Kea Webinar
DHCP introduction

Carsten Strotmann

16th September 2020

https://www.isc.org

Welcome

- Welcome to part one of our webinar series "the KEA DHCP Server"

About this Webinar

- short history of DHCP
- DHCPv4
- DHCPv6
- References

About DHCP

- DHCP is short for **Dynamic Host Configuration Protocol**
  - the Internet standard protocol to assign IP-addresses and network related configuration to TCP/IP connected machines
  - for IPv4: RFC 2131
  - for IPv6: RFC 8415

short history of DHCP (including ISC-DHCP and KEA)

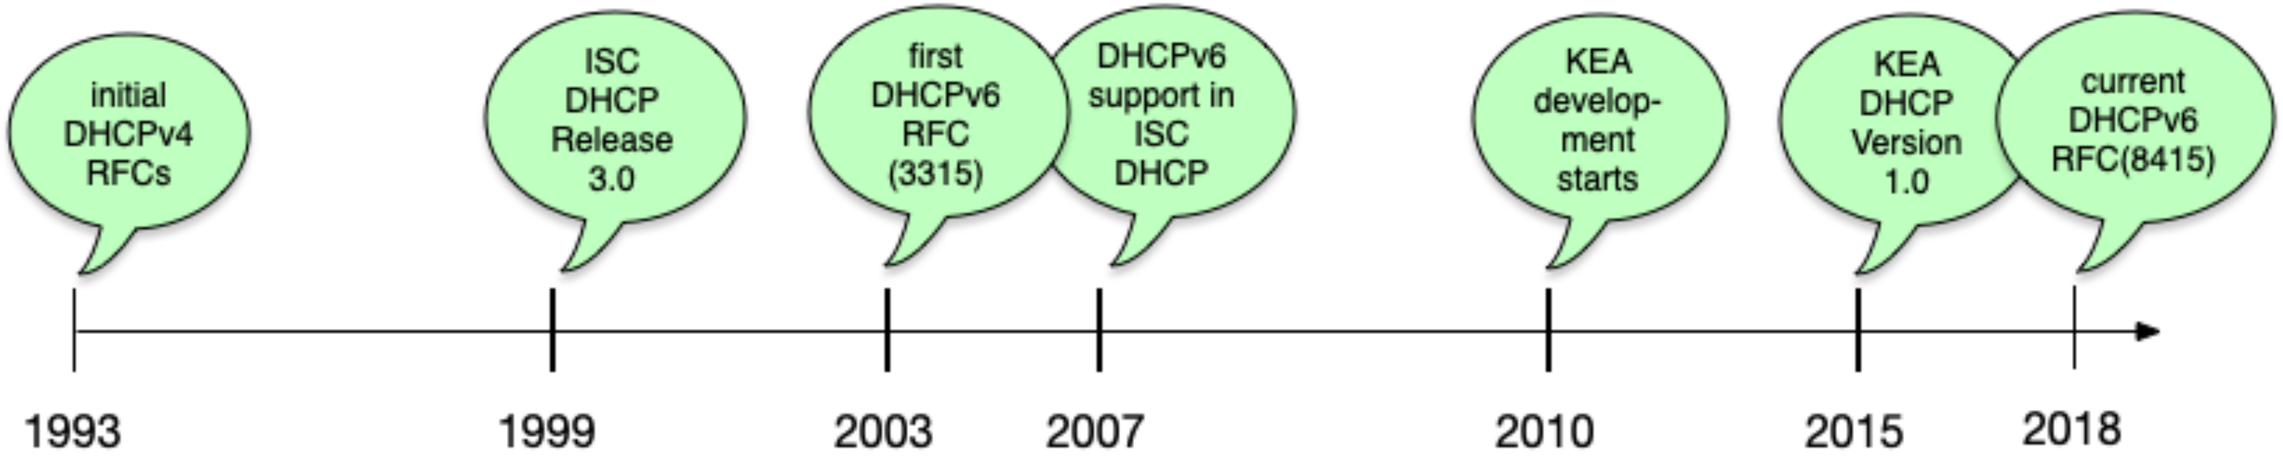

DHCPv4 overview

DHCPv4 Overview (1/2)

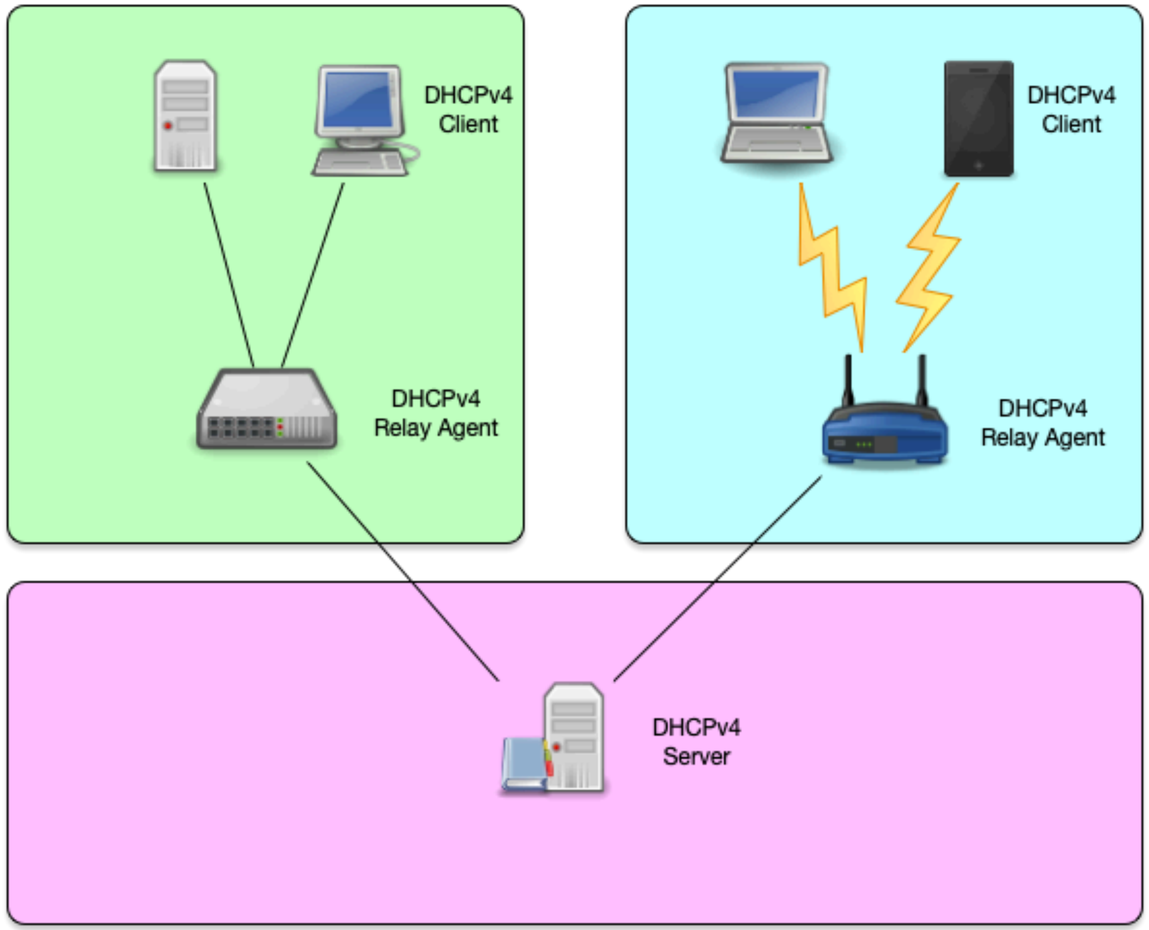

DHCPv4 Overview (2/2)

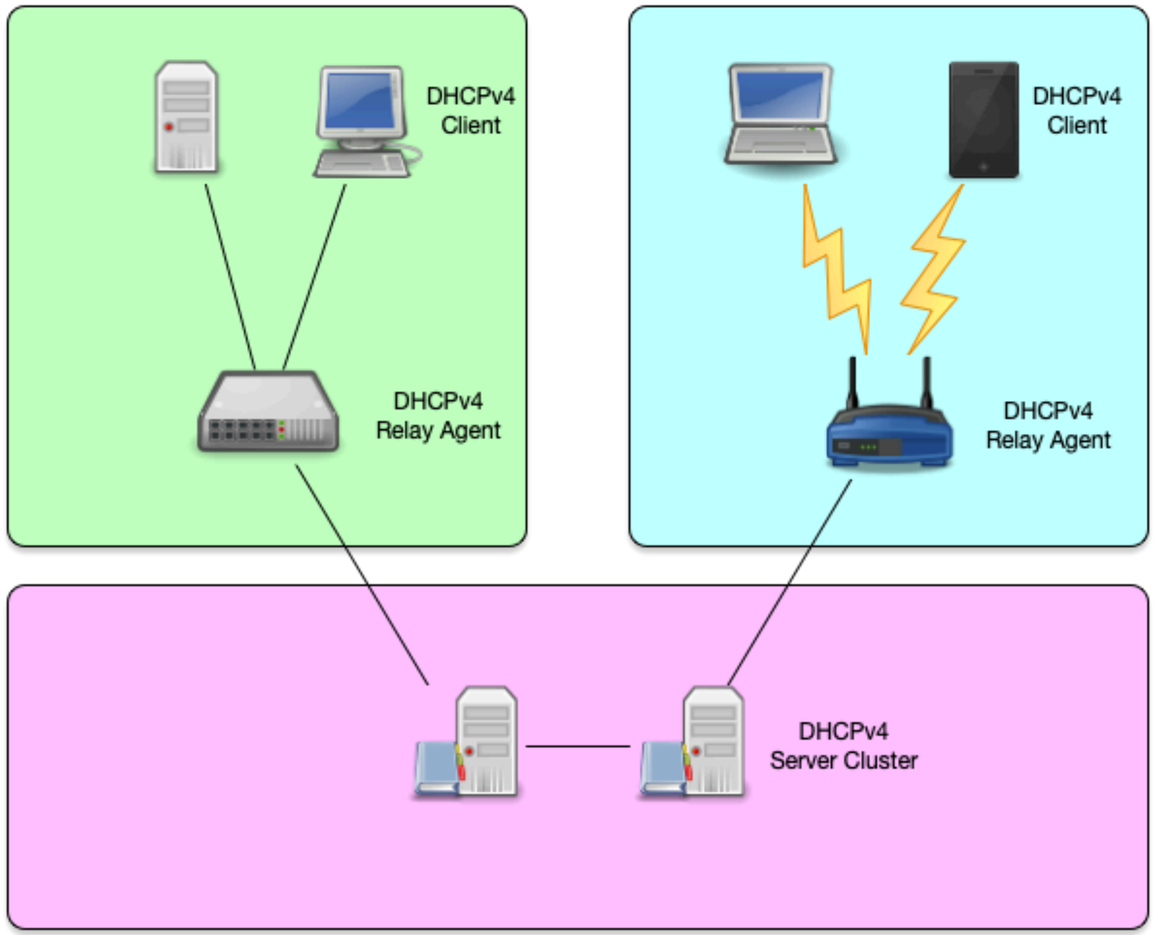

DHCPv4 protocol

- the DHCPv4 protocol uses UDP broadcast und (in some situations) unicast
  - the DHCPv4 server or relay agents listen on port 67
  - a DHCPv4 client listens on port 68 for messages from a server or relay agent
  - the initial request from a client requires layer 2 (Ethernet) communication

DHCPv4 Lease concept

DHCPv4 Lease (1/3)

- when using the DHCP protocol, a client can never keep an IP Address forever
  - each IP Address given out by a DHCP server has a "lease" time
  - this is the time in seconds that the client is allowed to use the IP Address

DHCPv4 Lease (3/3)

- according to the RFC, a DHCP server must store the lease information to permanent storage before giving the IP address to a client
  - this can be a performance bottleneck on a DHCP server
  - we will discuss the various options in one of the following webinars

DHCP Clients, Relays and Server

DHCP Clients, Relays and Server

- because a client machine without IP address can only communicate on the local link, the base DHCPv4 protocol is "link-local" only
- DHCPv4 relay-agents can be used to forward DHCPv4 requests to centralised DHCPv4 server
  - DHCP relay-agents are often found in network equipment (e.g. router)
  - dedicated "software based" relay-agents are available

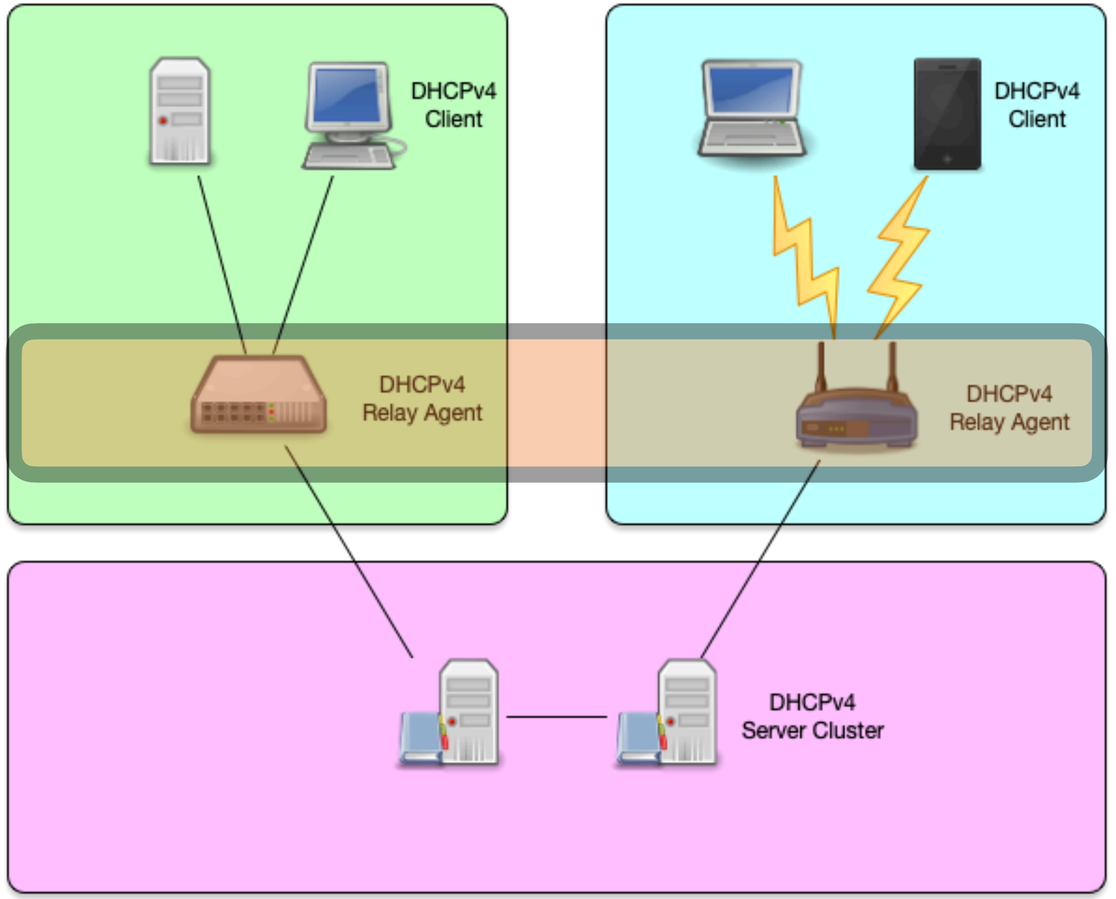

DHCP messages and client server communication

- **DHCPDISCOVER** (client asks: is there a DHCP server that can give me an address)
- **DHCPOFFER** (DHCP server offers an address to the client)
- **DHCPREQUEST** (client requests the IP address offered by the server)
- **DHCPACK** (server marks the IP address as leased and confirms that transaction)

DHCP messages (1)

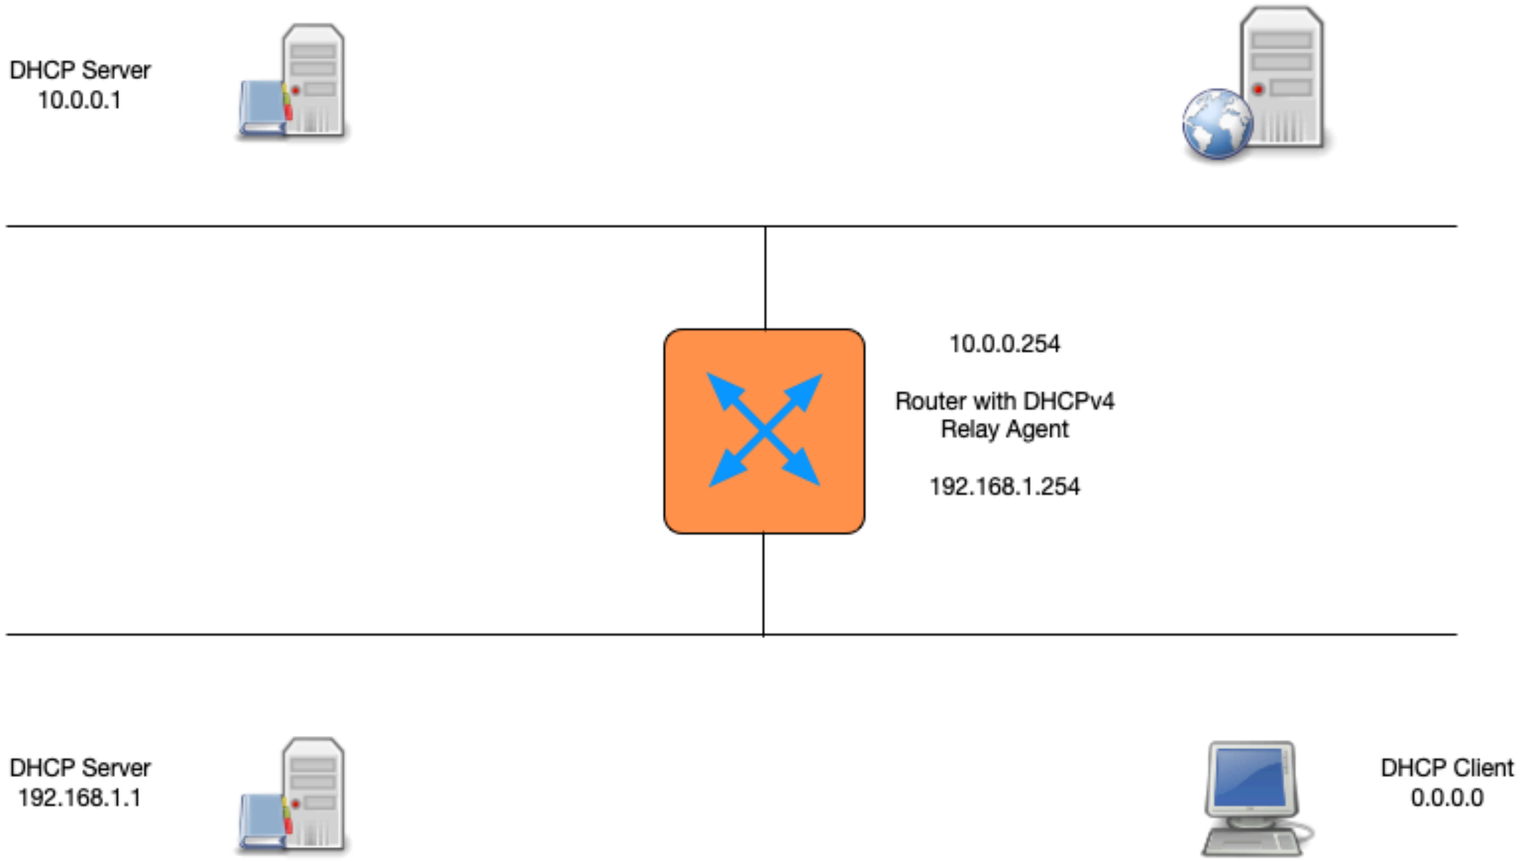

DHCP messages (2)

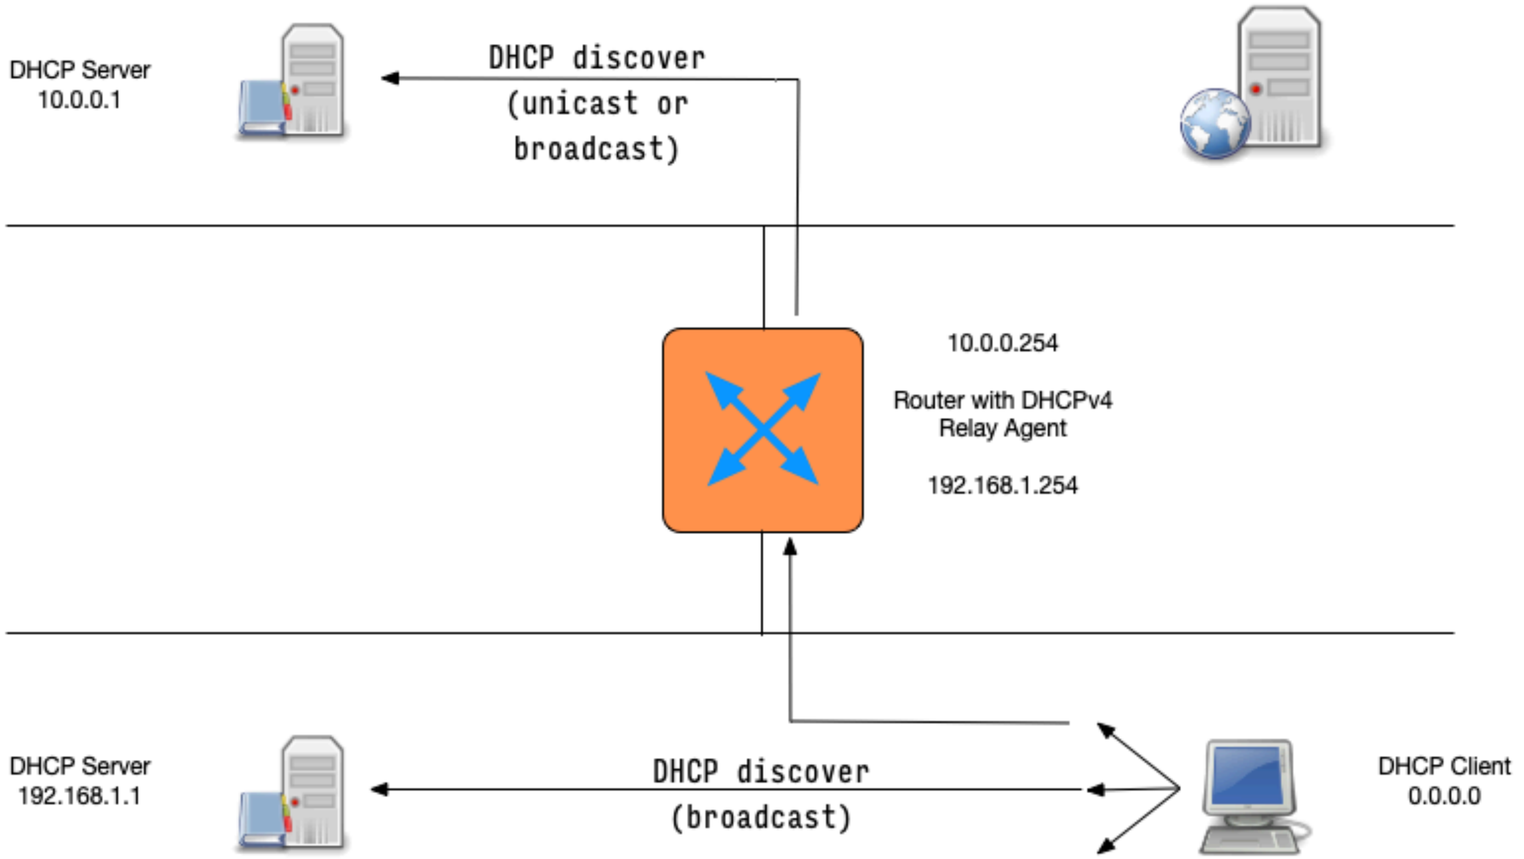

DHCP messages (3)

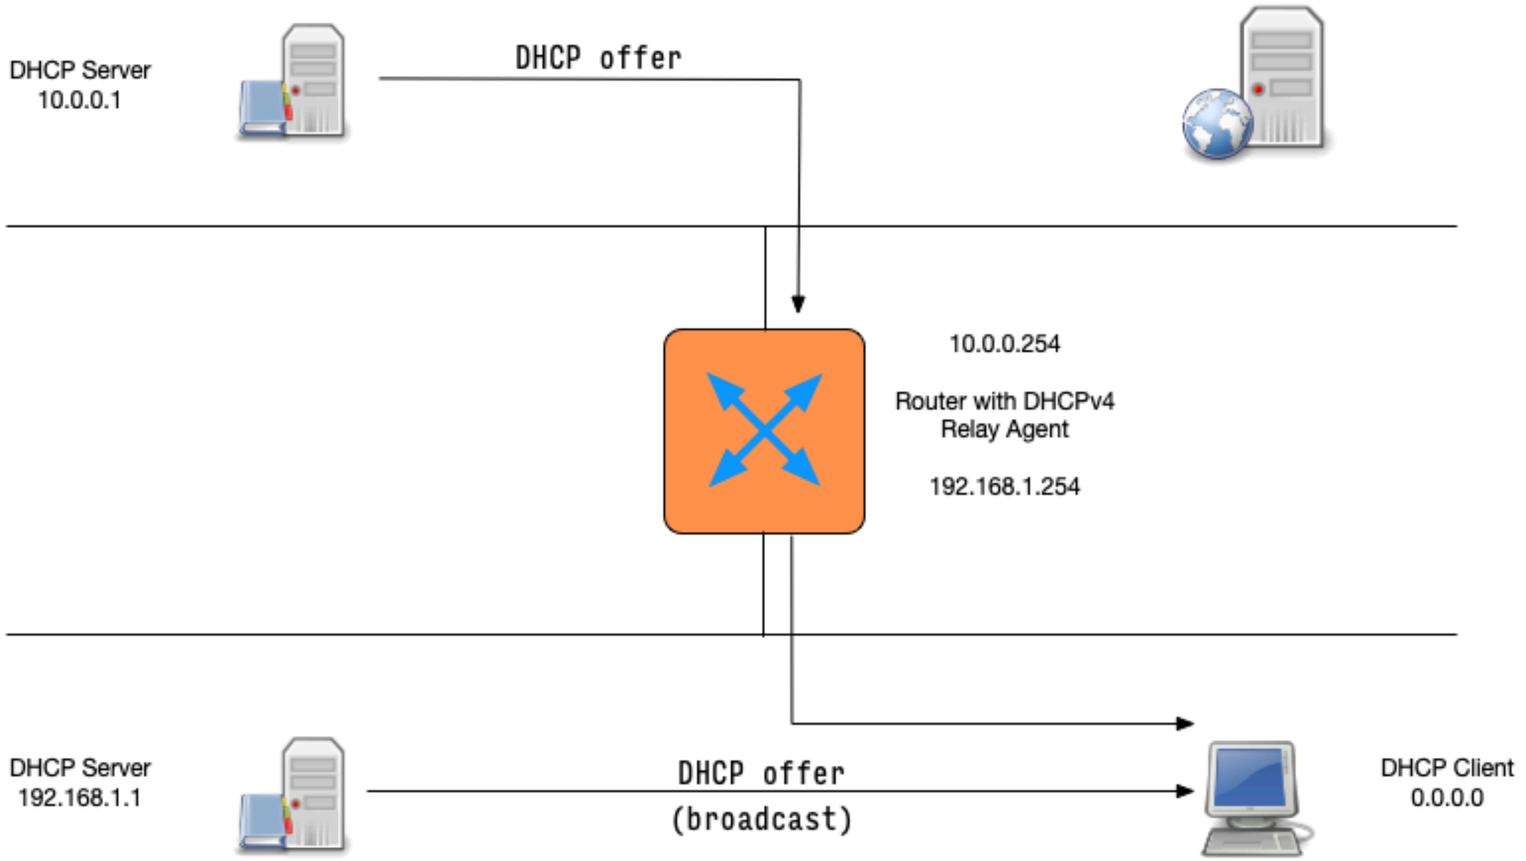

DHCP messages (4)

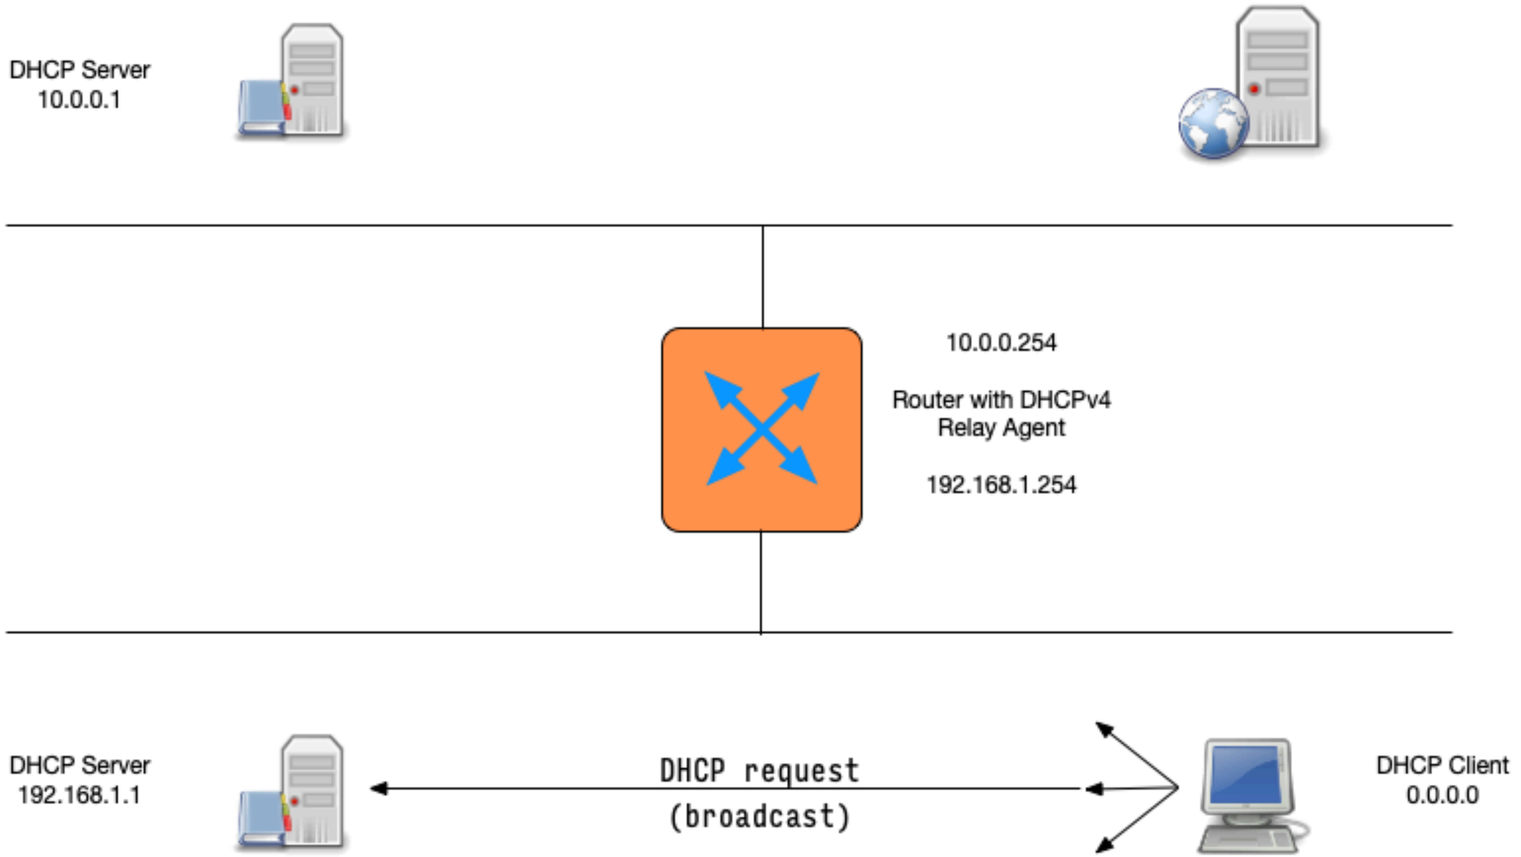

DHCP messages (5)

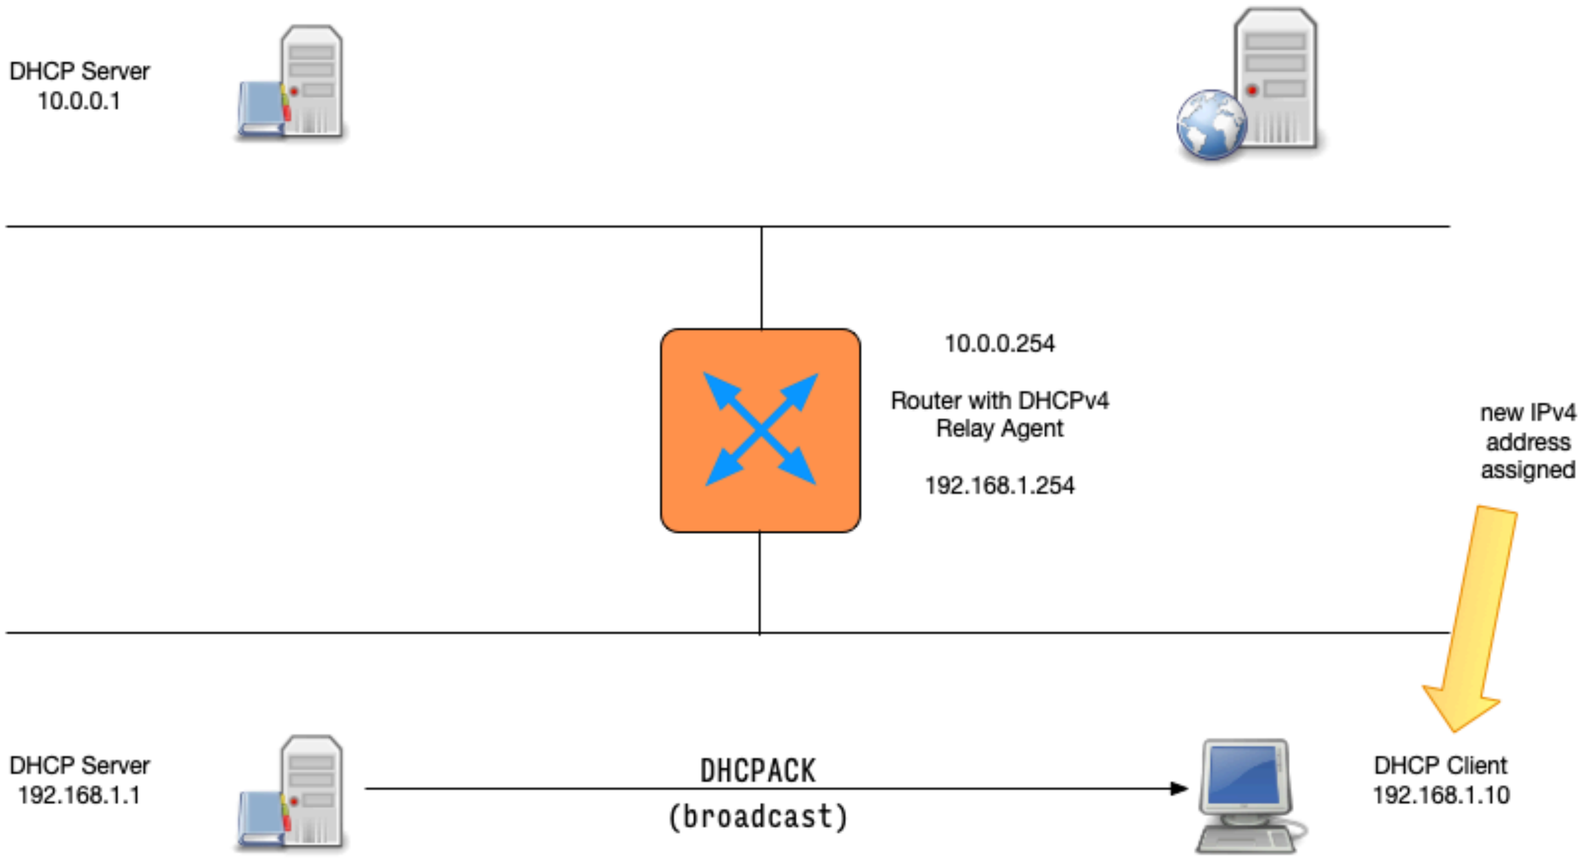

DHCPv4 client states

DHCPv4 client states

- A DHCP Client goes through a defined number of "states" when requesting or renewing a lease
  - `INIT-REBOOT`
  - `INIT`
  - `SELECTING`
  - `BOUND`
  - `RENEWING`
  - `REBINDING`

without IPv4 address (1/10)

without IPv4 address (2/10)

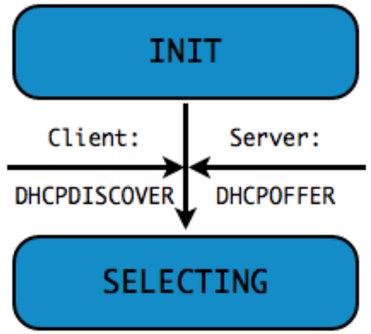

without IPv4 address (3/10)

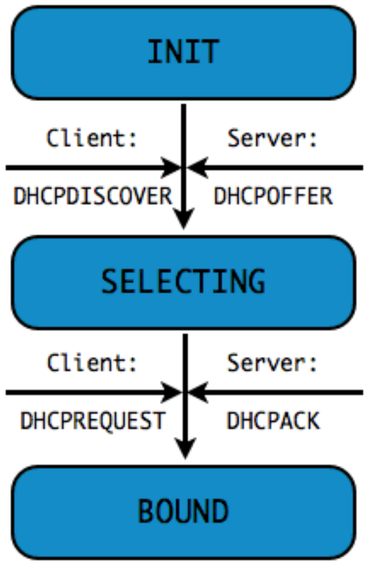

without IPv4 address (4/10)

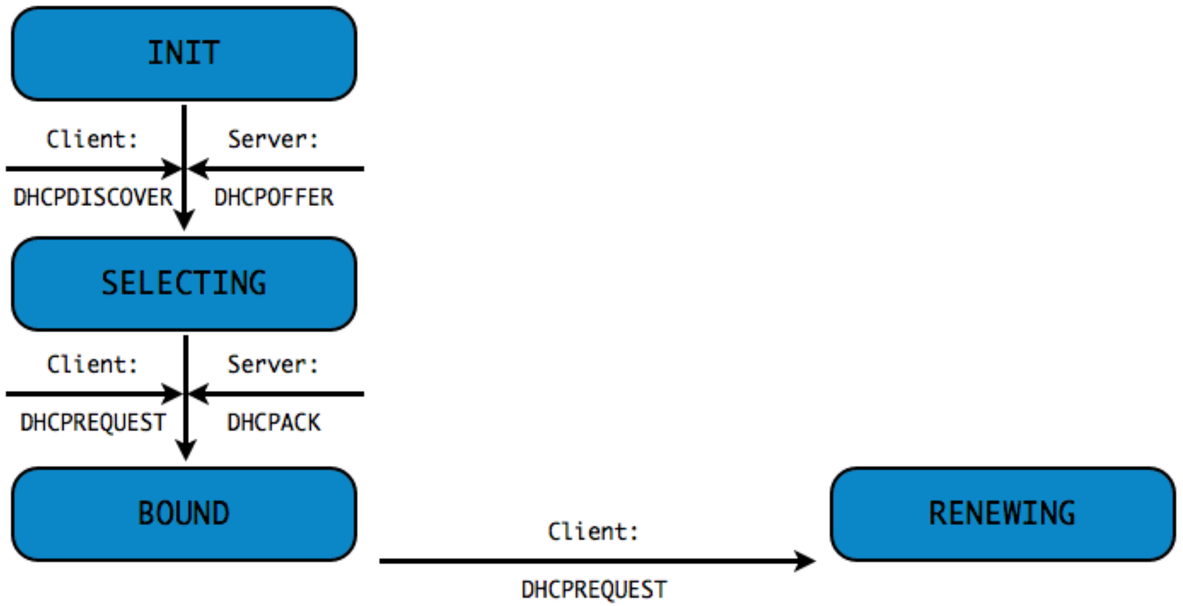

without IPv4 address (6/10)

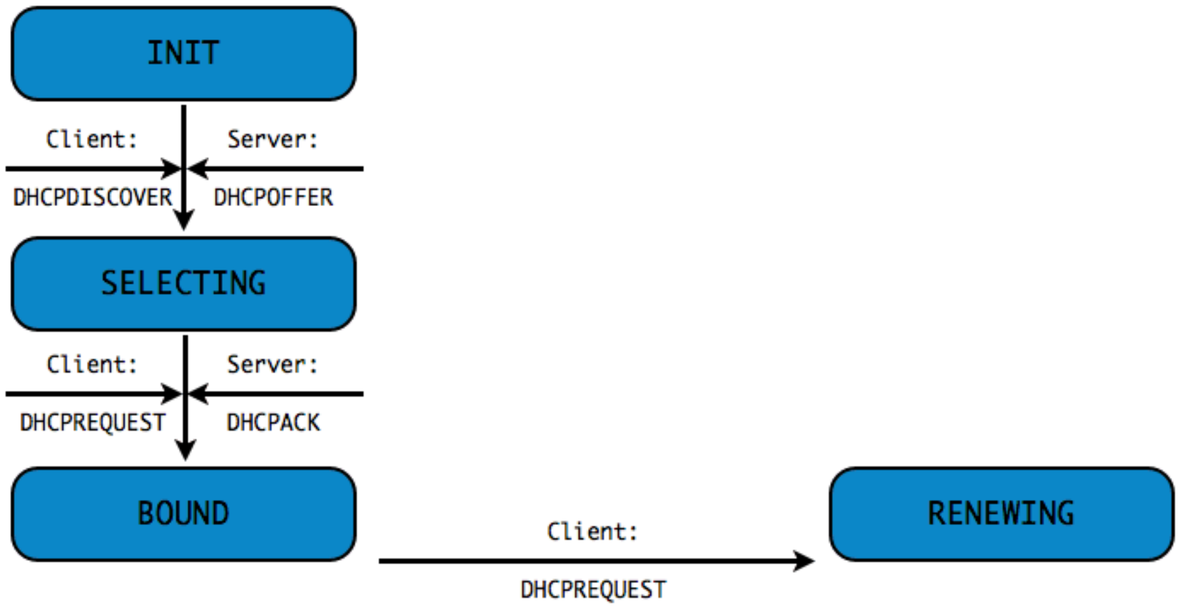

without IPv4 address (7/10)

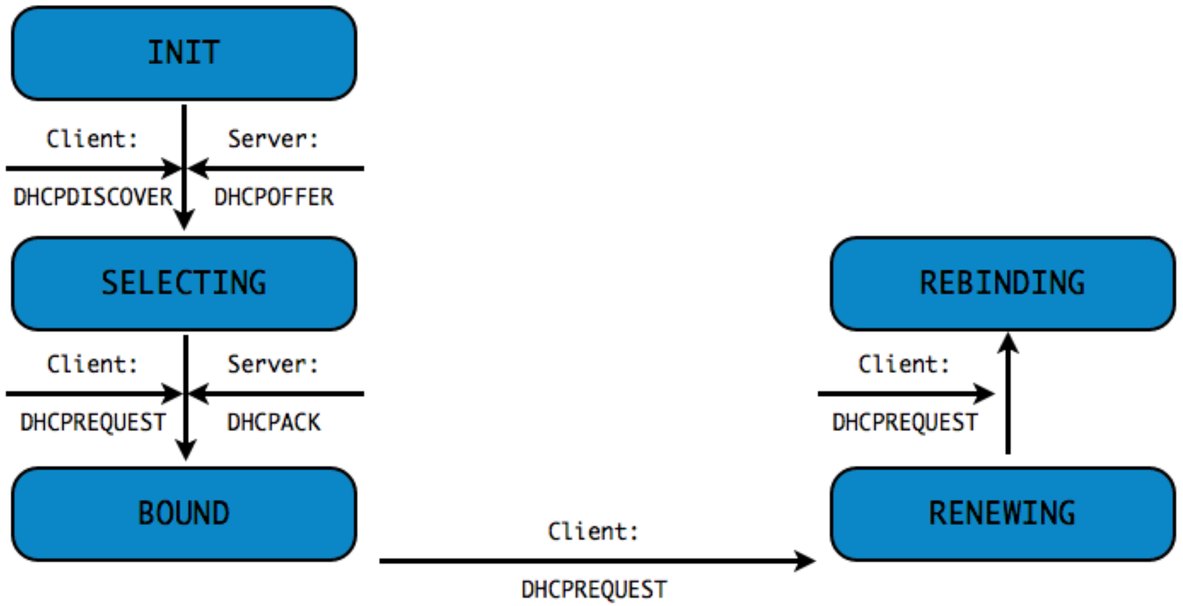

without IPv4 address (8/10)

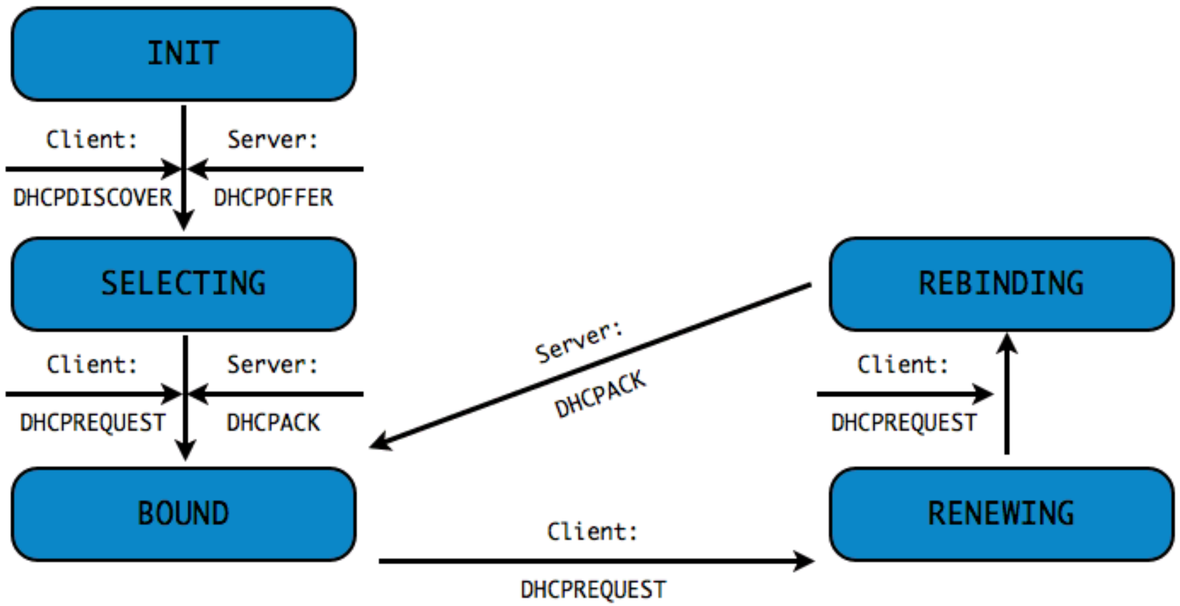

without IPv4 address (9/10)

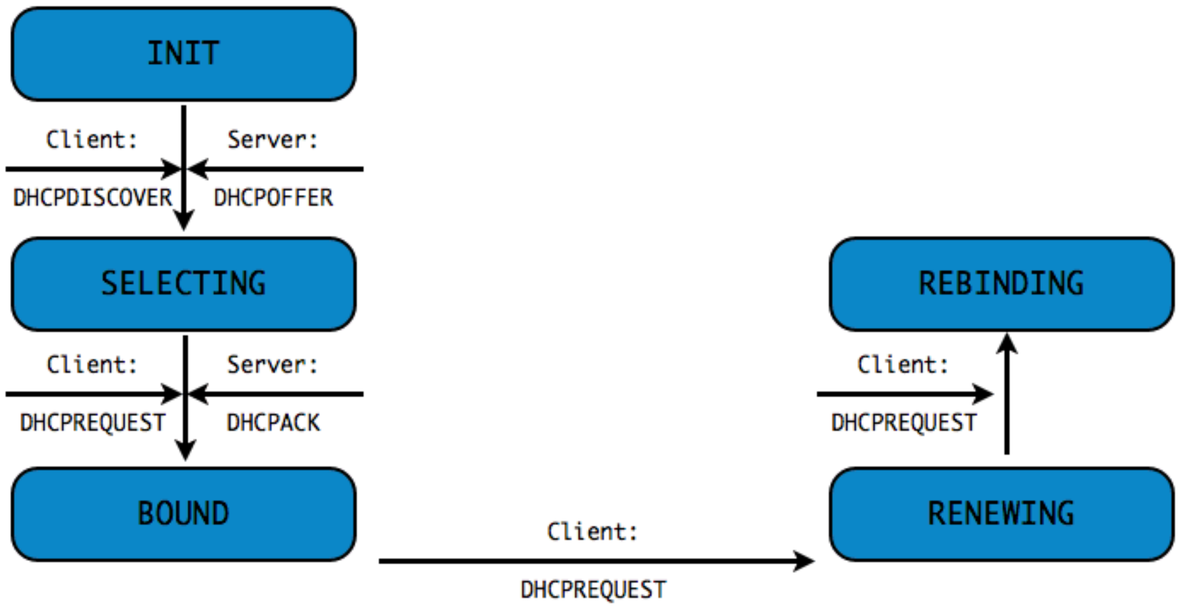

without IPv4 address (10/10)

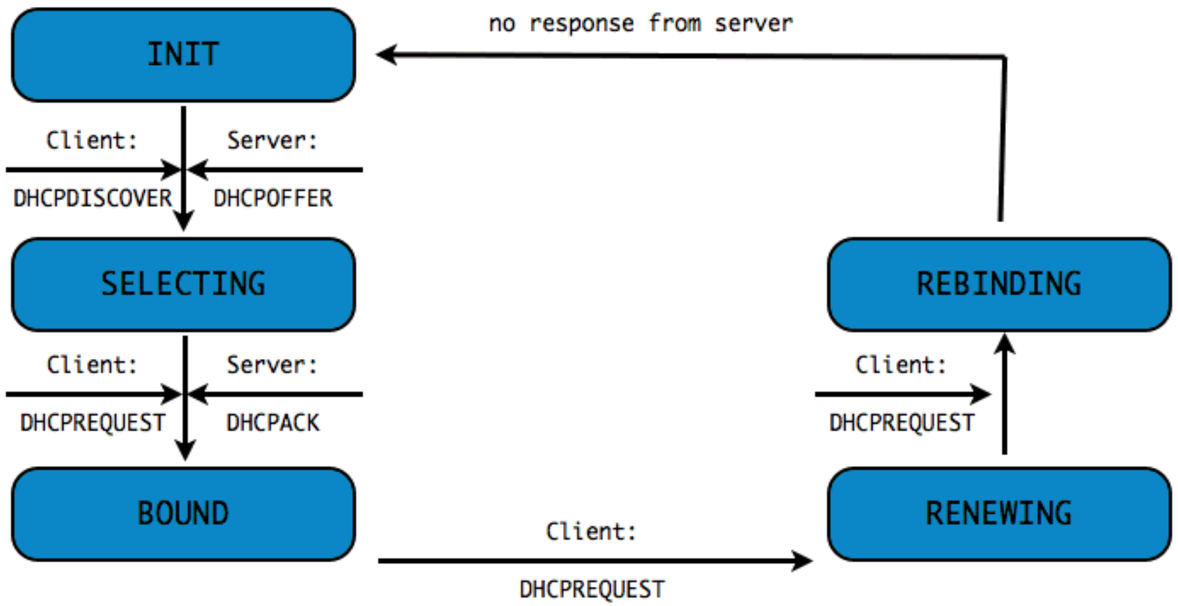

with IPv4 address (1/4)

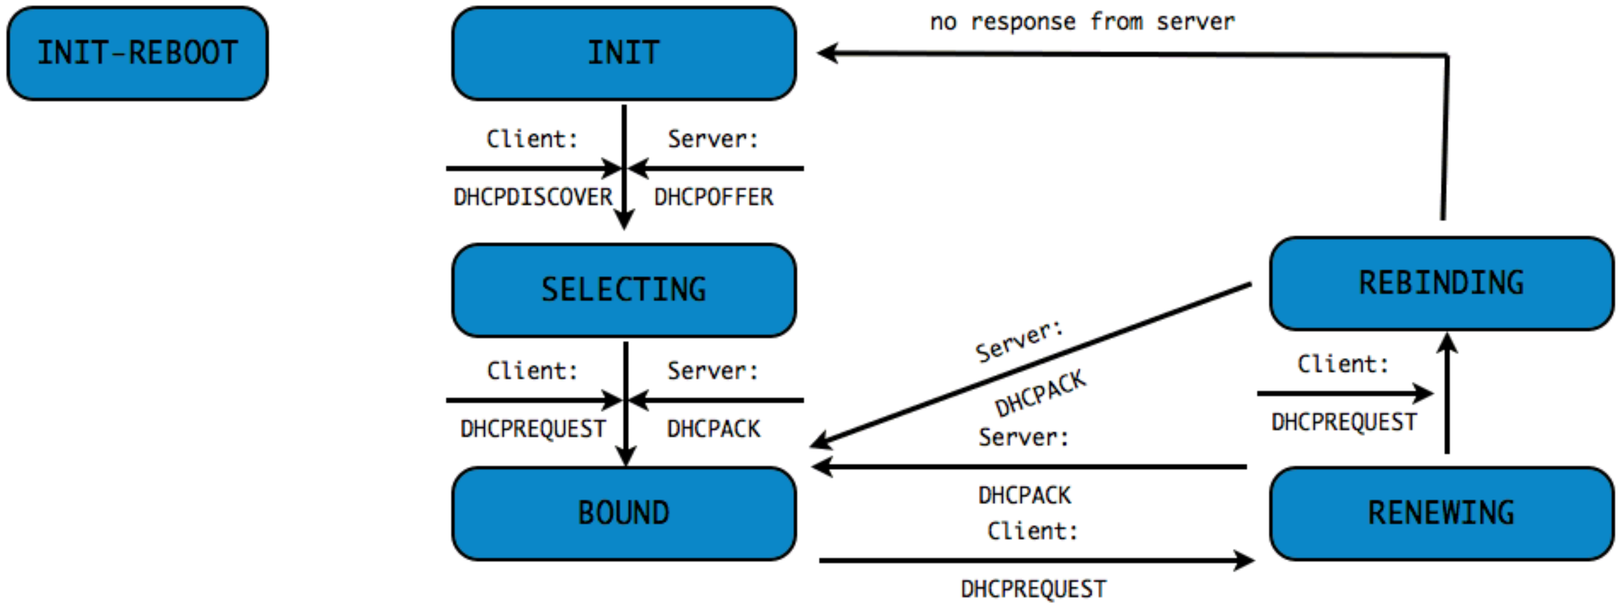

with IPv4 address (2/4)

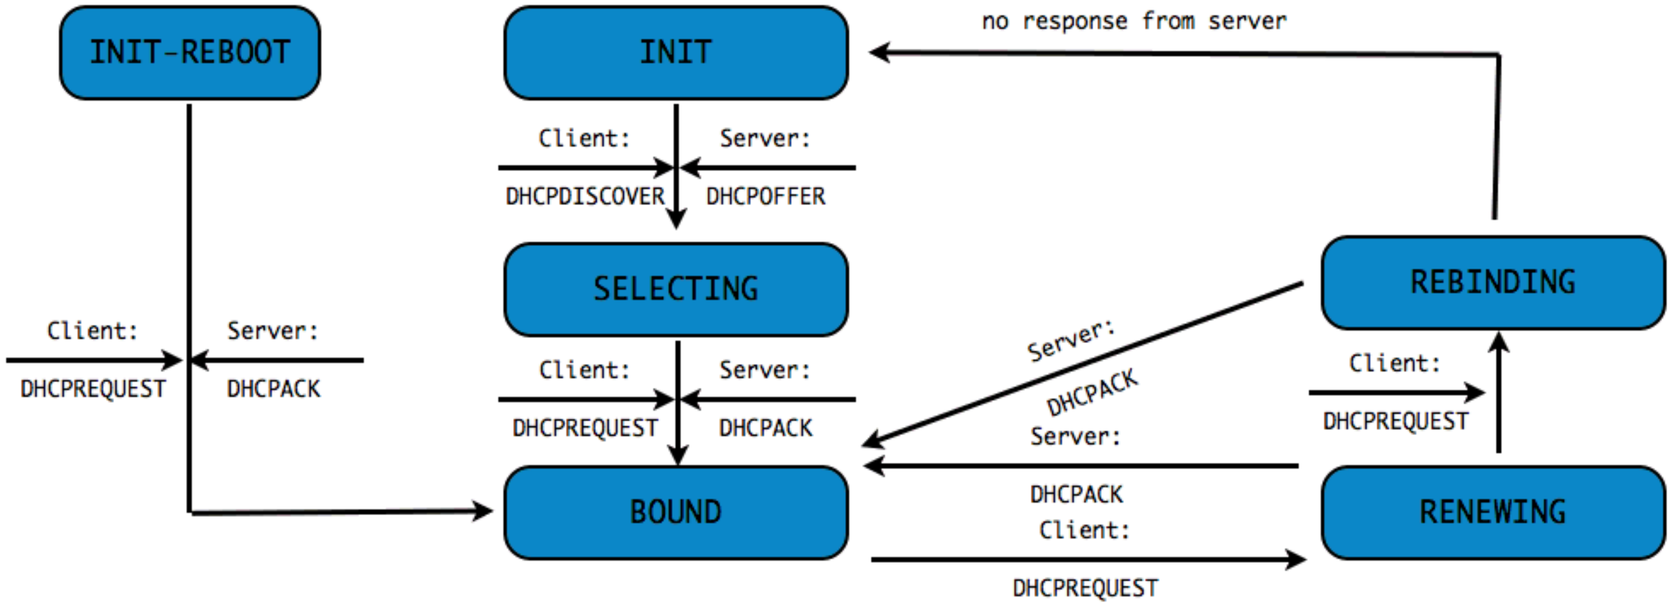

with IPv4 address (3/4)

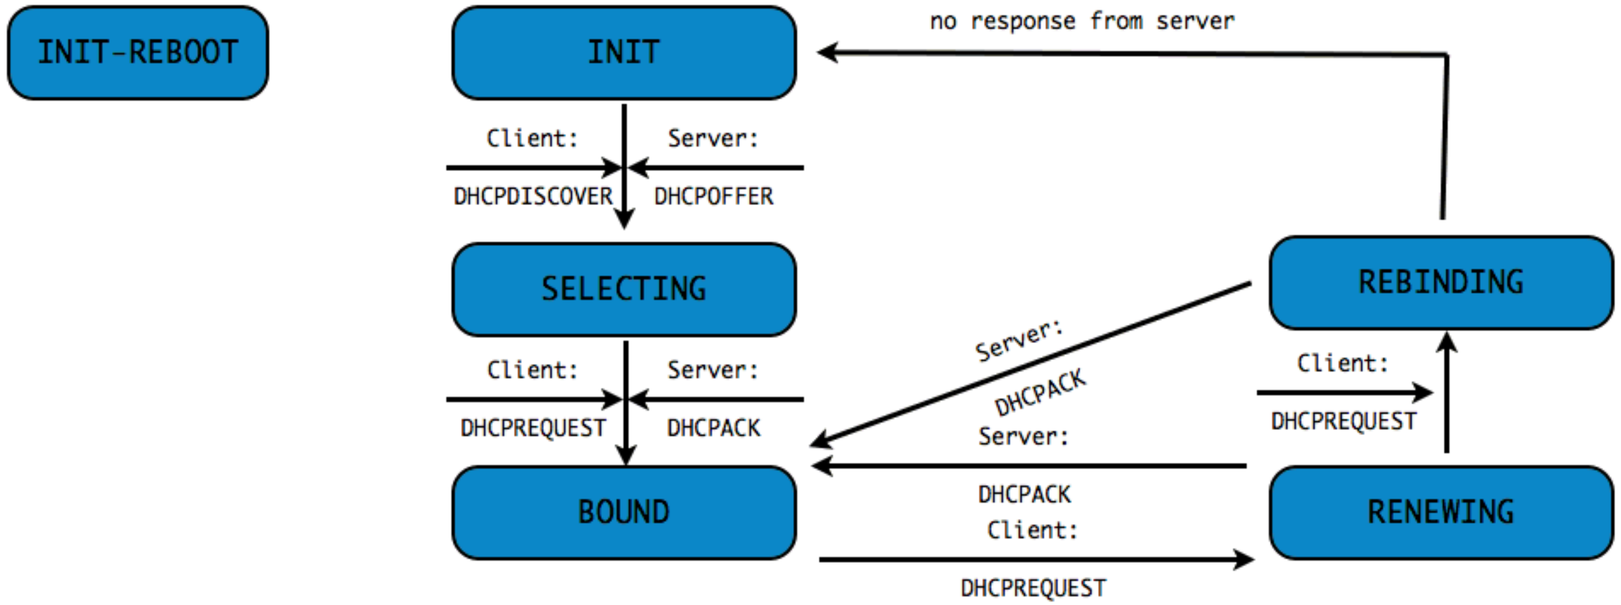

with IPv4 address (4/4)

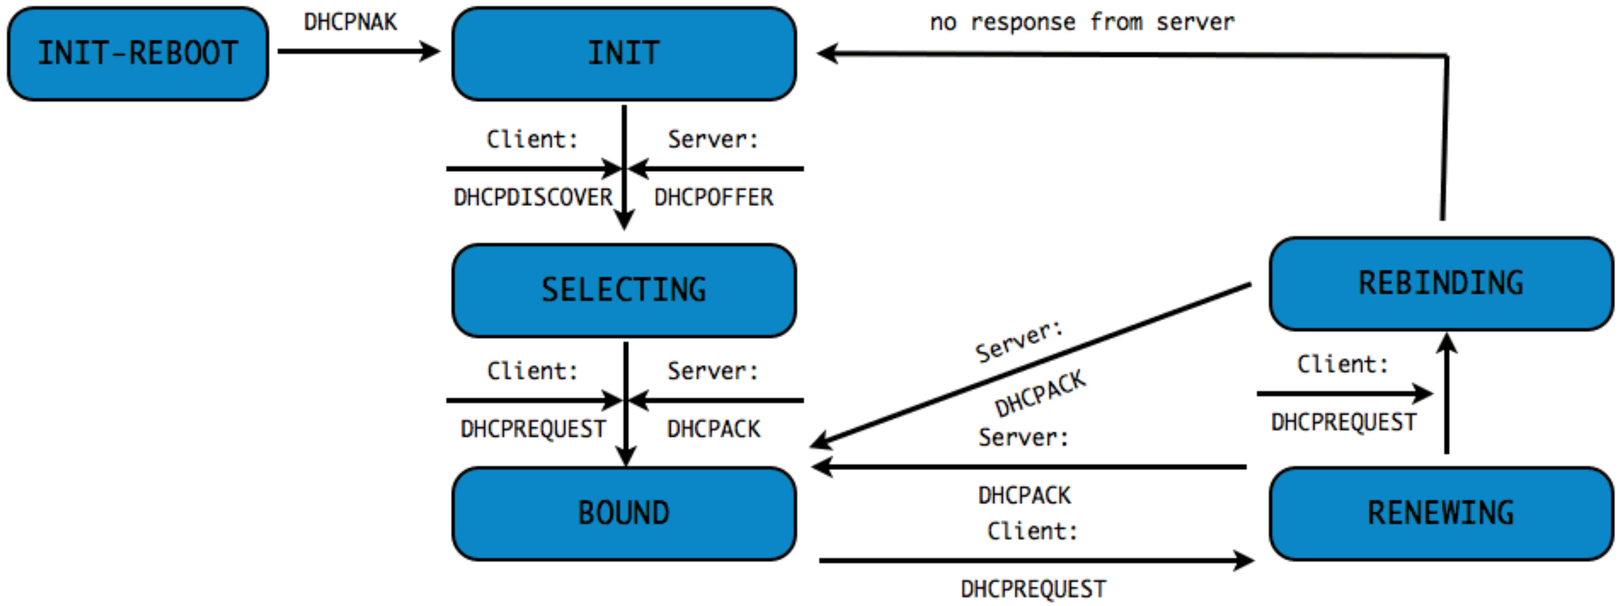

distributing network configuration with DHCP

BOOTP fields and DHCPv4 Options

- in addition to an IP address, DHCPv4 can be used to network configuration to a client
  - BOOTP configuration fields like next-server or boot-file-name
  - DHCPv4 options like domain-name-servers or domain-search
- we will cover the BOOTP fields and DHCPv4 options in our next webinar

Host reservations

Host reservation (1/3)

- sometimes a DHCP IP address should always be given to the same DHCP client machine
  - for example if that machine receives incoming connections (web-server, printer, database)
  - or if firewall rules define a security policy based on the IP address
- A **host reservation** binds a DHCP client via a client identifier (Ethernet MAC address) to a IP address

Host reservation (2/3)

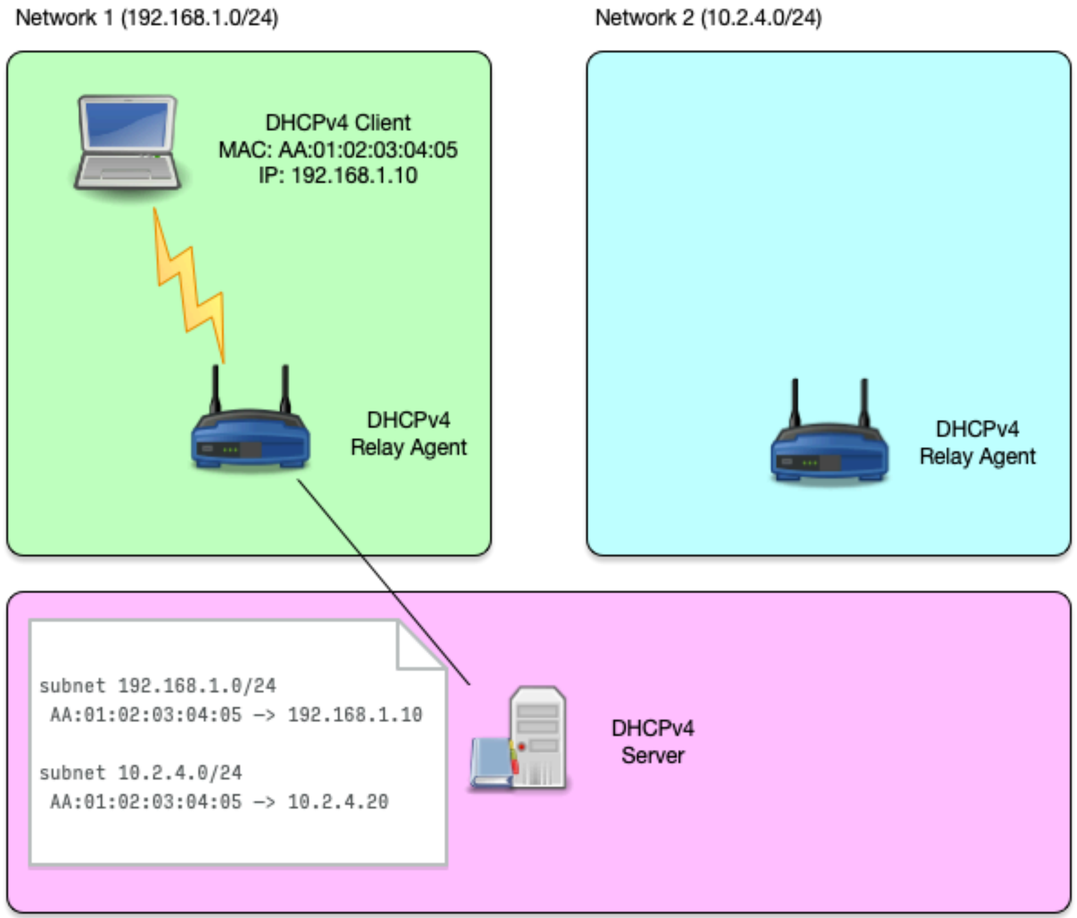

Host reservation (3/3)

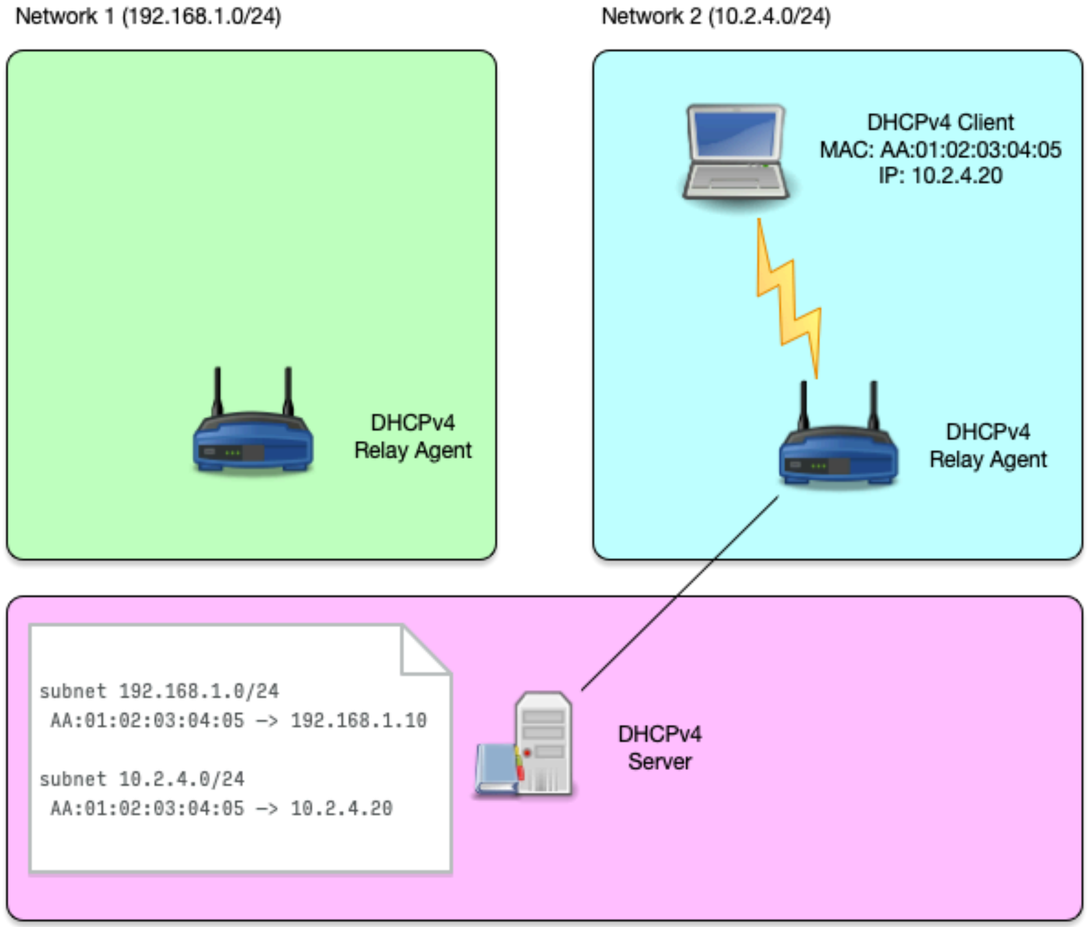

Shared Subnet

Shared Subnet

- A **shared subnet** is a physical network with more than one DHCPv4 managed subnet inside
  - shared subnet are sometimes created if a larger number of IP addresses are needed in a network, but because of IPv4 address shortage no continuous range of IPv4 addresses are available

Shared Subnet

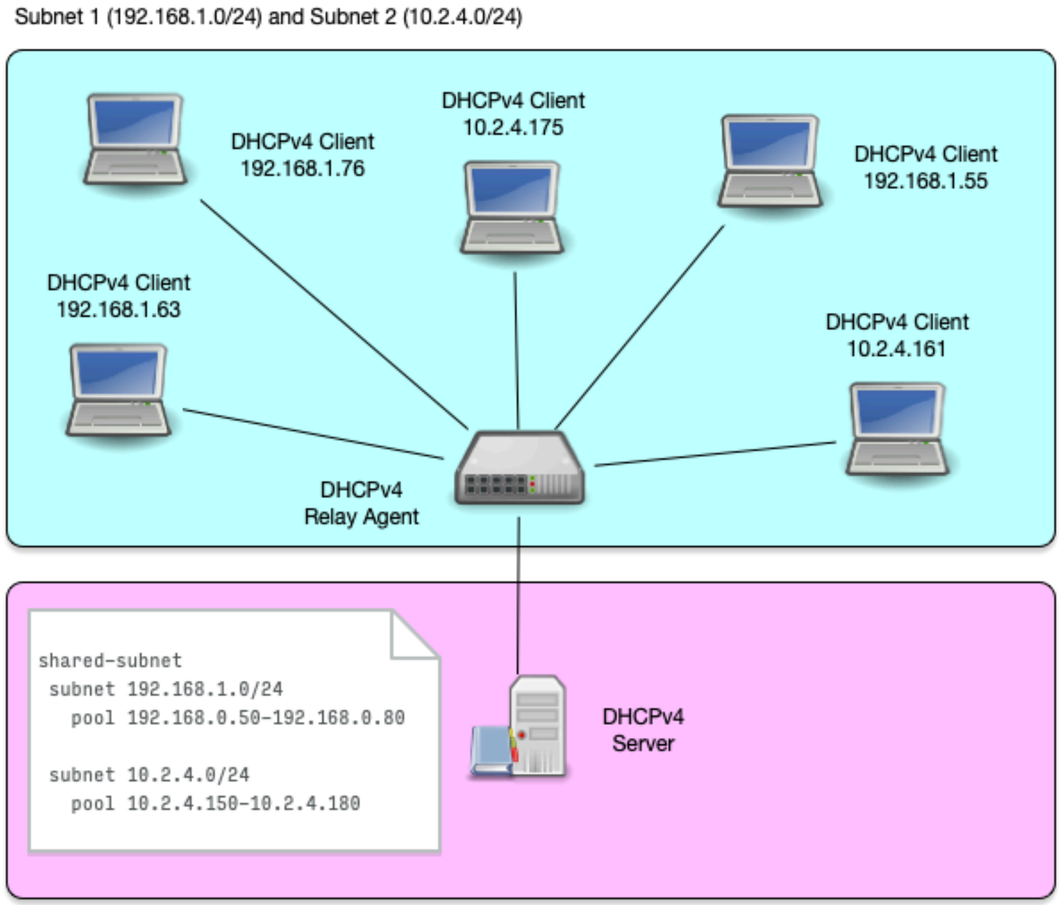

Shared Subnet

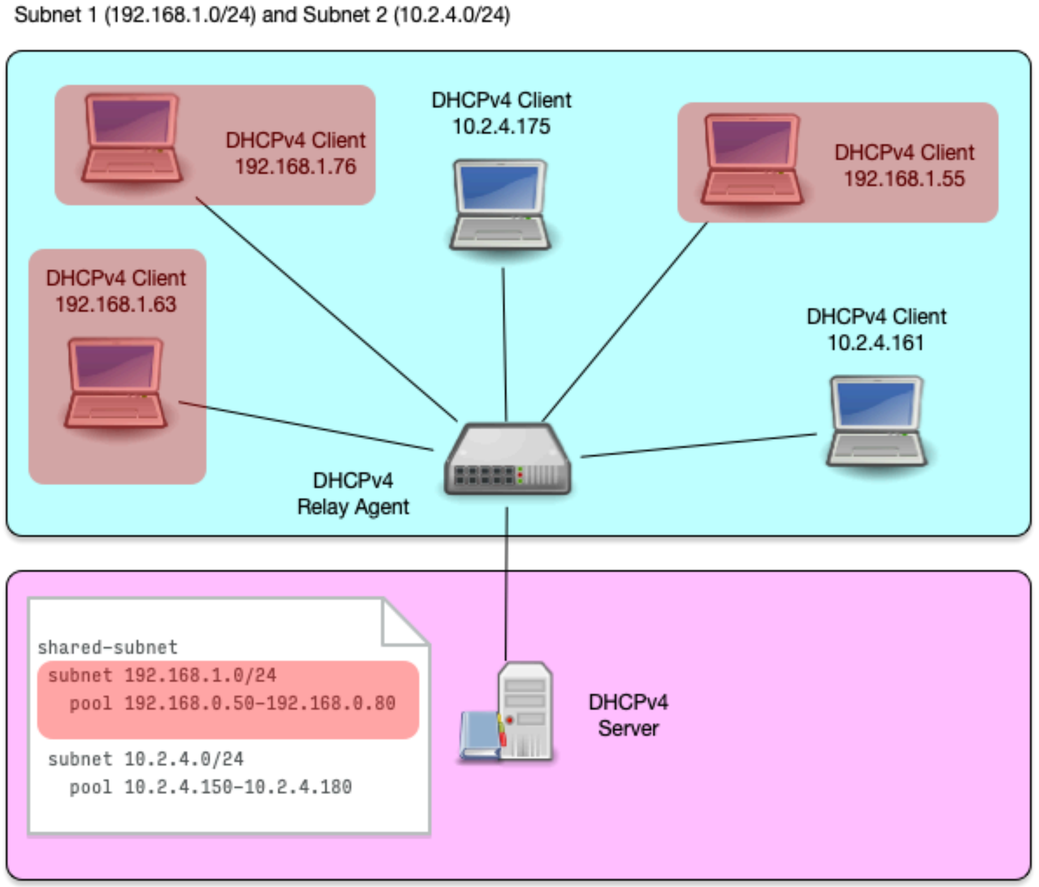

Shared Subnet

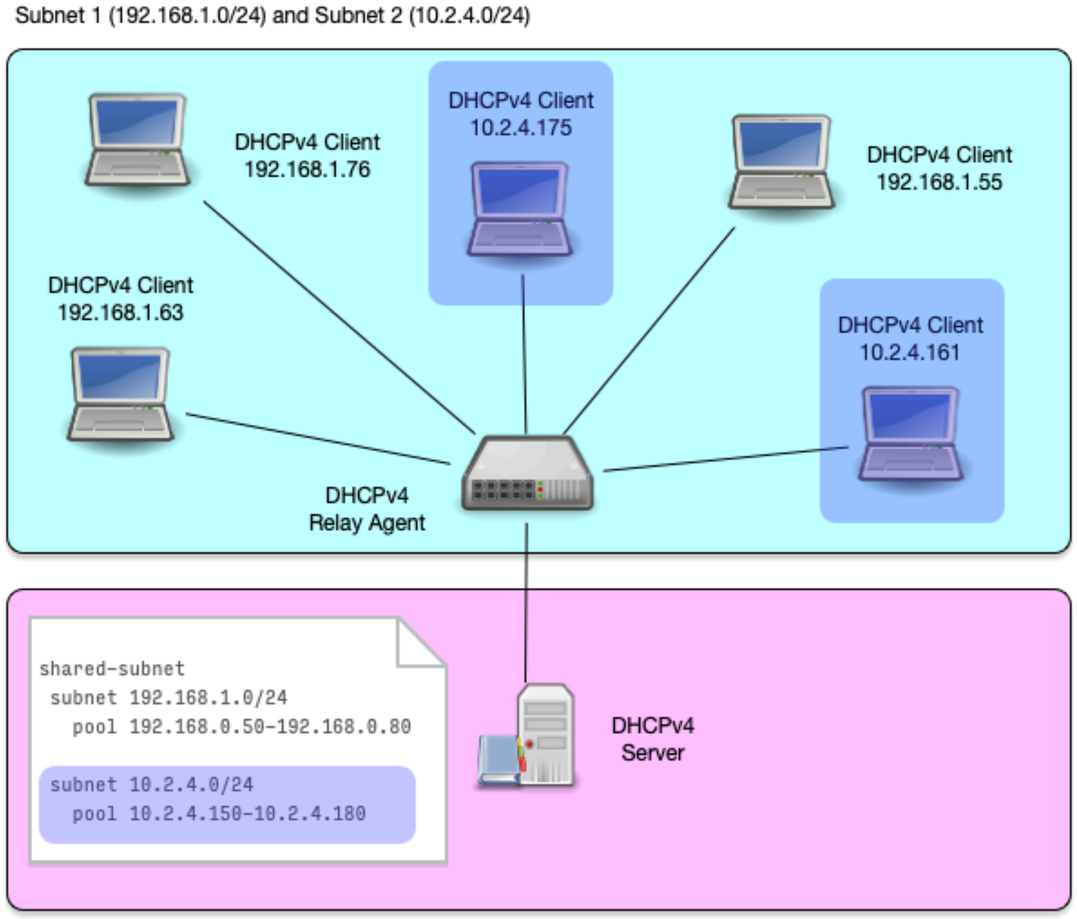

Shared Subnet

- another use case of shared subnets is a network where addresses from different IPv4 subnets (and possibly different network configuration) should be given to different network devices
  - cable modems and end user devices
  - printer, desktop and mobile devices
  - POS terminals and retail infrastructure devices (digital price tags)

Shared Subnet

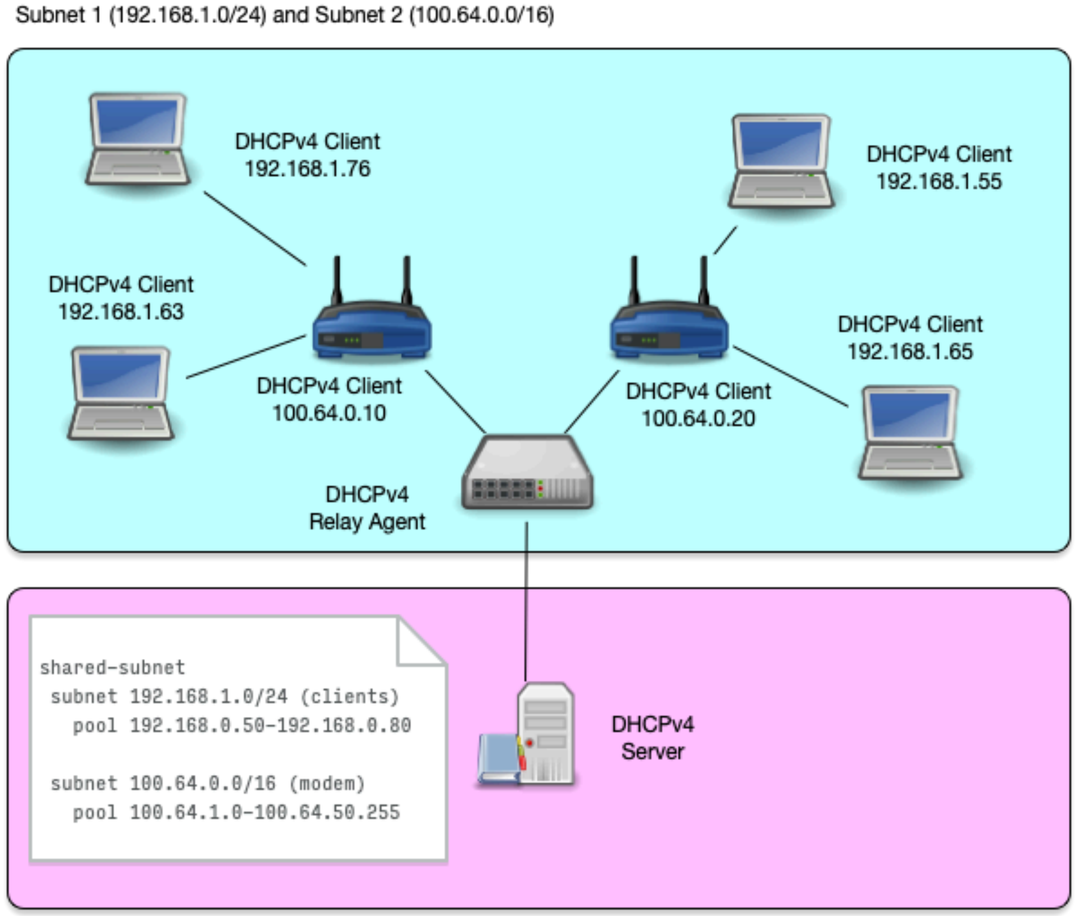

Shared Subnet

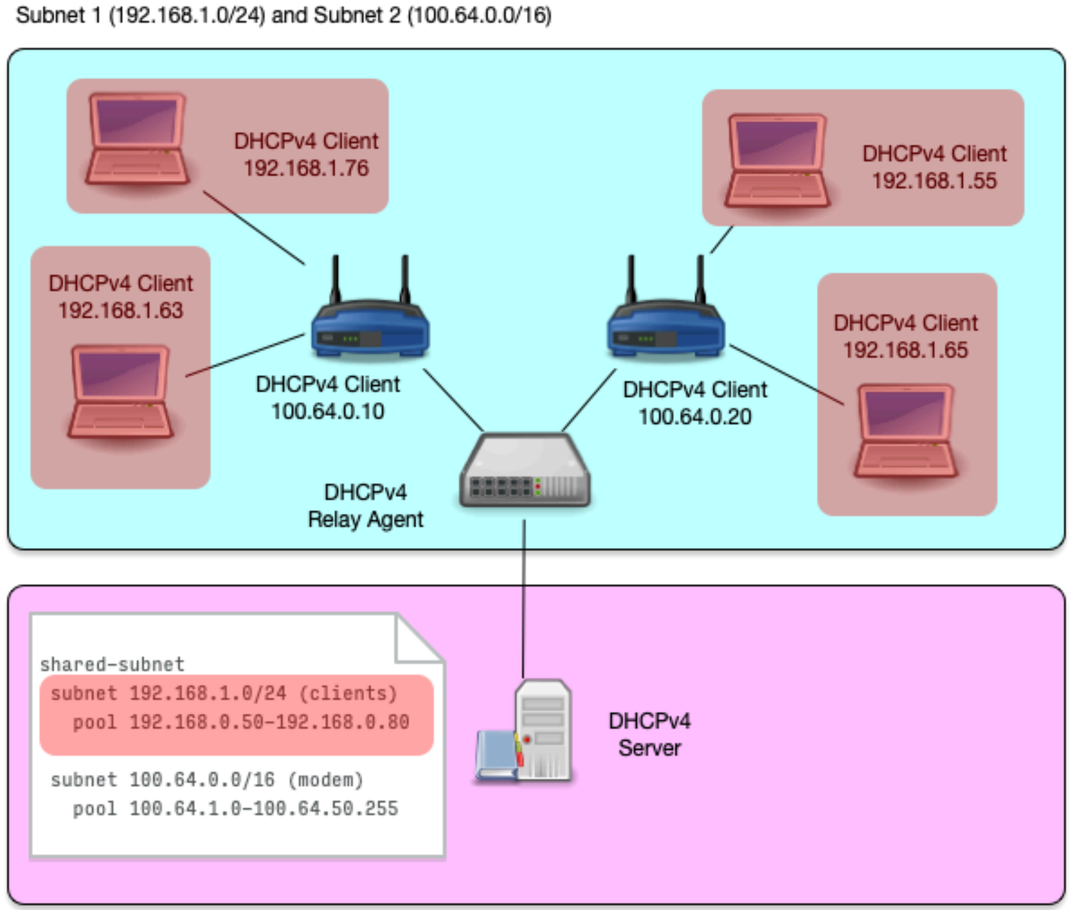

Shared Subnet

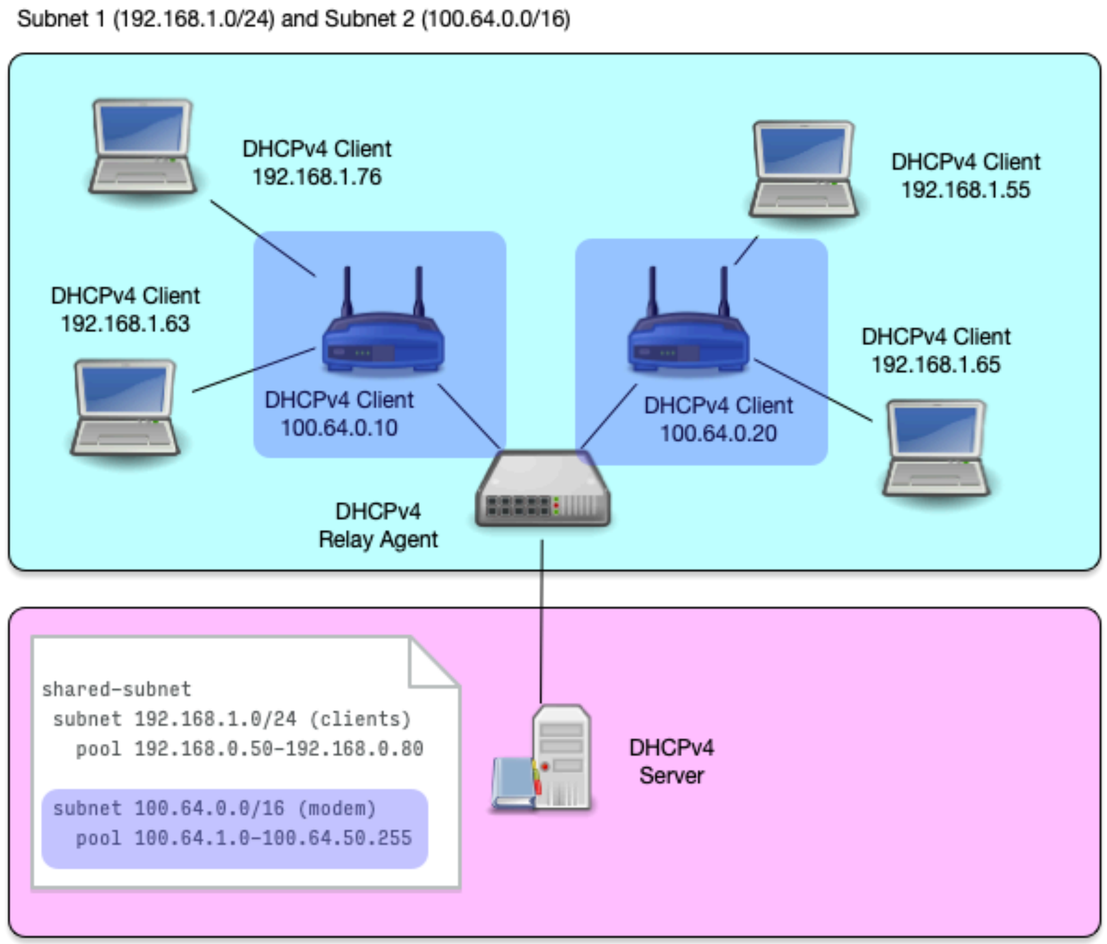

DHCPv6

DHCPv6

- from a birds-eye view, **DHCPv6** works the same way as DHCPv4
  - in the details, it is all different
  - **DHCPv6** is not an upgrade to DHCPv4, it is a protocol of its own

DHCPv6 IP based vs. DHCPv4 Layer2/ Ethernet based (use of link-local addresses)

- DHCPv6 is solely a Layer 3 protocol
  - a DHCPv6 client already has a working link-local IPv6 address (`fe80::`) when sending the first DHCPv6 request
  - no "low-level kernel trickery" required

DHCPv6 protocol (port numbers, communication)

- DHCPv6 Servers and Relay-Agents listen on Port 547 (UDPv6)
- DHCPv6 clients listen on Port 546 (UDPv6)

DHCPv6 multicast usage and addresses

- DHCPv6 clients communicate using link-local multicast addresses
  - All-DHCP-Relay-Agents-and-Servers (`ff02::1:2`)
  - All-DHCP-Servers (`ff05::1:3`)

the role of router in DHCPv6

- DHCPv6 has been designed to provide it's service in cooperation with the local router(s)
  - DHCPv6 must be enabled in the router configuration (`M-Flag` or `O-Flag`)
  - The Default-Gateway Address will be retrieved from a router and **not** from the DHCPv6 Server

Router in DHCPv6 (1/2)

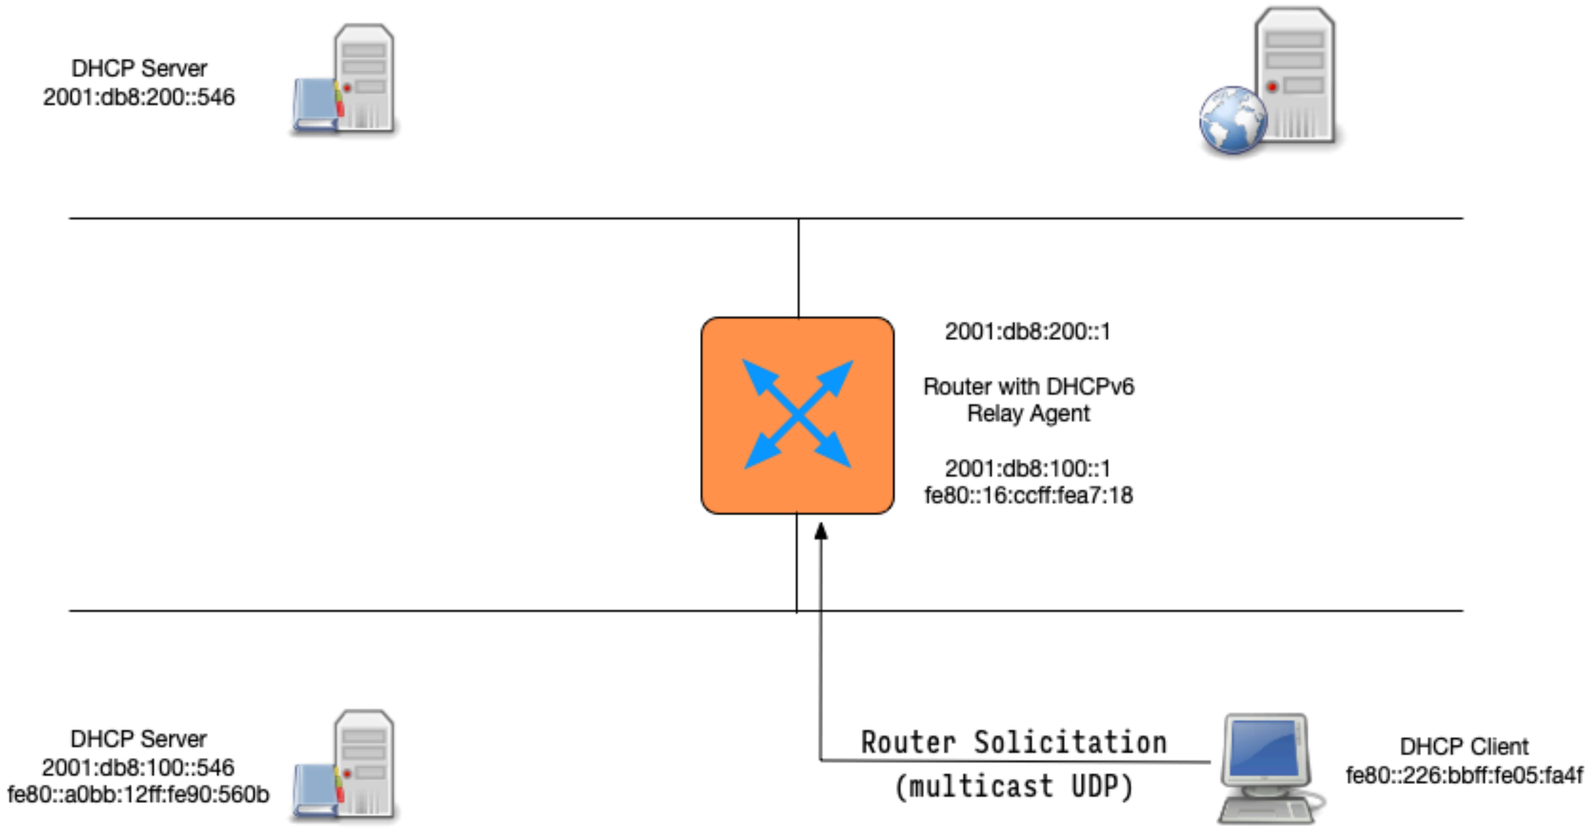

Router in DHCPv6 (1/2)

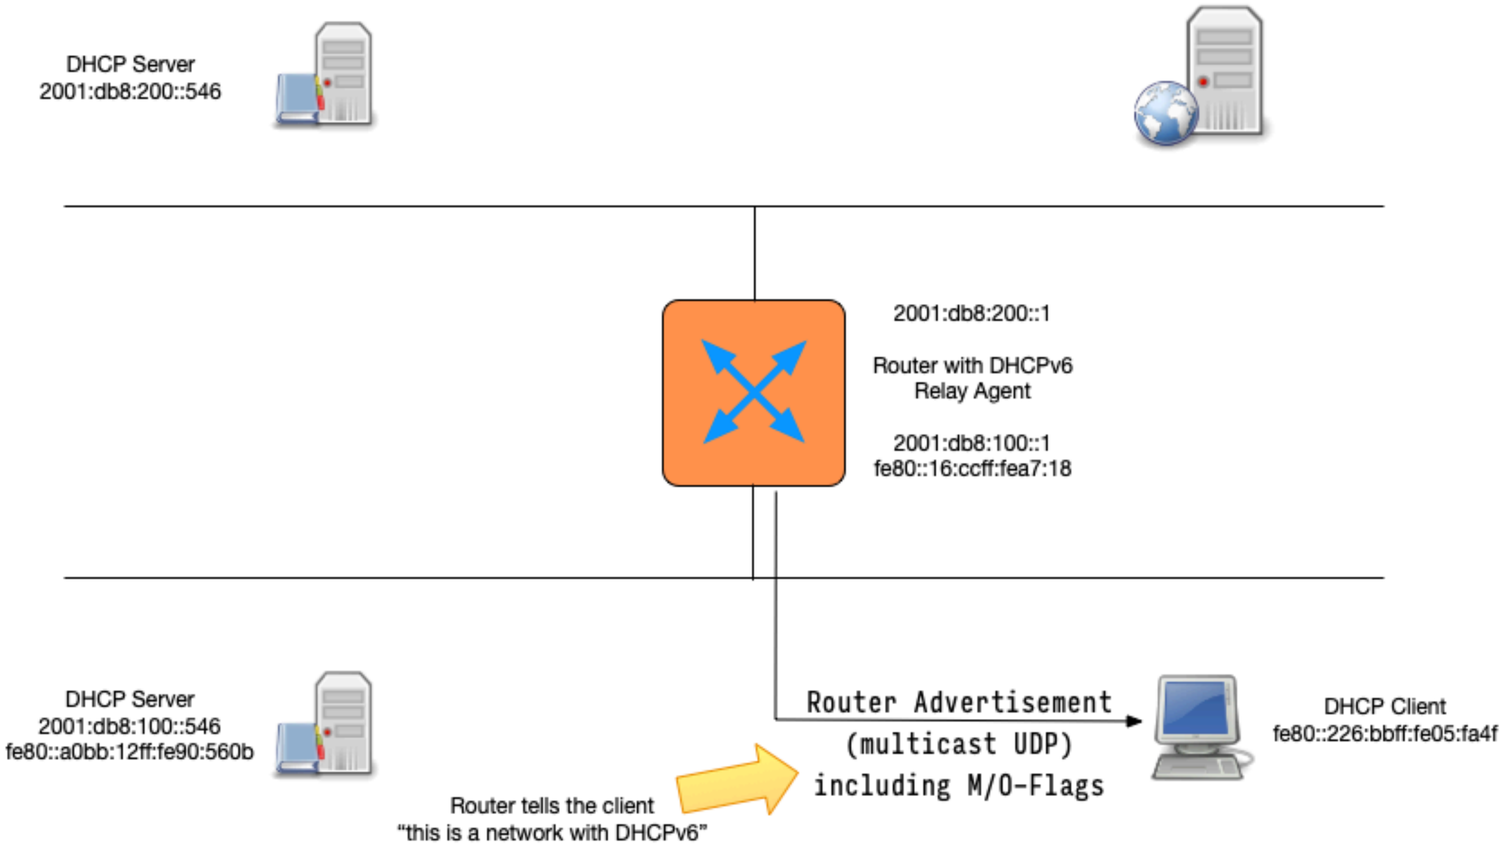

DHCPv6 address allocation vs. DHCPv4 address allocation

- DHCPv6 server must issue IP Addresses randomly from the available address pool
  - some DHCPv4 server products issue IP Addresses continuously
- The DHCPv6 scheme makes it harder to guess an IP Address or scan a network segment

DHCPv6 allocation types: non-temporary, temporary, multiple addresses, prefix-delegation

- a DHCPv6 client can send different kind of IP address requests
  - **non-temporary**: an address that will always be given to this client if available. This is similar to an DHCPv4 address request.
  - **temporary**: an address that will always change on each request and has a low lease time. Used by the client for outgoing connections for privacy reasons
  - **multiple**: a client can request multiple IP addresses from the DHCPv6 server
  - **prefix delegation**: a client that is a router and also a DHCPv6 server itself can request one or more IPv6 prefix networks from an upstream DHCPv6 server
    - this allows for an hierarchical DHCPv6 configuration

DHCPv6 in combination with SLAAC

- IPv6 support Stateless Automatic Address Configuration aka **SLAAC**
- **SLAAC** can be used as an alternative to **DHCPv6**
- **SLAAC** and **DHCPv6** can be combined

DHCPv6 - stateless vs. stateful

- There are two different ways to get an IPv6 address for a IPv6 enabled device
  - Stateless configuration
  - Stateful configuration

DHCPv6 - stateless vs. stateful

- **Stateless** configuration
  - The IPv6 address will be determined without a DHCP Server (IPv6 auto-configuration = SLAAC)
- **Stateful** configuration
  - The IPv6 address will be received from a DHCPv6 Server
- In both cases additional configuration parameters (DNS Server etc) can be retrieved by DHCPv6

DHCPv6 solicit message

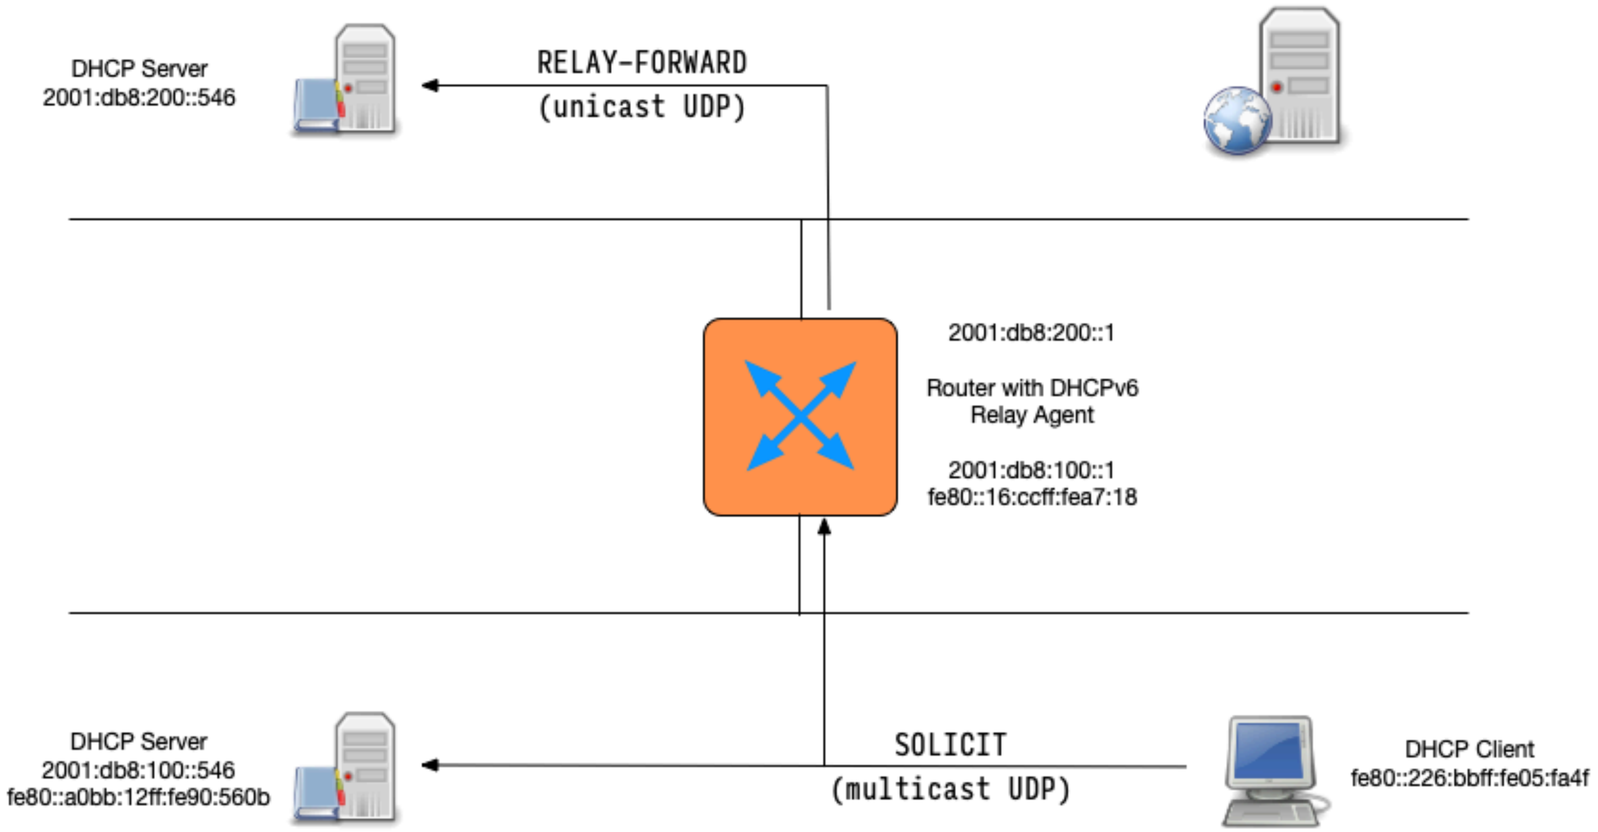

DHCPv6 advertise message

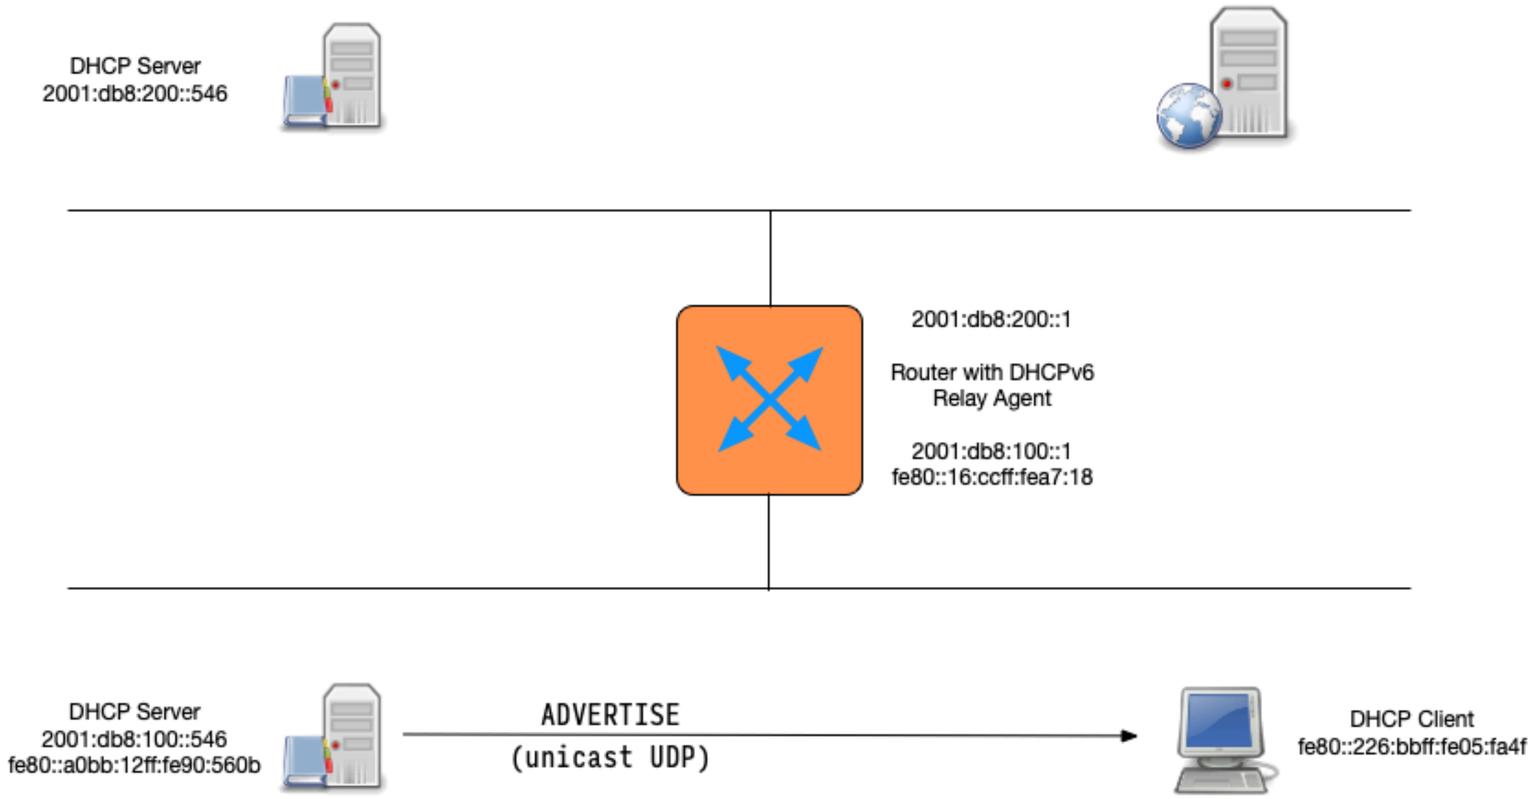

DHCPv6 advertise (via relay)

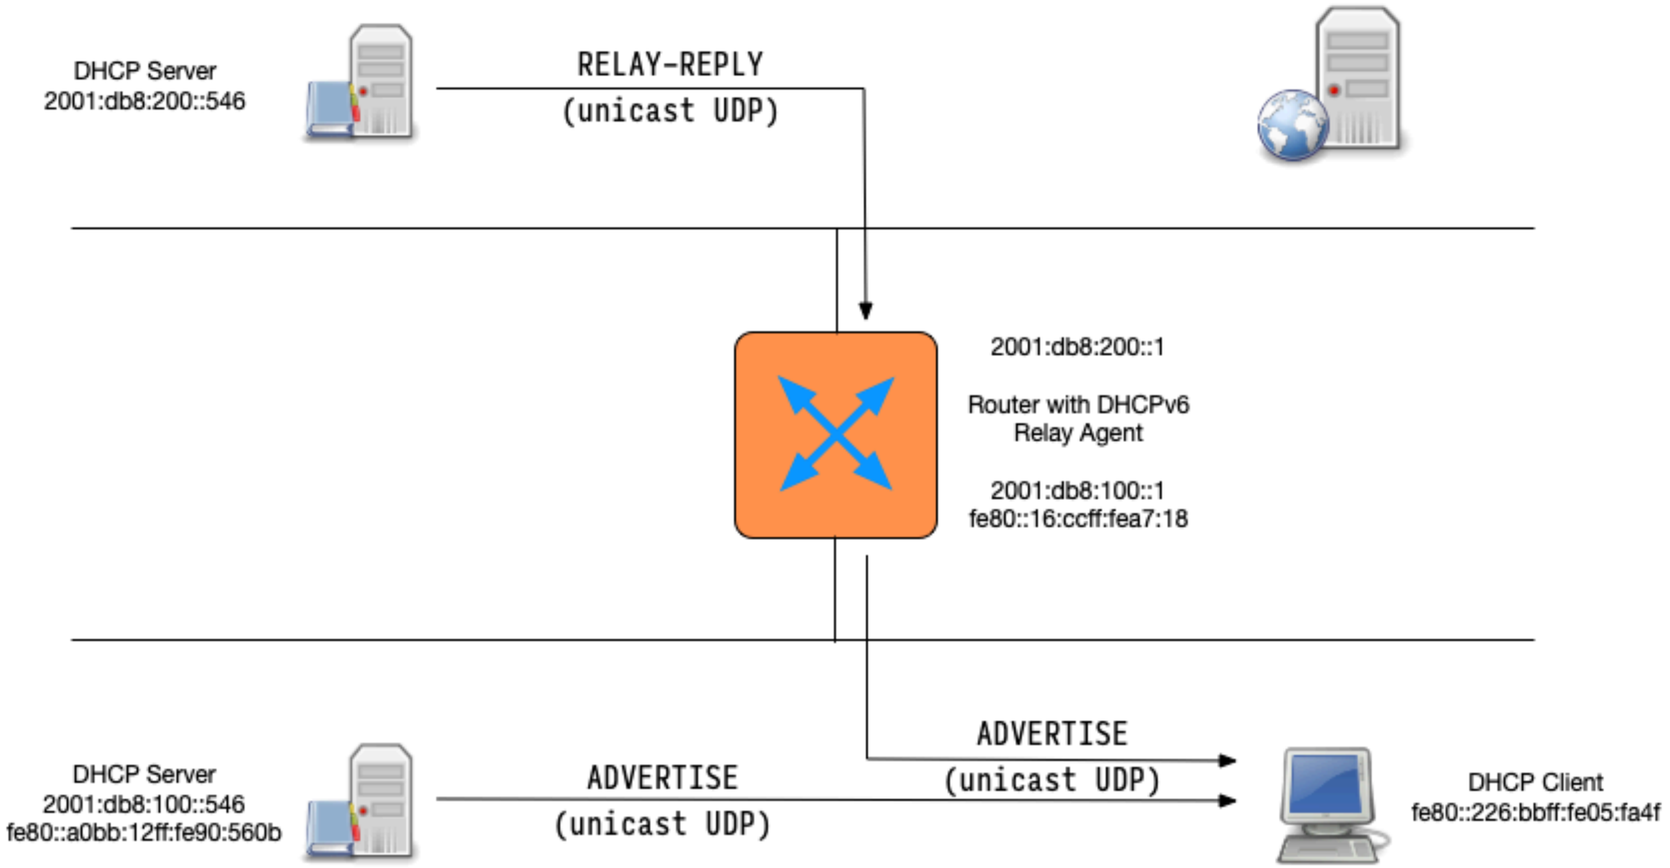

DHCPv6 request

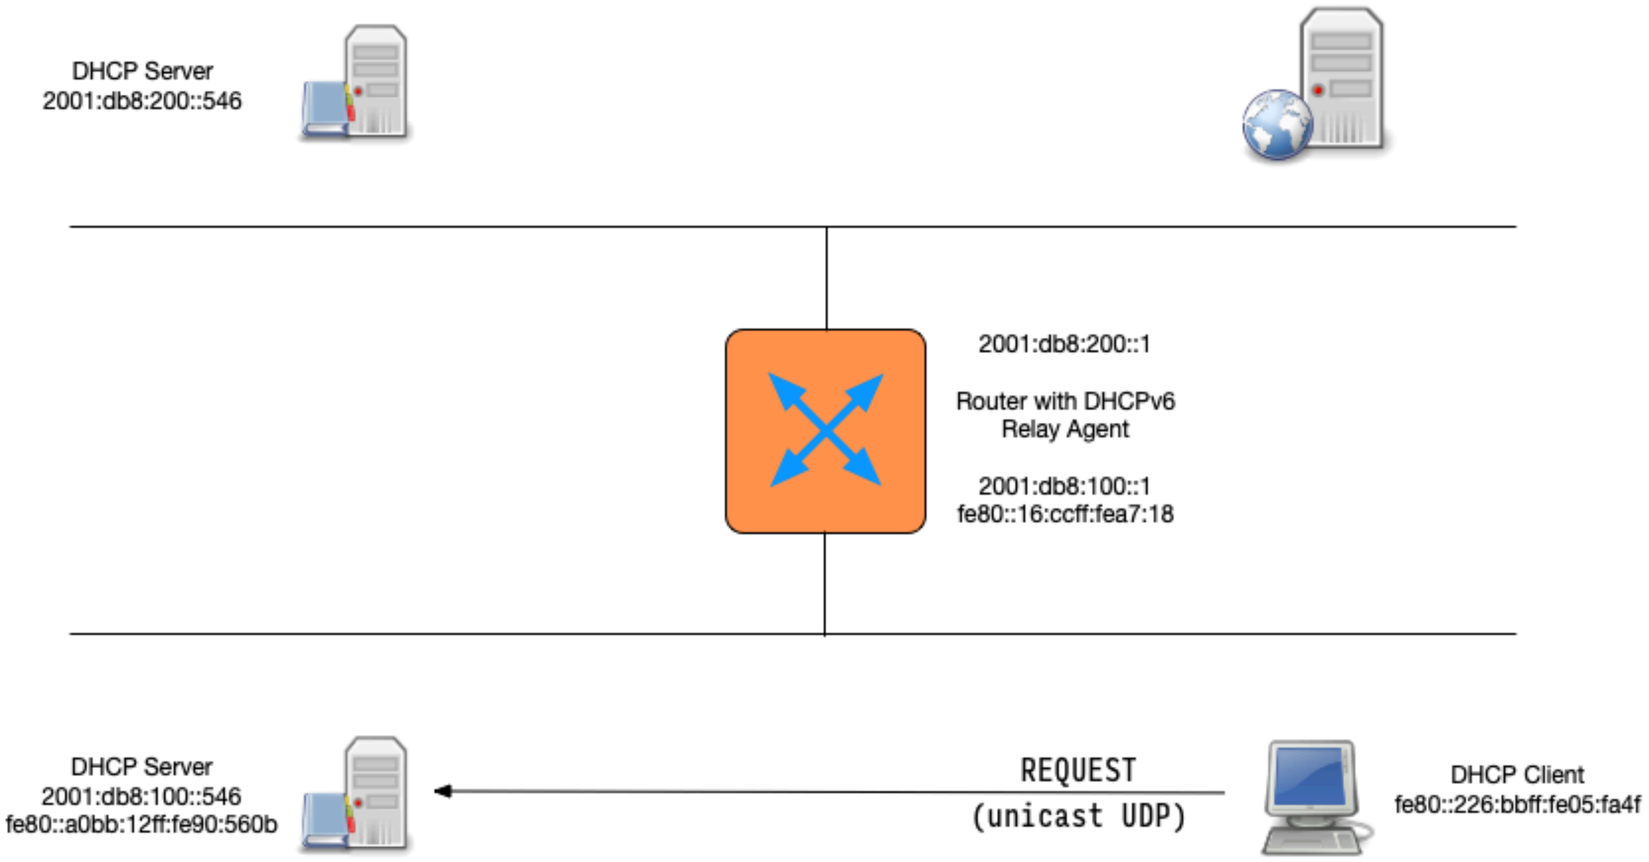

DHCPv6 reply

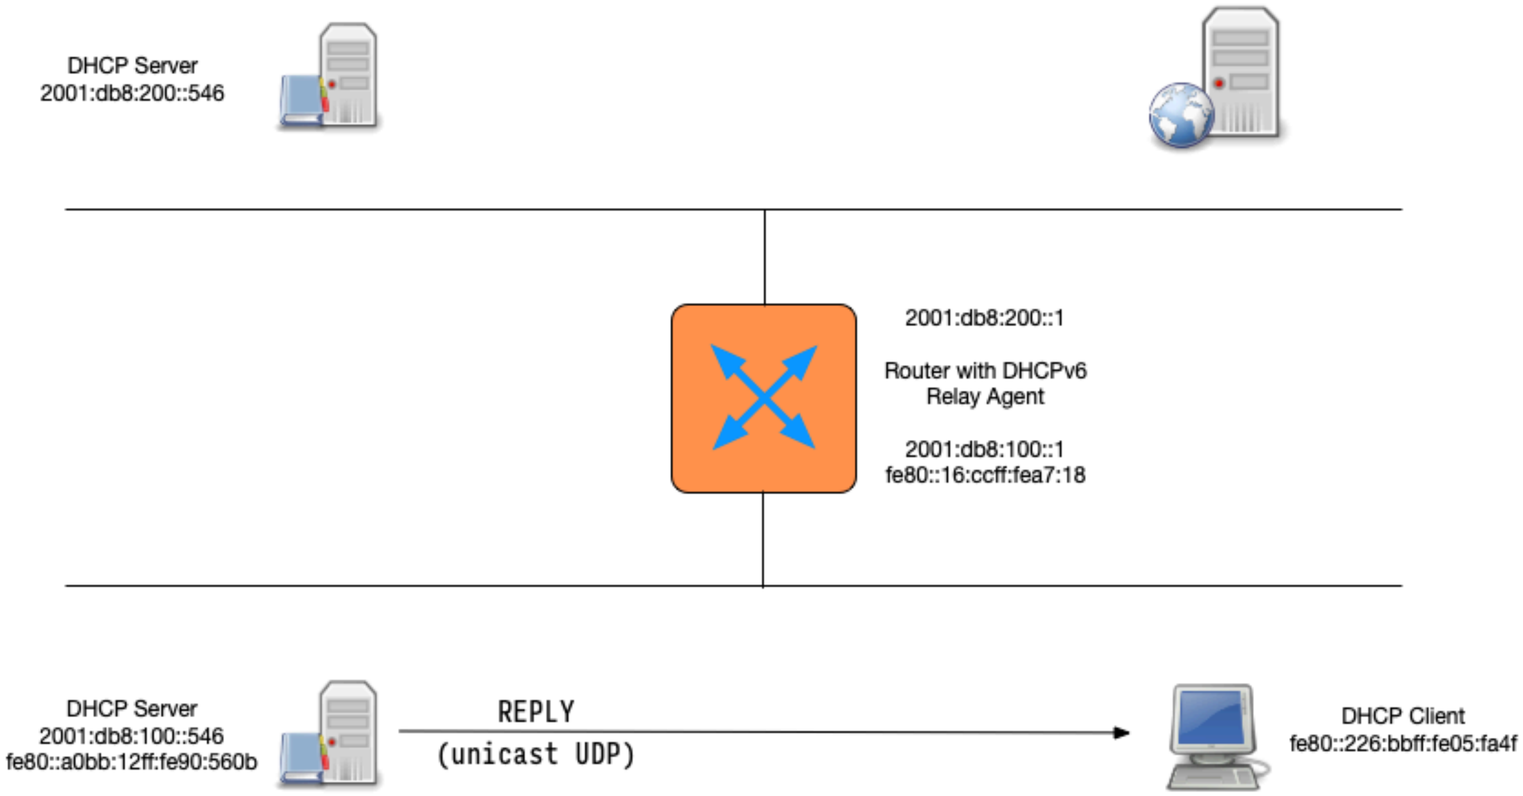

DHCPv6 client assigns a new IPv6 address

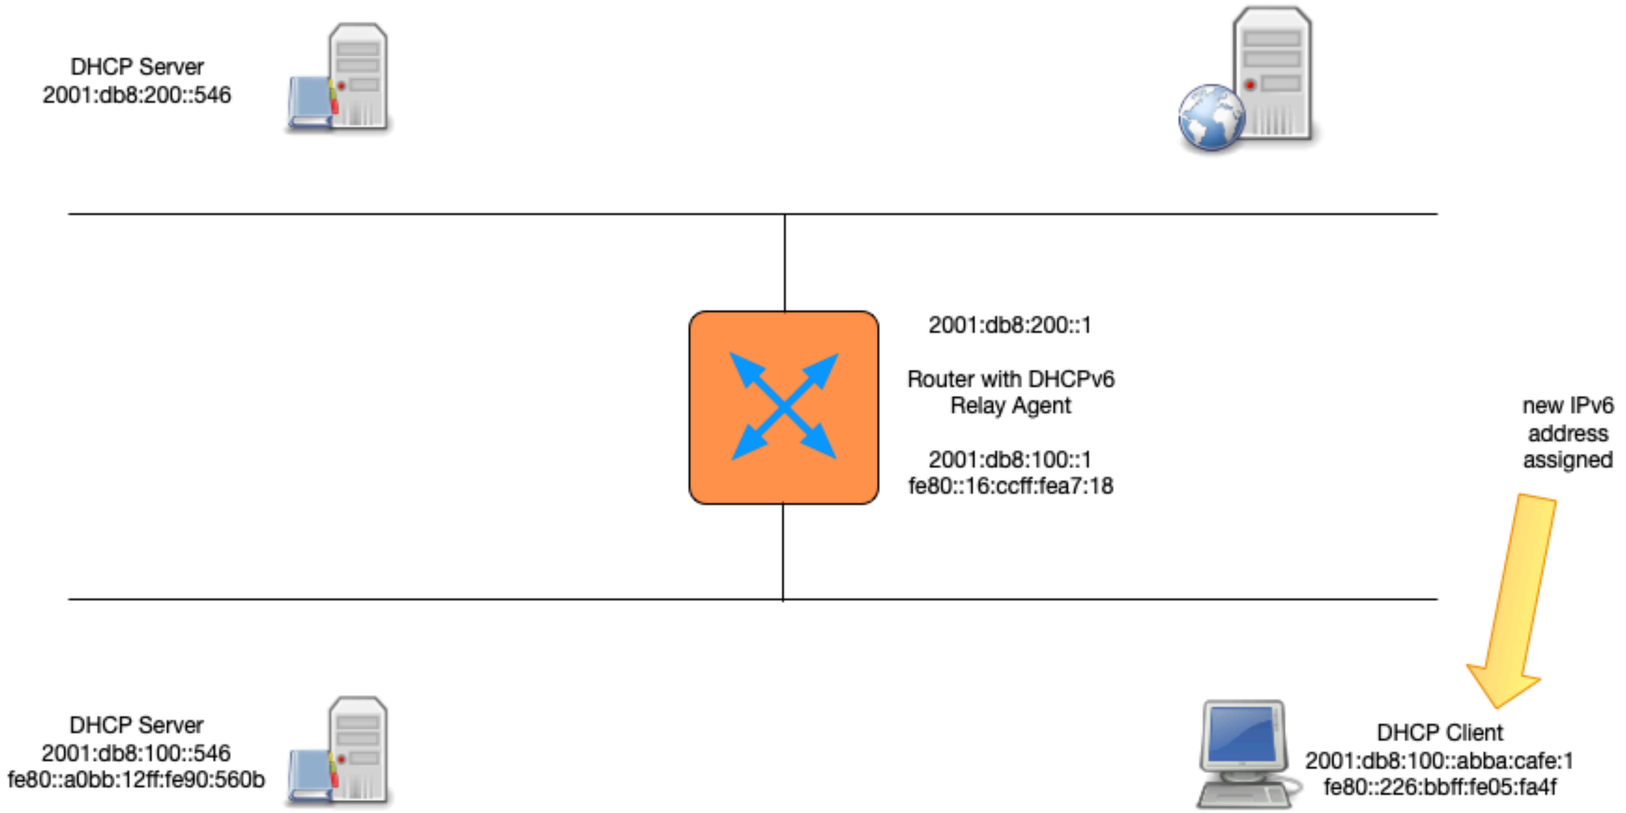

Identity Association (IA)

- An Identity Association (IA) is a construct through which a server and a client can identify, group, and manage a set of related IPv6 addresses (or delegated prefixes)
- Each IA consists of an IAID (Identity Association ID) and associated configuration information
- if a client has more than one network interface, every interface will be associated with one distinct IAID

Identity Association (IA)

- DHCPv6 clients can receive temporary and non- temporary addresses
  - temporary addresses are used for communication with outside, untrusted networks (like the Internet)
  - temporary addresses cannot be used to track a client, they are created randomly and change often
  - non-temporary addresses are stable and can be used to track a client machine (non-temporary addresses are used in trusted, internal networks)
- temporary and non-temporary (stable) IPv6 addresses are managed with the help of IAIDs

DHCPv6 client states

DHCPv6 client states

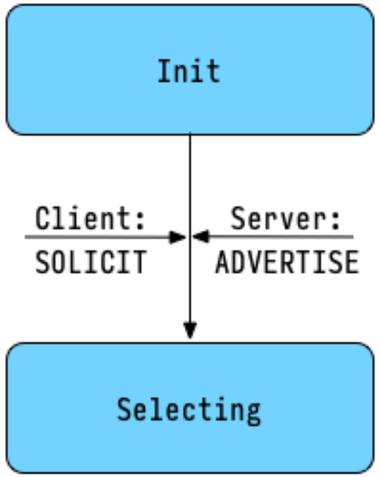

DHCPv6 client states

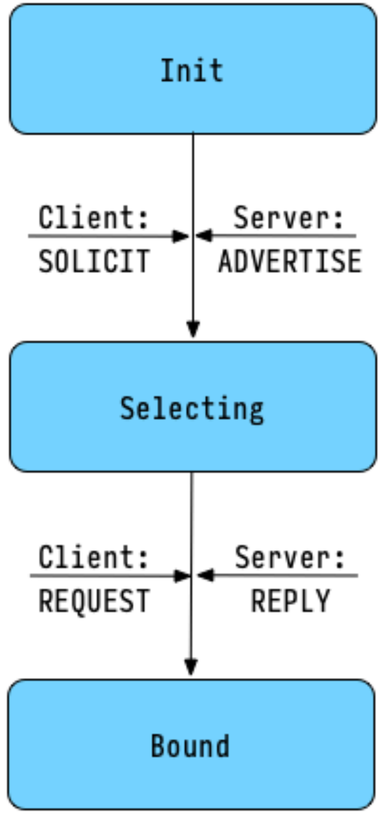

DHCPv6 client states

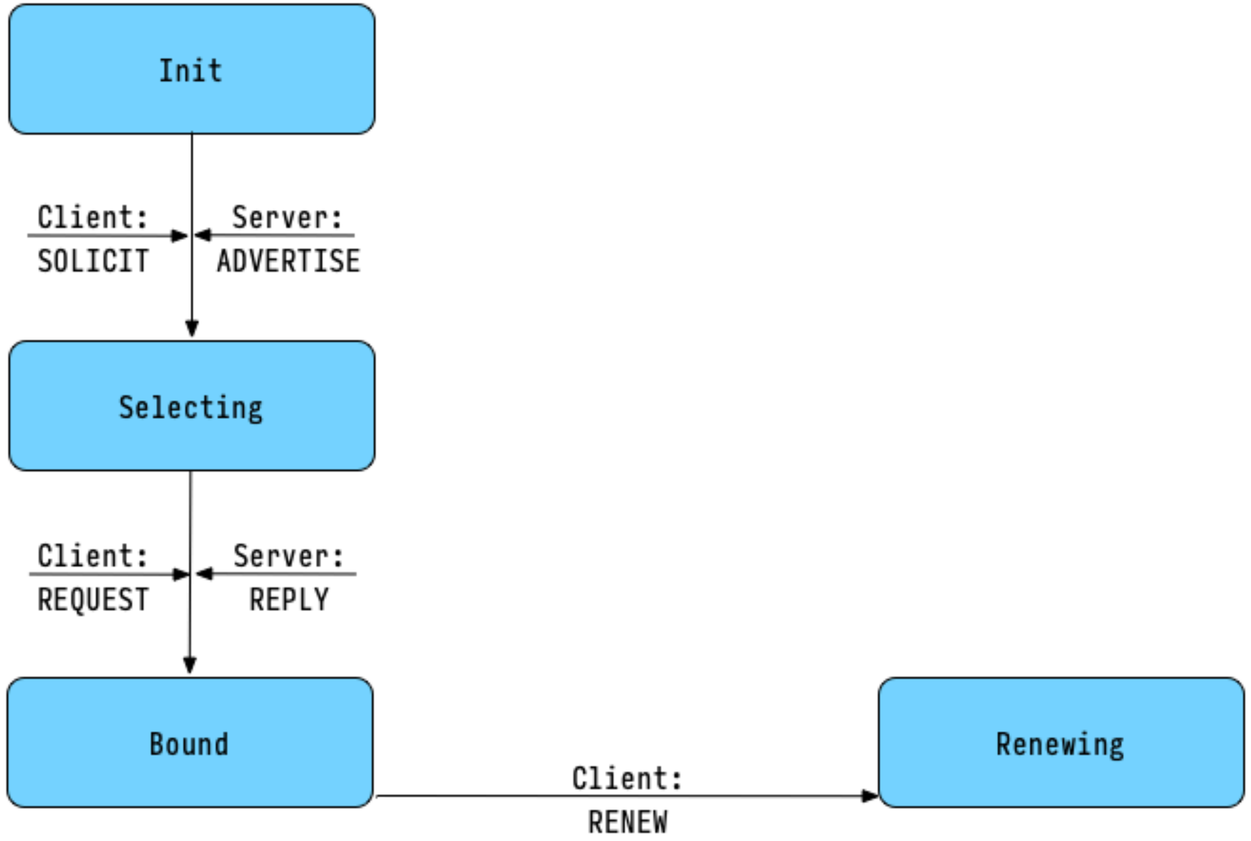

DHCPv6 client states

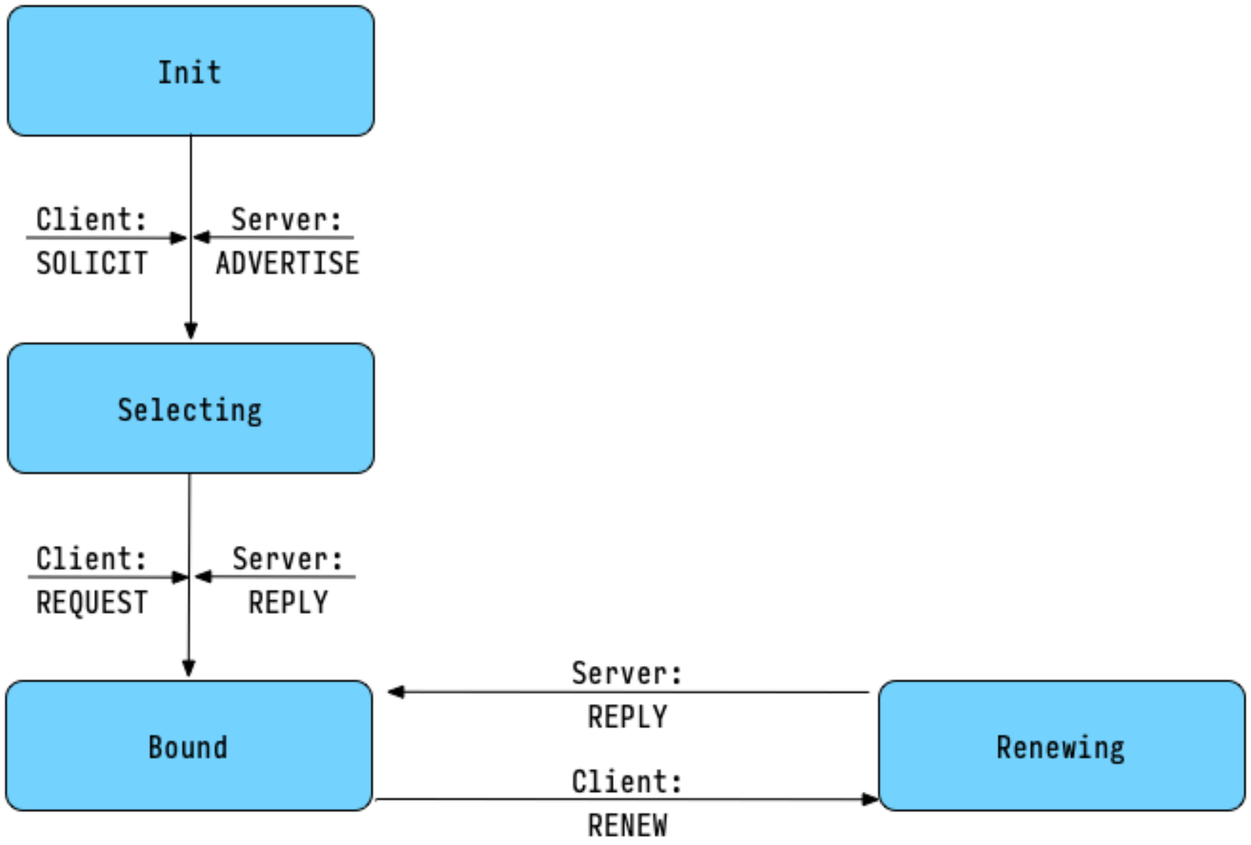

DHCPv6 client states

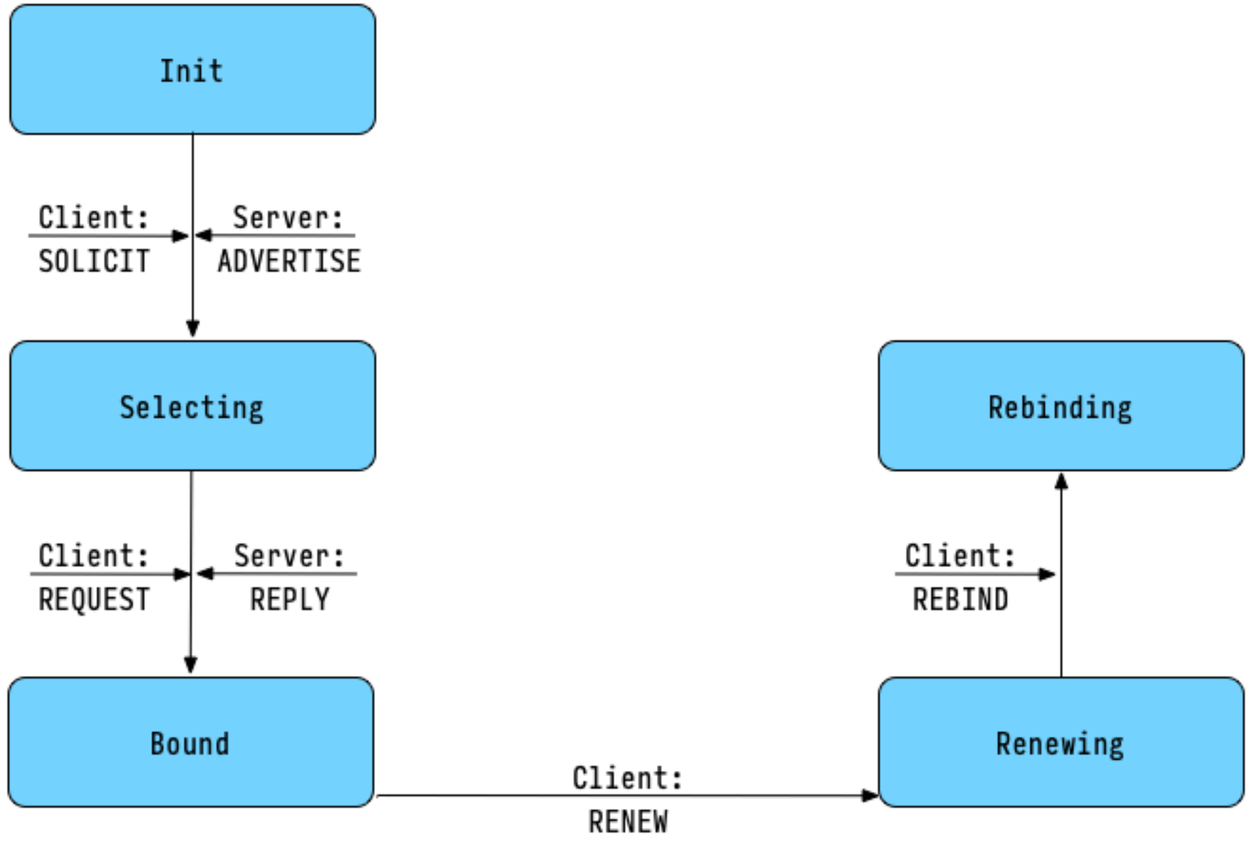

DHCPv6 client states

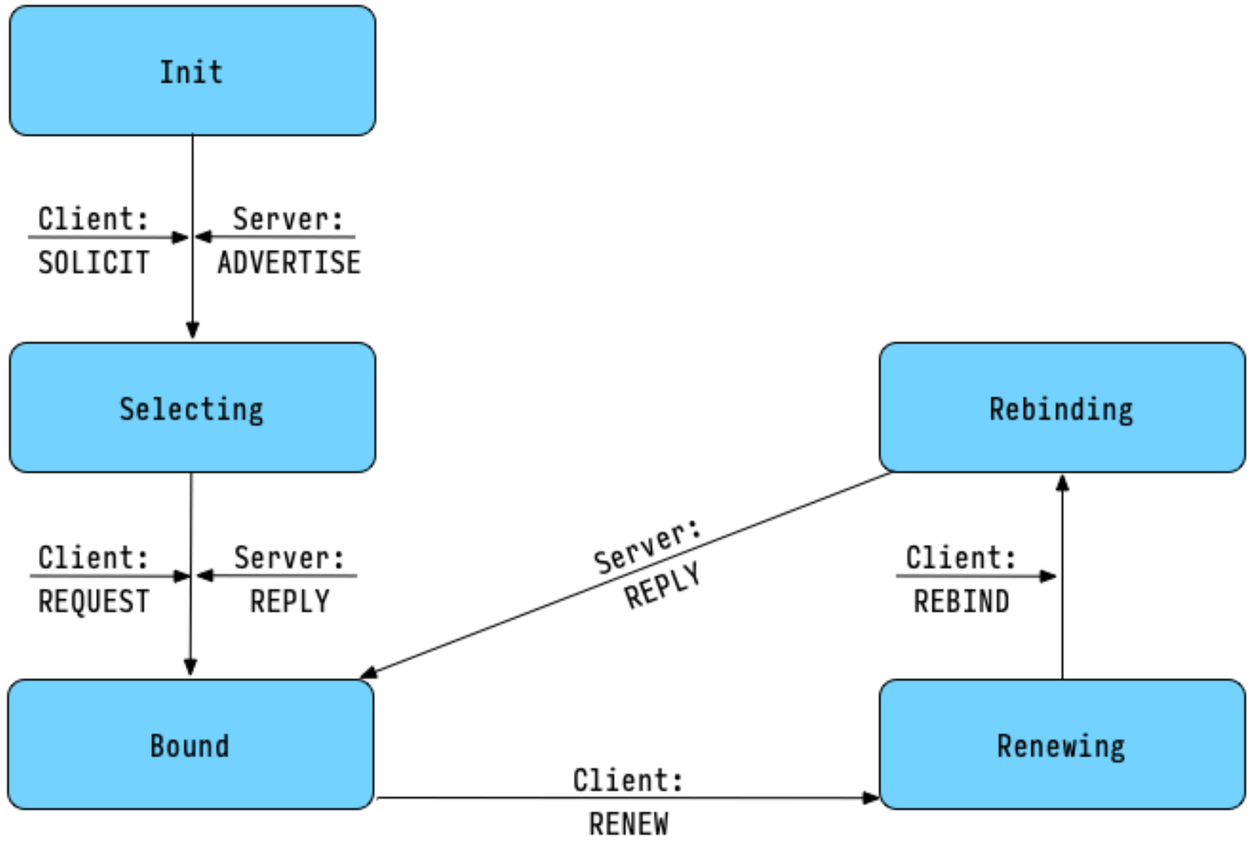

DHCPv6 client states

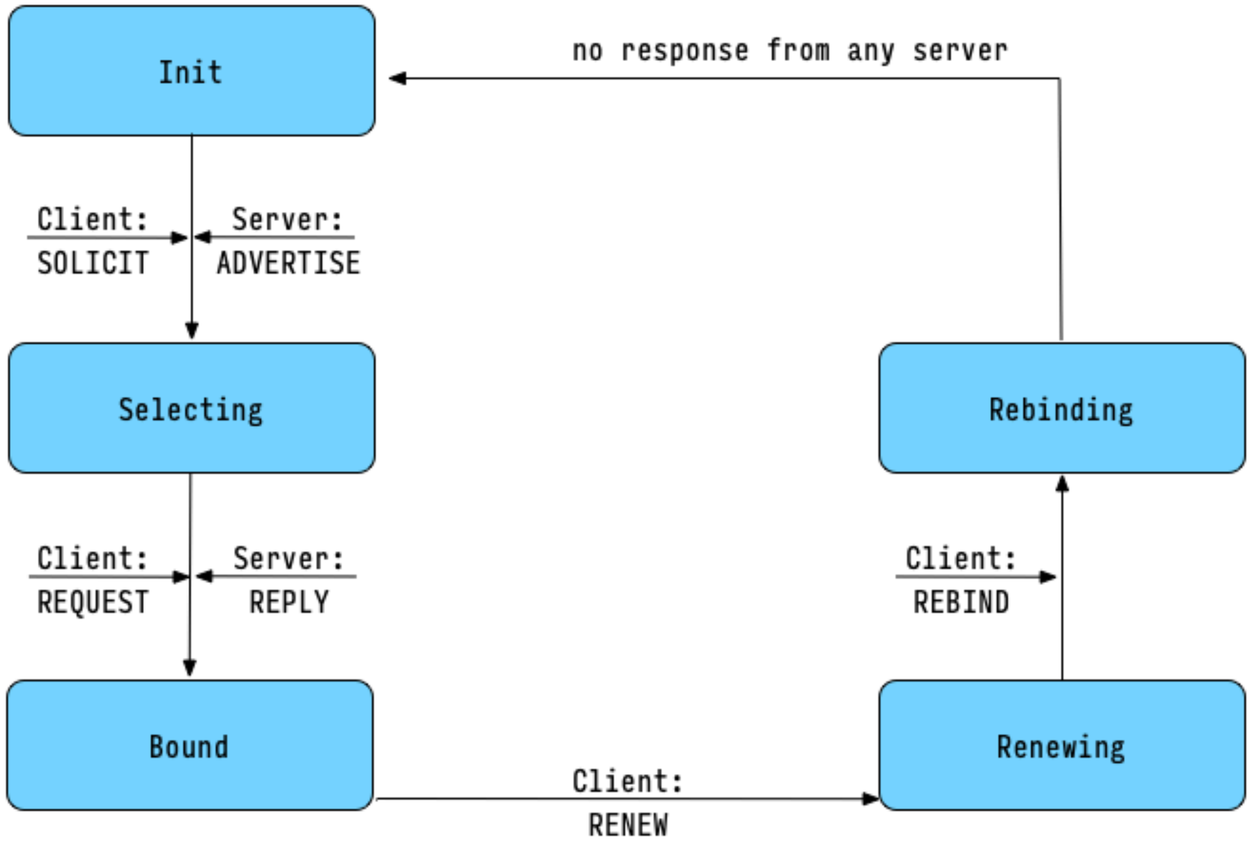

DHCPv6 rapid commit

- **rapid commit** speeds up the process of joining a network
- with "rapid commit" there is no information for the DHCPv6 server if the client is using the advertised IPv6 address
  - the DHCPv6 server must reserve the IPv6 address for the full lease time
    - this (temporary) squandering of IPv6 addresses is usually not a problem because of the large size of IPv6 subnets (/64 prefixes)

DHCPv6 rapid commit

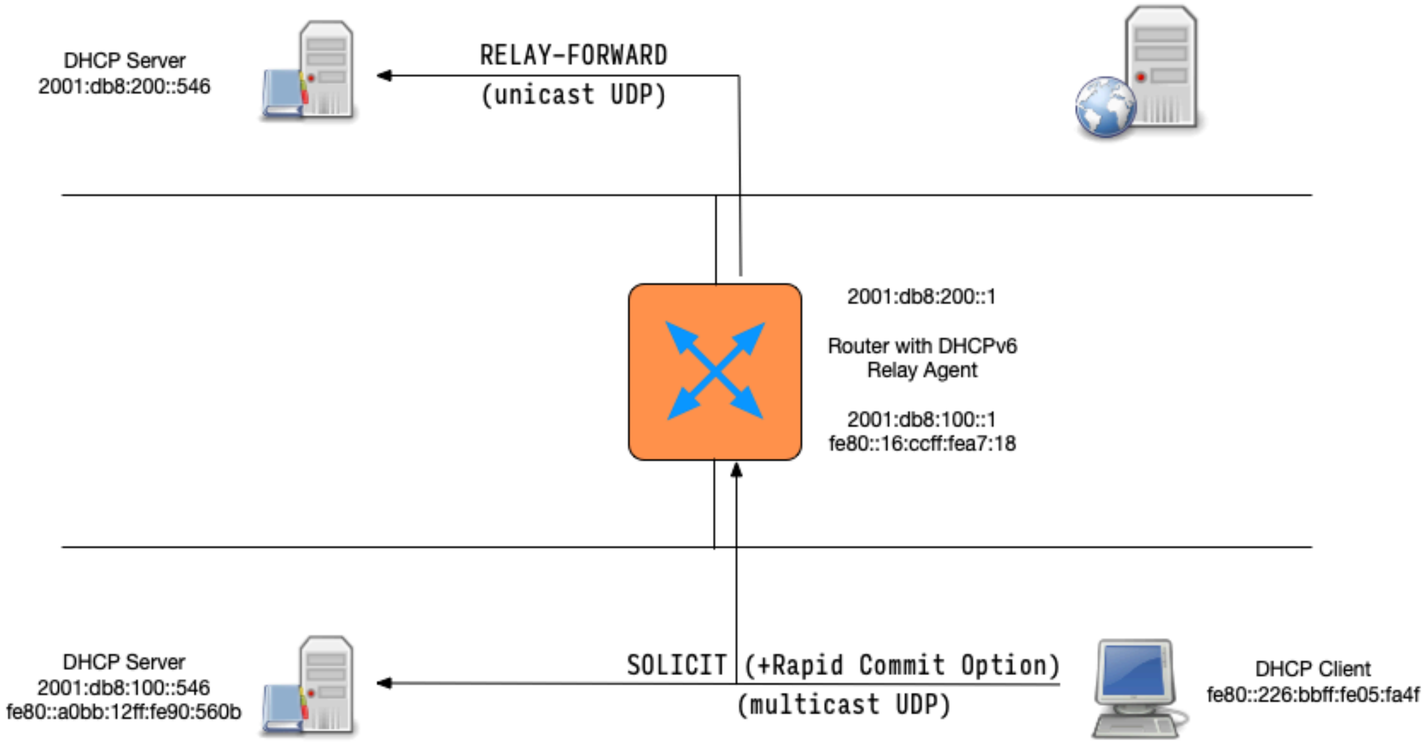

DHCPv6 rapid commit

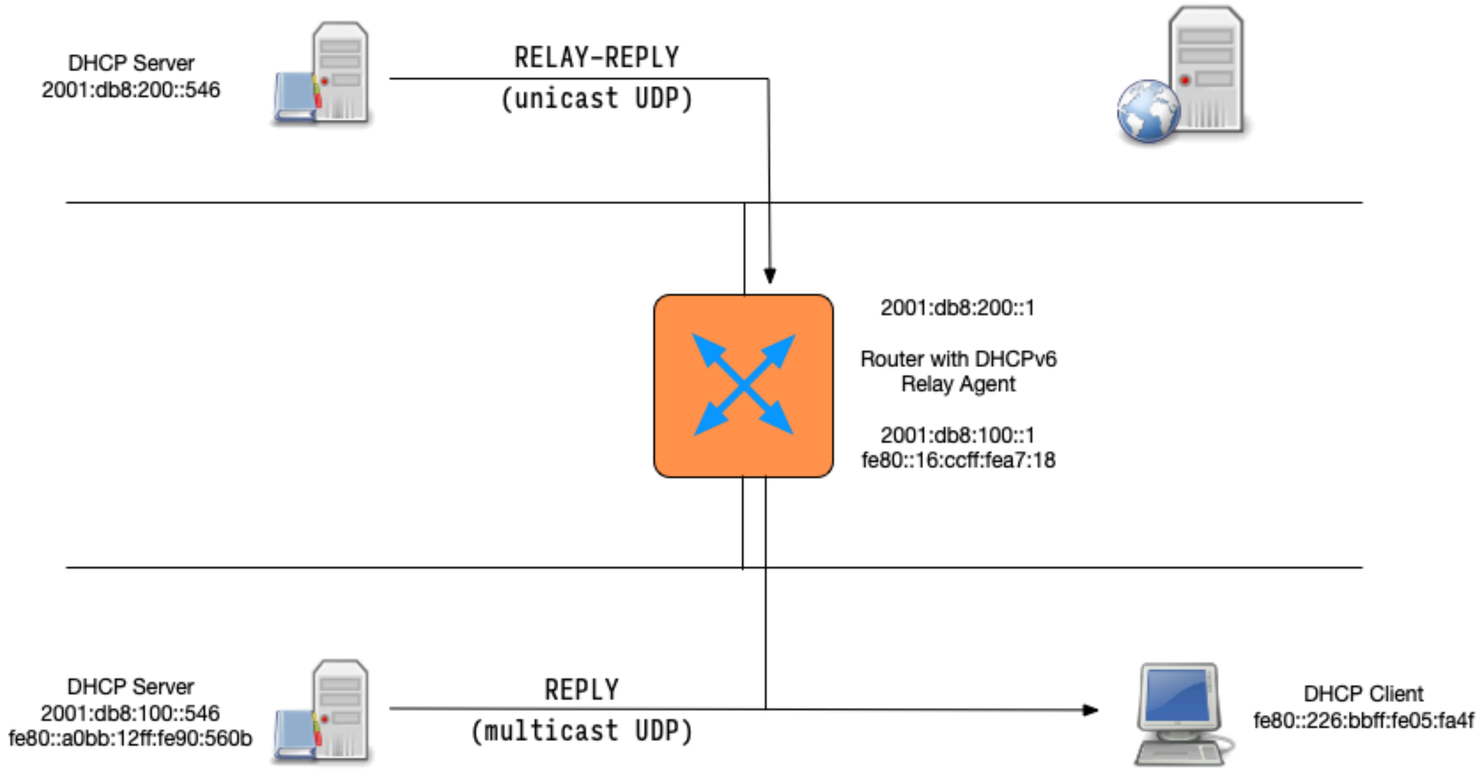

DHCPv6 rapid commit

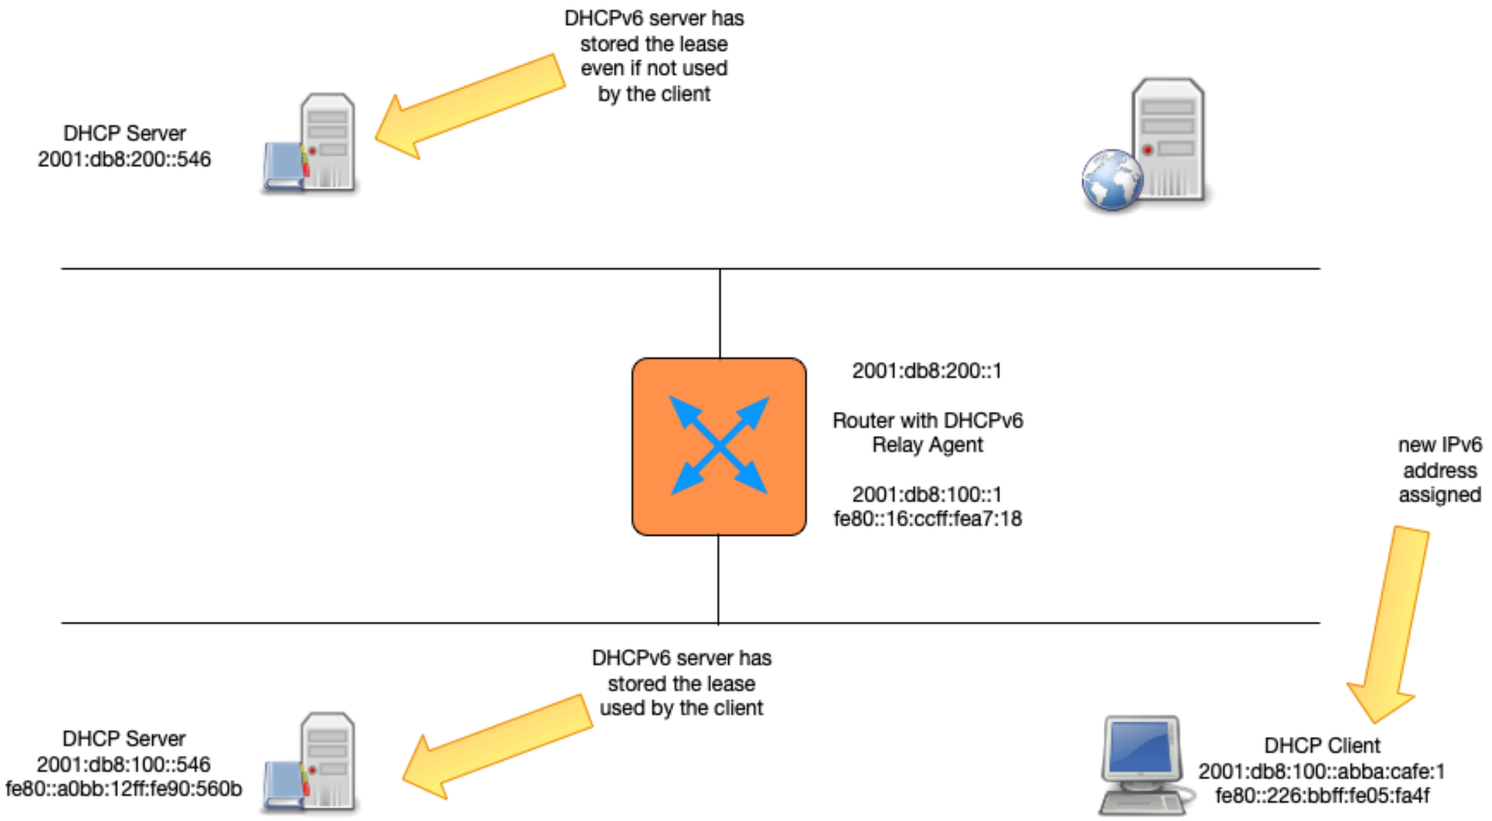

References: RFCs, Books, recommended Webpages

Internet Standards

- DHCPv4
  - RFC 2131 - DHCPv4 Dynamic Host Configuration Protocol
  - RFC 2132 - DHCP Options and BOOTP Vendor Extensions
  - RFC 3396 - Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)
  - RFC 4361 - Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)
  - RFC 6842 - Client Identifier Option in DHCP Server Replies
- DHCPv6
  - RFC 8415 - Dynamic Host Configuration Protocol for IPv6 (DHCPv6)

Books

- The DHCP Handbook - Understanding, Deploying, and Managing Automated Configuration Services (Ralph Droms, Ted Lemon) 1999
- IP Address Management - Principles and Practice (Timothy Rooney) 2011
- The TCP/IP Guide - A Comprehensive, Illustrated Internet Protocols Reference (Charles M. Kozierok) 2005
- Windows Server 2019 Inside Out (Orin Thomas)

Websites

- ISC Kea Documentation
  https://kea.readthedocs.io/en/latest/
- ISC Knowledgebase
  https://kb.isc.org/
- The TCP Guide
  http://www.tcpipguide.com/
- Microsoft - Dynamic Host Configuration Protocol (DHCP)
  https://docs.microsoft.com/en-us/windows-server/networking/technologies/dhcp/dhcp-top

Next Webinars

- 30th September - Kea DHCP - Installation and configuration
- 14th October - Kea DHCP - Lease allocation, client classification, and option assignment
- 28th October - Kea DHCP - High Availability and Database Backends
- 18th November - Kea DHCP - Monitoring, Logging, and Stork
- 2nd December - Kea DHCP - Migrating to Kea from ISC DHCP

Questions and Answers