

# The Domain Name System

tsaimh (2024-2026, CC-BY)

lwhsu (2020-2023, CC-BY)

? (?-2019)

國立陽明交通大學資工系資訊中心

Information Technology Center of Department of Computer Science, NYCU

# History of DNS

- What and Why is DNS?
  - IP is difficult to memorize, and IPv6 makes it worse
  - Domain Name ↔ IP Address(es)
- Before DNS
  - ARPANET
    - HOSTS.txt contains all the hosts' information (/etc/hosts)
    - Maintained by SRI's Network Information Center
      - Register → Distribute DB
  - Problems: Not scalable!
    - Traffic and Load
    - Name Collision
    - Consistency
- Domain Name System
  - **Administration decentralization**
  - Paul Mockapetris (University of Southern California)
    - RFC 882, 883 (1983) → 1034, 1035 (1987)




Paul Mockapetris (1948- )  
Inventor of DNS  
Internet Hall of Fame 2012  
ACM Fellow

# DNS Specification

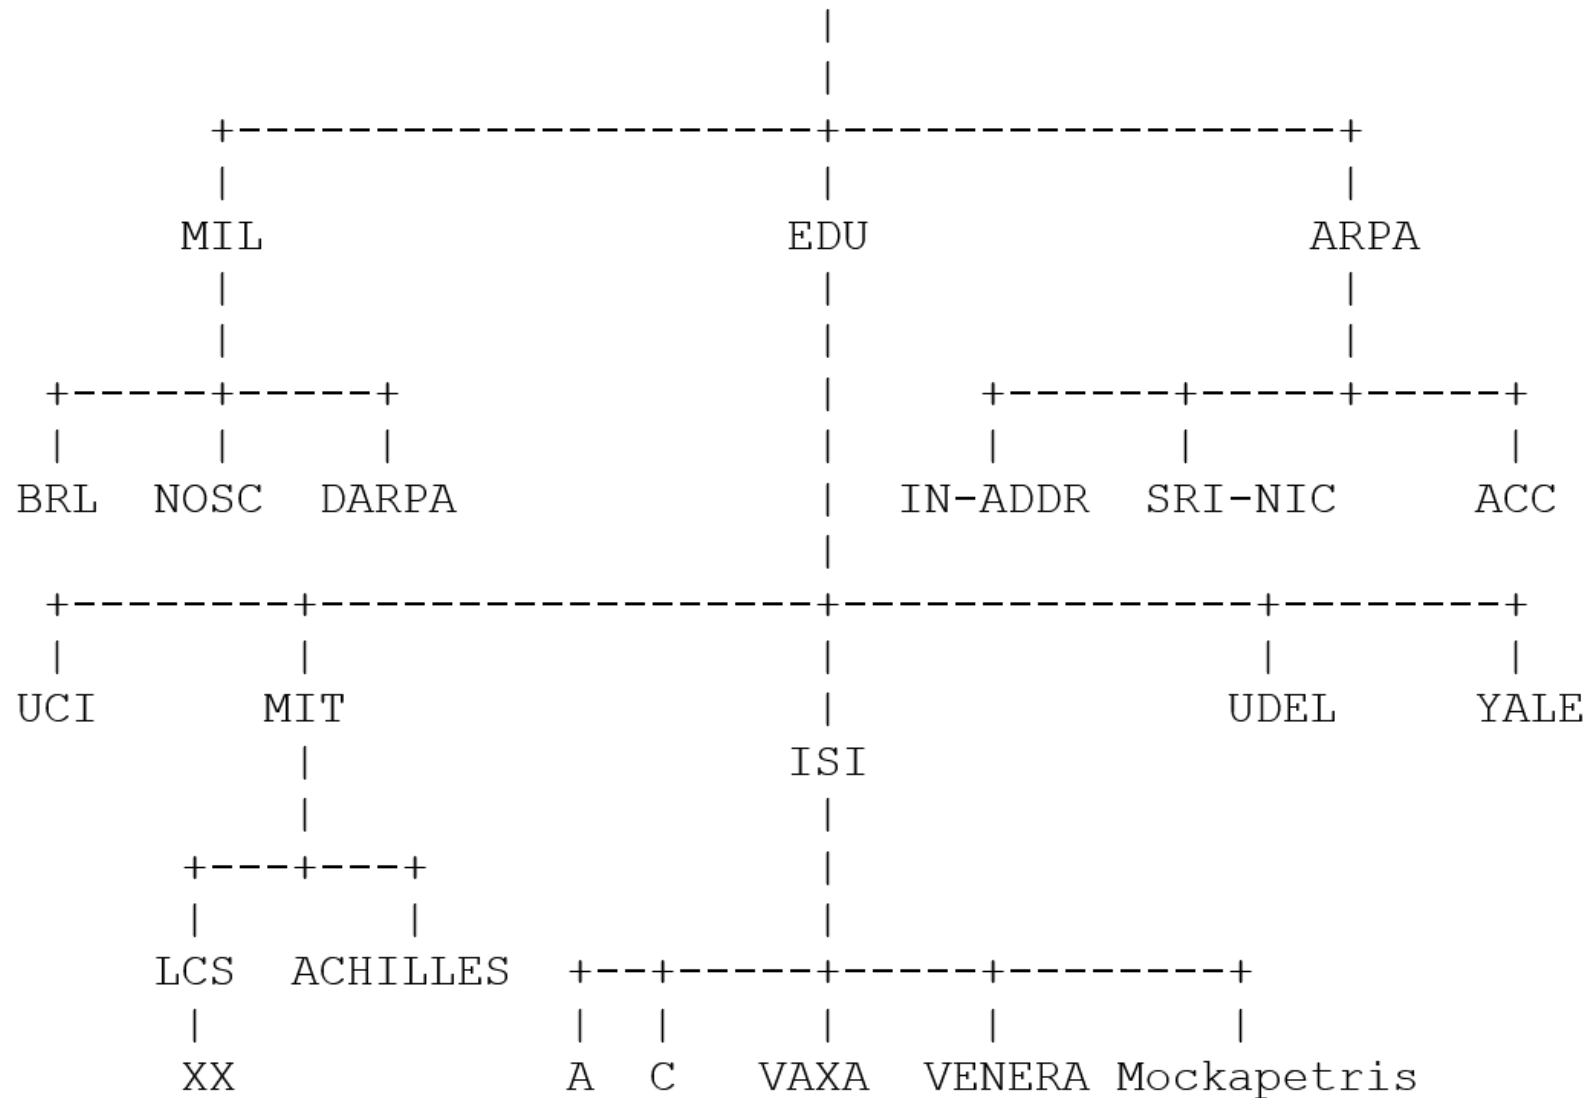
- **Tree architecture** – "domain" and "subdomain"
  - Divided into categories
  - Solves name collision
- **Distributed database**
  - Each site maintains a segment of the DB
  - Each site opens its information via network
- **Client-Server architecture**
  - Name servers provide information (Name Server)
  - Clients make queries to server (Resolver)

# The DNS Namespace – (1)

- Domain name is
  - An inverted tree (Rooted tree)
    - Root with label '.'
    - Root with label '' (Null)
- Domain and subdomain
  - Each domain has a "domain name" to identify its position in database
    - domain: nycu.edu.tw
    - subdomain: cs.nycu.edu.tw



# The DNS Namespace – (2)




# The DNS Namespace – (3)

- Domain name limitations (RFC1035: 2.3.4 “Size limits”)
  - Up to 63-octets in each label
  - Up to 255-octets in a full domain name
    - Up to 253 visible characters
  - What is the real maximum length of a DNS name?
    - <https://devblogs.microsoft.com/oldnewthing/20120412-00/?p=7873>

(63 letters).(63 letters).(63 letters).(62 letters) → Size limit exceeded!!

For example, `www.microsoft.com` is encoded as follows:



# Root Zone and Special Top Level Domains (TLDs) Maintained by IANA

- DNS **Root Zone** is the **upper-most part** of the DNS hierarchy, and involves **delegating** administrative responsibility of “**top-level domains**”.
- The **.int top-level domain**, designed for the sole use **of cross-national organizations**, such as treaty organizations (e.g., **NATO** and **WHO**)
- The **.arpa** domain (Address and Routing Parameter Area) is **used internally by Internet protocols**, such as for **reverse mapping of IP addresses**, and delivery of **ENUM phone number mapping**.

Source: <https://www.iana.org/domains/root/db>

IANA: Internet Assigned Numbers Authority

# Top Level Domains (TLDs)

- As of 2015, IANA distinguishes the following groups of top-level domains:
  - **Infrastructure** top-level domain (**ARPA**): It is managed by IANA on behalf of the IETF for various purposes specified in the RFCs.
  - **Generic** top-level domains (**gTLD**): Top-level domains with three or more characters
  - **Generic restricted** top-level domains (**grTLD**): These domains are managed **under official ICANN accredited registrars**. (**.biz .name .pro**)
  - **Sponsored** top-level domains (**sTLD**): These domains are proposed and sponsored by private agencies or organizations, and are managed **under official ICANN accredited registrars**.

# Top Level Domains (TLDs) (cont.)

- **country-code** top-level domains (**ccTLD**): Two-letter domains established for countries or territories. With some historical exceptions, the code for any territory is the same as its **two-letter ISO 3166** code.
- **Test** top-level domains (**tTLD**): These domains were installed under **.test** for **testing purposes in the IDN (Internationalization Domain Name) development process**; these domains are **not present in the root zone**.
- The following TLDs are reserved by RFCs:
  - **.example .invalid .localhost .test** (RFC 6761)
  - **.local** (RFC 6762)
  - **.onion** (RFC 7686)

# Generic TLD (gTLD)

- RFC 920 (1984) defines the first six generic TLDs
    - com: commercial organization, such as ibm.com
    - edu: educational organization, such as purdue.edu
    - gov: government organization, such as nasa.gov
    - mil: military organization, such as navy.mil
    - net: network infrastructure providing organization, such as hinet.net
    - org: noncommercial organization, such as x.org
  - In 1988, NATO requests for the TLD .int
    - int: International organization, such as nato.int
- Now Sponsored TLD
- 
- ```
graph LR; edu[edu] --> NowSponsored[Now Sponsored TLD]; gov[gov] --> NowSponsored; int[int] --> NowSponsored;
```

# Generic TLD (gTLD) (cont.)

- New gTLDs launched in year 2000:

- aero: for air-transport industry
- biz: for business
- coop: for cooperatives
- info: for all uses
- museum: for museum
- name: for individuals
- pro: for professionals

Now Sponsored TLD

Now generic restricted TLD

- <https://www.iana.org/domains/root/db>

# Sponsored TLD (sTLD)

- A **sponsored TLD** is a specialized top-level domain that **has a sponsor** representing a **specific community** served by the domain.

| TLD                     | Eligibility                                               | Sponsors                                                                                                     |
|-------------------------|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <a href="#">.jobs</a>   | Human resource managers                                   | <a href="#">The Society for Human Resource Management</a> (nonprofit organization)                           |
| <a href="#">.post</a>   | Postal services                                           | <a href="#">Universal Postal Union</a>                                                                       |
| <a href="#">.tel</a>    | For businesses and individuals to publish contact data    | Since 2008: <a href="#">Telnic Limited</a><br>Since 2017: <a href="#">Telnames Limited</a> (private company) |
| <a href="#">.travel</a> | Travel agents, airlines, hoteliers, tourism bureaus, etc. | Since 2020: Registry is <a href="#">Donuts Inc.</a> <sup>[10]</sup> (private company)                        |
| <a href="#">.xxx</a>    | Pornographic sites                                        | <a href="#">ICM Registry</a>                                                                                 |

# Country-code TLD (ccTLD)

- Country code extension applications **began in 1985**. The registered country code extensions in that year included **.us (United States)**, **.uk (United Kingdom)** and **.il (Israel)**.
- Creation and delegation of ccTLDs is described in **RFC 1591**, corresponding to **ISO 3166-1 alpha-2** country codes (i.e., all identifiers are two letters long).

| English short name                                         | French short name          | Alpha-2 code | Alpha-3 code | Numeric |
|------------------------------------------------------------|----------------------------|--------------|--------------|---------|
| China                                                      | Chine (la)                 | CN           | CHN          | 156     |
| Taiwan (Province of China) <small>QQ</small>               | Taiwan (Province de Chine) | TW           | TWN          | 158     |
| United Kingdom of Great Britain and Northern Ireland (the) |                            | GB           | GBR          | 826     |

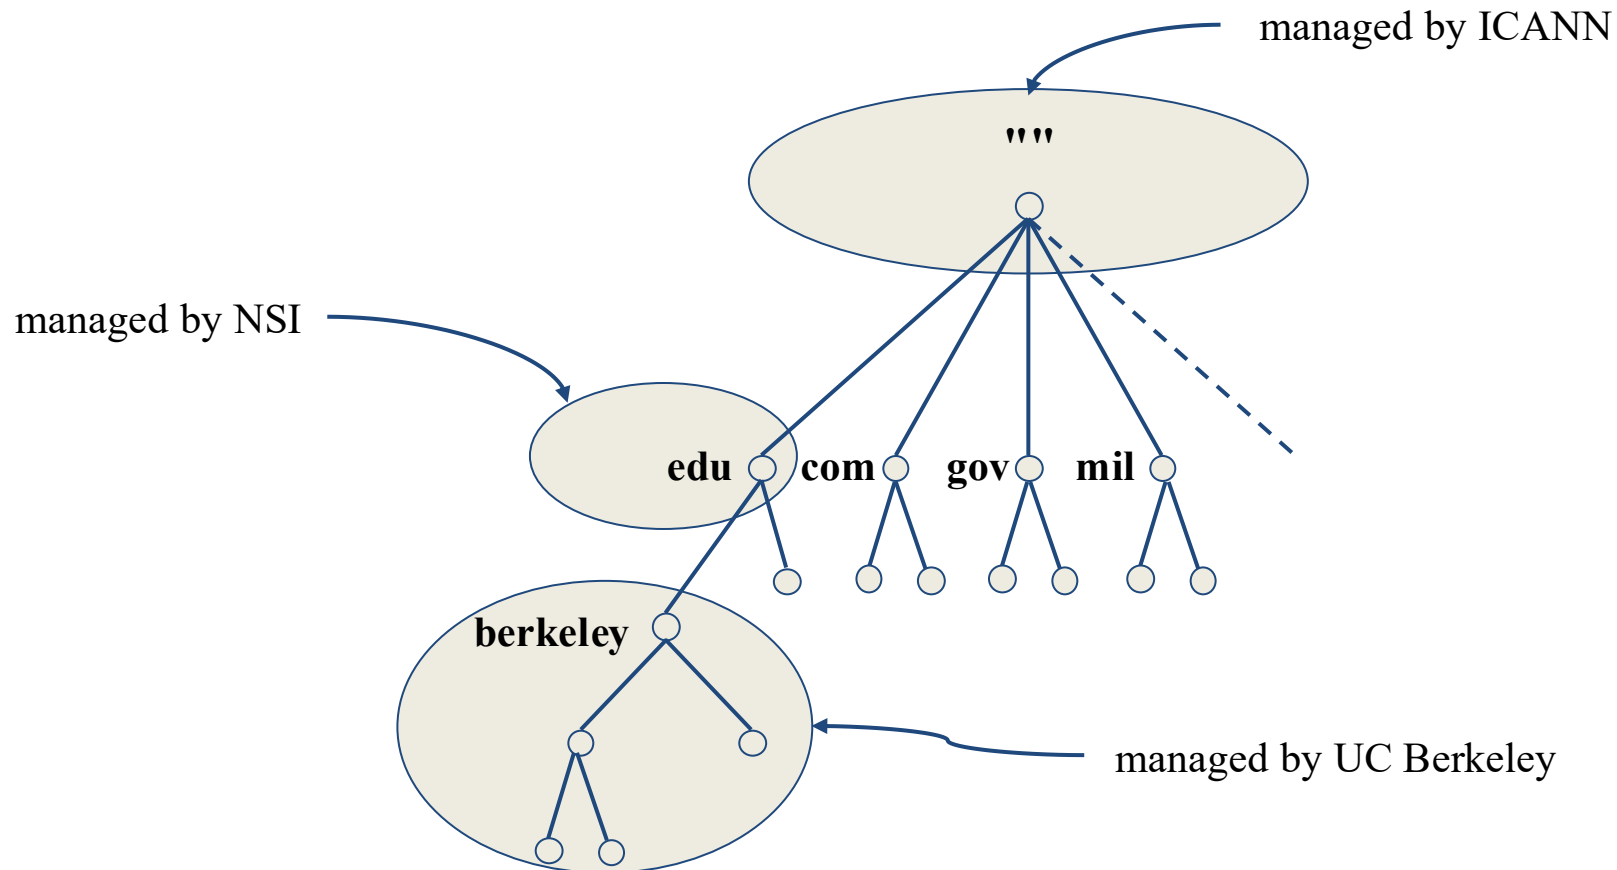
Source: <https://www.iso.org/obp/ui/#search>

# Country-code TLD (ccTLD)

- [ISO 3166](#), but just based on
  - Taiwan => tw (registered by MoE in July, 1989)
  - United Kingdom => uk (ISO3166 is GB)
  - European Union => eu
- Follow or not follow US-like scheme
  - US-like scheme example
    - edu.tw, com.tw, gov.tw
  - Other scheme
    - ac.jp, co.jp

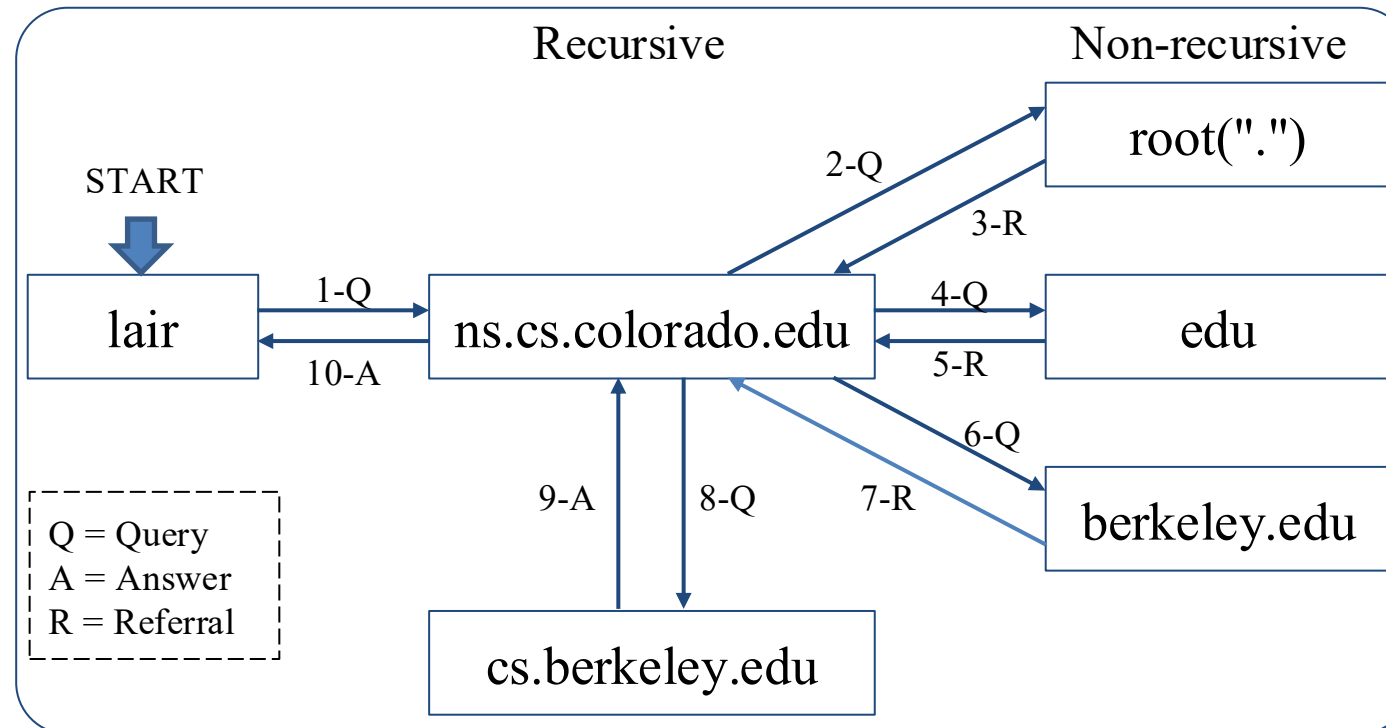
# How DNS Works – DNS Delegation

- Administration delegation
  - Each domain can delegate responsibility to subdomain
    - Specify name servers of subdomain



# How DNS Works – DNS query process

- Recursive query process
  - Ex: query lair.cs.colorado.edu => vangogh.cs.berkeley.edu, name server “ns.cs.colorado.edu” has no cache data




# DNS Delegation – Administered Zone

- Zone

- Autonomously administered piece of namespace

- Once the subdomain becomes a zone, it is independent to its parent

- Even parent contains NS's A record



# DNS Delegation – Administered Zone

- Two kinds of zone files
  - Forward Zone files
    - Hostname-to-Address mapping
    - Ex:
      - linux1.cs.nycu.edu.tw. IN A 140.113.17.11
  - Reverse Zone files
    - Address-to-Hostname mapping
    - Ex:
      - 11.17.113.140.in-addr.arpa. IN PTR linux1.cs.nycu.edu.tw.

# The Name Server Taxonomy (1)

- Categories of name servers
  - Based on the source of name server's data
    - **Authoritative**: official representative of a zone (master/slave)
      - **Master**: get zone data from disk
      - **Slave**: copy zone data from master
    - **Nonauthoritative**: answer a query from cache
      - **caching**: caches data from previous queries
  - Based on the type of answers handed out
    - **Recursive**: do query for you until it return an answer or error
    - **Nonrecursive**: refer you to the authoritative server
  - Based on the query path
    - **Forwarder**: performs queries on behalf of many clients with large cache
    - **Caching**: performs queries as a recursive name server

# The Name Server Taxonomy (2)


- Non-recursive referral
  - Hierarchical and **longest** known domain referral with cache data of other zone's name servers' addresses
  - Ex:
    - Query lair.cs.colorado.edu from a nonrecursive server
    - Whether cache has
      - IP of lair.cs.colorado.edu
      - Name servers of cs.colorado.edu
      - Name servers of colorado.edu
      - Name servers of edu
      - Name servers of root ("")
  - The resolver libraries do not understand referrals mostly. They expect the local name server to be recursive

# The Name Server Taxonomy (3)


- Caching
  - Positive cache (Long TTL)
  - Negative cache (Short TTL)
    - No host or domain matches the name queried
    - The type of data requested does not exist for this host
    - The server to ask is not responding
    - The server is unreachable of network problem
- Negative cache
  - 60% DNS queries are failed
  - To reduce the load of root servers, the authoritative negative answers must be cached

# The Name Server Taxonomy (4)

- Caching and forwarding DNS servers



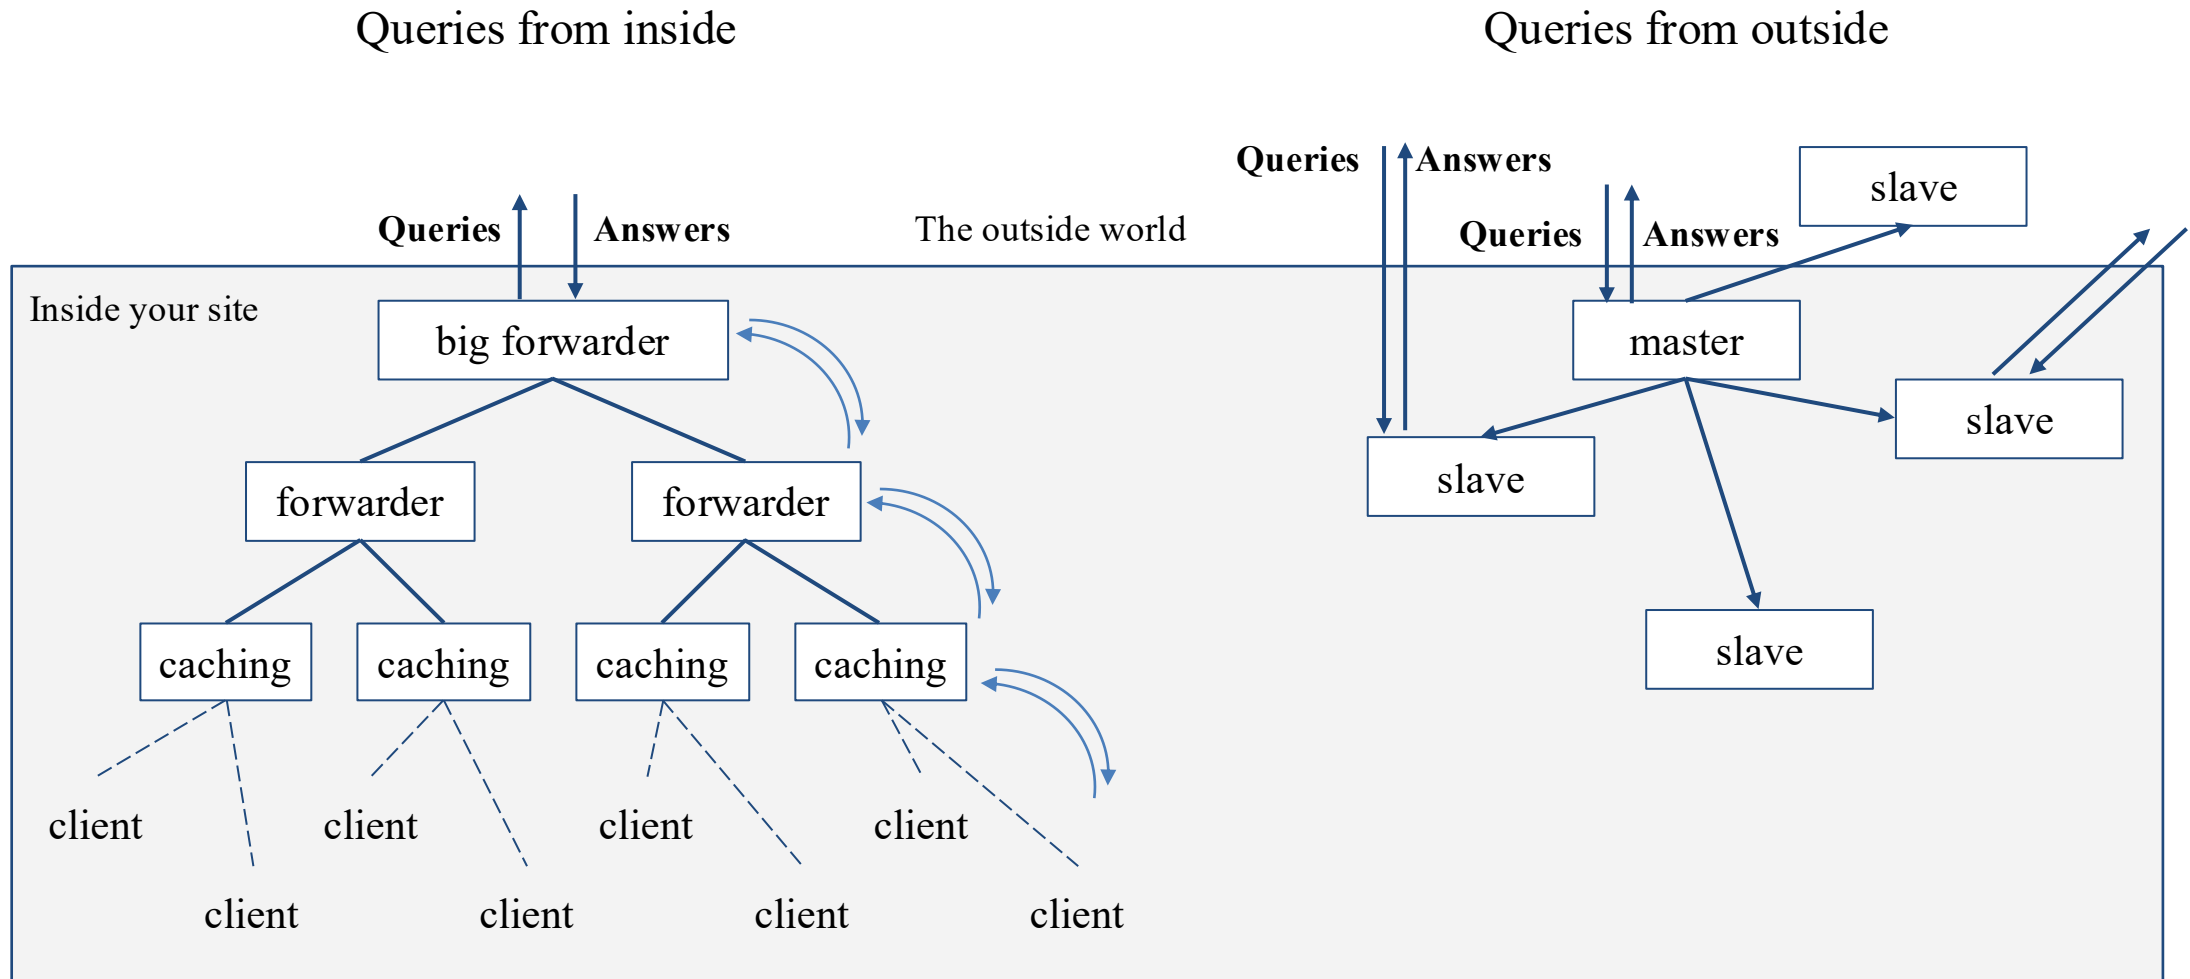
Caching



Forwarding

# The Name Server Taxonomy (5)

- How to arrange your DNS servers?
  - Ex:



# Root Name Servers

- Root name servers
  - In **named.root** file of BIND
  - <https://www.iana.org/domains/root/files>

|                     |         |    |      |                     |
|---------------------|---------|----|------|---------------------|
| .                   | 3600000 | IN | NS   | A.ROOT-SERVERS.NET. |
| A.ROOT-SERVERS.NET. | 3600000 |    | A    | 198.41.0.4          |
| A.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:503:ba3e::2:30 |
| .                   | 3600000 |    | NS   | B.ROOT-SERVERS.NET. |
| B.ROOT-SERVERS.NET. | 3600000 |    | A    | 199.9.14.201        |
| B.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:200::b     |
| .                   | 3600000 |    | NS   | C.ROOT-SERVERS.NET. |
| C.ROOT-SERVERS.NET. | 3600000 |    | A    | 192.33.4.12         |
| C.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:2::c       |
| .                   | 3600000 |    | NS   | D.ROOT-SERVERS.NET. |
| D.ROOT-SERVERS.NET. | 3600000 |    | A    | 199.7.91.13         |
| D.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:2d::d      |
| .                   | 3600000 |    | NS   | E.ROOT-SERVERS.NET. |
| E.ROOT-SERVERS.NET. | 3600000 |    | A    | 192.203.230.10      |
| E.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:a8::e      |
| .                   | 3600000 |    | NS   | F.ROOT-SERVERS.NET. |
| F.ROOT-SERVERS.NET. | 3600000 |    | A    | 192.5.5.241         |
| F.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:2f::f      |
| .                   | 3600000 |    | NS   | G.ROOT-SERVERS.NET. |
| G.ROOT-SERVERS.NET. | 3600000 |    | A    | 192.112.36.4        |
| G.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:12::d0d    |
| .                   | 3600000 |    | NS   | H.ROOT-SERVERS.NET. |
| H.ROOT-SERVERS.NET. | 3600000 |    | A    | 198.97.190.53       |
| H.ROOT-SERVERS.NET. | 3600000 |    | AAAA | 2001:500:1::53      |



# DNS Blocking (1)

- **DNS blocking** is a strategy for making it difficult for users to locate specific domains or websites on the Internet.
- It was first **introduced in 1997** as a means **to block spam email** from known malicious IP addresses.
- If an IP is on a block list, the DNS might reply that the **domain is unknown** or **with a different IP address that directs to a site** with a page stating that the requested domain is not permitted.
- Some public DNS Resolvers, like **Quad9** (for phishing and malicious domains) and **CleanBrowsing** (for adult contents), offer filters as part of their DNS.




# DNS Blocking (2)

```
$ dig www.meitufang.com @140.113.1.1
;; ANSWER SECTION:
www.meitufang.com.      300      IN      A
182.173.0.181

;; SERVER: 140.113.1.1#53(140.113.1.1) (UDP)

$ dig www.meitufang.com @8.8.8.8
;; ANSWER SECTION:
www.meitufang.com.      20866   IN      A
137.175.57.148

;; SERVER: 8.8.8.8#53(8.8.8.8) (UDP)
```



**此網域已經遭到封鎖**

**( This Domain Name Has Been Blocked )**

此網域涉違反 兒童及少年性剝削防制條例，經 衛生福利部 衛部  
護字第1131460390號 函聲請屏蔽  
若您對於本封鎖行為有疑義，請與 衛生福利部 聯繫

(If you have any questions about this termination, please contact the responsible  
government agency.)

# DNS Client Configurations

- /etc/resolv.conf

- nameserver
- domain
- search
- resolver(5), resolv.conf(5)

```
> cat /etc/resolv.conf
search it.cs.nycu.edu.tw cs.nycu.edu.tw
nameserver 10.1.1.1
nameserver 10.1.1.2
```

- /etc/hosts

- Format: **IP**      **FQDN**      **Aliases**
- **C:\Windows\system32\drivers\etc\hosts**
- hosts(5)

```
> cat /etc/hosts
::1          localhost localhost.my.domain
127.0.0.1   localhost localhost.my.domain
```

- /etc/nsswitch.conf

- **hosts: files (nis) (ldap) dns**
- nsswitch.conf(5)

```
> cat /etc/nsswitch.conf
hosts: files dns
```

# DNS Client Commands – host

- `$ host linux1.cs.nycu.edu.tw`  
`linux1.cs.nycu.edu.tw has address 140.113.17.11`
- `$ host 140.113.17.11`
- `11.17.113.140.in-addr.arpa domain name pointer IP-17-11.cs.nycu.edu.tw.`

# DNS Client Commands – nslookup

- `$ nslookup linux1.cs.nycu.edu.tw`

```
Server:          10.1.1.1
Address:         10.1.1.1#53
```

Non-authoritative answer:

```
Name:   linux1.cs.nycu.edu.tw
Address: 140.113.17.11
```

- `$ nslookup 140.113.17.11`

```
11.17.113.140.in-addr.arpa    name = IP-17-11.cs.nycu.edu.tw.
```

# DNS Client Commands – dig (1)

- `$ dig linux1.cs.nycu.edu.tw`

```
; <<>> DiG 9.18.44-1~deb12u1-Debian <<>> linux1.cs.nycu.edu.tw
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 42336
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;linux1.cs.nycu.edu.tw.          IN      A

;; ANSWER SECTION:
linux1.cs.nycu.edu.tw.  15      IN      A      140.113.17.11

;; Query time: 4 msec
;; SERVER: 10.1.1.1#53(10.1.1.1) (UDP)
;; WHEN: Mon Mar 09 09:04:22 CST 2026
;; MSG SIZE rcvd: 66
```

# DNS Client Commands – dig (2)

- `$ dig -x 140.113.17.11`

```
; <<>> DiG 9.18.44-1~deb12u1-Debian <<>> -x 140.113.17.11
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 1821
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;11.17.113.140.in-addr.arpa.      IN      PTR

;; ANSWER SECTION:
11.17.113.140.in-addr.arpa. 85969 IN      PTR      IP-17-11.cs.nycu.edu.tw.
```

# DNS Client Commands – drill

- Drop-in replacement of dig in unbound
- `$ drill -D www.cs.nctu.edu.tw`

```
;; ->>HEADER<<- opcode: QUERY, rcode: NOERROR, id: 46397
;; flags: qr rd ra ad ; QUERY: 1, ANSWER: 4, AUTHORITY: 0, ADDITIONAL: 0
;; QUESTION SECTION:
;; www.cs.nycu.edu.tw.   IN      A

;; ANSWER SECTION:
www.cs.nycu.edu.tw.      15      IN      CNAME   csproxy-itsc.cs.nycu.edu.tw.
csproxy-itsc.cs.nycu.edu.tw.  15      IN      A       140.113.96.55
www.cs.nycu.edu.tw.      15      IN      RRSIG   CNAME 13 5 3600 20260319000000
20260226000000 65052 cs.nycu.edu.tw.
u0dyVADm35ThwIHf3/8Enme0dL0t+2Ai411TPP4/vUm1P2q0lC/DvmjYx/3+tBb+oea3+eSV30ba+59
kwvlgZQ==
csproxy-itsc.cs.nycu.edu.tw.  15      IN      RRSIG   A 13 5 3600
20260319000000 20260226000000 65052 cs.nycu.edu.tw.
RrwUNLrqGZejU2Tpm/OxvunG5hLeof2ECJhH0BlefFm0GcWl zjFWz7ZjPz9VQnj0klNpxreL5udRXvf
fEiWZRQ==
```

# DNS Security

- DNSSEC

- Provide

- Origin authentication of DNS data
- Data integrity
- Authenticated denial of existence

- Not provide

- Confidentiality
- Availability

- `$ dig +dnssec linux1.cs.nycu.edu.tw`

```
;; ANSWER SECTION:
linux1.cs.nycu.edu.tw. 15      IN      A       140.113.17.11
linux1.cs.nycu.edu.tw. 15      IN      RRSIG   A 13 5 3600 20260319000000
20260226000000 65052 cs.nycu.edu.tw.
RQVPg4irDWzcUpJy7/NcZd6DBfSTYrtr4CNSDzEesIvjMpRRHG5XV4TW
DNGWRQ9T0Ow5ewPiNzWVc5TXpMKBSg==
```

# DNS Security (c.)

- DNS over TLS (DoT)
- DNS over HTTPS (DoH)
- DNS Amplification Attack
  - [http://www.cc.ntu.edu.tw/chinese/epaper/0028/20140320\\_2808.html](http://www.cc.ntu.edu.tw/chinese/epaper/0028/20140320_2808.html)

# DNS Server Software


- BIND <https://www.isc.org/bind/>
  - Complete DNS Server solution
- Name Server Daemon (NSD) <https://www.nlnetlabs.nl/projects/nsd/about/>
  - Authoritative DNS Server
    - No recursion, No caching
    - DNSSEC
- Unbound <https://www.nlnetlabs.nl/projects/unbound/about/>
  - Local resolver
    - Validating, Recursive, Caching
    - DoH, DoT

Ref: [https://en.wikipedia.org/wiki/Comparison\\_of\\_DNS\\_server\\_software](https://en.wikipedia.org/wiki/Comparison_of_DNS_server_software)

# Misc. (1)

- Internationalized Domain Name (IDN)
  - Punycode
    - A representation of Unicode with ASCII
    - .台灣 <-> .xn--kpry57d
    - <https://en.wikipedia.org/wiki/Punycode>
- Public & cloud services
  - Hurricane Electric Free DNS Hosting
    - <https://dns.he.net/>
  - AWS Route53
    - <https://aws.amazon.com/route53/>
- GeoDNS
  - Different DNS answers based on client's geographical location

# Misc. (2)



- DNS for fun
  - <https://www.dns.toys/>
- DNS Key Value Storage
  - <https://dnskv.com/>
- Tunnel
  - [net/iodine](https://net/iodine)
- Config
- FOSDEM 2023: Bizarre and Unusual Uses of DNS
  - Rule 53: If you can think of it, someone's done it in the DNS
  - [https://fosdem.org/2023/schedule/event/dns\\_bizarre\\_and\\_unusual\\_uses\\_of\\_dns/](https://fosdem.org/2023/schedule/event/dns_bizarre_and_unusual_uses_of_dns/)

## Useful utilities and services over DNS

dns.toys is a DNS server that takes creative liberties with the DNS protocol to offer handy utilities and services that are easily accessible via the command line. Copy and run the below commands to try it out.

| Service                                                                                                                                  | Usage (Click to copy)                                            |
|------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| <b>World time</b><br>Get current time for cities. Pass city names without spaces suffixed with .time. Optional two letter country codes. | <code>dig mumbai.time @dns.toys </code>                          |
| <b>Timezone conversion</b><br>Convert time between cities using format YYYY-MM-DDTHH:MM-\$fromCity-\$toCity                              | <code>dig 2023-05-28T14:00-mumbai-paris/fr.time @dns.toys</code> |
| <b>Weather</b><br>Get weather for cities. Pass city names without spaces. Optional country codes.                                        | <code>dig mumbai.weather @dns.toys</code>                        |

# References

- What are root name servers?
  - <https://www.netnod.se/i-root/what-are-root-name-servers>
- Ukraine asked the internet's governing body to remove Russian sites
  - <https://www.cnbc.com/2022/03/01/ukraine-asked-icann-to-revoke-russian-domains-shut-dns-servers.html>
- ICANN asked to suspend .ru and .su domains
  - <https://brandsec.com.au/icann-to-consider-suspending-ru-and-su-domains/>
- 紮根過去 放眼未來 - 網域名稱的歷史脈絡
  - <https://nccnews.com.tw/202103/ch3a.html>
- 臺北專屬域名「.taipei」現況與發展專題報導
  - <https://ws.ndc.gov.tw/Download.ashx?u=LzAwMS9hZG1pbmlzdHJhdG9yLzEwL1JlbEZpbGUvNTU2Ni83MDAwLzAwNjE0MDlfMS5wZGY%3D&n=cGFydDIucGRm&icon=..pdf>