

# IP-related security issues + DNS + DHCP

Michael Tsai  
2016/04/13

# Common IP-related security issues

- IP forwarding - pretending to be a router
- ICMP redirect (please send your packet here)  
/ source routing (please follow the specified route)
- Broadcast pings (“smurf attack”)  
(directed broadcast - from “far-far-away”)
- IP spoofing -  
send a packet with somebody else’s IP as the source  
address  
“unicast reverse path forwarding” (uRPF)

Download the VM  
for today's exercise  
lubuntu.ova  
<http://goo.gl/Bax8B8>

# DNS + DHCP

- DNS: domain name  $\longleftrightarrow$  IP address
- DHCP:  
gives you a IP + configuration when you joins a new network

DHCP =  
Dynamic Host  
Configuration Protocol

# DHCP (Dynamic Host Configuration Protocol)

---

- ▶ 每個地方有自己的subnet及IP設定
- ▶ 到一個新的地方，一開始怎麼取得此一subnet的IP呢?
- ▶ 通常同一個subnet中會設置一台DHCP server
- ▶ 此server將負責“接待”新來的機器，分發未使用的IP給它們
- ▶ 想像全系如果都需要手動設定IP, 會發生什麼事情?
  - ▶ 網管需要分配IP給所有電腦 (全系有多少電腦???)
  - ▶ IP衝突 (同樣的IP被不同的電腦使用)




# DHCP 4部曲

DHCP Discover: 請問有人可以發IP給我嗎?

Src: 0.0.0.0, 68  
Dest: 255.255.255.255, 67  
DHCPDISCOVER  
Yiaddr: 0.0.0.0  
Transaction ID: 654  
Request:  
Subnet Mask, Router, Domain Name Server

DHCP Offer: 我這邊有一組IP看看你要不要用.

Src: 192.168.55.254, 67  
Dest: 255.255.255.255, 68  
DHCPOFFER  
Yiaddr: 192.168.48.15  
DHCP server ID: 192.168.55.254  
Transaction ID: 654  
Lifetime: 4 hrs  
Netmask: 255.255.248.0  
Router: 192.168.55.254  
DNS: 140.112.30.21, 140.112.254.4



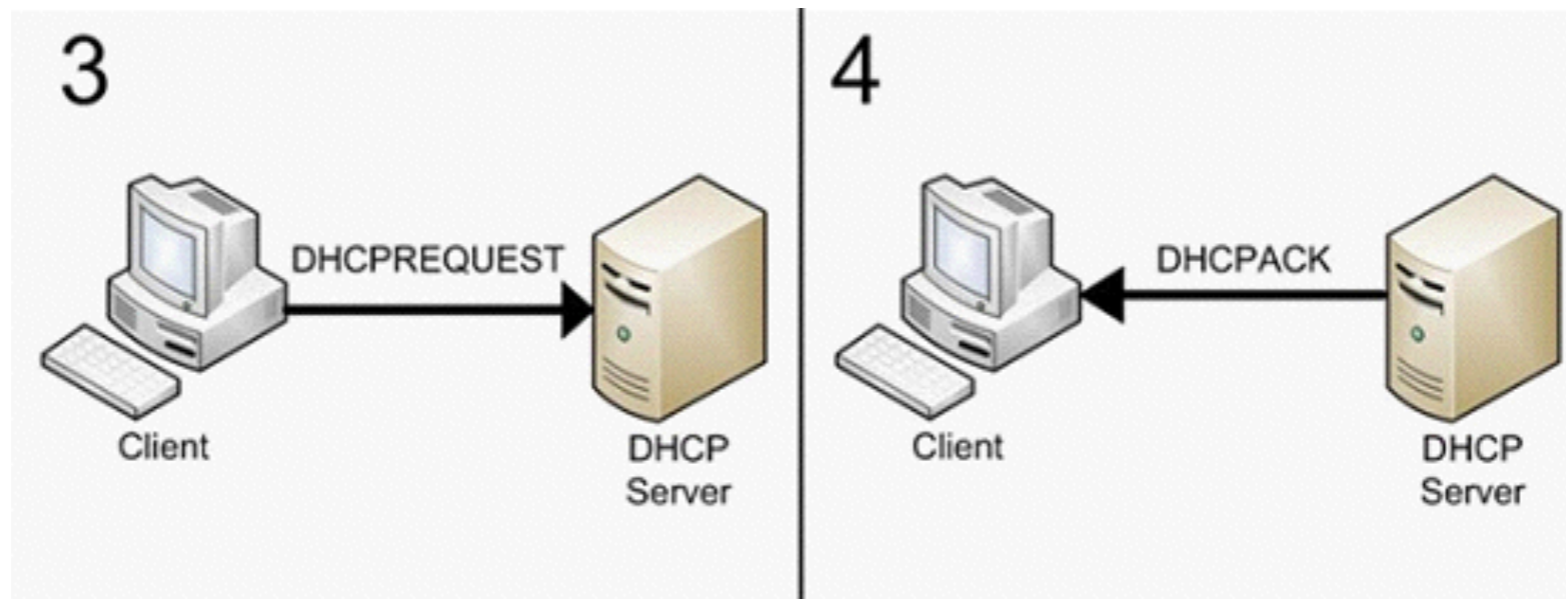
# DHCP 4部曲

DHCP Request:那我要把這組IP拿走囉!

Src: 0.0.0.0, 68  
Dest: 255.255.255.255, 67  
DHCPREQUEST  
Yiaddr: 192.168.48.15  
Transaction ID: 655  
DHCP server ID: 192.168.55.254  
Lifetime: 4 hrs

DHCP Ack: 沒問題. 請用.

Src: 192.168.55.254, 67  
Dest: 255.255.255.255, 68  
DHCPACK  
Yiaddr: 192.168.48.15  
DHCP server ID: 192.168.55.254  
Transaction ID: 655  
Lifetime: 4 hrs  
Netmask: 255.255.248.0  
Router: 192.168.55.254  
DNS: 140.112.30.21, 140.112.254.4





# Some additional facts

---

- ▶ Originally designed as BOOTP - for diskless workstations.
- ▶ Today, You can still use PXEBOOT on various network interface card for diskless operation.
- ▶ A server keeps track of “lease” (of IPs), which would expire after a predefined time period
- ▶ A client usually “renews” the lease when the time is half over
- ▶ The lease information must survive reboot for network stability

- ▶ A typical lease entry:

```
lease 192.168.20.4 {  
    starts 6 2009/06/27 00:40:00;  
    ends 6 2009/06/27 12:40:00;  
    hardware ethernet 00:00:00:00:00:00;  
    uid 00:00:00:00:00:00;  
    client-hostname "example-workstation1";
```

```
}  
▶
```

# DHCP 的細節

---

- ▶ 一個subnet上可能有多個DHCP server. 因此發出DHCPREQUEST之後，可能收到多個DHCPPOFFER。
- ▶ Client可以要求使用之前使用過的IP，但DHCP server可以拒絕(可能根本已經不在同一個網段，或是已經被別的client使用中)
- ▶ Authoritative & non-authoritative: 有主管權的DHCP server可以發出“拒絕”client使用某IP的要求，而沒有主管權的DHCP server則會忽略該要求(沒有回應)
- ▶ 想想看: DHCP server的安全漏洞. 如果有人接在系上網路上且開啟DHCP server，會發生什麼事情?



# A typical DHCP server configuration file

---

- ▶ The following dhcpd.conf is used by ISC DHCPD

```
default-lease-time 600;
max-lease-time 7200;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option routers 192.168.1.254;
option domain-name-servers 192.168.1.1, 192.168.1.2;
option domain-search "example.com";
subnet 192.168.1.0 netmask 255.255.255.0 {
    range 192.168.1.10 192.168.1.100;
}

host apex {
    option host-name "apex.example.com";
    hardware ethernet 00:A0:78:8E:9E:AA;
    fixed-address 192.168.1.4;
}
```

DNS =

Domain Name System

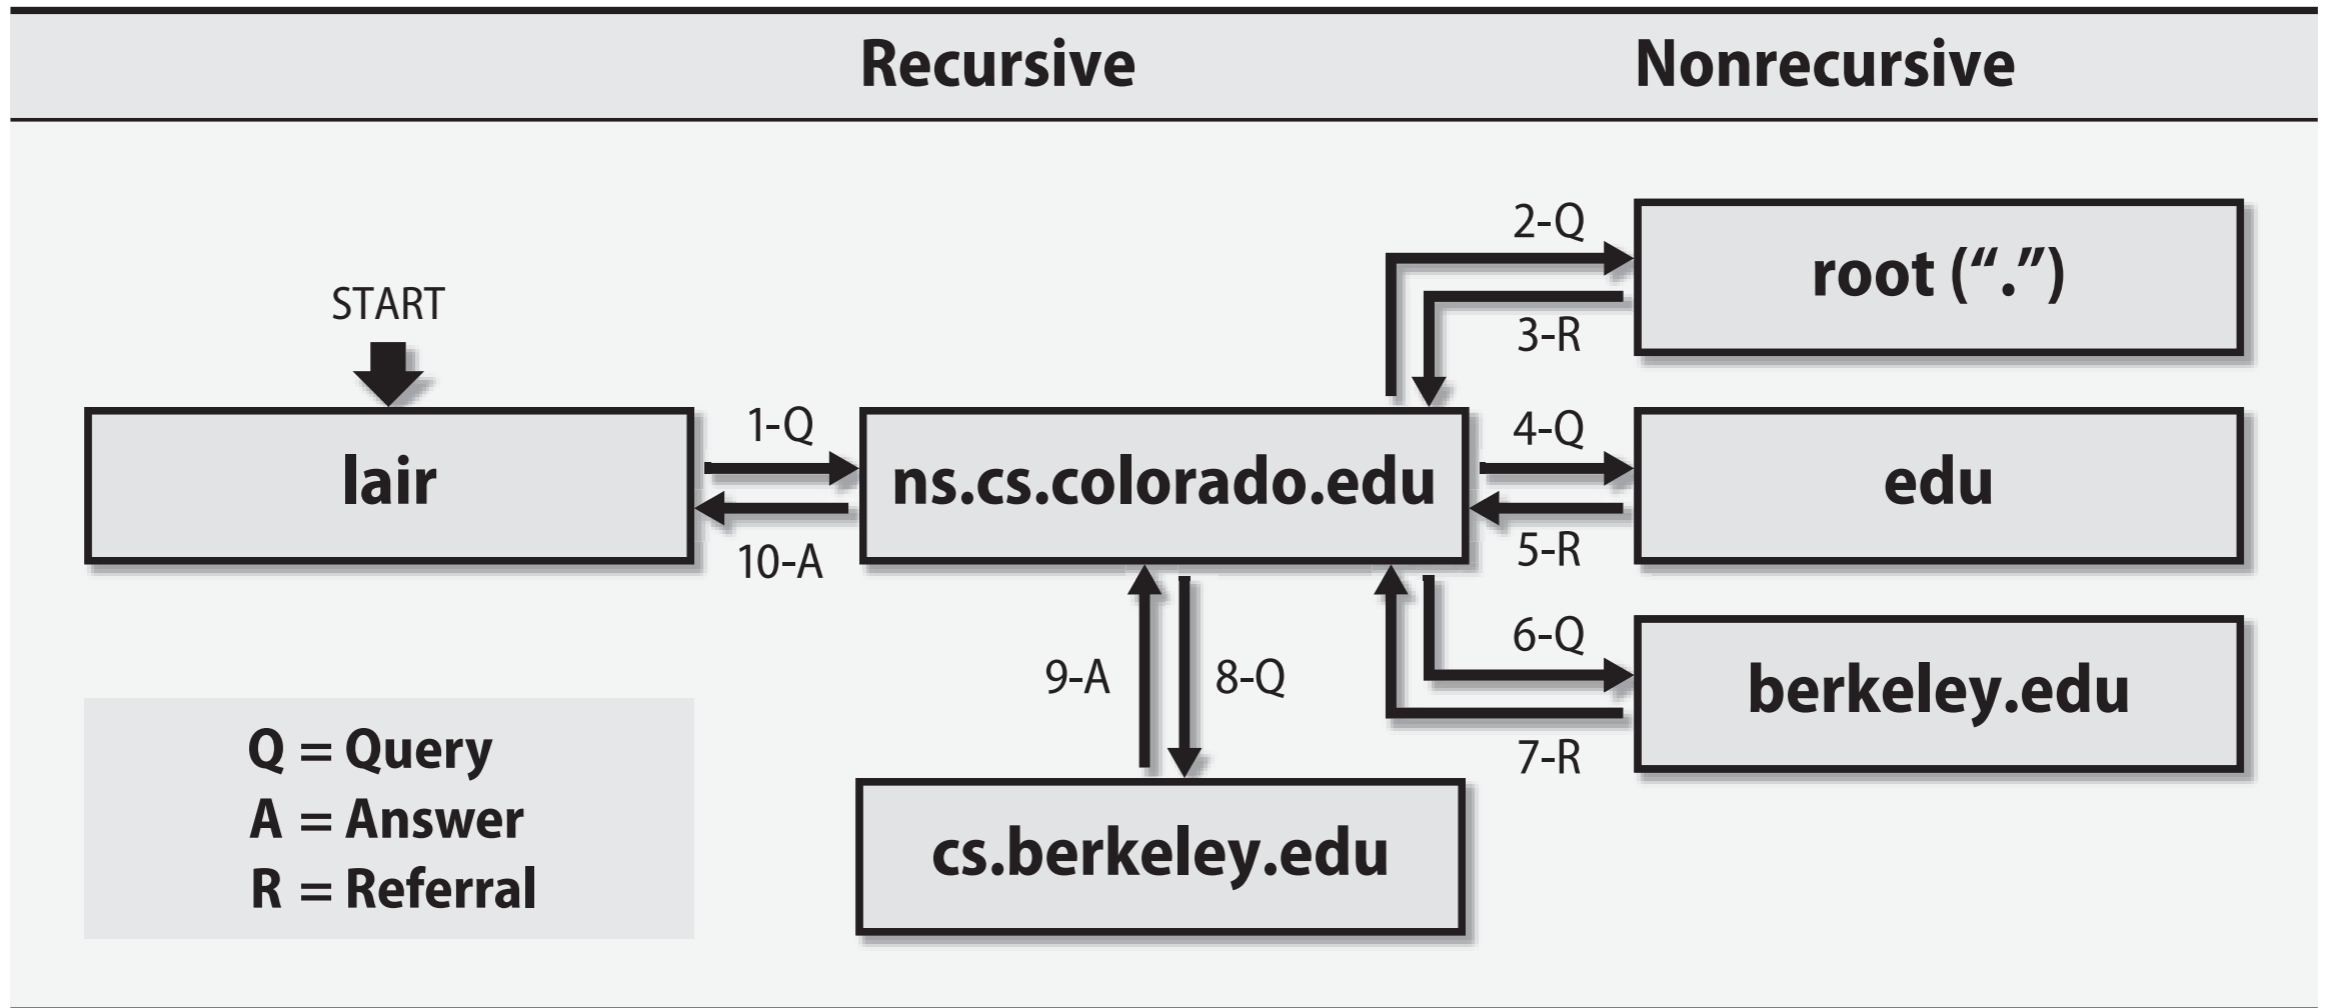
# DNS (Domain Name Service)

---

- ▶ 一言以蔽之: 將名稱轉為IP的服務
- ▶ 常見的轉換種類:
  - ▶ Domain name -> IP (type A):  
ntucsv.csie.ntu.edu.tw -> 140.112.30.28
  - ▶ @domainname的mail server (type MX):  
csie.ntu.edu.tw -> ms.csie.ntu.edu.tw
  - ▶ Domain name -> domain name (type CNAME):  
www.csie.ntu.edu.tw -> ntucsv.csie.ntu.edu.tw
  - ▶ IP -> domain name (type PTR)  
140.112.30.21 -> csman.csie.ntu.edu.tw
- ▶ 可以多重宣告: 增加可靠度或分散性.
  - ▶ 例如www.google.com的A指到了6個IP!



# Delegation = distributed



# DNS Basics

- **Forward mapping:** domain name → IP
- **Reverse mapping:** IP → domain name
- **Master (primary) server:**  
holds the “reference copy” of the domain data
- **Slave (secondary) server:**  
copy the data from the master server
- **Caching-only server:**  
holds the domain data only temporarily in cache
- **Zone transfer:**  
the “copying” or “propagation” of DNS data from master to slave

# DNS Basics

- **Authoritative**: from the server responsible for that “zone”
- **Non-authoritative**: from caching server (your local DNS)
- **TTL (time-to-live)**: How long can the caching server save the answer in the (temporary) database
- Example:
  - roots: 42 days
  - edu: 2 days
  - berkeley.edu: 2 days
  - vangogh.cs.berkeley.edu: 1 day
- Example: moving CSIE servers to BL —>  
set TTL to be 1 hr or less



# Client Setting

- In `/etc/resolv.conf`  
(but most OS have similar fields)
- **search**: search domain name with these postfix
- **nameserver**: the **IP addresses** of the nameserver  
first will be used  
if not working (timeout), the next one will be used
- **options**: additional options can be specified  
(e.g., `options rotate timeout:2 attempts:2`)  
try every server, times out in 2 sec, query each server at  
most 2 times

# 常用DNS指令

- Examples:
  - `dig @8.8.8.8 -t MX csie.ntu.edu.tw`
  - `dig @140.112.30.21 www.csie.ntu.edu.tw`

```
;; ANSWER SECTION:
www.csie.ntu.edu.tw.      600      IN       A       140.112.30.28

;; AUTHORITY SECTION:
csie.ntu.edu.tw.        86400    IN       NS      csman2.csie.ntu.edu.tw.
csie.ntu.edu.tw.        86400    IN       NS      ntuns.ntu.edu.tw.
csie.ntu.edu.tw.        86400    IN       NS      csman.csie.ntu.edu.tw.

;; ADDITIONAL SECTION:
csman.csie.ntu.edu.tw.  600      IN       A       140.112.30.21
ntuns.ntu.edu.tw.       85489    IN       A       140.112.254.6
csman2.csie.ntu.edu.tw. 600      IN       A       140.112.30.12
```

# Resource Records

				Mail server
mail	IN	MX 10	mail	
	IN	A	140.112.91.209	
www	IN	CNAME	<u>mvnl.csie.ntu.edu.tw.</u>	Alias
mvnl	IN	NS	ns-mvnl	Sub-domain
ns-mvnl	IN	A	140.112.91.208	Forward
208	IN	PTR	<u>mvnl.csie.ntu.edu.tw</u>	Reverse

# Resource Records

- ; comments
- @ the current zone name
- () allows data to span line
- **IN** Internet (default)
- Formal syntax: zone [ttl] [IN] record\_type record\_content

# Fire up your VM!

1. Install bind9 in your ubuntu VM (<http://goo.gl/Bax8B8>)  
(sudo apt-get update ; sudo apt-get install bind9)
2. Add forward setting —> point to 140.112.30.21  
(in /etc/bind/named.conf.options  
“service bind9 reload” after you’re done)
3. Make sure your ubuntu is using this nameserver  
(check /etc/resolv.conf)
4. Check if you can get the answers from your own DNS  
server

# Considerations in DNS service

- No name service = no service at all (mostly)
- Diversity: Geographically, software, etc.
- Standalone (not many users)
- Uninterruptible power supply
- Two slaves, one is off-site
- Separate authoritative and caching-only servers (security)

# DNS

## can be used for attacks!

- DNS spoofing (cache poisoning)
  - Redirect the user to a malicious server instead of the official server
- DNS hijacking
- DNS DDoS (reflection / amplification)
  - Request (30 bytes) versus response (3297 bytes): **100x !**
  - Send requests to a large number of DNS servers **on behalf of the host to be attacked** (no check on the source IP)
  - Mitigation (DNS server): limit the hosts (with IP in certain ranges) which can query your DNS

# SOA record

```
; start of the authority record for
nasa2015.com
@    IN    SOA    ns.nasa2015.com.
root.nasa2015.com. (
                2015042701 ; serial number
                10800      ; refresh (3 hr)
                1200       ; retry (20 min)
                3600000     ; expire (40+ day)
                3600 )     ; minimum ( 1 hr)
```



# 課堂作業

- Create a new zone “[your ID].com” in your DNS server!
  1. In /etc/bind/named.conf.local, add  
zone “[your ID].com” {  
    type master;  
    file “/etc/bind/db.[your ID].com”;  
};
  2. Copy /etc/bind/db.local to /etc/bind/db.[your ID].com  
And modify the record.
  3. Edit /etc/bind/db.[your ID].com and add a few entries.  
(try adding an A record and test it with dig)  
(remember to do “service bind9 reload”)

# 課堂作業

- Show to one of the TAs the following queries are successful: (please check with dig)
  1. `www.[your ID].com` is mapped to multiple IPs (for load-balancing, for example)
  2. Map `www-csie.[your ID].com` to `www.csie.ntu.edu.tw` (not its IP)
  3. Secure your DNS and make sure that only the IPs on your list can use your DNS (as a non-authoritative / local DNS)