DHCP & NAT

tsaimh (2024-2025, CC BY-SA)
wangth (2018-2023, CC BY-SA)
? (2009-2017)

國立陽明交通大學資工系資訊中心
Information Technology Center, Department of Computer Science, NYCU

DHCP – Dynamic Host Configuration Protocol

BOOTP (Bootstrap Protocol)

- BOOTP (Bootstrap Protocol) was originally defined in RFC 951 in 1985, as a replacement for RARP (defined in RFC 903 in 1984).
- BOOTP was used to automatically assign an IP address to network devices from a configuration server.
- To allow network booting, BOOTP has also been used for Unix-like diskless workstations to obtain the network location of their boot image.
- Unlike RARP defined in link layer, BOOTP uses UDP (port 67 by server and port 68 by client), which allows that one central BOOTP server could serve hosts on many subnets.
- BOOTP supports IPv4 only.
- Problems of BOOTP
  - BOOTP was limited to static IP assignment, lacking dynamic allocation capabilities.
  - BOOTP lacks the lease mechanism essential for efficient IP utilization.

DHCP Introduction

- An increasing set of BOOTP vendor information extensions was defined to supply relevant information about the network, like default gateway, name server IP address, the domain name, etc.
- The BOOTP vendor information extensions were incorporated as DHCP option fields, such that DHCP server can also serve BOOTP clients (kind of backward compatible).
- DHCP was first introduced in RFC 1541 in 1993, which was obsoleted by RFC 2131 in 1997. (DHCP Options are defined in RFC 2132, obsoleting RFC 1533)
- DHCPv6 was first introduced in RFC 3315 in 2003, which was obsoleted by RFC 8415 in 2018.

Client-Server Model

- A DHCP client typically queries the information immediately after booting, and periodically thereafter before the expiration of the information.
- Any DHCP server on the network may service the request.
- On large networks that consist of multiple links, a single DHCP server may service the entire network when aided by DHCP relay agents located on the interconnecting routers.

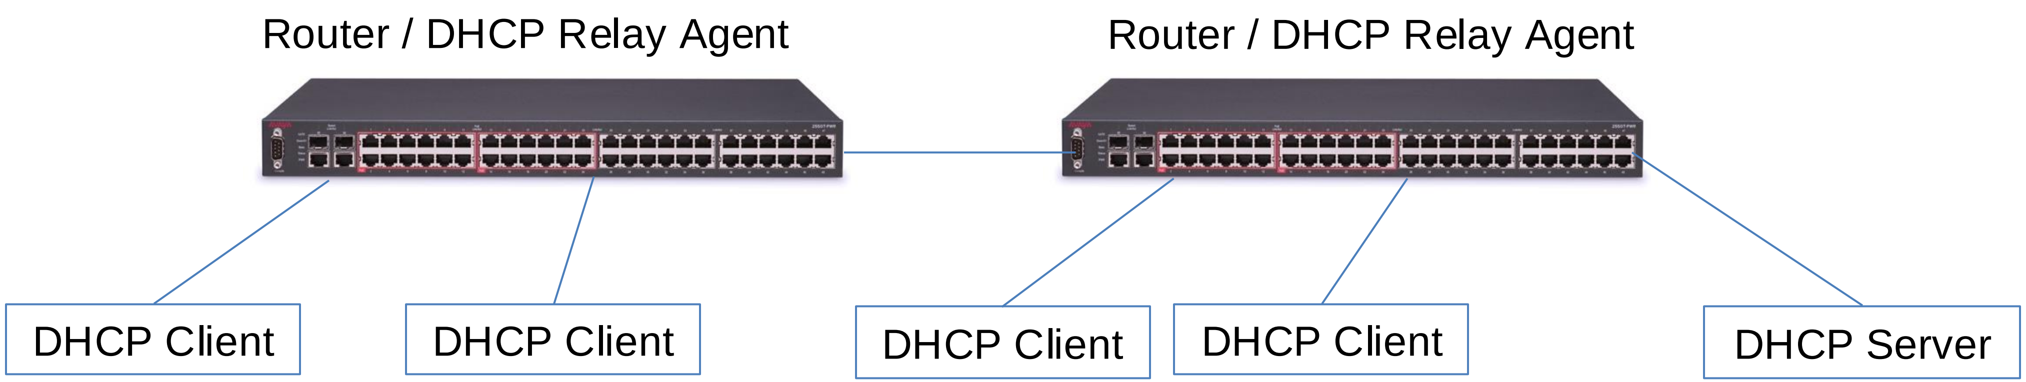

DHCP Address Assignment

- Address allocation mechanisms
  - Dynamic allocation (with lease time)
    - The request-and-grant process uses a lease concept with a controllable time period, allowing the DHCP server to reclaim and then reallocate IP addresses that are not renewed.
  - Automatic allocation (without lease time)
    - Like dynamic allocation, but the DHCP server keeps a table of past IP address assignments, so that it can preferentially assign to a client the same IP address that the client previously had.
  - Manual allocation (compatible with bootp)
    - Each IP address is pre-allocated to a single device.

Dynamic allocation

- Benefits for dynamic allocation
  - Automation
    - No intervention for an administrator
  - Centralized management
    - An administrator can easily look to see which devices are using which addresses
  - Address reuse and sharing
  - Portability and universality
    - Do NOT require DHCP server know the identity of each client
    - Support mobile devices
  - Conflict avoidance

DHCP Leases

- Dynamic address allocation is by far the most popular
  - Hosts are said to “lease” an address instead of “own” one
- Lease Time Management
  - Short Lease (e.g., 1 hour - 1 day):
    - Suitable for mobile devices or frequently changing network environments.
    - Allows quick reallocation of IPs when devices disconnect.
  - Long Lease (e.g., several days - permanent):
    - Used for static workstations, servers, and infrastructure devices.
    - Reduces frequent renewal traffic but can lead to inefficient IP utilization if devices leave the network.

DHCP Protocol – The DORA model

- DHCP Discover
  - The client sends a broadcast request to find available DHCP servers.
- DHCP Offer
  - DHCP servers respond with an available IP and other network settings.
- DHCP Request
  - The client selects one of the offered IPs and requests to use it.
- DHCP Acknowledge
  - The server confirms and officially assigns the IP lease.

※ Question
  - Why not use the IP after DHCP offer?

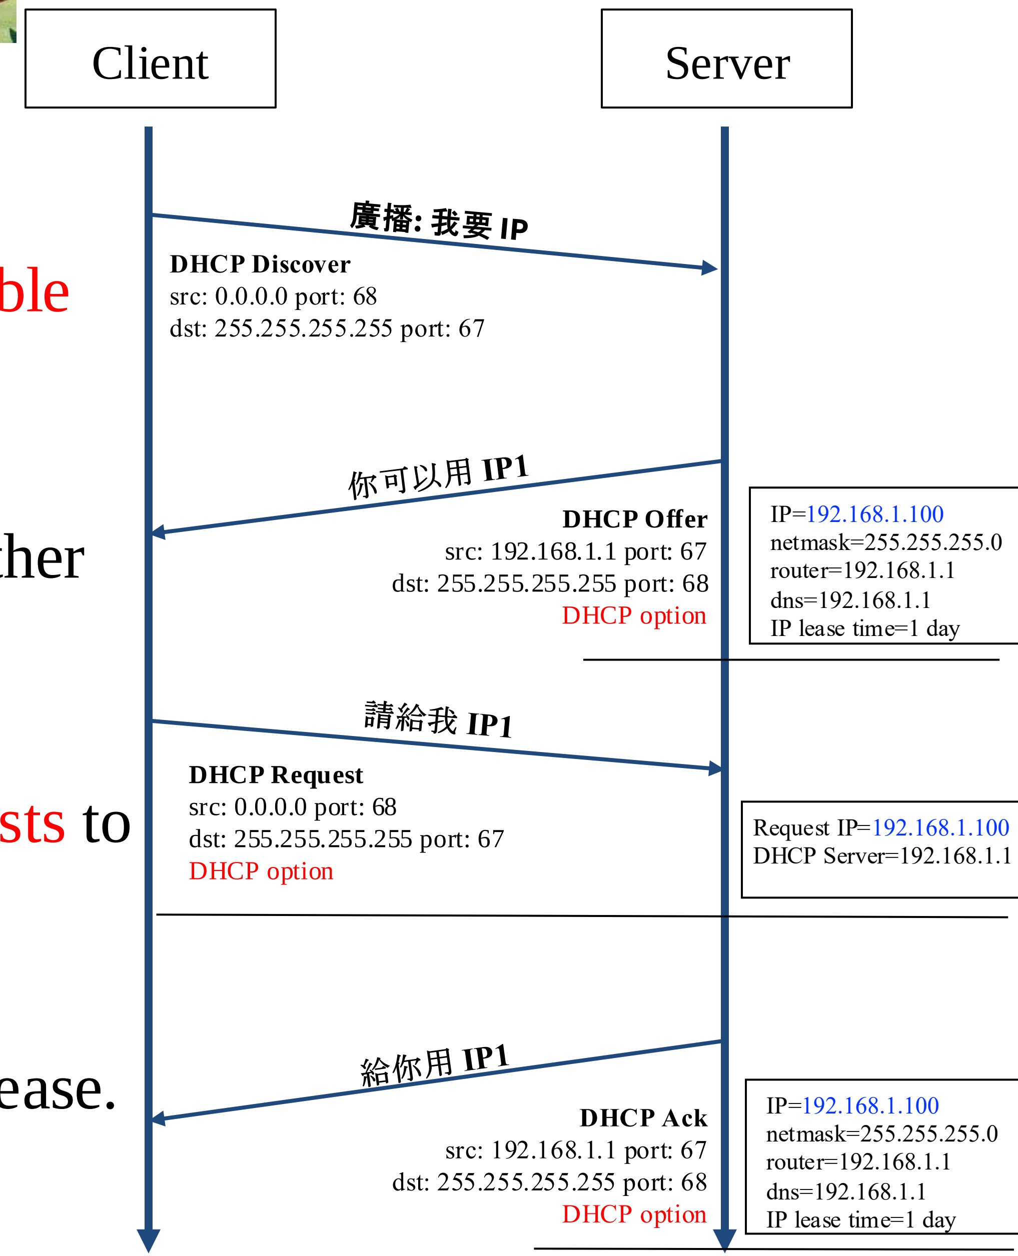

DHCP Protocol – Inform and Release

- DHCP Inform
  - Request more information than the server sent
  - Repeat data for a particular application
    - ex. browsers request web proxy settings from server
  - It does not refresh the IP expiry time in server's database
- DHCP Release
  - Client send this request to server to release the IP, and the client will un-configure this IP
  - Not mandatory

DHCP Protocol – dhclient

- You can use dhclient -v to observe DHCP behavior.

```
wilicw@switch ~> sudo dhclient -v
Internet Systems Consortium DHCP Client 4.4.1
Copyright 2004-2018 Internet Systems Consortium.
All rights reserved.
For info, please visit https://www.isc.org/software/dhcp/

Listening on LPF/veth4d659537/be:41:28:54:55:77
Sending on   LPF/veth4d659537/be:41:28:54:55:77
Listening on LPF/lxdbr0/00:16:3e:32:40:b6
Sending on   LPF/lxdbr0/00:16:3e:32:40:b6
Listening on LPF/docker0/02:42:e2:dc:e4:0e
Sending on   LPF/docker0/02:42:e2:dc:e4:0e
Listening on LPF/br-a8b53c9721b7/02:42:d2:70:41:09
Sending on   LPF/br-a8b53c9721b7/02:42:d2:70:41:09
Listening on LPF/LAN/f6:77:b2:cc:d0:3a
Sending on   LPF/LAN/f6:77:b2:cc:d0:3a
Listening on LPF/ens160/00:0c:29:46:98:06
Sending on   LPF/ens160/00:0c:29:46:98:06
Sending on   Socket/fallback
DHCPDISCOVER on veth4d659537 to 255.255.255.255 port 67 interval 3 (xid=0xa32aa863)
DHCPDISCOVER on lxdbr0 to 255.255.255.255 port 67 interval 3 (xid=0xb4cc7872)
DHCPDISCOVER on docker0 to 255.255.255.255 port 67 interval 3 (xid=0xccd40263)
DHCPDISCOVER on br-a8b53c9721b7 to 255.255.255.255 port 67 interval 3 (xid=0x766e7e13)
DHCPDISCOVER on LAN to 255.255.255.255 port 67 interval 3 (xid=0xd88773)
DHCPREQUEST for 172.16.249.131 on ens160 to 255.255.255.255 port 67 (xid=0x3583c91d)
DHCPACK of 172.16.249.131 from 172.16.249.254 (xid=0x1dc98335)
RTNETLINK answers: File exists
bound to 172.16.249.131 -- renewal in 841 seconds.
```

DHCP Lease Address Pools

- Each DHCP server maintains a set of IP addresses
  - Used to allocate leases to clients
    - Most of clients are equals
      - A range of addresses is normally handled as a single group defined for a particular network

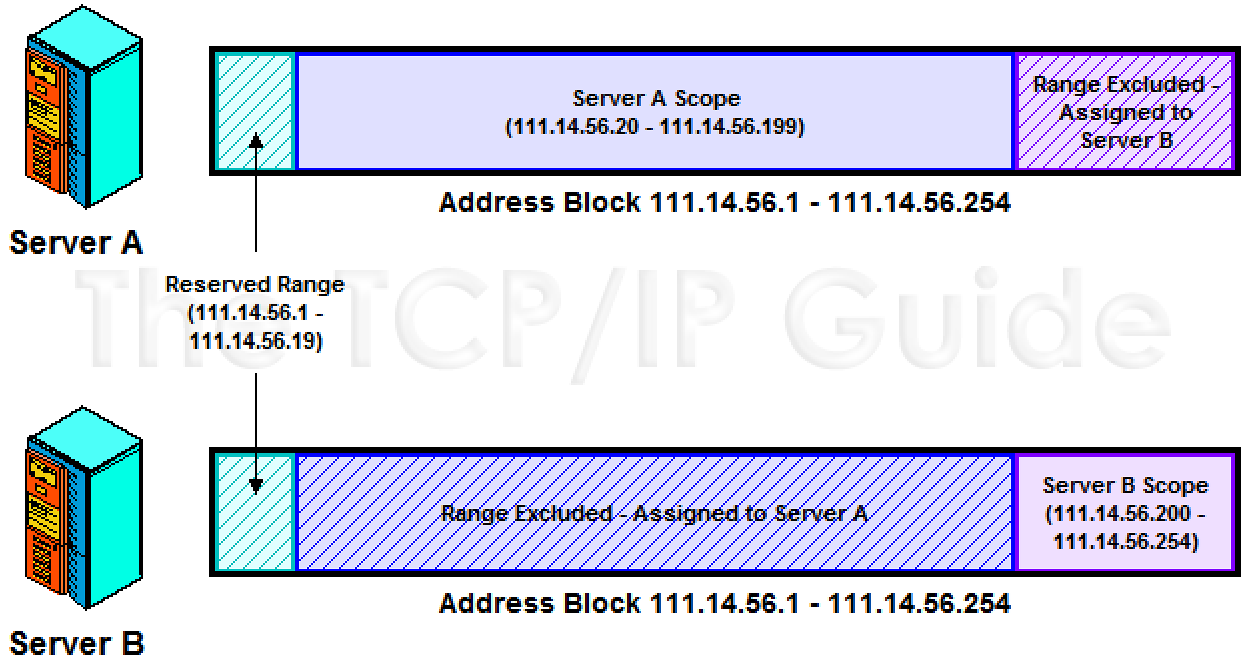

[The TCP/IP Guide - DHCP Lease Address Pools, Ranges (Scopes) and Address Management]

DHCP Lease “Life Cycle” (1)

- Allocation
  - A client without active lease acquires a lease through allocation.
- Reallocation
  - When a client reboots, it asks the DHCP to confirm the lease and acquires operating parameters.
- Normal operation
  - Once a lease is active, the client functions normally.
- Renewal
  - After a certain portion of the lease time has expired, the client attempts to contact the server to renew the lease.
- Rebinding
  - If the server did not response to renewal, the client tries to rebind to any active DHCP server.
- Release
  - The client may decide at any time that it no longer wishes to use the IP address.

Ref: [DHCP Lease "Life Cycle" Overview (Allocation, Reallocation, Renewal, Rebinding and Release) and Lease Timers]

DHCP Lease "Life Cycle" (2)

- The client unicasts the request at Day 5 to renew the lease.

  (not shown in the figure)

- In the rebinding period, the client broadcasts its request instead of unicasting.
  (the request message will reach all available DHCP servers)

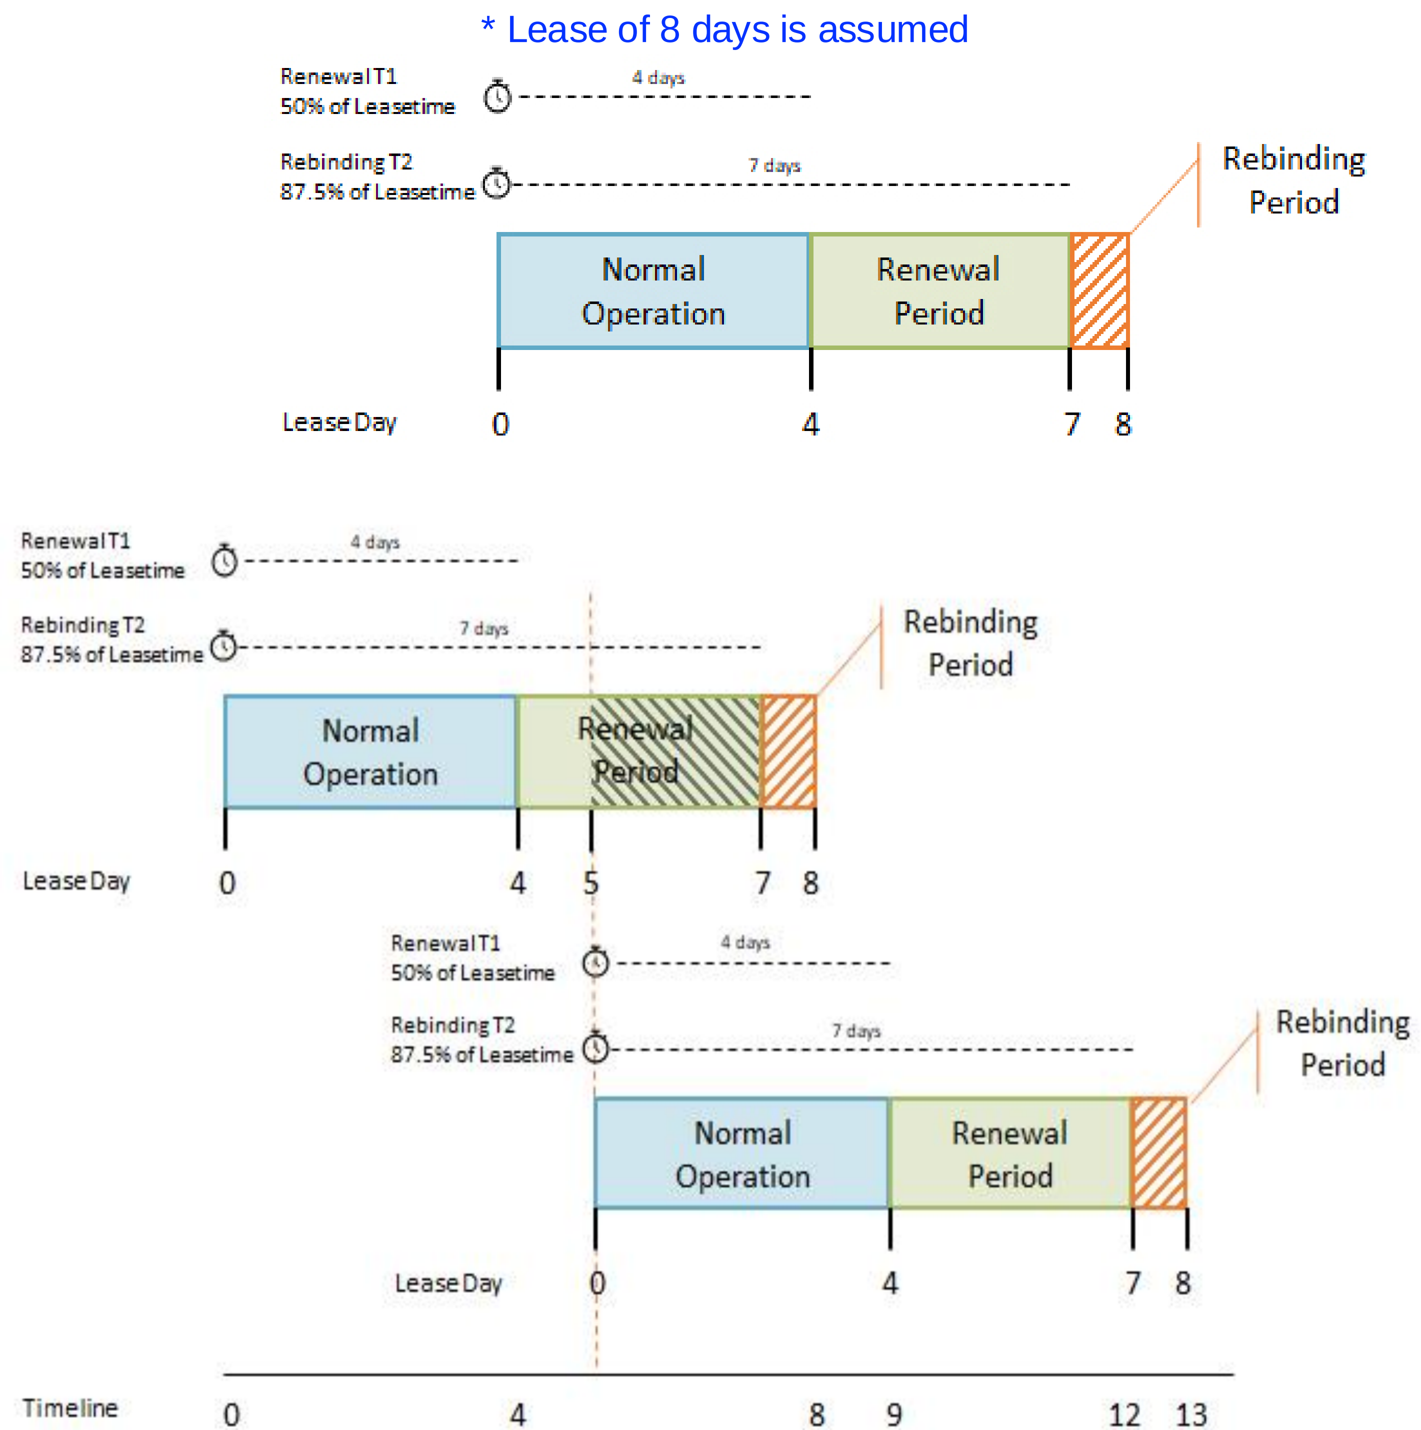

[DHCP Lease Time - What is it and How does it work? — LazyAdmin]

ISC DHCP Server (1)

- The ISC DHCP software system was originally written for Internet Systems Consortium by Ted Lemon and Vixie Enterprises.
- https://www.isc.org/dhcp/
- FreeBSD
  - /usr/ports/net/isc-dhcp44-server/
  - pkg install isc-dhcp44-server
- Linux
  - apt install isc-dhcp-server

ISC DHCP Server (2)

- In 2019, ISC released a [Kea Migration utility](), essentially a modified version of ISC DHCP written by Francis Dupont
- Helps users migrate a configuration file from ISC DHCP to Kea by translating the common elements to Kea configuration syntax.

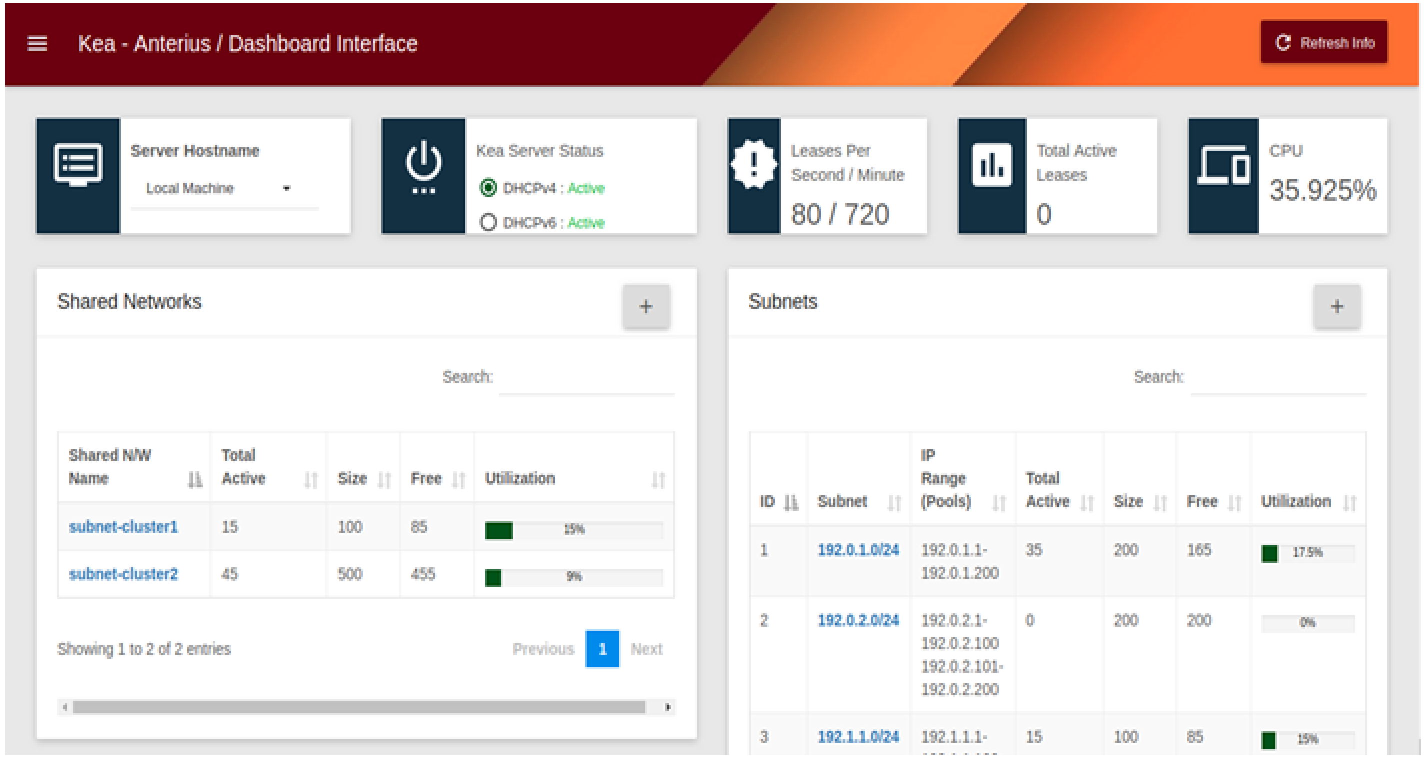

DHCP Server on FreeBSD (1)

- Kernel support
  - device bpf
- Install DHCP server
  - /usr/ports/net/isc-dhcp44-server/
  - pkg install isc-dhcp44-server
- Enable DHCP server in /etc/rc.conf
  - dhcpd_enable="YES"
  - dhcpd_flags="-q"
  - dhcpd_conf="/usr/local/etc/dhcpd.conf"
  - dhcpd_ifaces=""
  - dhcpd_withumask="022"

DHCP Server on FreeBSD (2)

Three-way handshake

- Option definitions

```
option domain-name "cs.nycu.edu.tw";
option domain-name-servers 140.113.235.107, 140.113.1.1;

default-lease-time 600;
max-lease-time 7200;
ddns-update-style none;
log-facility local7;
```

/etc/syslogd.conf
/etc/newsyslog.conf

DHCP Server on FreeBSD (3)

- Subnet definition

```
subnet 192.168.1.0 netmask 255.255.255.0 {
        range 192.168.1.101 192.168.1.200;
        option domain-name "cs.nycu.edu.tw";
        option routers 192.168.1.254;
        option broadcast-address 192.168.1.255;
        option domain-name-servers 140.113.17.5, 140.113.1.1;
        default-lease-time 3600;
        max-lease-time 21600;
}
```

- Host definition

```
host fantasia {
        hardware ethernet 08:00:07:26:c0:a5;
        fixed-address 192.168.1.30;
}
host denyClient {
        hardware ethernet 00:07:95:fd:12:13;
        deny booting;
}
```

DHCP Server on FreeBSD (4)

- Important files
  - /usr/local/sbin/dhcpd
  - /usr/local/etc/dhcpd.conf
  - /var/db/dhcpd.leases (leases issued)
  - /usr/local/etc/rc.d/isc-dhcpd

NAT –
Network Address Translation

Ref: https://en.wikipedia.org/wiki/IPv4_address_exhaustion
https://en.wikipedia.org/wiki/List_of_countries_by_IPv4_address_allocation

IPv4 Address Crisis

- Because the original Internet architecture had fewer than 4.3 billion addresses available, depletion has been anticipated since the late 1980s when the Internet started experiencing dramatic growth.
- The anticipated shortage has been the driving factor in creating and adopting several new technologies, including Classless Inter-Domain Routing (CIDR) in 1993 (RFC 1518), network address translation (NAT) in 1994 (RFC 1631), and IPv6 in 1998 (RFC 2460).

Network Address Translation (NAT)

- Network Address Translation (NAT) is a method of mapping an IP address space into another by modifying network address information in the IP header of packets while they are in transit across a routing device.
- NAT is introduced in RFC 1631 in May 1994 as a ”short-term solution” to the two most compelling problems at that time: IP address depletion and scaling in routing.
- NAT is also known as IP masquerading, which hides an entire IP address space (usually private IP addresses) behind a single IP address (usually public address).
- In 1999, RFC 2663 introduced Network Address and Port Translation (NAPT), which expanded the translation of addresses to include port numbers.

Private Address Space

- Private addresses are not allocated to any specific organization, such that anyone may use these addresses without approval.
- Private addresses are often seen as enhancing network security for the internal network since use of private addresses internally makes it difficult for an external host to initiate a connection to an internal system.
- Private addresses space defined by RFC1918
    - 24-bit block (Class A): 10.0.0.0
    - 20-bit block (16 Class B): 172.16.0.0 ~ 172.31.255.255
    - 16-bit block (256 Class C): 192.168.0.0 ~ 192.168.255.255

Carrier-Grade NAT

Ref: https://en.wikipedia.org/wiki/Carrier-grade_NAT

Ref: https://datatracker.ietf.org/doc/html/rfc6598

- In April 2012, IANA allocated the 100.64.0.0/10 block of IPv4 addresses specifically for use in carrier-grade NAT scenarios.

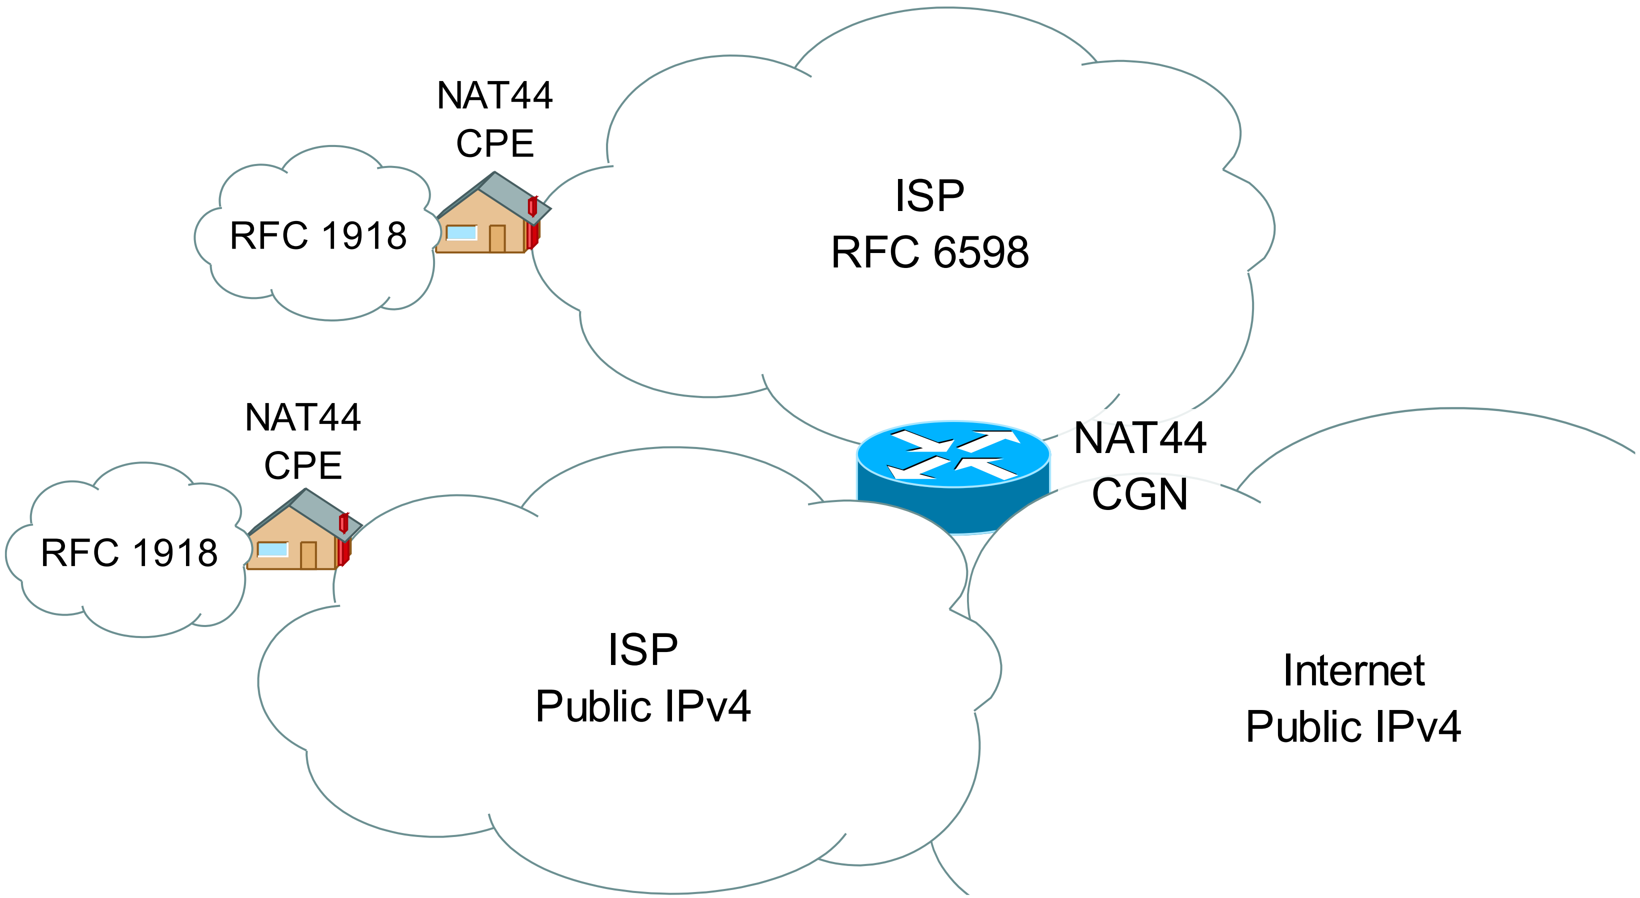

Types of NAT

- Source NAT vs. Destination NAT
- Uni-directional NAT vs. Bi-directional NAT
- Full-Cone vs. Restricted Cone vs. Port-Restricted Cone vs. Symmetric
- …

Source NAT (SNAT)

- Rewrite the source IP and/or Port.
- The rewritten packet looks like one sent by the NAT server.

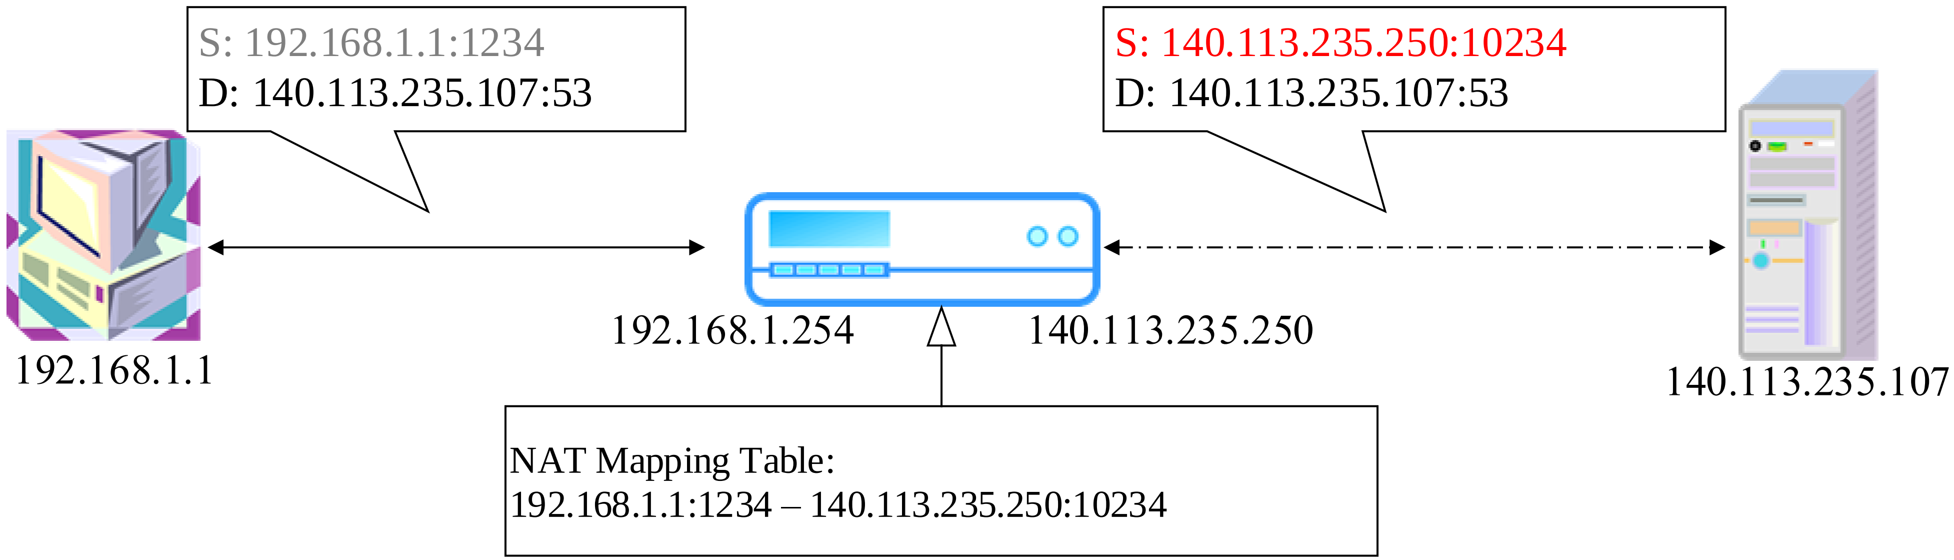

Destination NAT (DNAT)

- DNAT is commonly used to publish a service located in a private network on a publicly accessible IP address.
  - Rewrite the destination IP and/or Port.
  - The rewritten packet will be redirected to another IP address when it pass through NAT server.

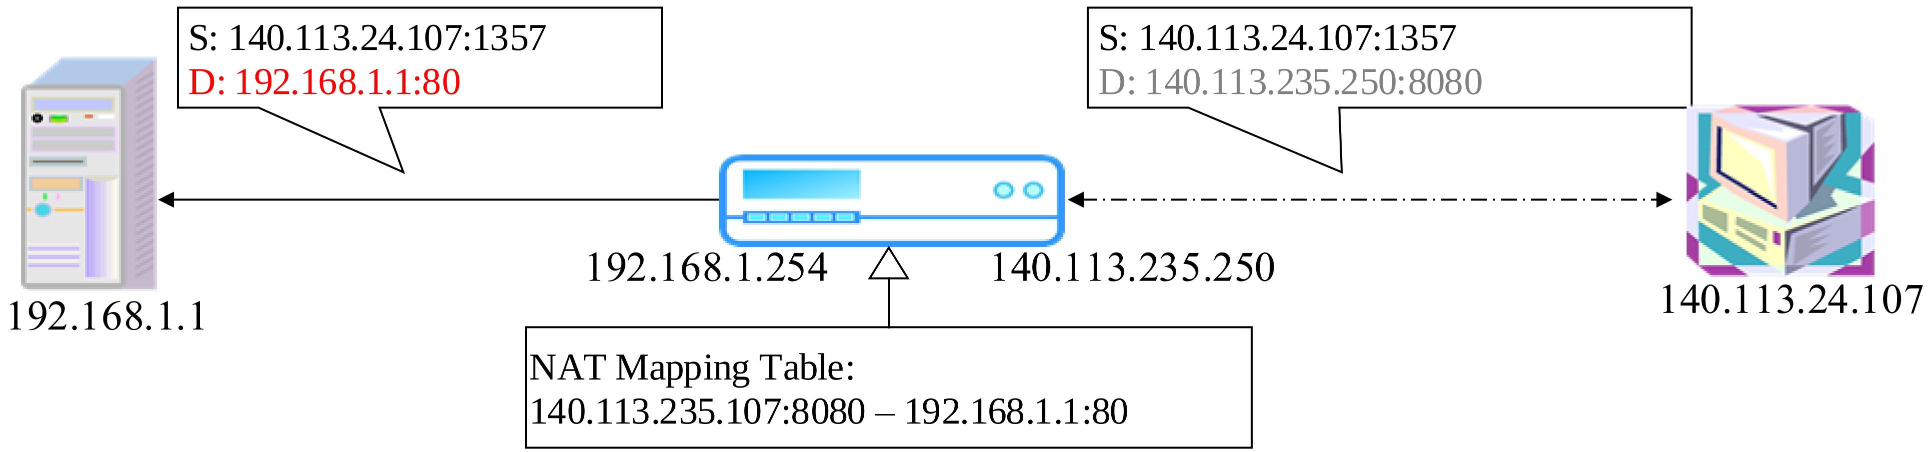

- Both SNAT and DNAT are usually used together in coordination for two-way communication (bi-directional NAT).

NAT Unidirectional Operation

- NAT Unidirectional Operation
  - Traditional/Outbound operation
  - The original variety of NAT in RFC 1631
    - The simplest NAT
    - The client sends request from the inside to outside network

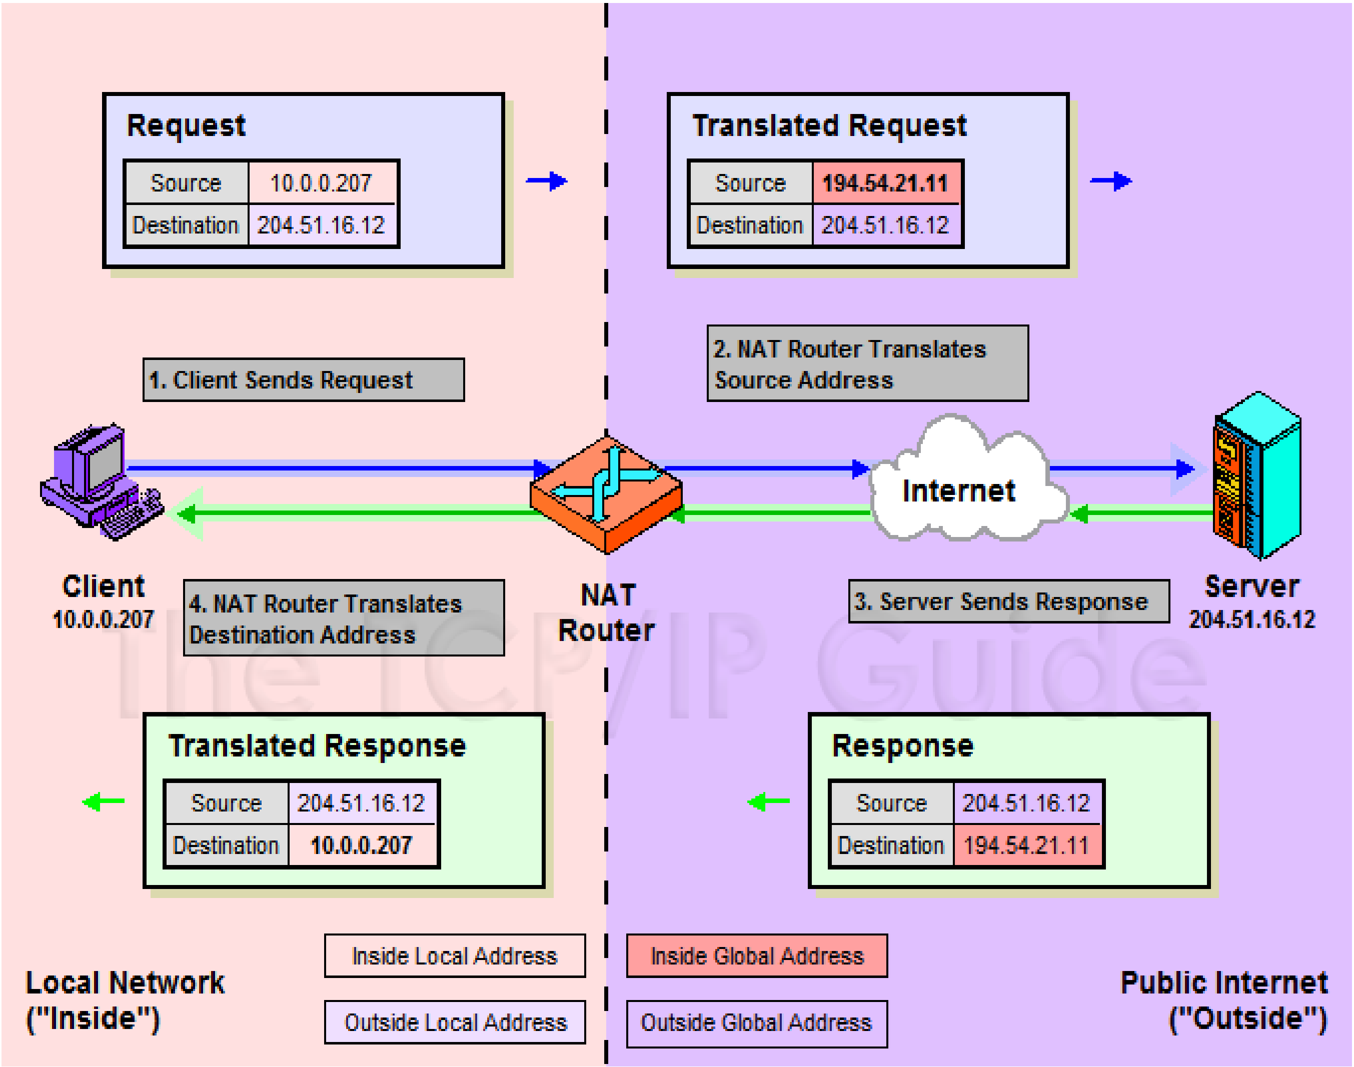

The TCP/IP Guide - IP NAT Unidirectional (Traditional/Outbound) Operation

NAT Bidirectional Operation

- Two-way / Inbound
- SNAT + DNAT

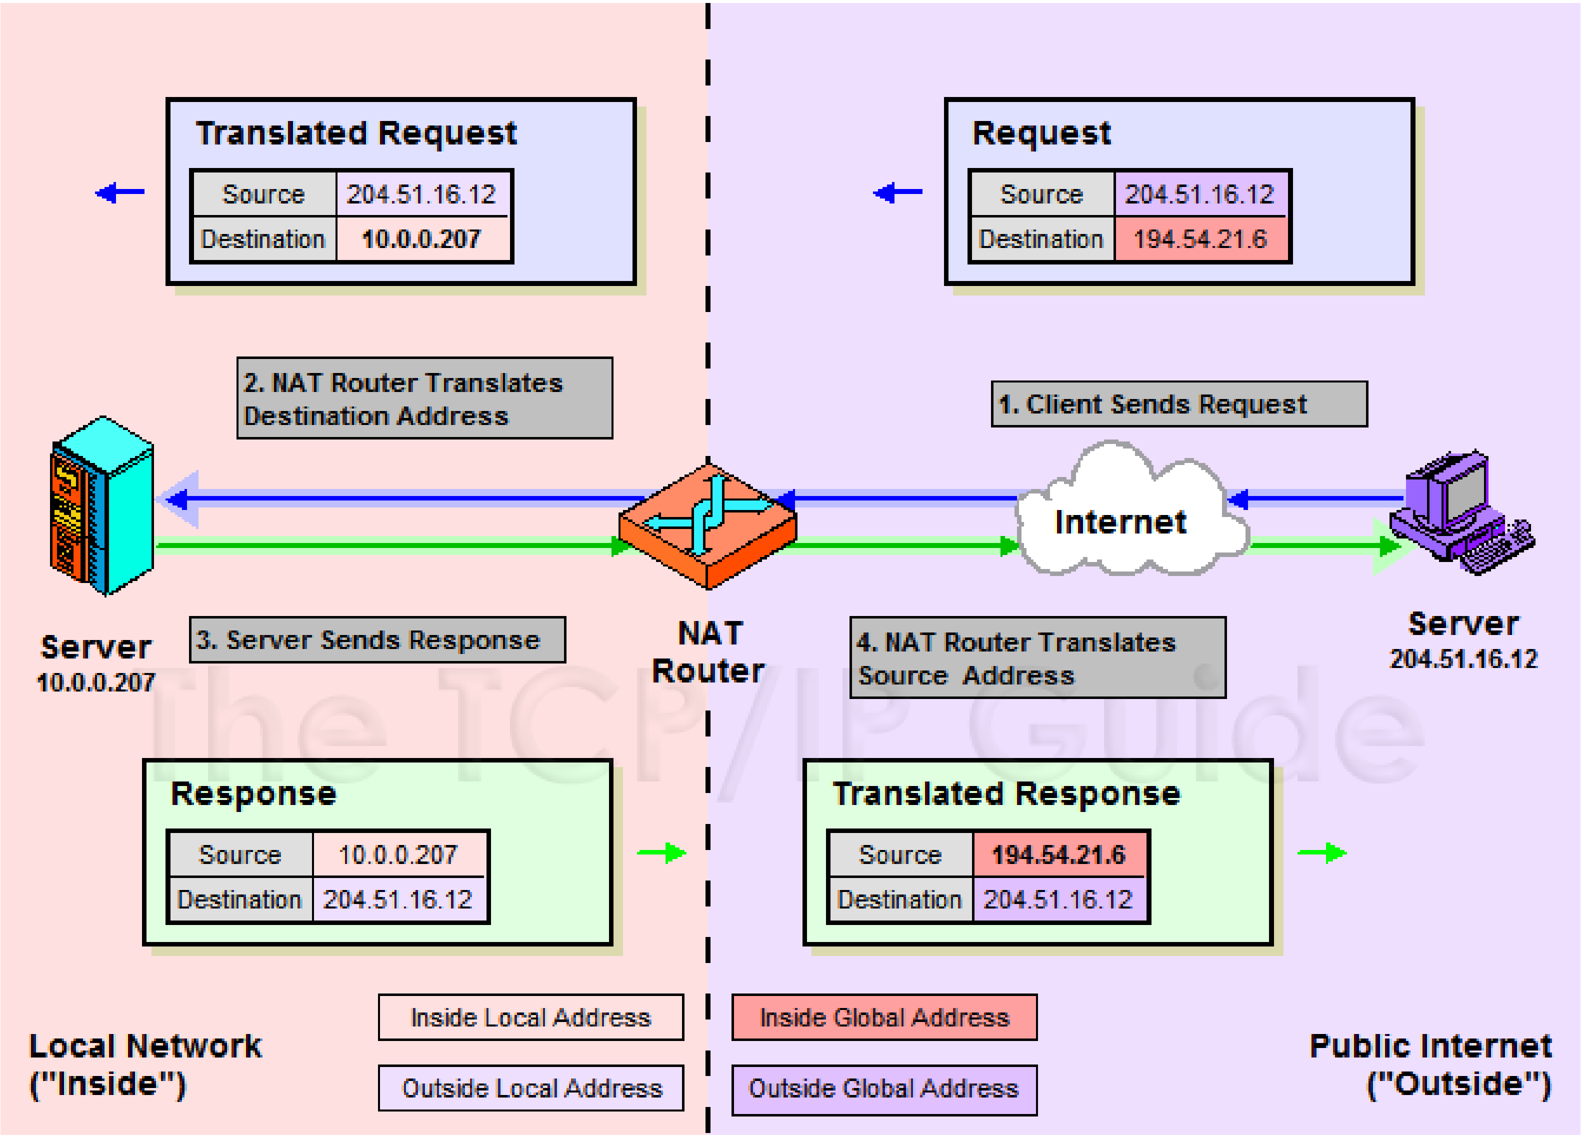

The TCP/IP Guide - IP NAT Bidirectional (Two-Way/Inbound) Operation

NAT Port-Based Operation (1)

- NAT example:

NAT mapping table

Orig	Alias	Remote
192.168.1.1:1029	140.113.235.219:1029	⇔ 140.113.235.72:23
192.168.1.1:1030	140.113.235.219:1030	⇔ 140.113.235.72:23
192.168.1.2:1029	140.113.235.219:30029	⇔ 140.113.235.72:21
192.168.1.2:1030	140.113.235.219:30030	⇔ 140.113.235.72:21

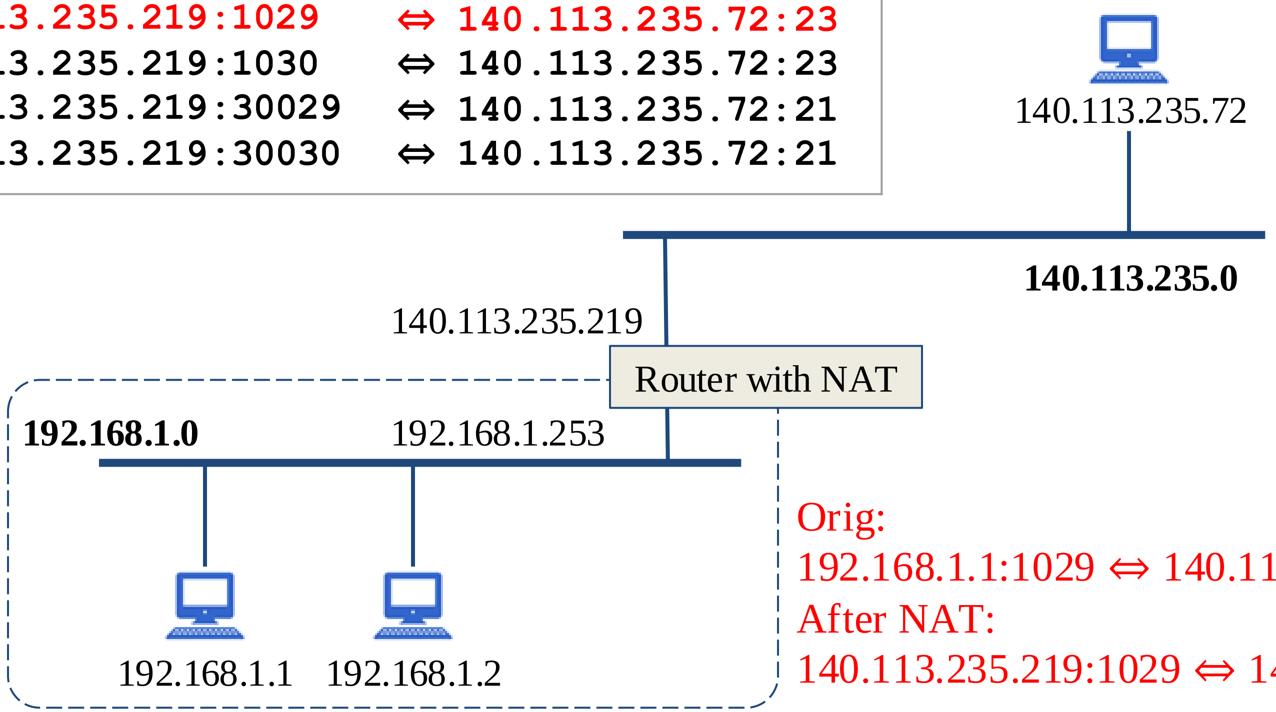

Orig:
192.168.1.1:1029 ⇔ 140.113.235.72:23
After NAT:
140.113.235.219:1029 ⇔ 140.113.235.72:23

NAT Port-Based Operation (2)

```
% ifconfig en0
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
	options=400<CHANNEL_IO>
	ether f0:18:98:5e:a7:b6
	inet6 fe80::14fb:c233:f28e:1c00%en0 prefixlen 64 secured scopeid 0x6
	inet 192.168.0.104 netmask 0xffffff00 broadcast 192.168.0.255
	inet6 fdfe:b150:6c38:699d:8ad:82e7:438c:2e50 prefixlen 64 autoconf secured
	nd6 options=201<PERFORMNUD,DAD>
	media: autoselect
	status: active
% curl ifconfig.me
140.113.210.231
```

NAT Compatibility Issues

- It is NOT possible for NAT to be completely transparent to the hosts that use it
  - ICMP Manipulations
  - Applications that embed IP address
    - FTP
  - Additional issues with port translation
    - The issues applying to addresses now apply to ports as well
  - Problems with IPSec

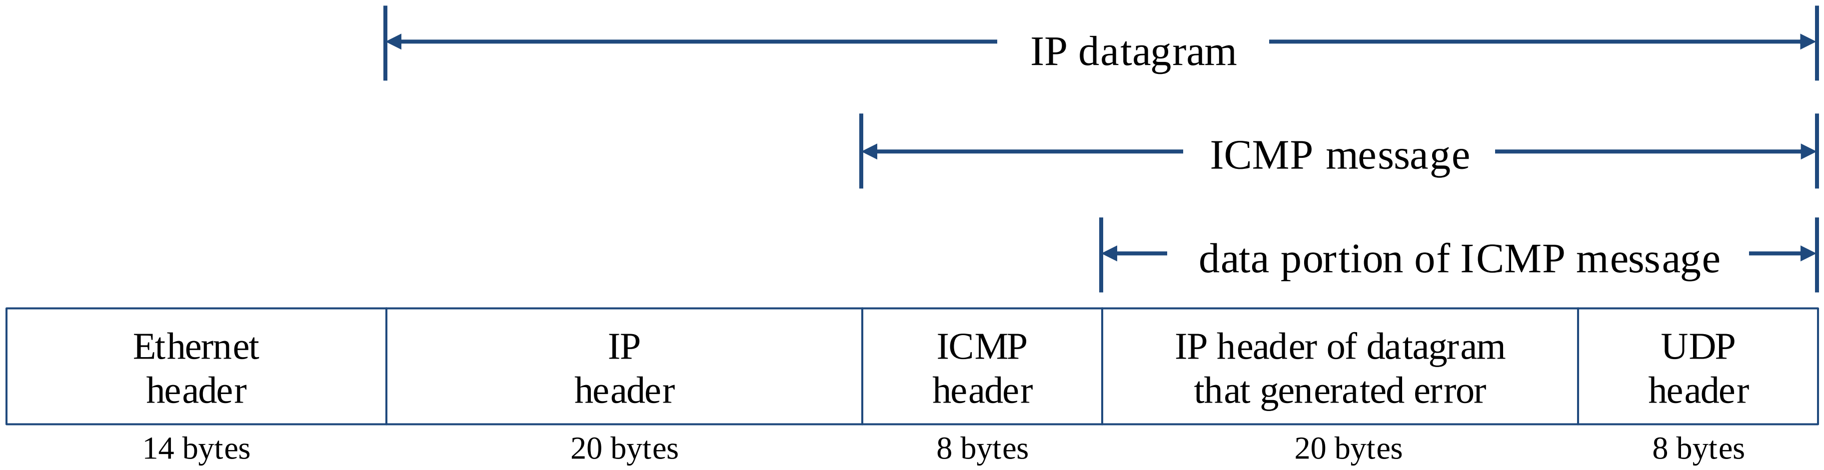

Set up NAT on FreeBSD

- Check IP configuration and enable Packeg Filter (PF) in /etc/rc.conf

```
ifconfig_fxp0="inet 140.113.235.4  netmask 255.255.255.0 media autoselect"
ifconfig_fxp1="inet 192.168.1.254  netmask 255.255.255.0 media autoselect"
defaultrouter="140.113.235.254"
gatewy_enable="YES"
pf_enable="YES"
pflog_enable="YES"
```

- Three types of translation in /etc/pf.conf
  - **nat**: normal uni-directional NAT
  - **rdr** (redirect): redirected to another destination and possibly a different port.
  - **binat** (bi-directional nat): a bidirectional mapping between an external IP netblock and an internal IP netblock.

NAT on FreeBSD – Example 1

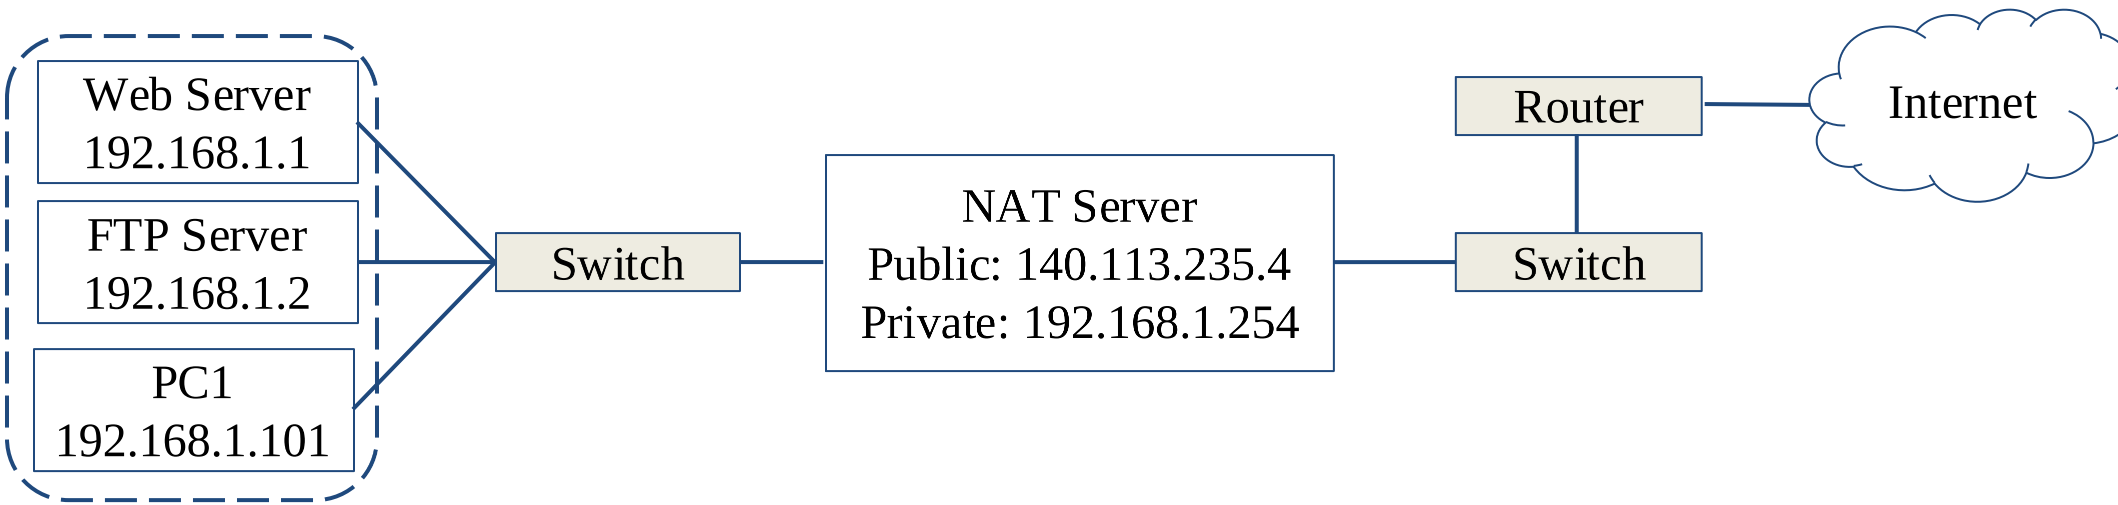

/etc/pf.conf

```
# macro definitions
extdev='fxp0'
intranet='192.168.1.0/24'
webserver='192.168.1.1'
ftpserver='192.168.1.2'
pc1='192.168.1.101'
# nat rules
nat on $extdev inet from $intranet to any -> $extdev
rdr on $extdev inet proto tcp to port 80 -> $webserver port 80
rdr on $extdev inet proto tcp to port 443 -> $webserver port 443
rdr on $extdev inet proto tcp to port 21 -> $ftpserver port 21
```

NAT on FreeBSD – Example 2

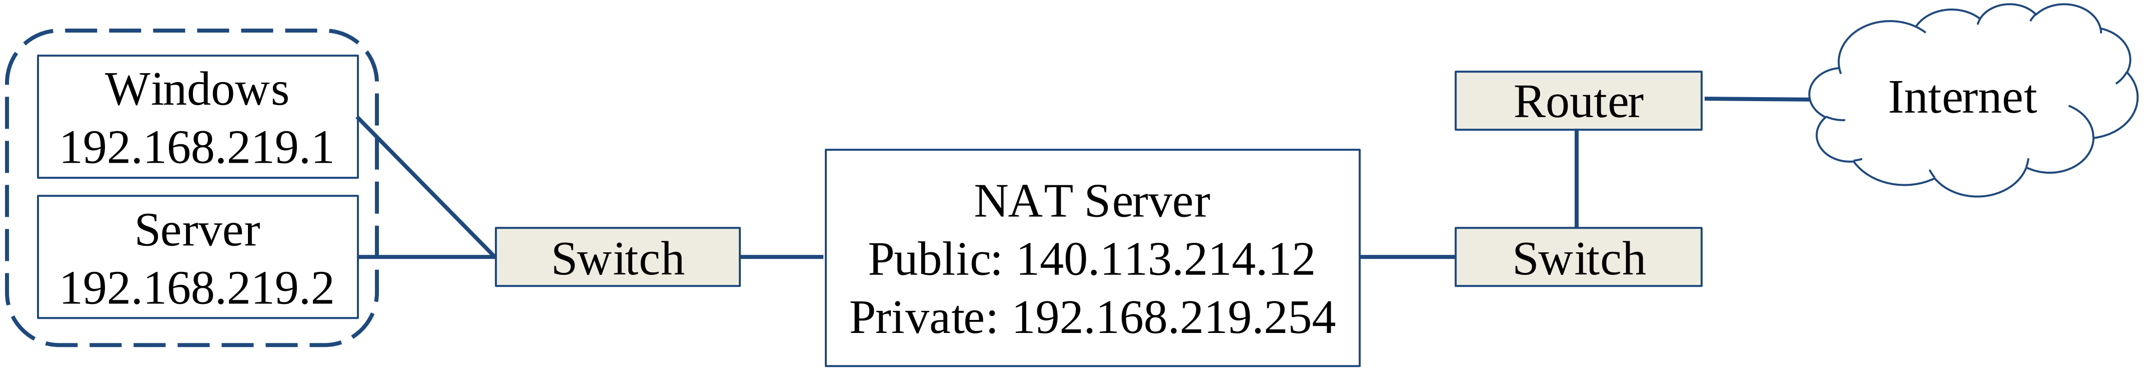


```
# macro definitions
extdev='fxp0'
intranet='192.168.219.0/24'
windows='192.168.219.1'
server_int='192.168.219.2'
server_ext='140.113.214.13'

# nat rules
nat on $extdev inet from $intranet to any -> $extdev
rdr on $extdev inet proto tcp to port 3389 -> $windows port 3389
binat on $extdev inet from $server_int to any -> $server_ext
```

References

- Reverse Address Resolution Protocol (RARP), RFC 903 (1984)
- Bootstrap Protocol (BOOTP), RFC 951 (1985)
- Dynamic Host Configuration Protocol (DHCP), RFC 2131 (1997)
- DHCP Options and BOOTP Vendor Extensions, RFC 2132 (1997) (updated by RFCs 3442, 3942, 4361, 4833 and 5494)

- The IP Network Address Translator (NAT), RFC 1631 (1994)
- IP Network Address Translator (NAT) Terminology and Considerations , RFC 2663 (1999)
- Address Allocation for Private Internets, RFC 1918 (1996)
- IANA-Reserved IPv4 Prefix for Shared Address Space, RFC 6598 (2012)