DNS Fundamentals

ONLINE TUTORIAL

28 JUL 2021

Overview

- DNS Technical Overview
- DNS Operations
- Recursive DNS
  - Lab: Recursive DNS
- Authoritative DNS
  - Lab: Primary and Secondary DNS
- DNS Troubleshooting
- Reverse DNS
  - Lab: Reverse DNS
- Secure Zone Transfer
  - Lab: RNDC
- DNS Anycast or DNS Privacy

Housekeeping

- Please mute your microphone when the Instructor is presenting to avoid disruptions to the class;
- For Q&As please use the Shared document;
- After every 30 minutes we will discuss the Q&A;
- The chat section will be used to share information, URLs, etc;
- If you raise your hand be sure to lower it again;
- If something goes wrong, re-join the meeting;
- If you have any trouble, please email training@apnic.net

Resources

- Course Materials
  - https://wiki.apnictraining.net/dns-20210728-online/

Ask Questions

- Shared Doc

- Email the mailing list
  - CommunityTrainers@apnic.net

DNS Overview

Module 1

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

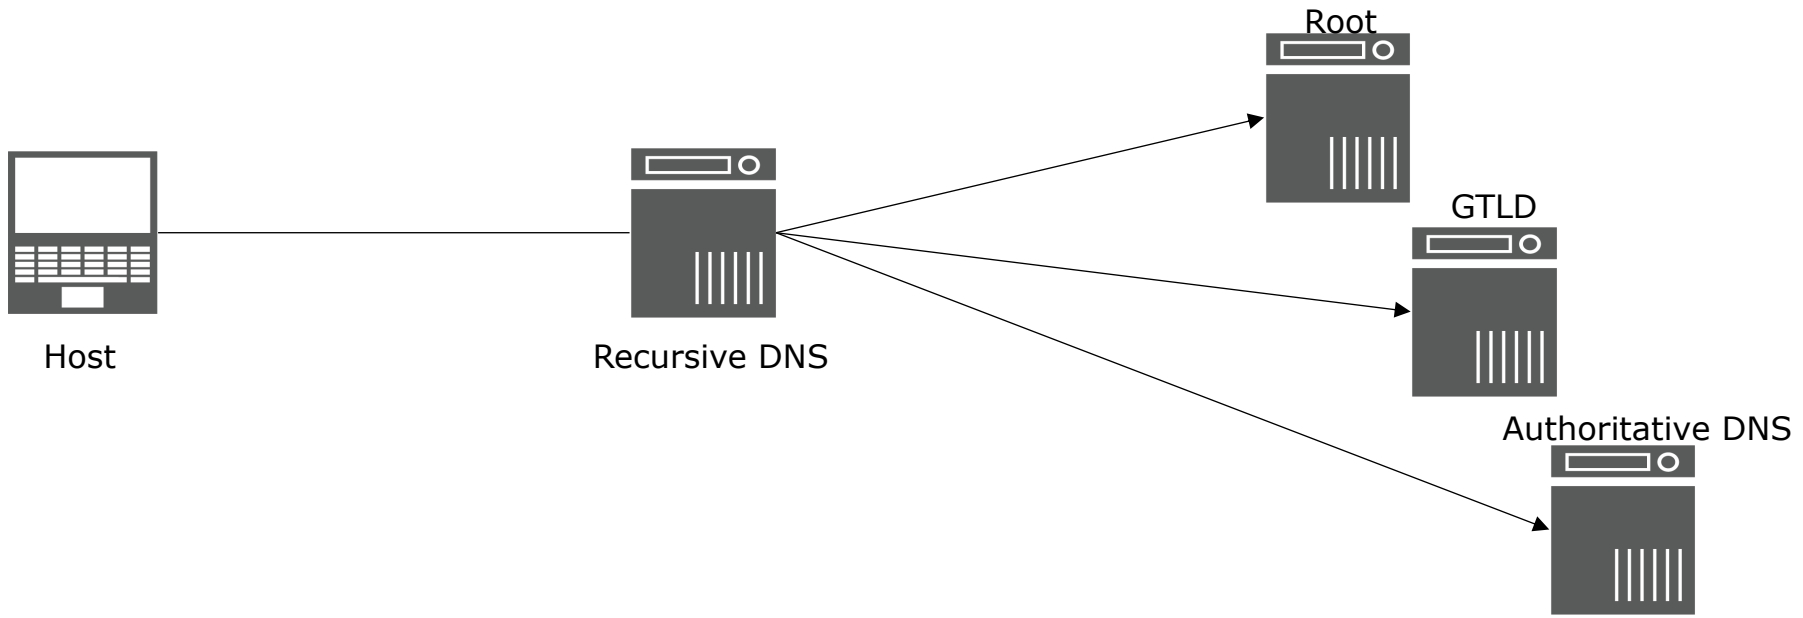

DNS is a critical piece of the Internet infrastructure

IP Addresses vs Domain Names

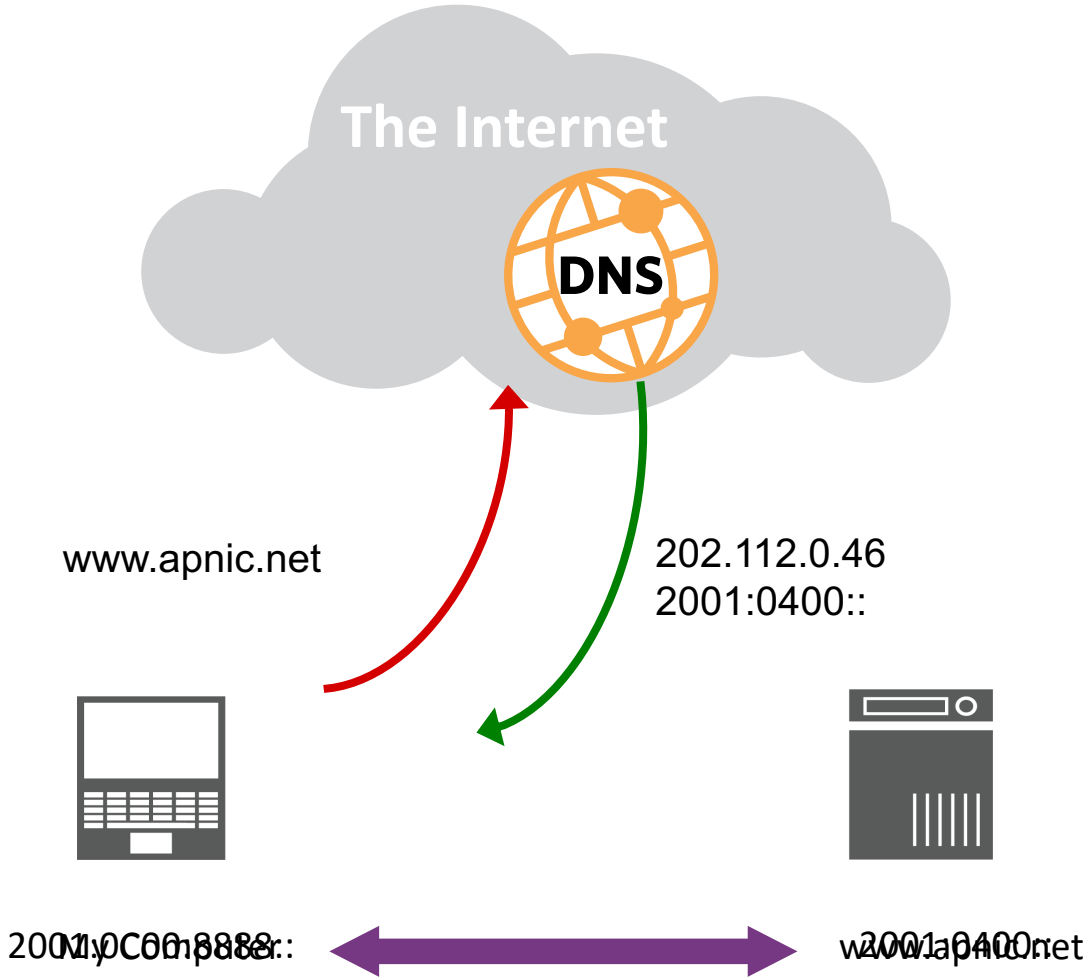

DNS Features

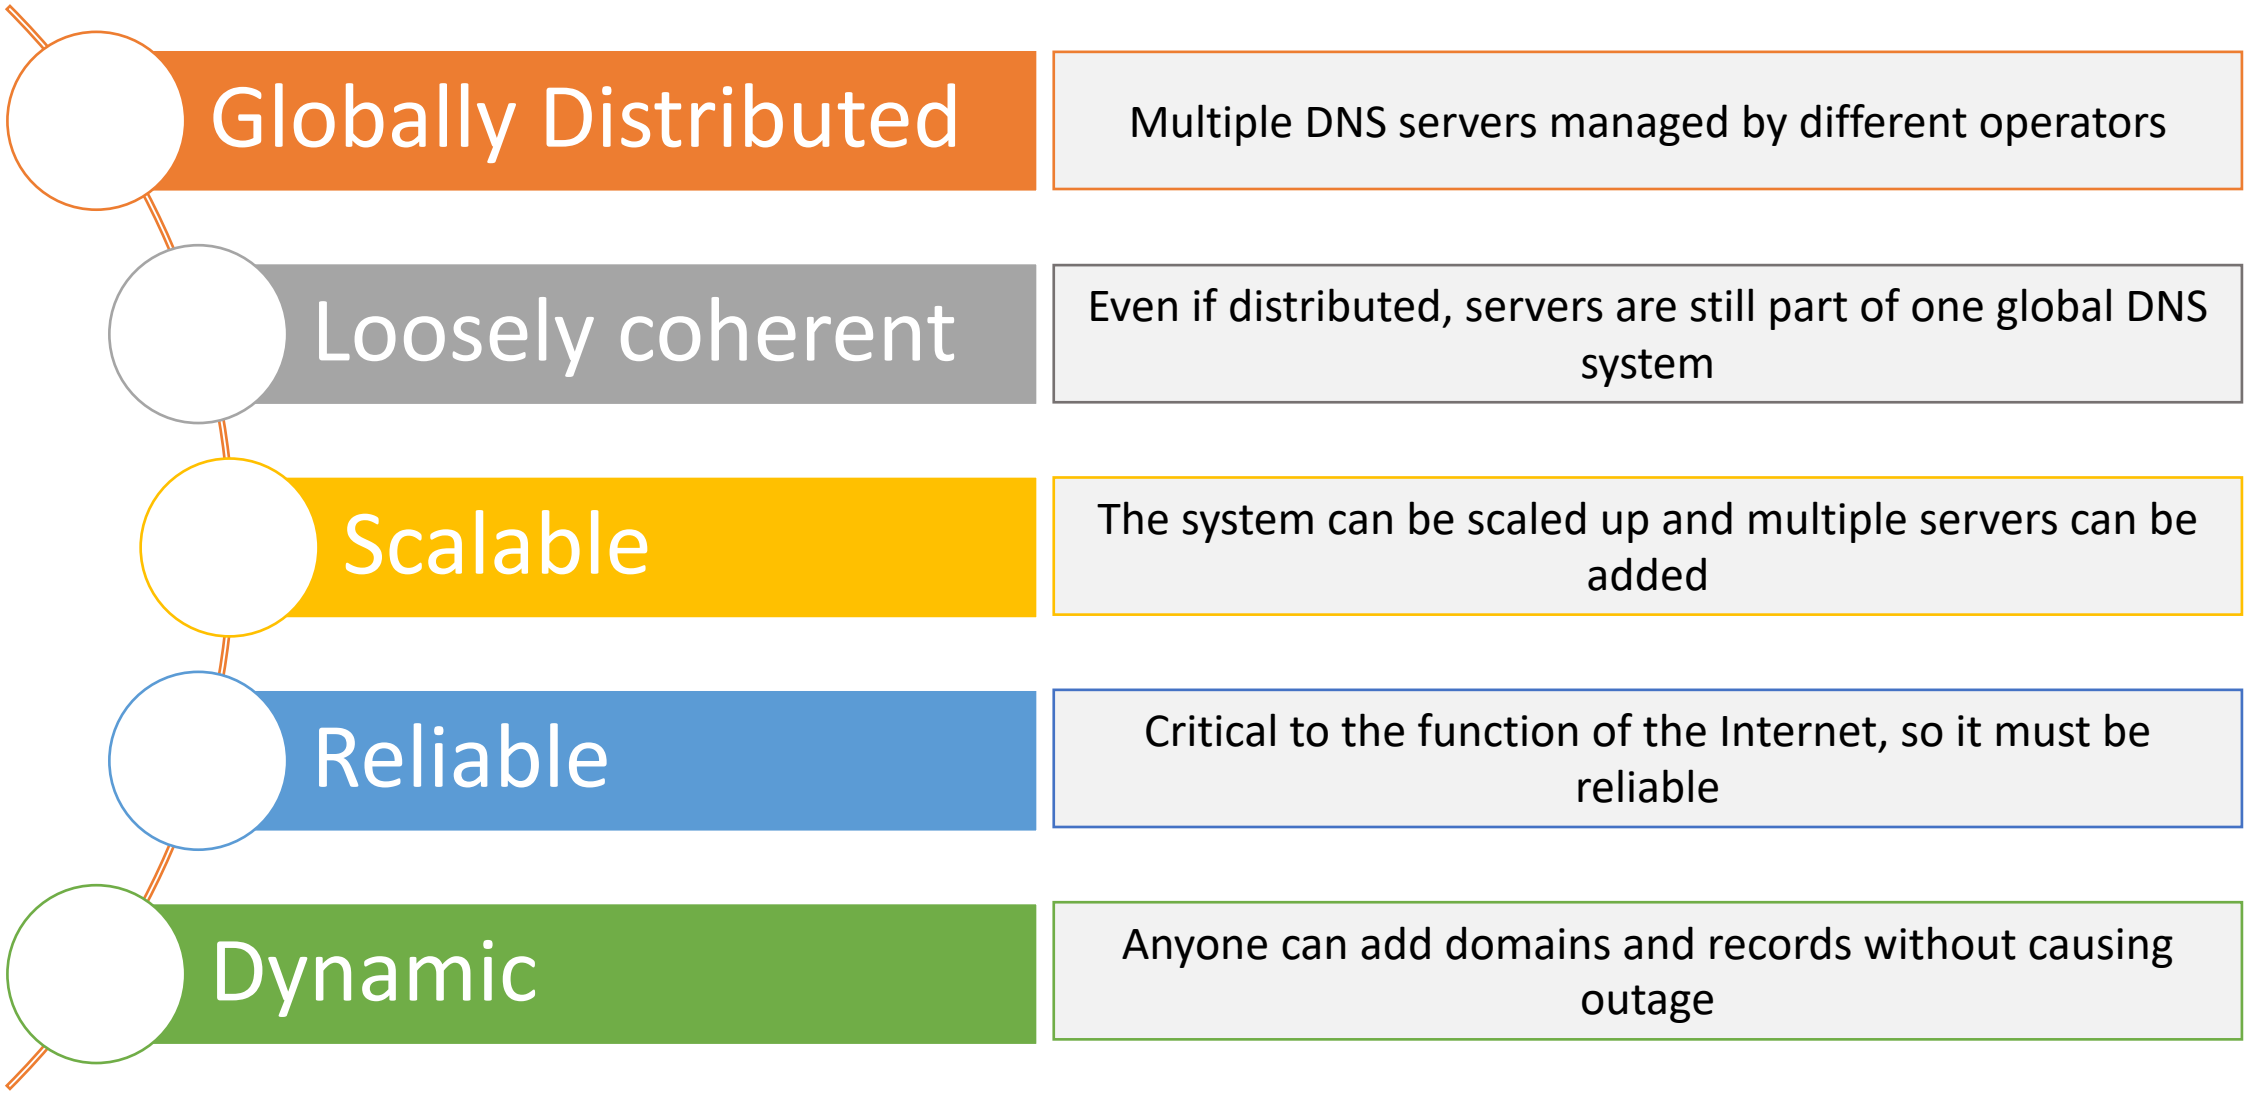

More about DNS

- 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

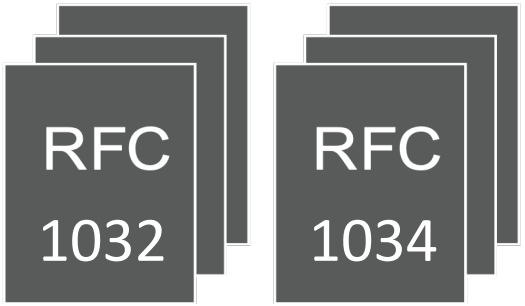

What is DNS?

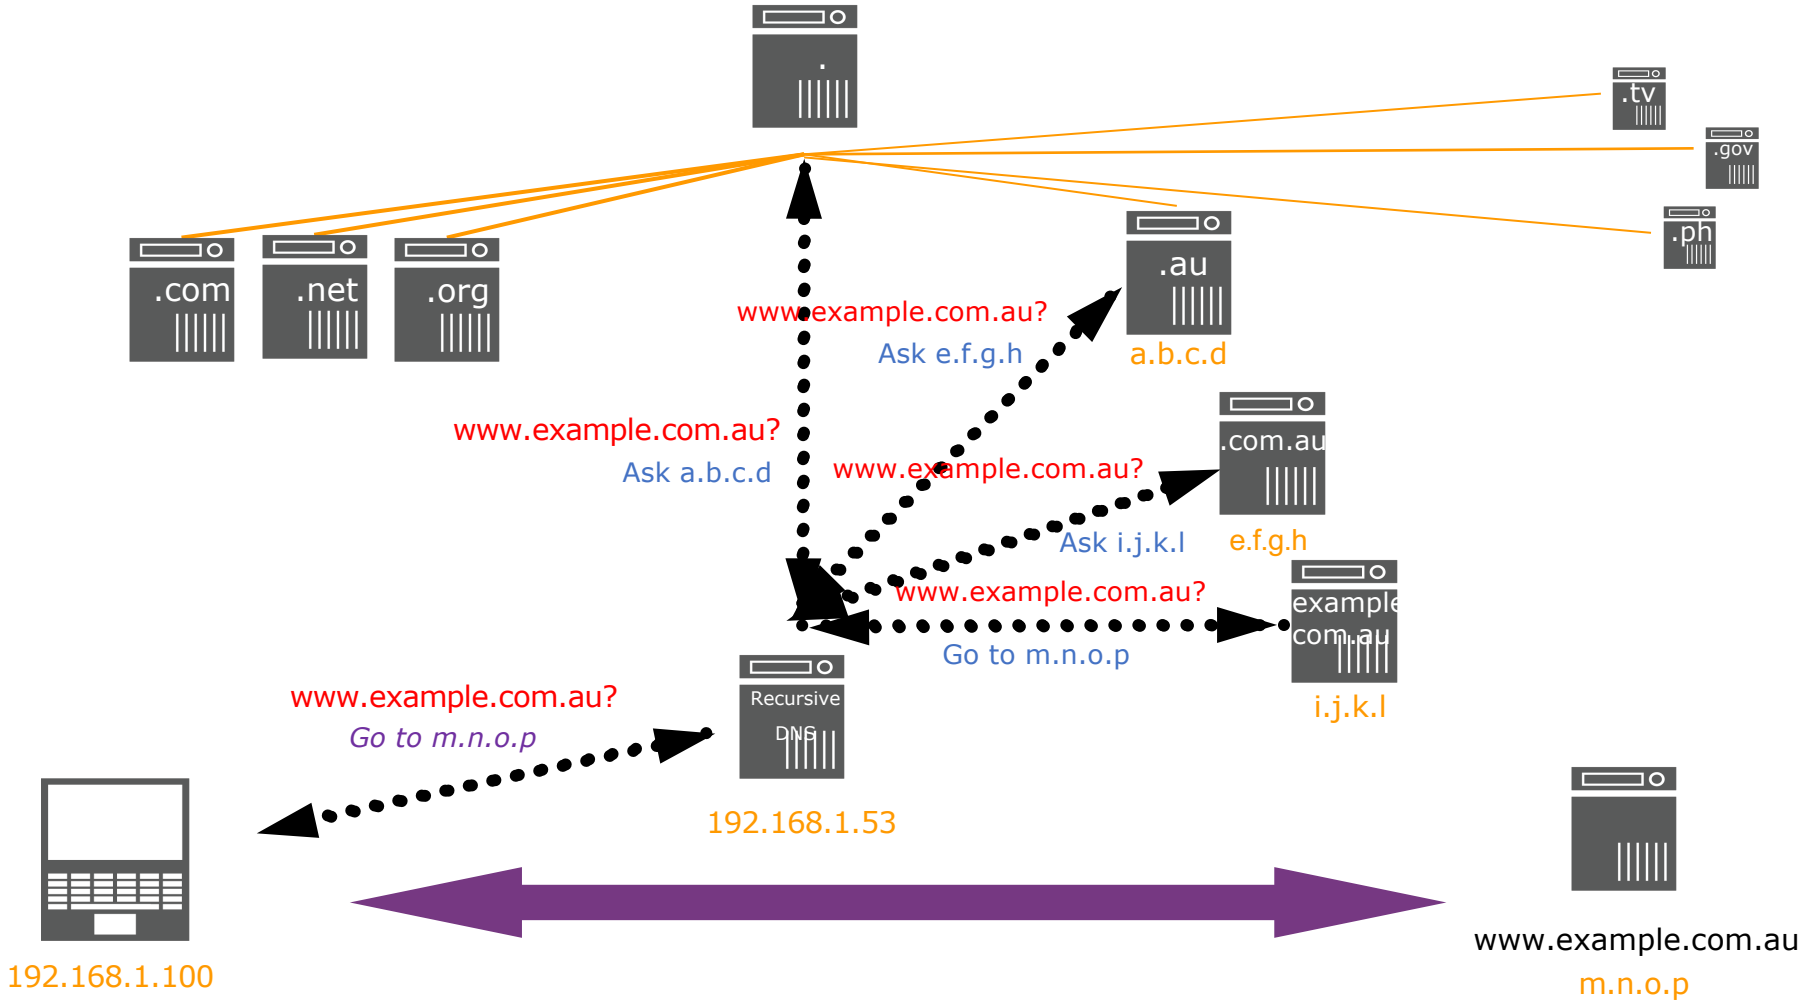
.
.com
.net
.org
.tv
.gov
.ph
.au
a.b.c.d
www.example.com.au?
Ask e.f.g.h
www.example.com.au?
Ask a.b.c.d
.com.au
e.f.g.h
www.example.com.au?
Ask i.j.k.l
www.example.com.au?
Go to m.n.o.p
example
com.au
i.j.k.l
www.example.com.au?
Go to m.n.o.p
Recursive
DNS
192.168.1.53
192.168.1.100
www.example.com.au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

DNS Hierarchy Tree

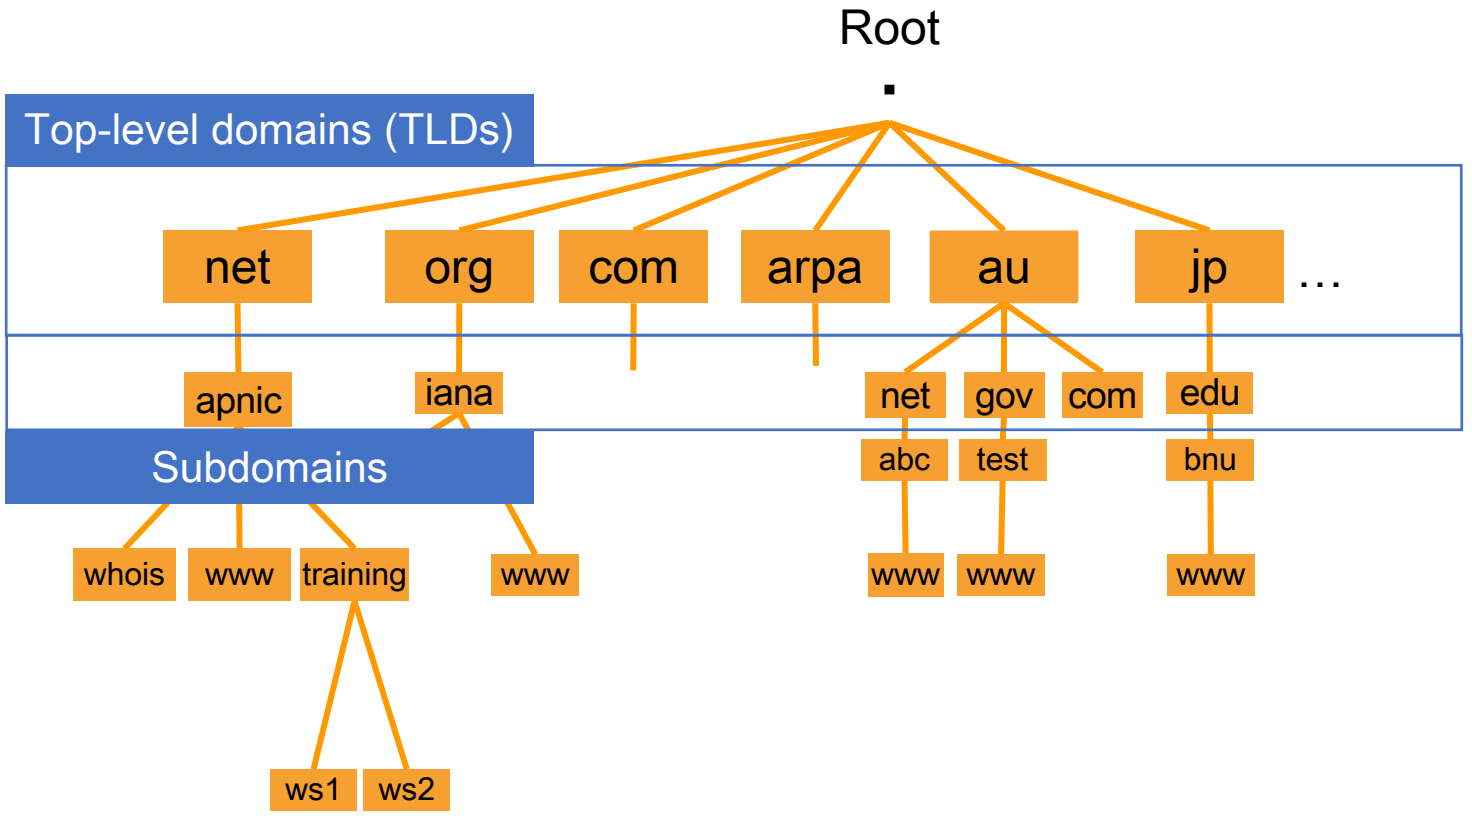

Country-code TLDs (ccTLDs)

Generic TLDs (gTLDs)

Infrastructure TLD

Internationalized TLDs (IDN)

FQDN: ws1.training.apnic.net.

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

DNS Components

Namespace	Domains and zones
Nameserver	Makes the namespace available
Resolvers or clients	Query the nameserver for records in the namespace

Domains

Domains are “namespaces”

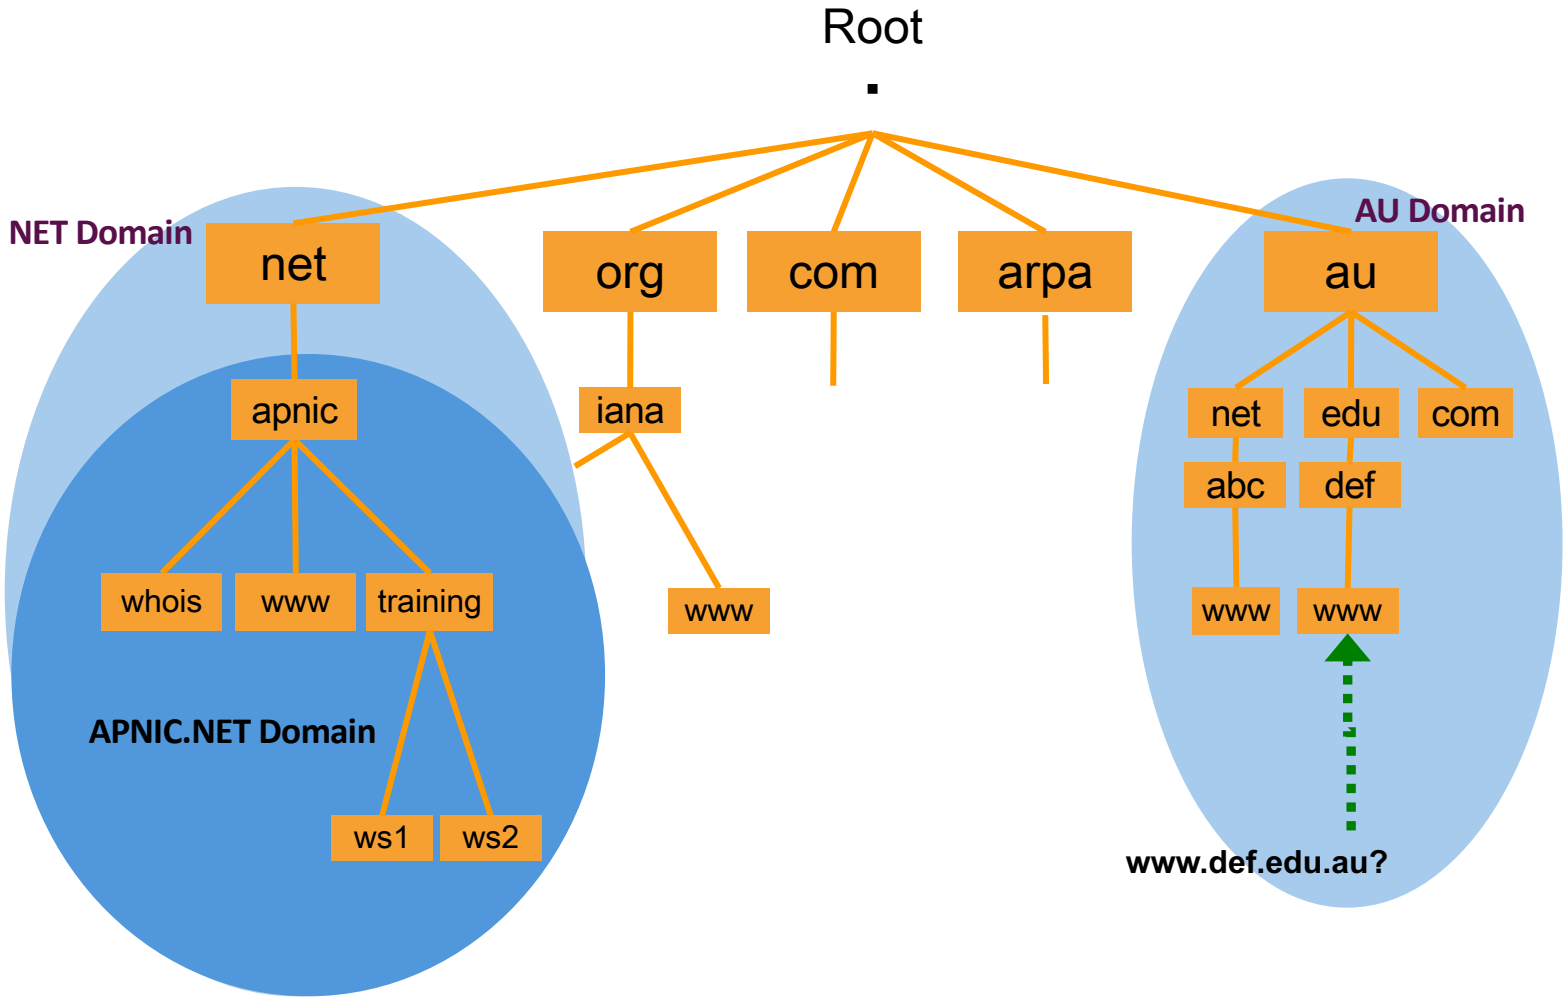

Delegation

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

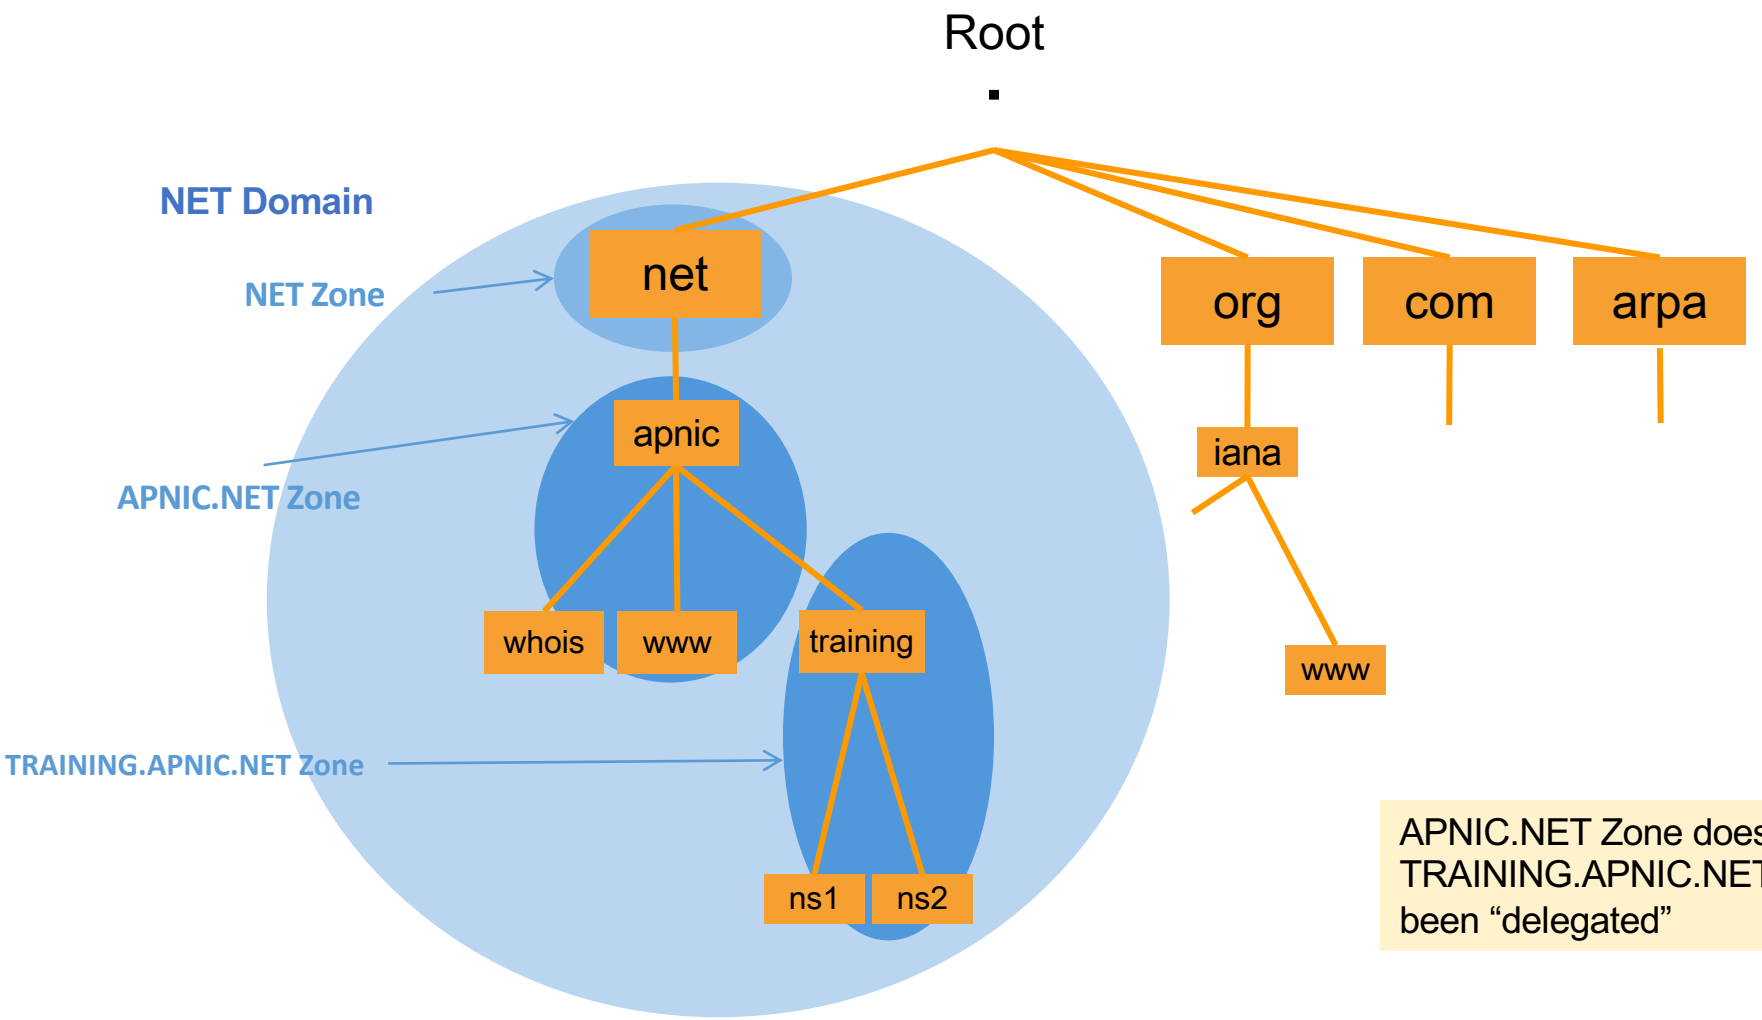

APNIC.NET Zone doesn't include TRAINING.APNIC.NET since it has been "delegated"

Name Servers

- Name servers answer DNS questions

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

- Mixture of functions

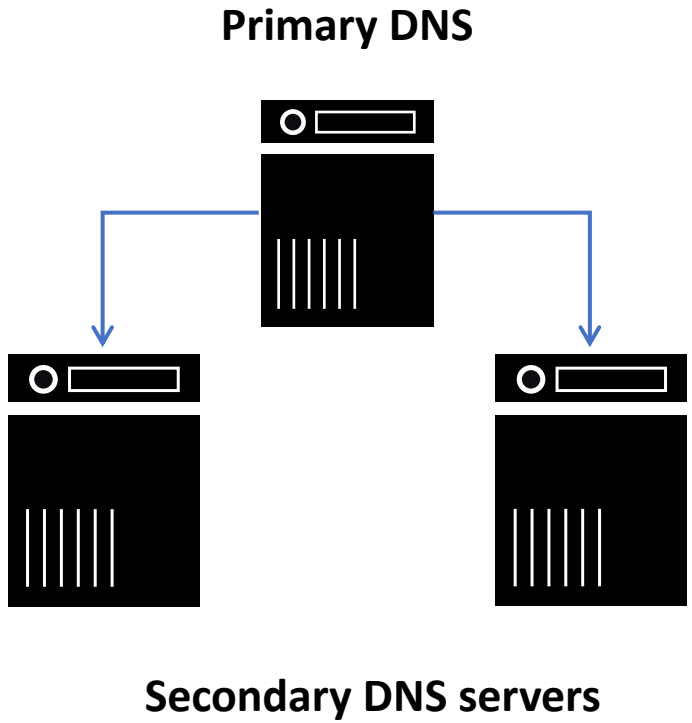

Authoritative Nameserver

- 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

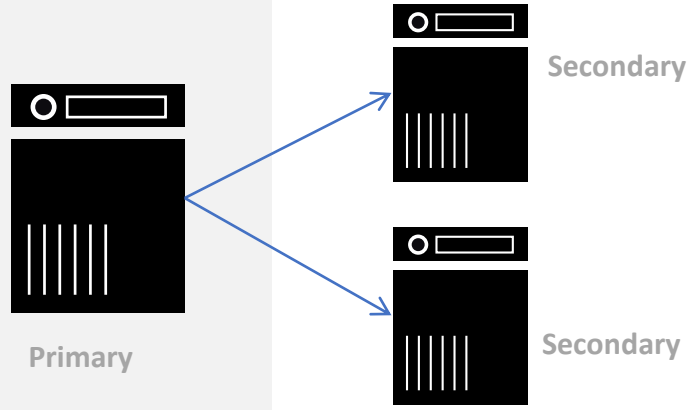

Recursive Nameserver

- 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

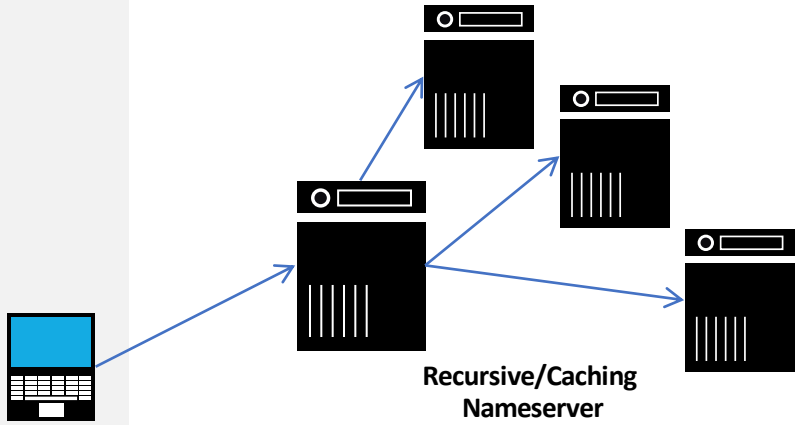

Root Servers

- 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

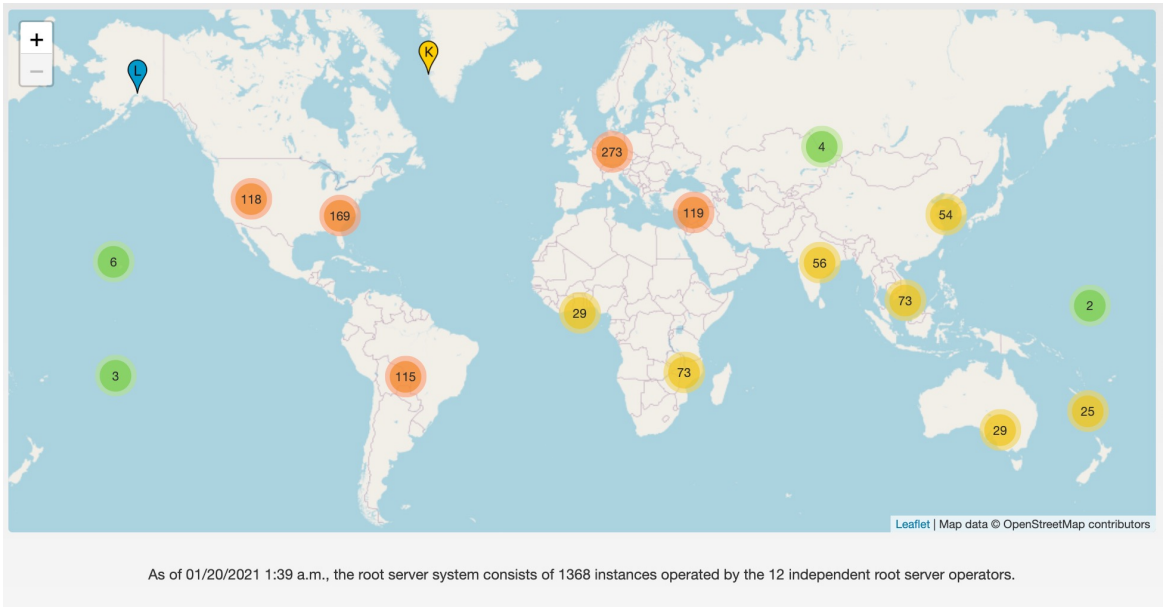

As of 01/20/2021 1:39 a.m., the root server system consists of 1368 instances operated by the 12 independent root server operators.

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

```
training.apnic.net.   86400      IN      A      192.168.1.100
```

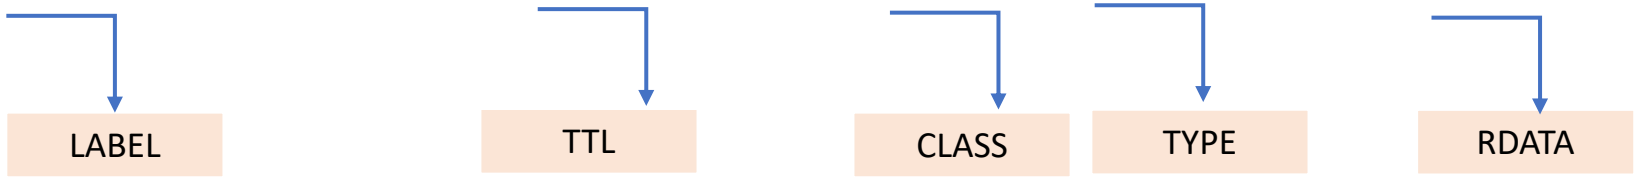

Common Resource Record Types

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

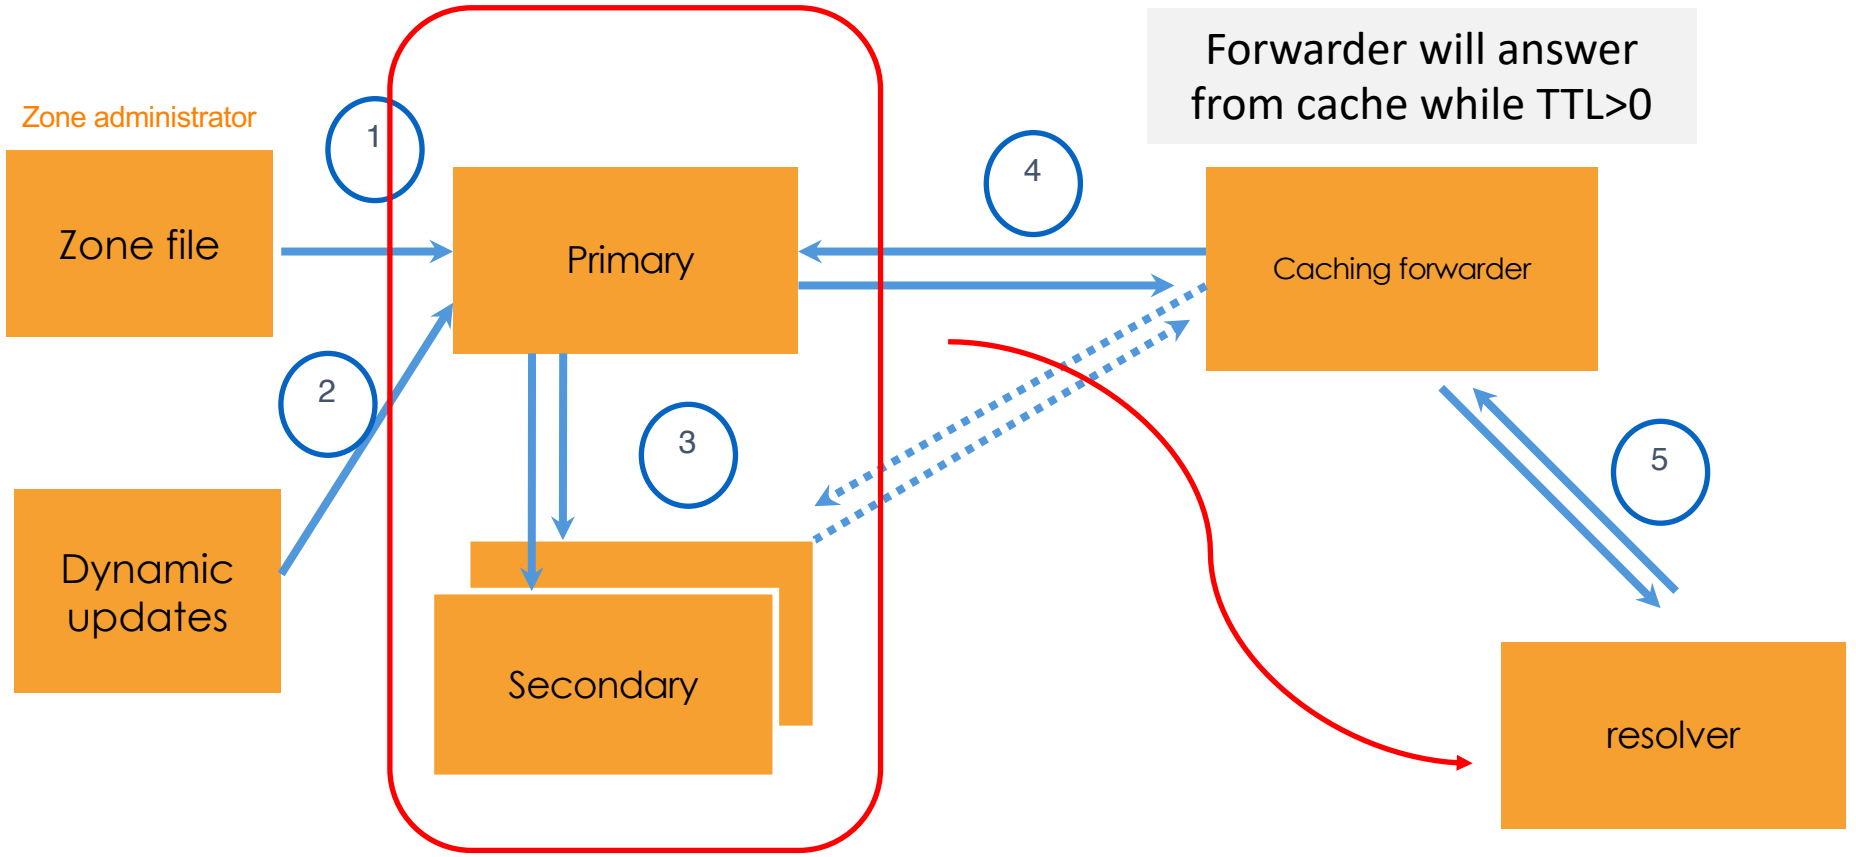

Delegating a Zone

Delegation is done by adding NS records.

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

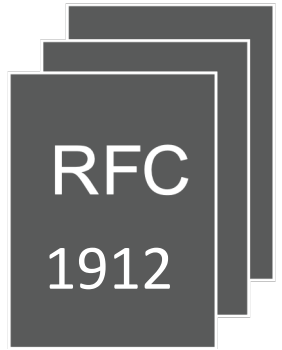


```
;From apnic.net zone, add these records:

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.

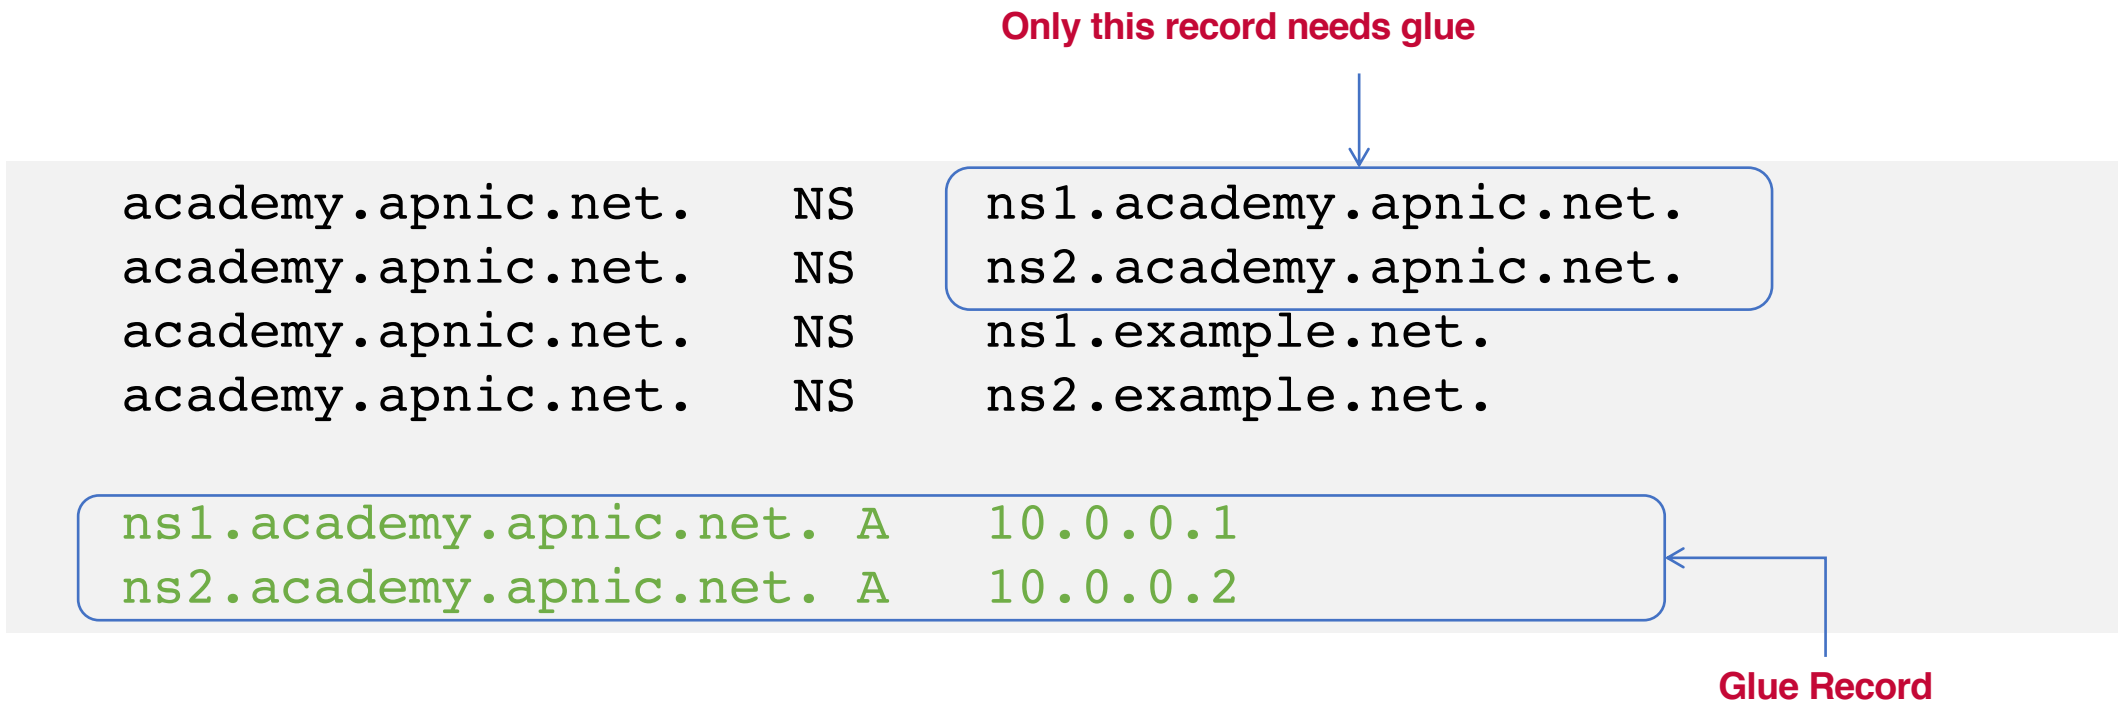

Delegation Example

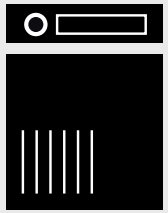

ns.apnic.net

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

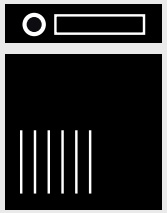

ns.academy.apnic.net

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
```

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

Performance Considerations

- How many DNS servers?
- How many zones are expected to load?
- How large are the zones?
- How often are zone transfers?
- Where are the DNS servers located?
- What is the expected bandwidth?
- Are these servers multihomed?
- How many interfaces are to be enabled for listening?
- How many queries are expected to receive?
- Is the server caching?
- Is the server doing recursion?
- Is dynamic updates allowed?

Questions

Thank You!