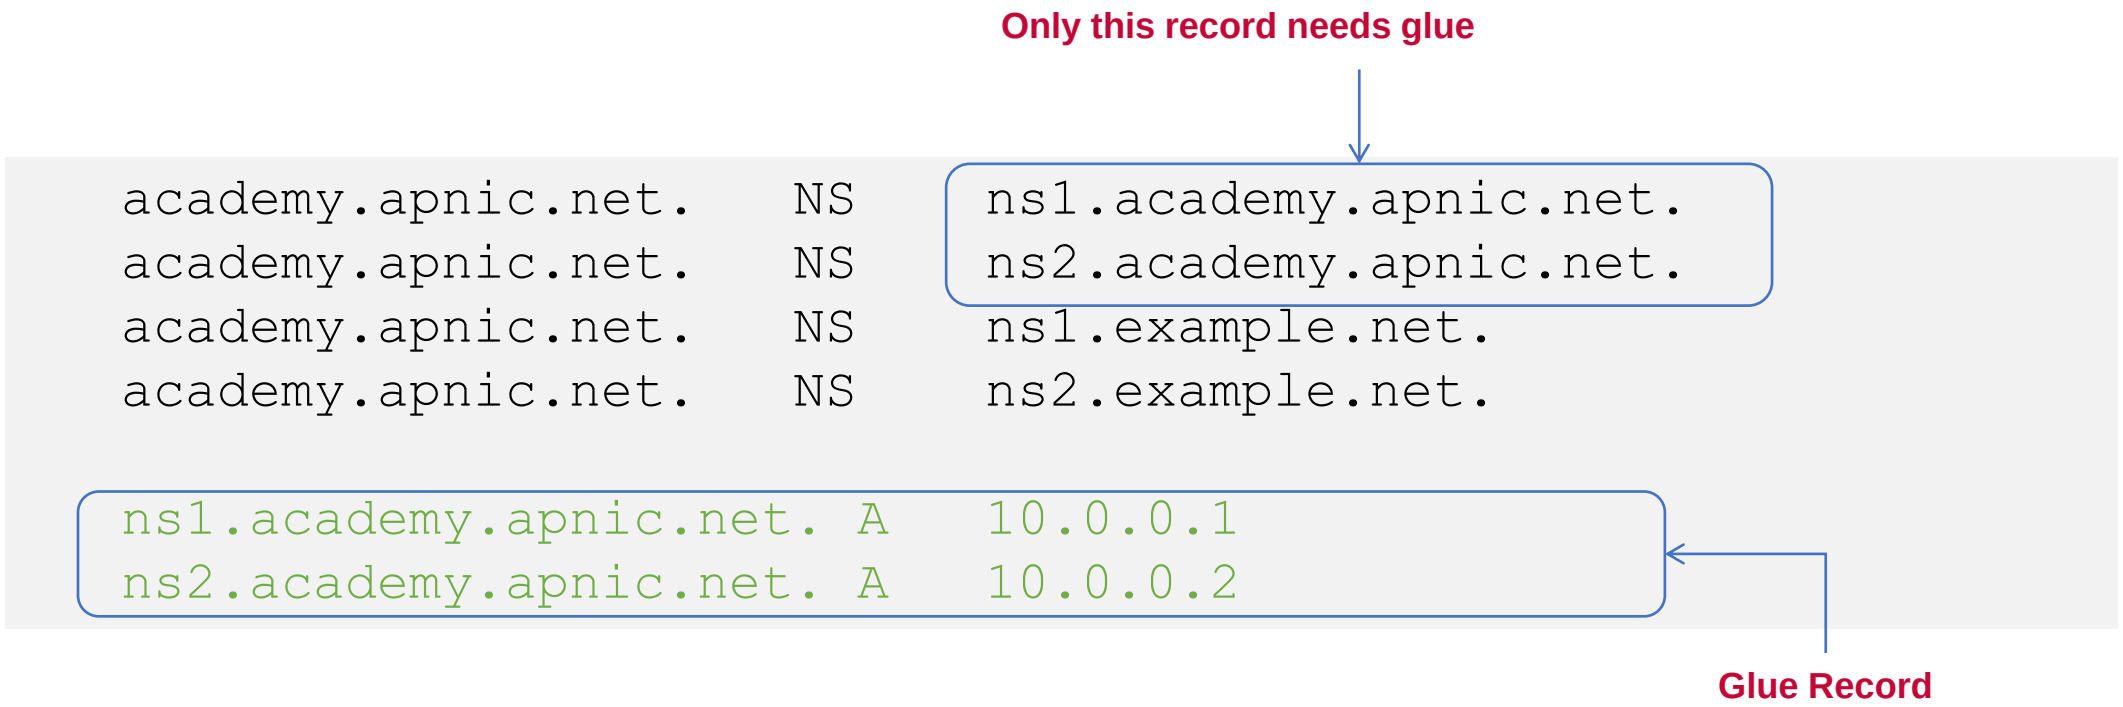

Delegation Example

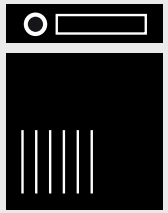

ns.apnic.net

1. Add NS records and glue
2. Make sure there is no other data from the academy.apnic.net. zone in the zone file

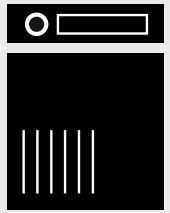

ns.academy.apnic.net

1. Setup minimum two servers
2. Create zone file with NS records
3. Add all academy.apnic.net data in its own zonefile.

Resolvers

A piece of software (usually in the operating system) which formats the DNS request into UDP packets

A stub resolver is a minimal resolver that forwards all requests to a local recursive nameserver

Every host needs a resolver

- In Linux, this is in /etc/resolv.conf
- Configure to use more than one DNS server

DNS Query – dig

What is the IP address of academy.apnic.net?

```
dig academy.apnic.net

; <<>> DiG 9.14.10 <<>> academy.apnic.net
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 60912
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;academy.apnic.net.                IN      A

;; ANSWER SECTION:
academy.apnic.net.        86400   IN      A       203.119.101.88

;; Query time: 17 msec
;; SERVER: 202.12.29.236#53(202.12.29.236)
;; WHEN: Wed Jan 20 10:58:42 AEST 2021
;; MSG SIZE  rcvd: 62
```

DNS Query – drill

```
drill academy.apnic.net

;; ->>HEADER<<- opcode: QUERY, rcode: NOERROR, id: 62275
;; flags: qr rd ra ; QUERY: 1, ANSWER: 1, AUTHORITY: 4, ADDITIONAL: 6
;; QUESTION SECTION:
;; academy.apnic.net.         IN        A

;; ANSWER SECTION:
academy.apnic.net.  86400      IN        A         203.119.101.88

;; AUTHORITY SECTION:
apnic.net.          3600       IN        NS        ns4.apnic.net.
apnic.net.          3600       IN        NS        netnod.apnic.net.
apnic.net.          3600       IN        NS        ns2.apnic.net.
apnic.net.          3600       IN        NS        apnic.authdns.ripe.net.

;; ADDITIONAL SECTION:
ns2.apnic.net.      2547       IN        A         203.119.95.53
ns4.apnic.net.      2547       IN        A         202.12.31.53
netnod.apnic.net.   2575       IN        A         194.146.106.106
ns2.apnic.net.      2547       IN        AAAA      2001:ddd::53
ns4.apnic.net.      2547       IN        AAAA      2001:dd8:12::53
netnod.apnic.net.   2575       IN        AAAA      2001:67c:1010:27::53

;; Query time: 107 msec
;; SERVER: 203.119.110.16
;; WHEN: Mon Jan 25 15:34:07 2021
;; MSG SIZE  rcvd: 273
```

DNS Cloud Providers

- There have been an increase of third-party cloud DNS providers over the years.

- Why we use them?
  - It's free and generally fast
  - Avoid surveillance and blocking
  - Lack of trust in the current provider
  - Focus on privacy

Public DNS Providers	
Google	8.8.8.8 8.8.4.4
Cloudflare DNS	1.1.1.1
Quad9	9.9.9.9
OpenDNS	208.67.222.222 208.67.220.220

Public DNS Resolvers

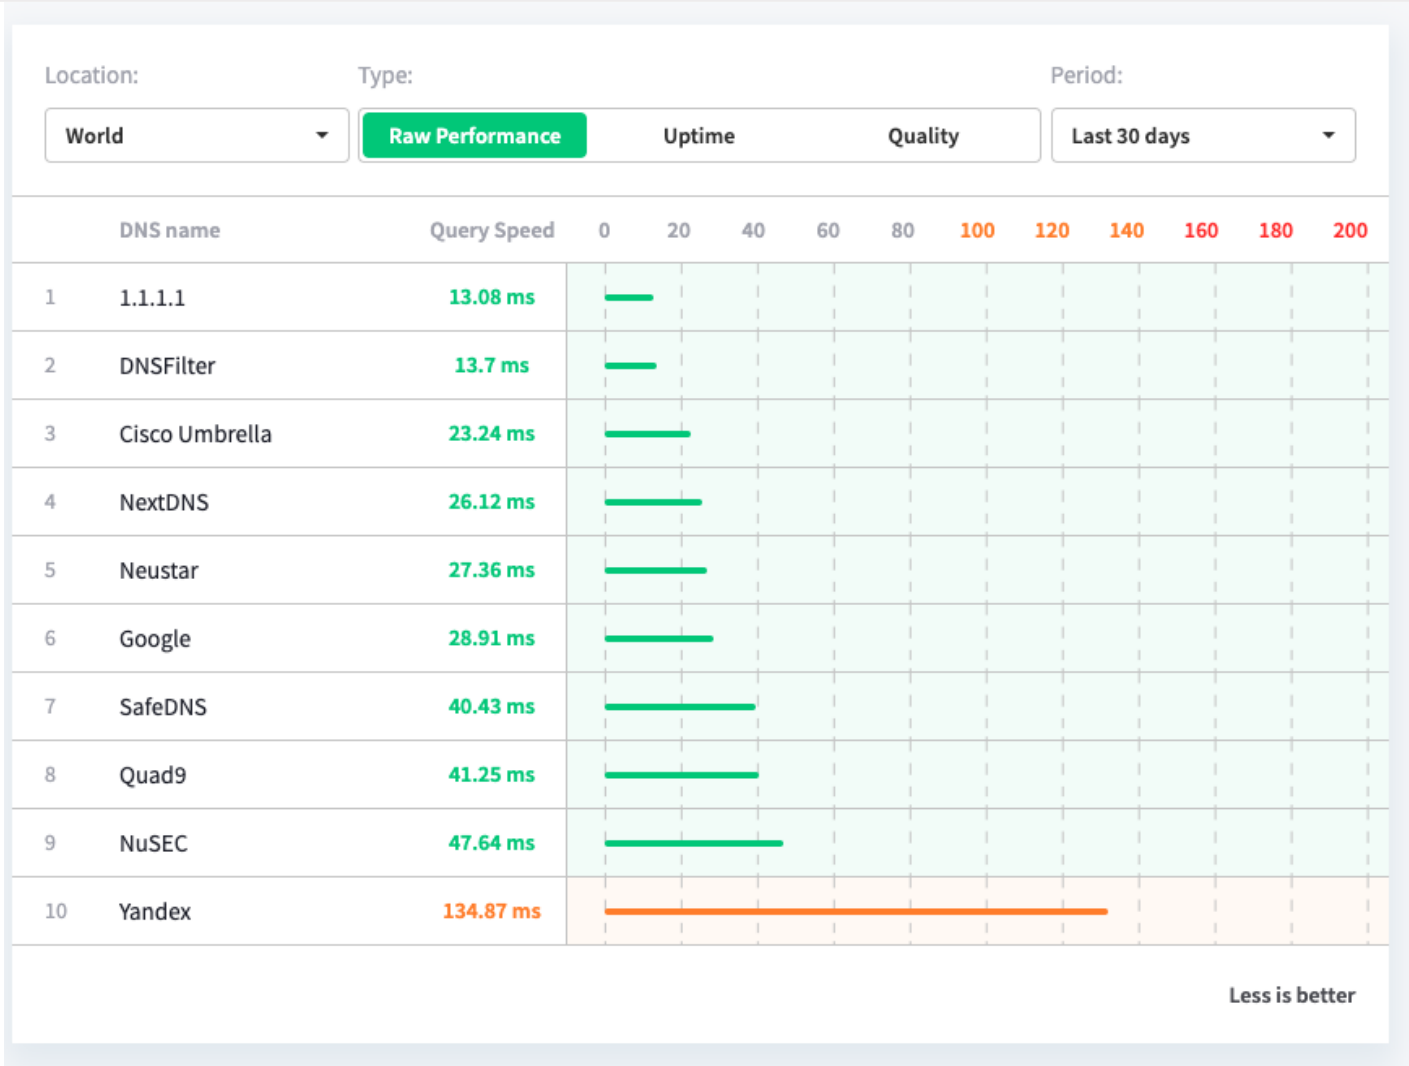

https://www.dnsperf.com/#!dns-resolvers

Cloudflare DNS (1.1.1.1)

- Part of a research agreement between APNIC and Cloudflare (link)

```
inetnum:        1.0.0.0 - 1.0.0.255
netname:        APNIC-LABS
descr:          APNIC and Cloudflare DNS Resolver project
descr:          Routed globally by AS13335/Cloudflare
descr:          Research prefix for APNIC Labs
country:        AU
org:            ORG-ARAD1-AP
admin-c:        AR302-AP
tech-c:         AR302-AP
abuse-c:        AA1412-AP
status:         ASSIGNED PORTABLE
remarks:        ---------------
remarks:        All Cloudflare abuse reporting can be done via
remarks:        resolver-abuse@cloudflare.com
remarks:        ---------------
mnt-by:         APNIC-HM
mnt-routes:     MAINT-AU-APNIC-GM85-AP
mnt-irt:        IRT-APNICRANDNET-AU
last-modified:  2020-07-15T13:10:57Z
source:         APNIC
```

```
route:          1.0.0.0/24
origin:         AS13335
descr:          APNIC Research and Development
                6 Cordelia St
mnt-by:         MAINT-AU-APNIC-GM85-AP
last-modified:  2018-03-16T16:58:27Z
source:         APNIC
```

Remember …

Deploy multiple authoritative servers to distribute load and risk

Use cache to reduce load to authoritative servers and response times

SOA timers and TTL need to be tuned to the needs of the zone

Thank You!

- Any questions?