

# DNS Protocol and Attacks

## (1) DNS (Domain Name Service)

2004 presidential debate between vice president Dick Cheney and John Edward:

**Cheney:** Well, the reason they keep mentioning Halliburton is because they're trying to throw up a smokescreen. They know the charges are false. They know that if you go, for example, to **FactCheck.com**, an independent Web site sponsored by the University of Pennsylvania, you can get the specific details with respect to Halliburton.

The actual web site should be **FactCheck.org** (a politically neutral web size), not **FactCheck.com**. George Soro (who does not like President Bush) immediately capitalized on this mistake by somehow (he might paid the owner of FactCheck.com for doing so) redirecting all the traffic to **FackCheck.com** to his own site, where the top item is an article by Soros entitled "*Why we must not Re-Elect President Bush*".

- ❖ Motivation
  - Human prefer pronounceable names rather numeric IP addresses
- ❖ Original Naming Scheme
  - The original names formed a flat namespace without structure
  - A central site, the Network Information Center (NIC), administered the namespace
  - Later, the NIC was replaced by the INTERNET Network Information Center.
  - Advantage: names are convenient and short
  - Disadvantage: a flat namespace cannot generalize to large sets of machines for both technical and administrative reasons.
    - Potential conflict
    - Names are assigned by a center server
    - Maintaining correct copies of the entire list at each site is difficult
- ❖ Hierarchical Names
  - Decentralizing the naming mechanism: delegating authority and distributing responsibility
  - A hierarchical naming scheme works like the management of a large organization.
    - The namespace is partitioned at the top level
    - The authority for names in subdivisions is passed to designated agents
- ❖ DNS
  - Specifies the name syntax and rules for delegating authority over names
  - Specifies the implementation of distributed computing system that efficiently map names to addresses.
- ❖ DNS Syntax
  - Set of labels separated by delimiter character (period)

- Example: `ecs.syr.edu`
- `syr.edu` is also a domain
- The top-level domain is `edu`

❖ Original Top-Level Domains


Domain Name	Meaning
<b>COM</b>	<b>Commercial organizations</b>
<b>EDU</b>	<b>Educational institutions (4-year)</b>
<b>GOV</b>	<b>Government institutions</b>
<b>MIL</b>	<b>Military groups</b>
<b>NET</b>	<b>Major network support centers</b>
<b>ORG</b>	<b>Organizations other than those above</b>
<b>ARPA</b>	<b>Temporary ARPANET domain (obsolete)</b>
<b>INT</b>	<b>International organizations</b>
<i>country code</i>	<b>Each country (geographic scheme)</b>

❖ New Top-Level Domains

Domain Name	Assigned To
<b>aero</b>	<b>Air-Transport Industry</b>
<b>biz</b>	<b>Businesses</b>
<b>coop</b>	<b>Non-Profit Cooperatives</b>
<b>info</b>	<b>Unrestricted</b>
<b>museum</b>	<b>Museums</b>
<b>name</b>	<b>Individuals</b>
<b>pro</b>	<b>Professionals (accountants, lawyers, physicians)</b>

❖ Mapping Domain Names to Addresses

- Name server: supplies name-to-address translation
- Client: uses one or more name servers when translating a name
- DNS uses a set of on-line servers
- Servers arranged in tree
- Given server can handle entire subtree. For example, ECS manages domain names within the `ecs.syr.edu` domain.



#### ❖ Root Servers

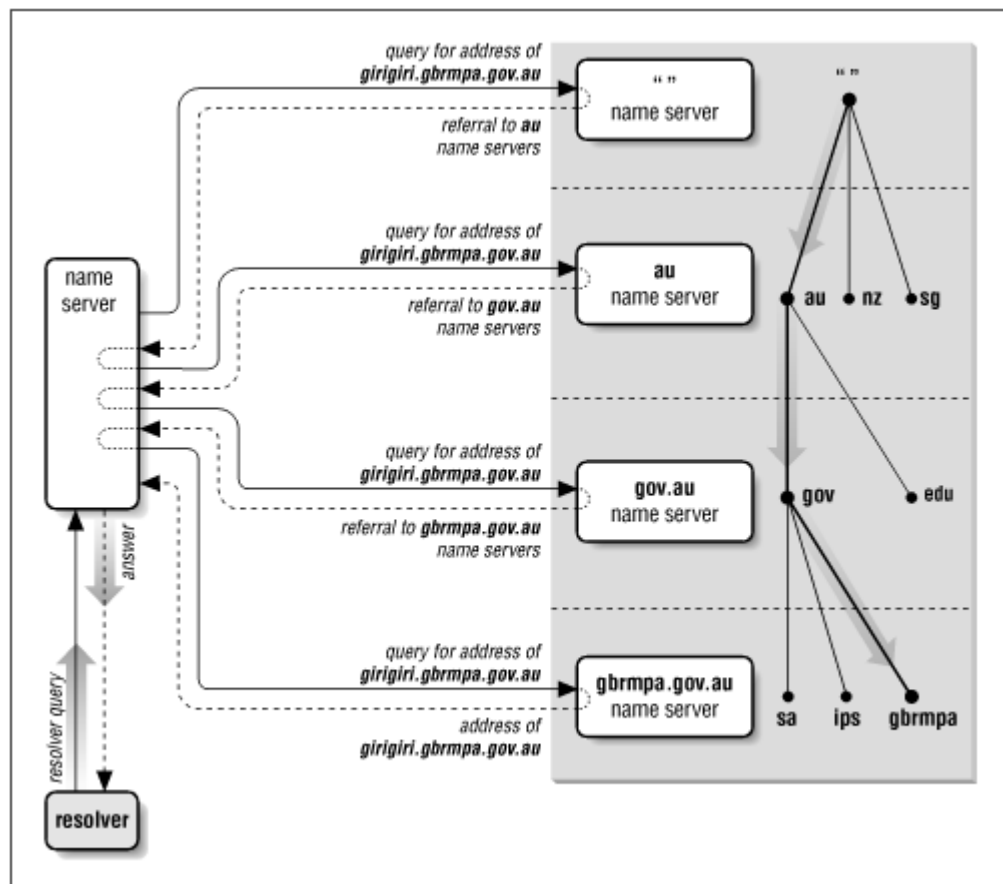
- There are 13 DNS root servers from A.ROOT-SERVERS.NET, B.ROOT-SERVERS.NET to M.ROOT-SERVERS.NET.
- These root servers have well known IP addresses, and should be used when configuring DNS servers. For example, the A server's IP address is 198.41.0.4, operated by VeriSign and the B server's IP address is 192.5.5.241, operated by ISC, etc.
- Currently (2007), 6 of the 13 root servers are not single machines. They represent several physical servers each in multiple geographical locations. For example, F-root answers queries over IPv4 on 192.5.5.241, and over IPv6 on 2001:500::1035. Service for f.root-servers.net is provided by a distributed collection of nameserver nodes located in many places, using a [Hierarchical Anycast technique](#) running [ISC BIND 9](#).
- Until mid-2000, the root servers also handled all requests for the generic top level domains. Due to the potential denial of service attacks, this responsibility was later removed from the root servers and led to the creation of dedicated Top-Level Domain Servers to handle .com, net, .org, country code, etc.

#### ❖ Efficient Translation

- Facts
  - Most lookups refer to local names
  - Name-to-address bindings change infrequently
  - User is likely to repeat the same lookup
- Initial contact begins with the local DNS server
- Caching: every server caches answers
  - Local server maintain caches
  - Bindings change very infrequently
  - Time to Live for each entry: set by the authoritative server


#### ❖ Types of queries

- Recursive: often used by the client
  - Iterative: often used by the local DNS server
- ❖ Recursive query:
- A resolver sends a *recursive query* to a name server for information about a particular domain name. The queried name server is then obliged to respond with the requested data or with an error stating that data of the requested type don't exist or that the domain name specified doesn't exist.
  - If the queried name server isn't authoritative for the data requested, it will have to query other name servers to find the answer. It could send recursive queries to those name servers, thereby obliging them to find the answer and return it (and passing the buck). Or it could send iterative queries and possibly be referred to other name servers "closer" to the domain name it's looking for. Current implementations are polite and do the latter, following the referrals until an answer is found.
- ❖ Iterative process
- *Iteration*, or *iterative resolution*, refers to the resolution process used by a name server when it receives iterative queries.
  - The figure is from "DNS and BIND" by Paul Albitz & Cricket Liu



## ❖ Inverse Mappings

- Implemented by a separate, parallel tree, keyed by IP address.
- 222.33.44.3 --> 3.44.33.222.in-addr.arpa
- The Internet root domain servers maintain a database of valid IP addresses along with information about domain name servers that can resolve each address.



## ❖ DNS Port

- The DNS uses TCP and UDP on port 53 to serve requests.
- Almost all DNS queries consist of a single UDP request from the client followed by a single UDP reply from the server.
- TCP typically comes into play only when the response data size exceeds 512 bytes, or for such tasks as zone transfer.

## ❖ DNS Caching

```
;; ANSWER SECTION:
apollo.ecs.syr.edu.      1D IN CNAME   apollo.syr.edu.
apollo.syr.edu.         1D IN A       128.230.208.46

;; AUTHORITY SECTION:
syr.edu.                1D IN NS      ns3.broadwing.net.
syr.edu.                1D IN NS      ns4.broadwing.net.
syr.edu.                1D IN NS      lurch.cns.syr.edu.
syr.edu.                1D IN NS      icarus.syr.edu.

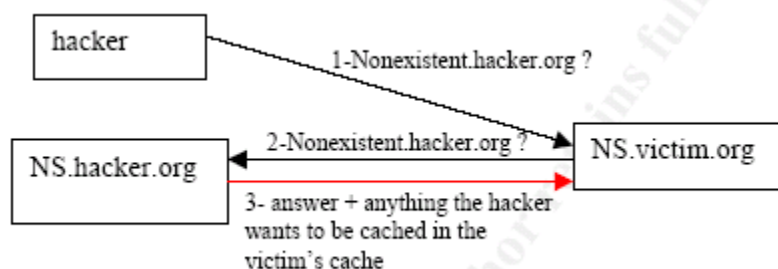
;; ADDITIONAL SECTION:
lurch.cns.syr.edu.      1D IN A       128.230.12.5
icarus.syr.edu.         1D IN A       128.230.1.49

;; Total query time: 1 msec
;; FROM: nyx to SERVER: default -- 128.230.12.5
```

- ❖ DNS software
  - BIND (Berkeley Internet Name Domain)
  - Microsoft DNS (in the server editions of Windows 2000 and Windows 2003)
  - DNS-oriented utilities
    - dig: the "domain information groper"

## (2) Attacking DNS Protocols

- ❖ DNS cache poisoning
  - Victim DNS server asks other DNS servers for mappings if it doesn't have them. It then caches the mappings.
  - DNS cache poisoning is to poison the client's cache
- ❖ Unrelated Data Attack
  - To improve performance, DNS servers can send back more information than what the client has asked for. For example, if the client asks for the IP address of `www.mysite.com`, the DNS server can also send back the IP addresses for `ftp.mysite.com` and `mail.mysite.com` to avoid another likely DNS lookup.
  - In the older version of DNS servers, the validity of the extra information is not verified. The actual information does not even need to be related to the original query. Therefore, a malicious DNS server from `mysite.com` can send back the faked IP addresses for `Citibank.com`, tricking users to go to the malicious site when they try to connect to Citibank.
  - This problem has been fixed in BIND, by forbidding anything that is not related to the original request to be cached.



- ❖ Related Data Attack
  - The process is the same as the unrelated data attack
  - The hacker has to make the "extra" information related to the original query
    - MX: mail server for a domain
    - CNAME: canonical name for an alias
    - NS: DNS servers for a domain
  - The above information is "related" to the original request, but they can point to totally different information the hacker wants to be cached.
  - The problem has also been fixed in BIND, by rejecting all the "out of zone" information.

- ❖ DNS Spoofing:
  - Answer DNS queries intended for another server.
  - Difficulties: Transaction ID (16 bits) and source UDP port must be guessed.
  
- ❖ Reverse DNS attacks
  - Assume `ulysses.ecs.syr.edu` and `apollo.ecs.syr.edu` trust each other, and the trust is based on name, not the IP address.
  - Also assume that you are from `hacker.com`, and you have the control of your own DNS server (which means IP to name and name to IP for machines in `hacker.com`).
  - How can you exploit the trust relationship between Ulysses and Apollo?
  - Q: How to prevent this?
    - Use IP address for the trust relationship
    - A second (forward) DNS look up.
  
- ❖ **DNS Rebinding Attacks**
  
- ❖ **Pharming Attacks:** aiming to redirect a website's traffic to another, bogus website. Many techniques can be used, and mostly target DNS.
  - *Insider attack:* corrupted insiders can modify the local DNS servers to mislead users to bogus websites.
  - *Corrupted hosts:* if a host is already corrupted, the `/etc/hosts` file can be modified; moreover, the DNS lookup process can also be modified. These can cause users to go to bogus websites.
  - *Domain Registration Attacks:* when a domain registration expires and the owner forgot to renew it, attackers can buy that domain legally and hijack the domain. Attackers can also buy the domain names that are similar to their targets. If users misspell the domain names, they might become victims.
  - *Corrupting DNS servers.* Change the mappings on the corrupted DNS servers.
  - *DNS cache Poisoning and DNS Spoofing.*