

The Domain Name System

lwhsu (2020-2021, CC-BY)

? (?-2019)

國立陽明交通大學資工系資訊中心

Computer 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)

DNS Specification

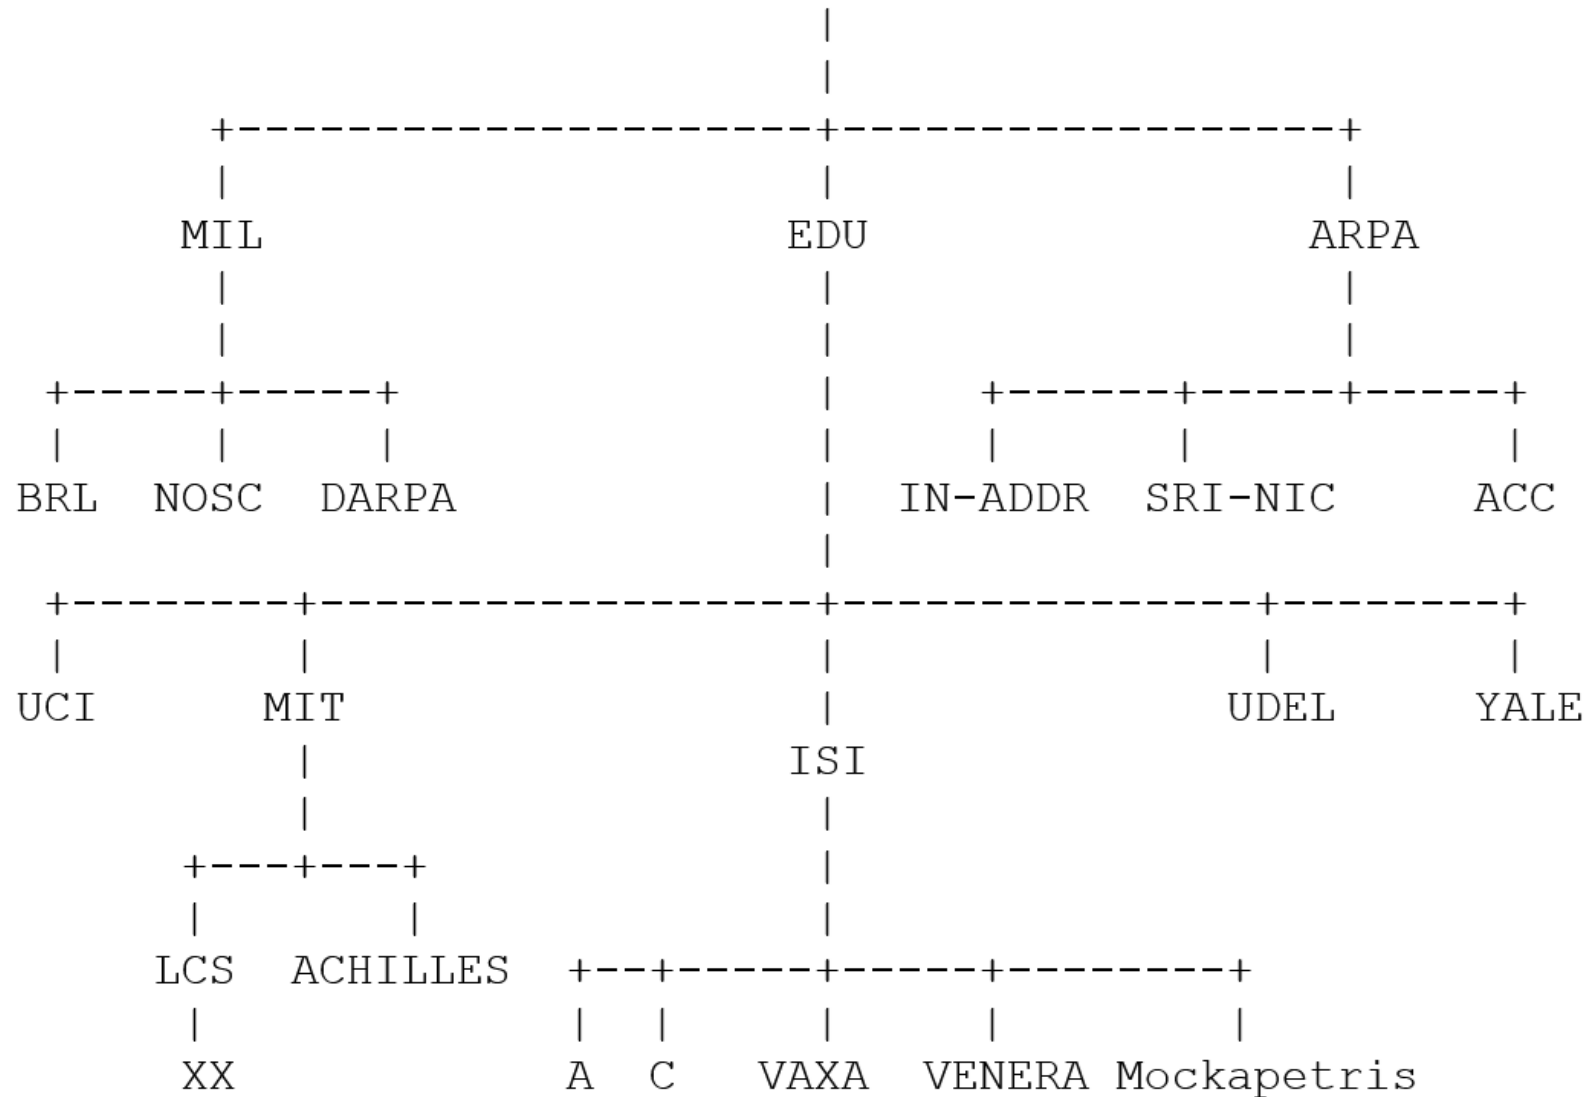
- **Tree architecture** – “**domain**” and “**subdomain**”
 - Divided into categories
 - Solves name collision
- **Distributed database**
 - Each site maintains a segment of the DB
 - Each site opens its information via network
- **Client-Server architecture**
 - Name servers provide information (Name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

The DNS Namespace – (1)

- Domain name is
 - A 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 level
 - Top-level / First level
 - Direct child of “root”
 - Maintained by ICANN (Internet Corporation for Assigned Names and Numbers)
 - Second-level
 - Child of a Top-level domain
- 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
 - 253 visible characters and 2 length bytes
 - What is the real maximum length of a DNS name?
 - <https://devblogs.microsoft.com/oldnewthing/20120412-00/?p=7873>

The DNS Namespace – (4)

- gTLDs (generic Top-Level Domains)
- 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
- int: International organization, such as nato.int

The DNS Namespace – (5)

- 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

 - xxx: for adult entertainment industry (sTLD)
 - On March 18st , 2011


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

The DNS Namespace – (6)

- Other than US, ccTLD (country code TLD)
 - ISO 3166, but just based on
 - Taiwan => tw
 - Japan => jp
 - United States => us
 - 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

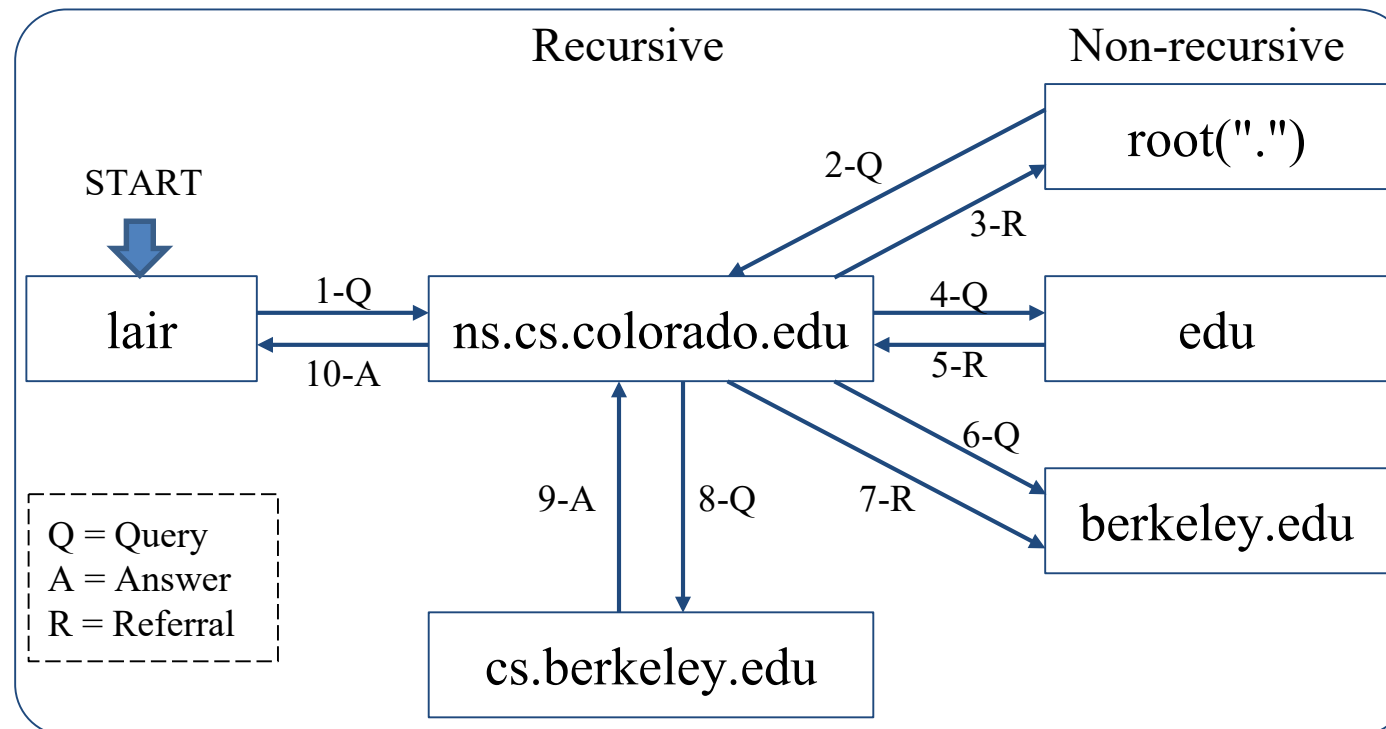
How DNS Works – DNS Delegation

- 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:
 - bsd1.cs.nctu.edu.tw. IN A 140.113.235.131
 - Reverse Zone files
 - Address-to-Hostname mapping
 - Ex:
 - 131.235.113.140.in-addr.arpa. IN PTR bsd1.cs.nctu.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)


- Nonrecursive 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



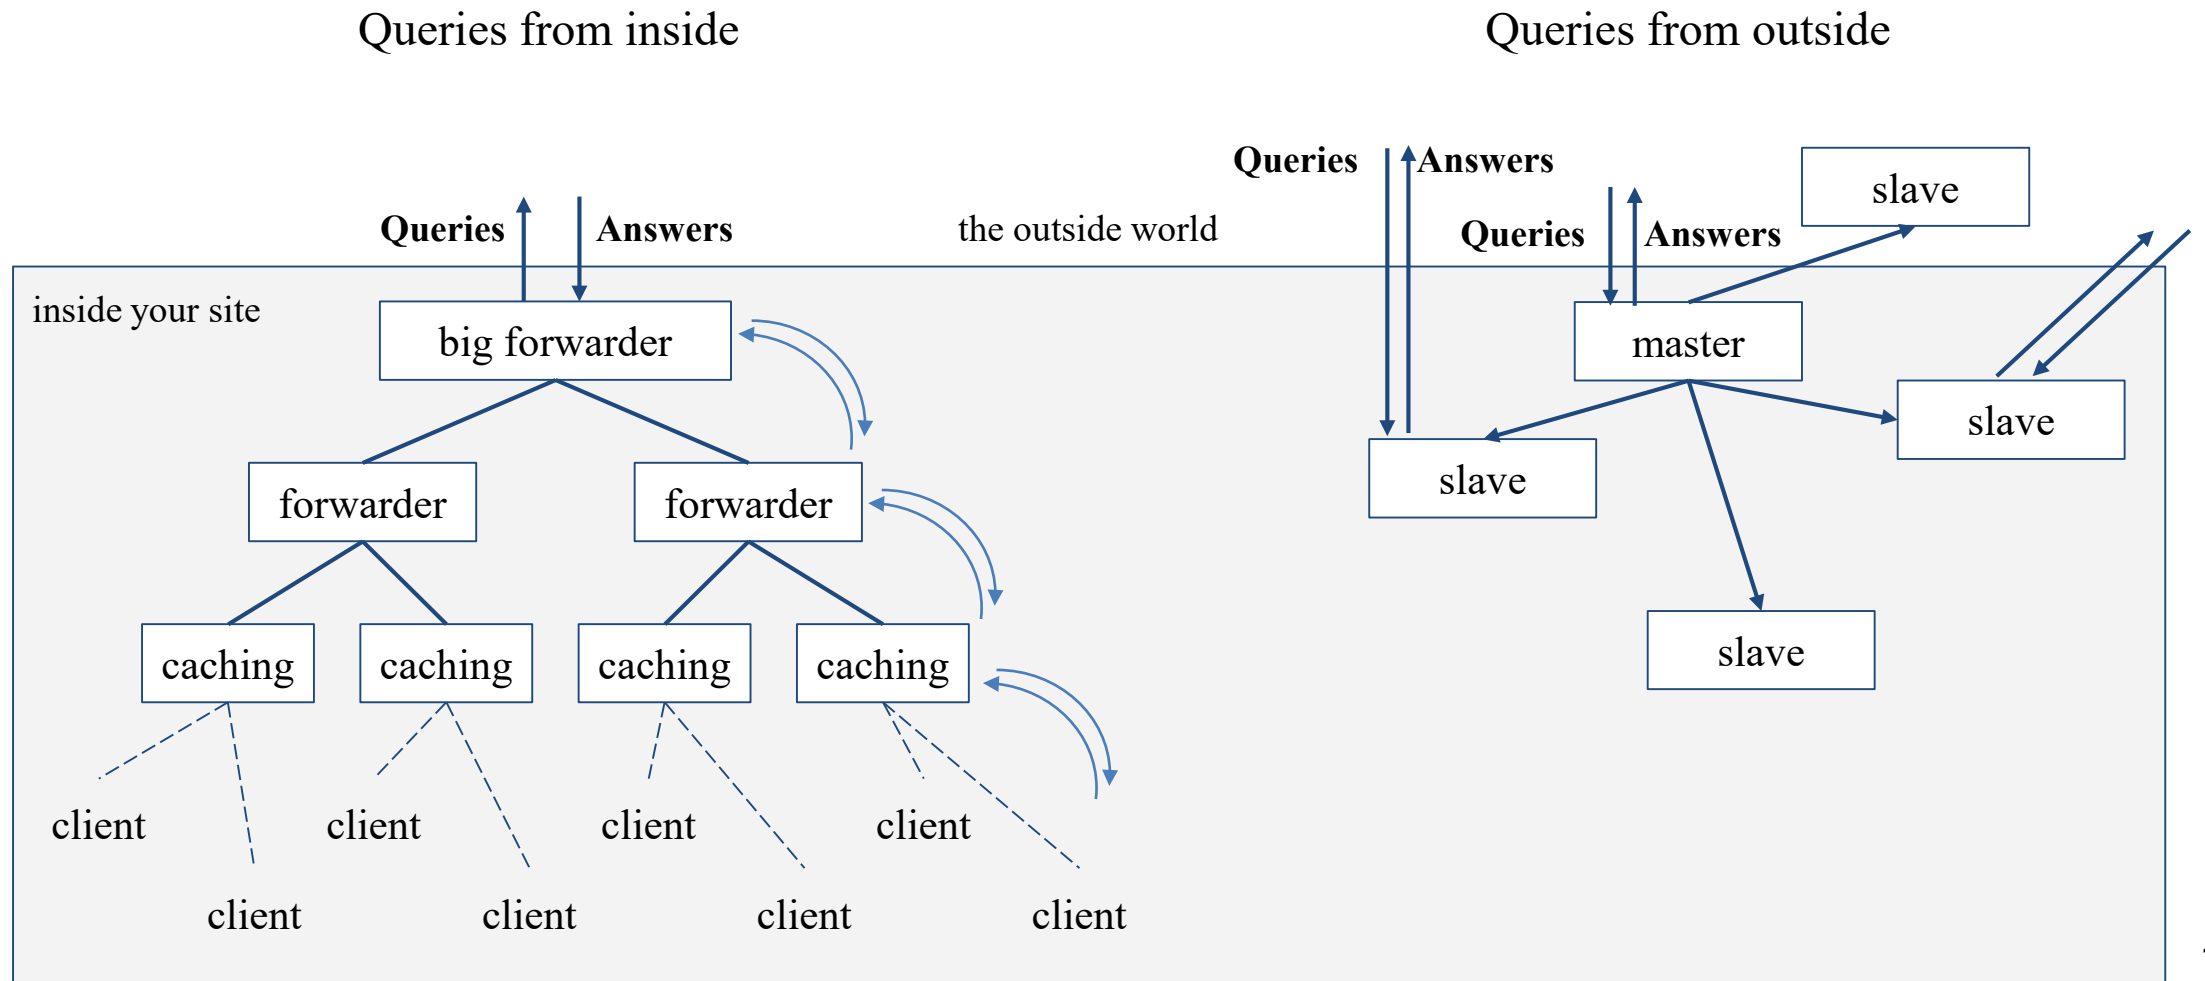
Caching



Forwarding

The Name Server Taxonomy (5)

- How to arrange your DNS servers?
 - Ex:



The Name Server Taxonomy (6)

- 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 Client Configurations

- `/etc/resolv.conf`
 - `nameserver`
 - `domain`
 - `search`
 - `resolver(5)`, `resolverconf(8)`
- `/etc/hosts`
 - IP FQDN Aliases
 - `C:\Windows\system32\drivers\etc\hosts`
 - `hosts(5)`
- `/etc/nsswitch.conf`
 - `hosts: files (nis) (ldap) dns`
 - `nsswitch.conf(5)`

DNS Client Commands – host

- `$ host nasa.cs.nctu.edu.tw`
 - `nasa.cs.nctu.edu.tw` has address `140.113.17.32`
- `$ host 140.113.17.32`
 - `32.17.113.140.in-addr.arpa` domain name pointer `nasa.cs.nctu.edu.tw`.

DNS Client Commands – nslookup

- `$ nslookup nasa.cs.nctu.edu.tw`
 - Server: 140.113.235.1
 - Address: 140.113.235.1#53
 - Name: nasa.cs.nctu.edu.tw
 - Address: 140.113.17.32
- `$ nslookup 140.113.17.225`
 - Server: 140.113.235.1
 - Address: 140.113.235.1#53
 - 32.17.113.140.in-addr.arpa name = nasa.cs.nctu.edu.tw.

DNS Client Commands – dig (1)

- \$ dig nasa.cs.nctu.edu.tw

```
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 47883
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3

;; QUESTION SECTION:
;nasa.cs.nctu.edu.tw.          IN      A

;; ANSWER SECTION:
nasa.cs.nctu.edu.tw.         3600    IN      A      140.113.17.32

.....
```

DNS Client Commands – dig (2)

- \$ dig -x nasa.cs.nctu.edu.tw

```
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 5514
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3

;; QUESTION SECTION:
;32.17.113.140.in-addr.arpa.      IN      PTR

;; ANSWER SECTION:
32.17.113.140.in-addr.arpa. 86400 IN      PTR      nasa.cs.nctu.edu.tw.

.....
```


DNS Security

- DNSSEC
 - Provide
 - Origin authentication of DNS data
 - Data integrity
 - Authenticated denial of existence
 - Not provide
 - Confidentiality
 - Availability
 - \$ dig +dnssec bsd1.cs.nctu.edu.tw

```
;; ANSWER SECTION:
bsd1.cs.nctu.edu.tw. 3600 IN A 140.113.235.131
bsd1.cs.nctu.edu.tw. 3600 IN RRSIG A 7 5 3600 ...
```

RRSIG: Resource Record Signature

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

DNS Server Software

- BIND
 - Complete DNS Server solution
- NSD
 - Authoritative DNS Server
 - No recursion, No caching
 - DNSSEC
- Unbound
 - Local resolver
 - Validating, Recursive, Caching
 - DoH, DoT
- https://en.wikipedia.org/wiki/Comparison_of_DNS_server_software

Misc.

- 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