

Introduction to Computer Networks

# **NAT, IPv6, and Multicast**

<https://pages.cs.wisc.edu/~mgliu/CS640/S26/index.html>

**Ming Liu**

**mgliu@cs.wisc.edu**

# Outline

- Last
  - Inter-domain Routing
- Today
  - NAT, IPv6, and Multicast
- Announcements
  - Lab3 due on 03/27/2026

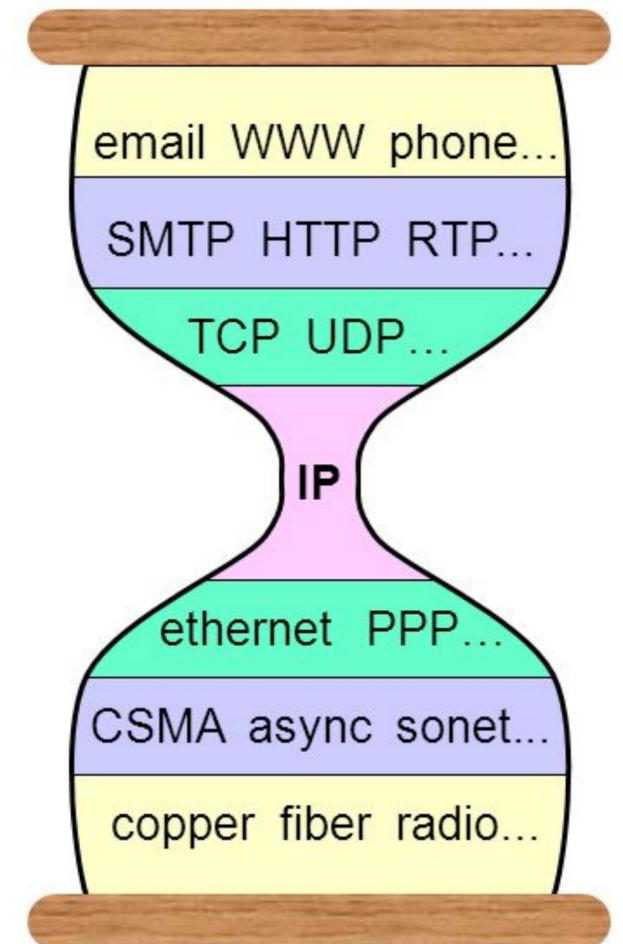
# Recap: BGP

- Key Questions:
  - How can we interconnect domains at scale?
- Terminology
  - Autonomous System (AS)
  - BGP (Border Gateway Protocol)
  - Policy-driven Routing
  - Learn and Speak
  - Peering and Customer-Provider

**IP is powerful,**

# Recap: The Hour Glass Model

- Hide underlying L2 technologies
  - E.g., Ethernet, Wi-Fi, etc.
- Support many different types of apps
  - E.g., Email, browsing, streaming, etc.
- Provide minimal functionalities
  - Addressing, forwarding, and routing



**IP is powerful, but the Internet is rising rapidly!**

# The Scalability Challenge

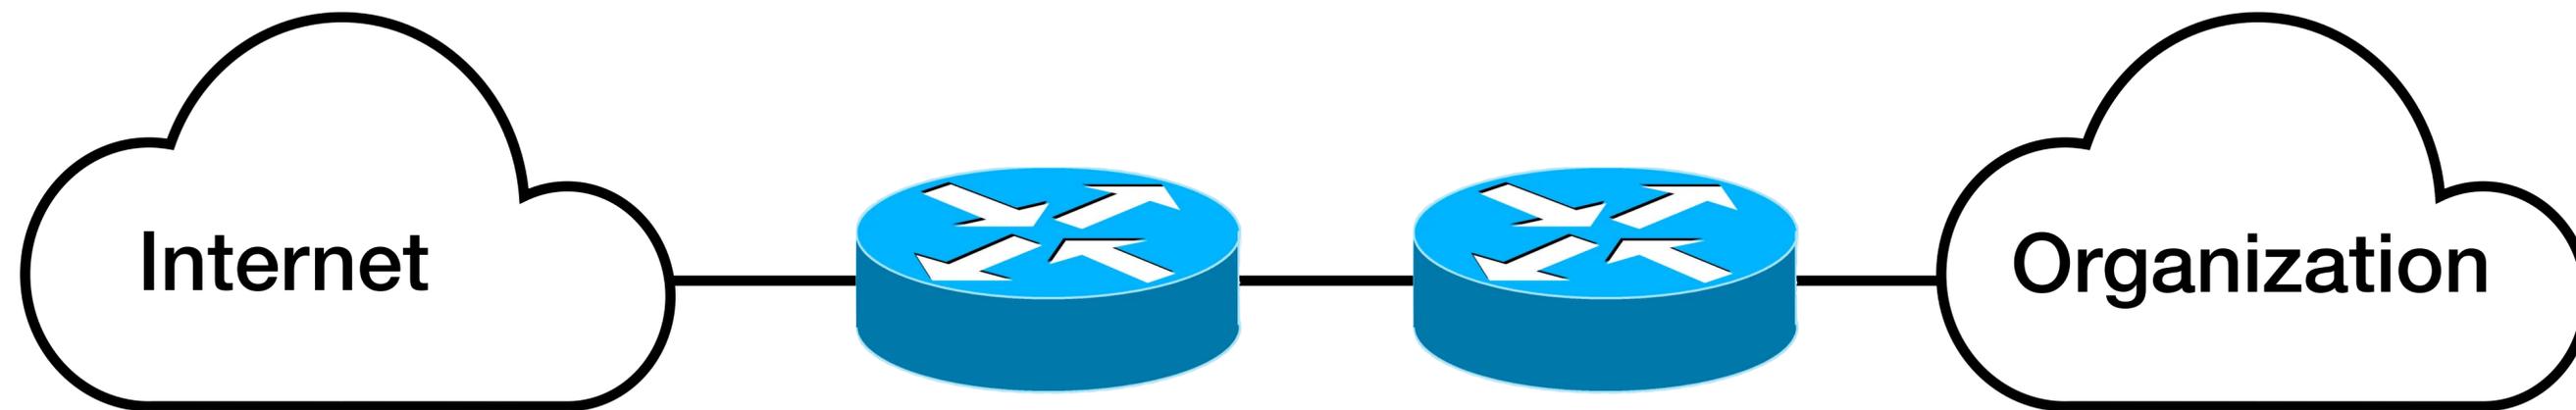
- Q1: How can we support a rising number of hosts?
- Q2: How can we design the routing scheme to reduce traffic?
- Q3: How can we build the underlying infrastructure to provide enough routing capacity?

# The Scalability Challenge

- Q1: How can we support a rising number of hosts?
- Q2: How can we design the routing scheme to reduce traffic?
- *Q3: How can we build the underlying infrastructure to provide enough routing capacity?*

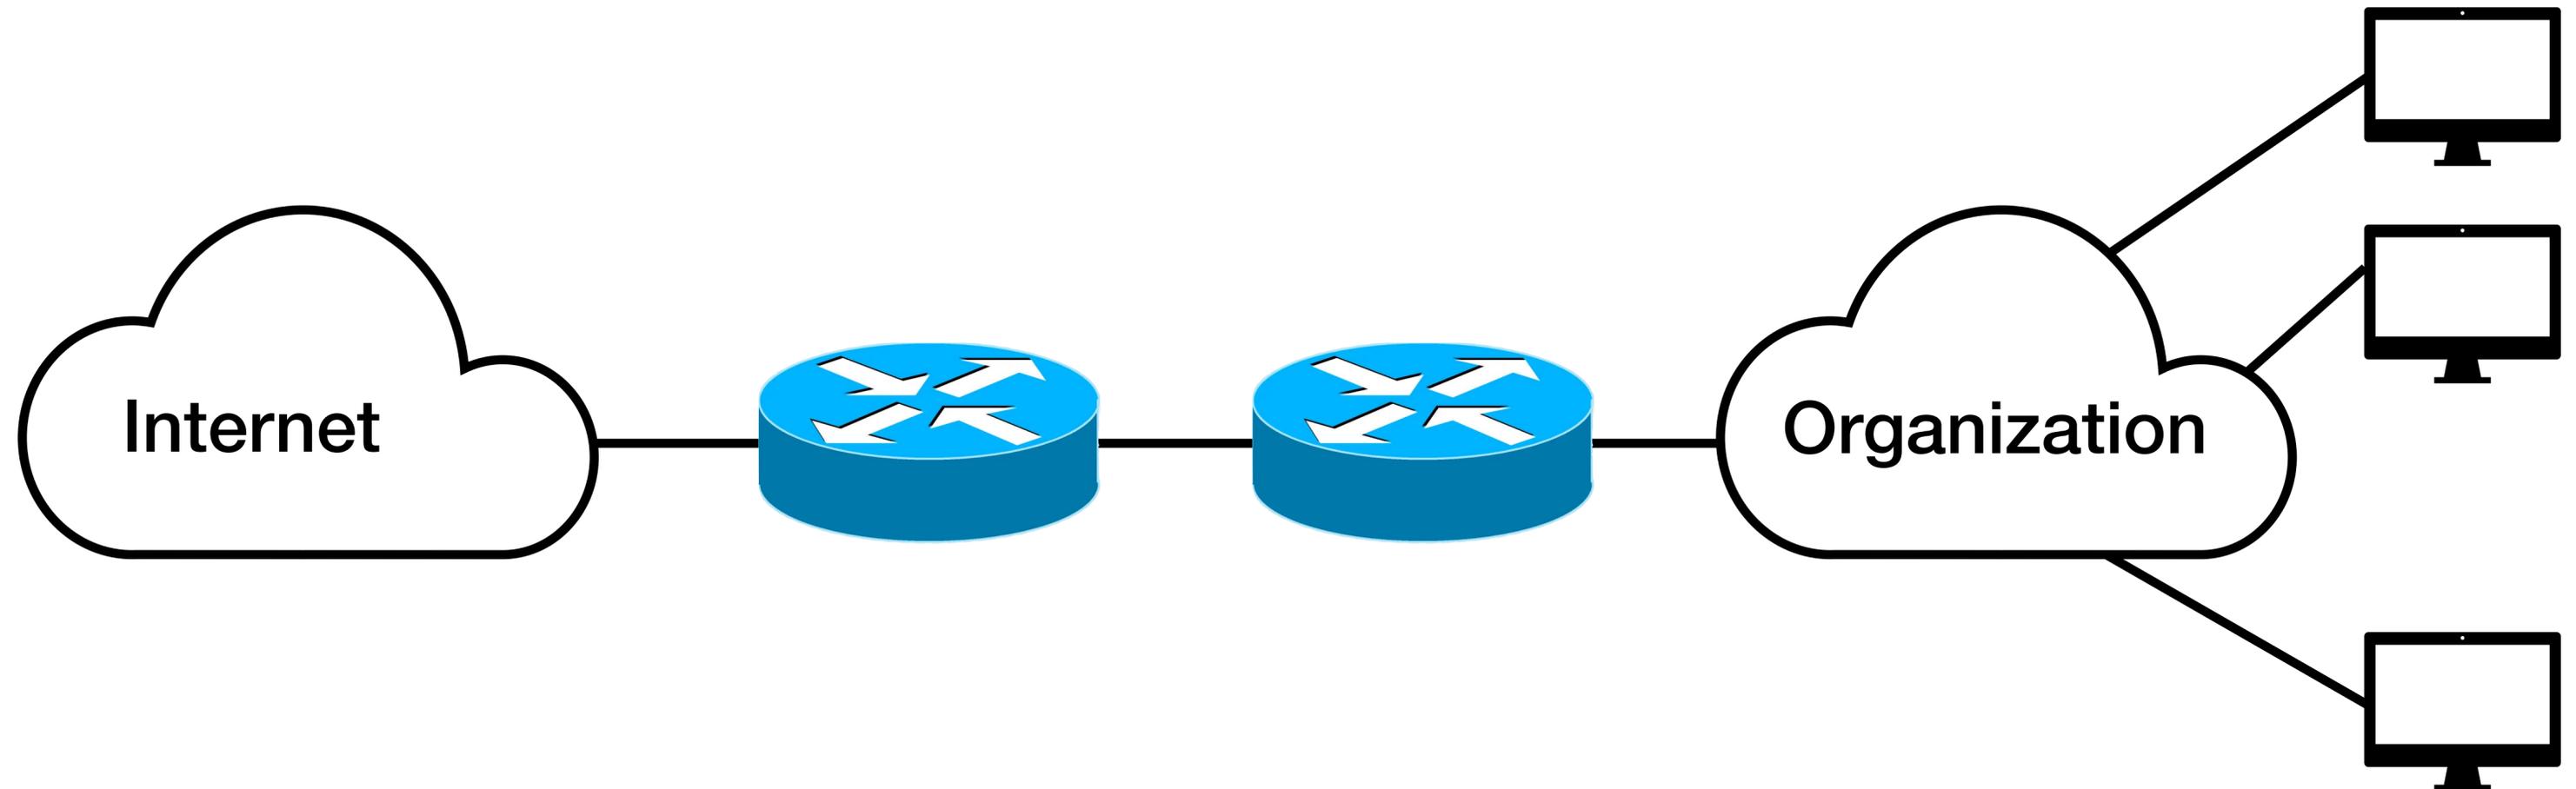
# Private Network

- The proliferation of LANs at small offices, home offices,...



# Private Network

- The proliferation of LANs at small offices, home offices,...
- Work with ISP to allocate a contiguous IP address range



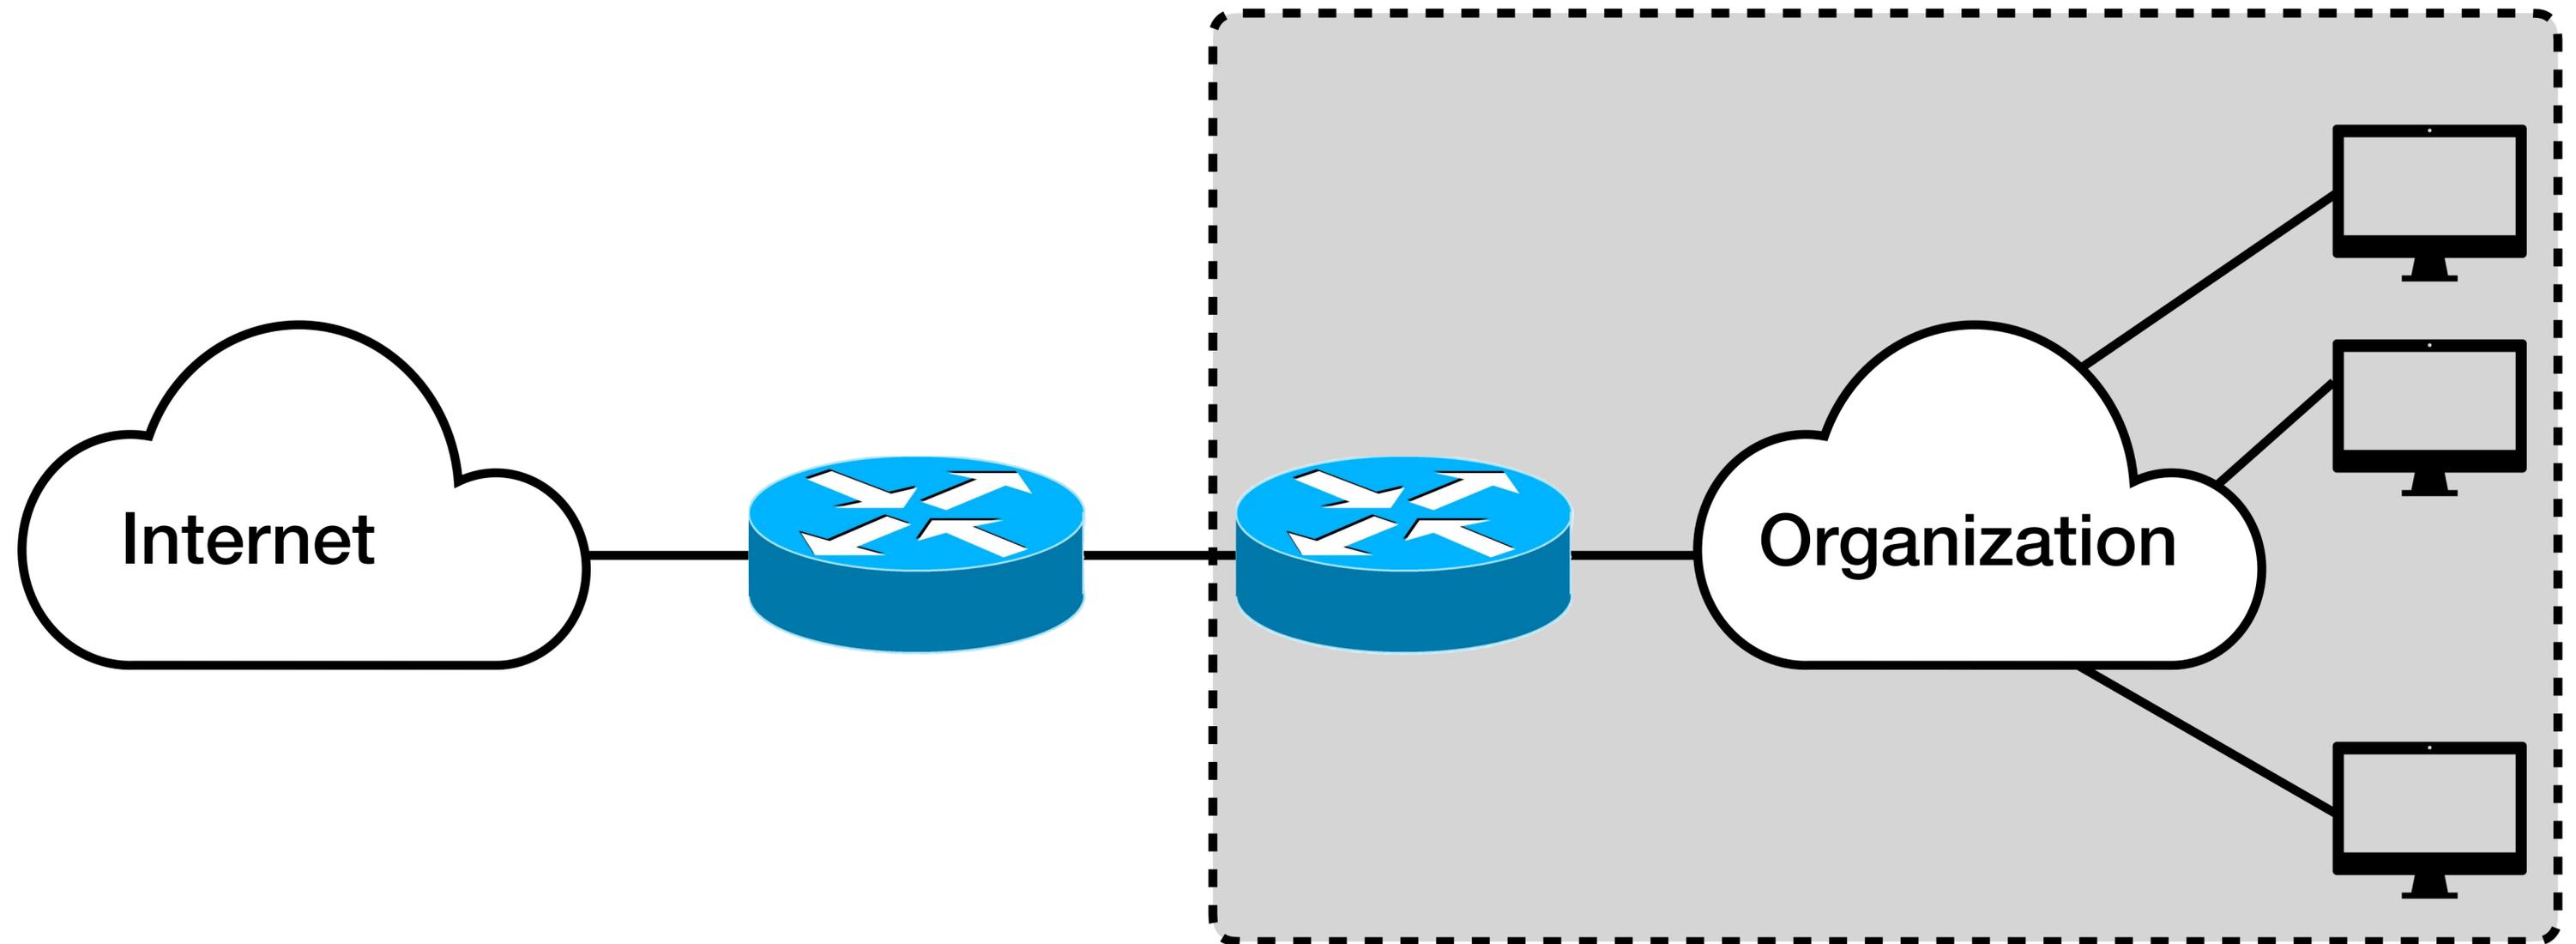
# Private Network

- The proliferation of LANs at small offices, home offices,...
- Work with ISP to allocate a contiguous IP address range

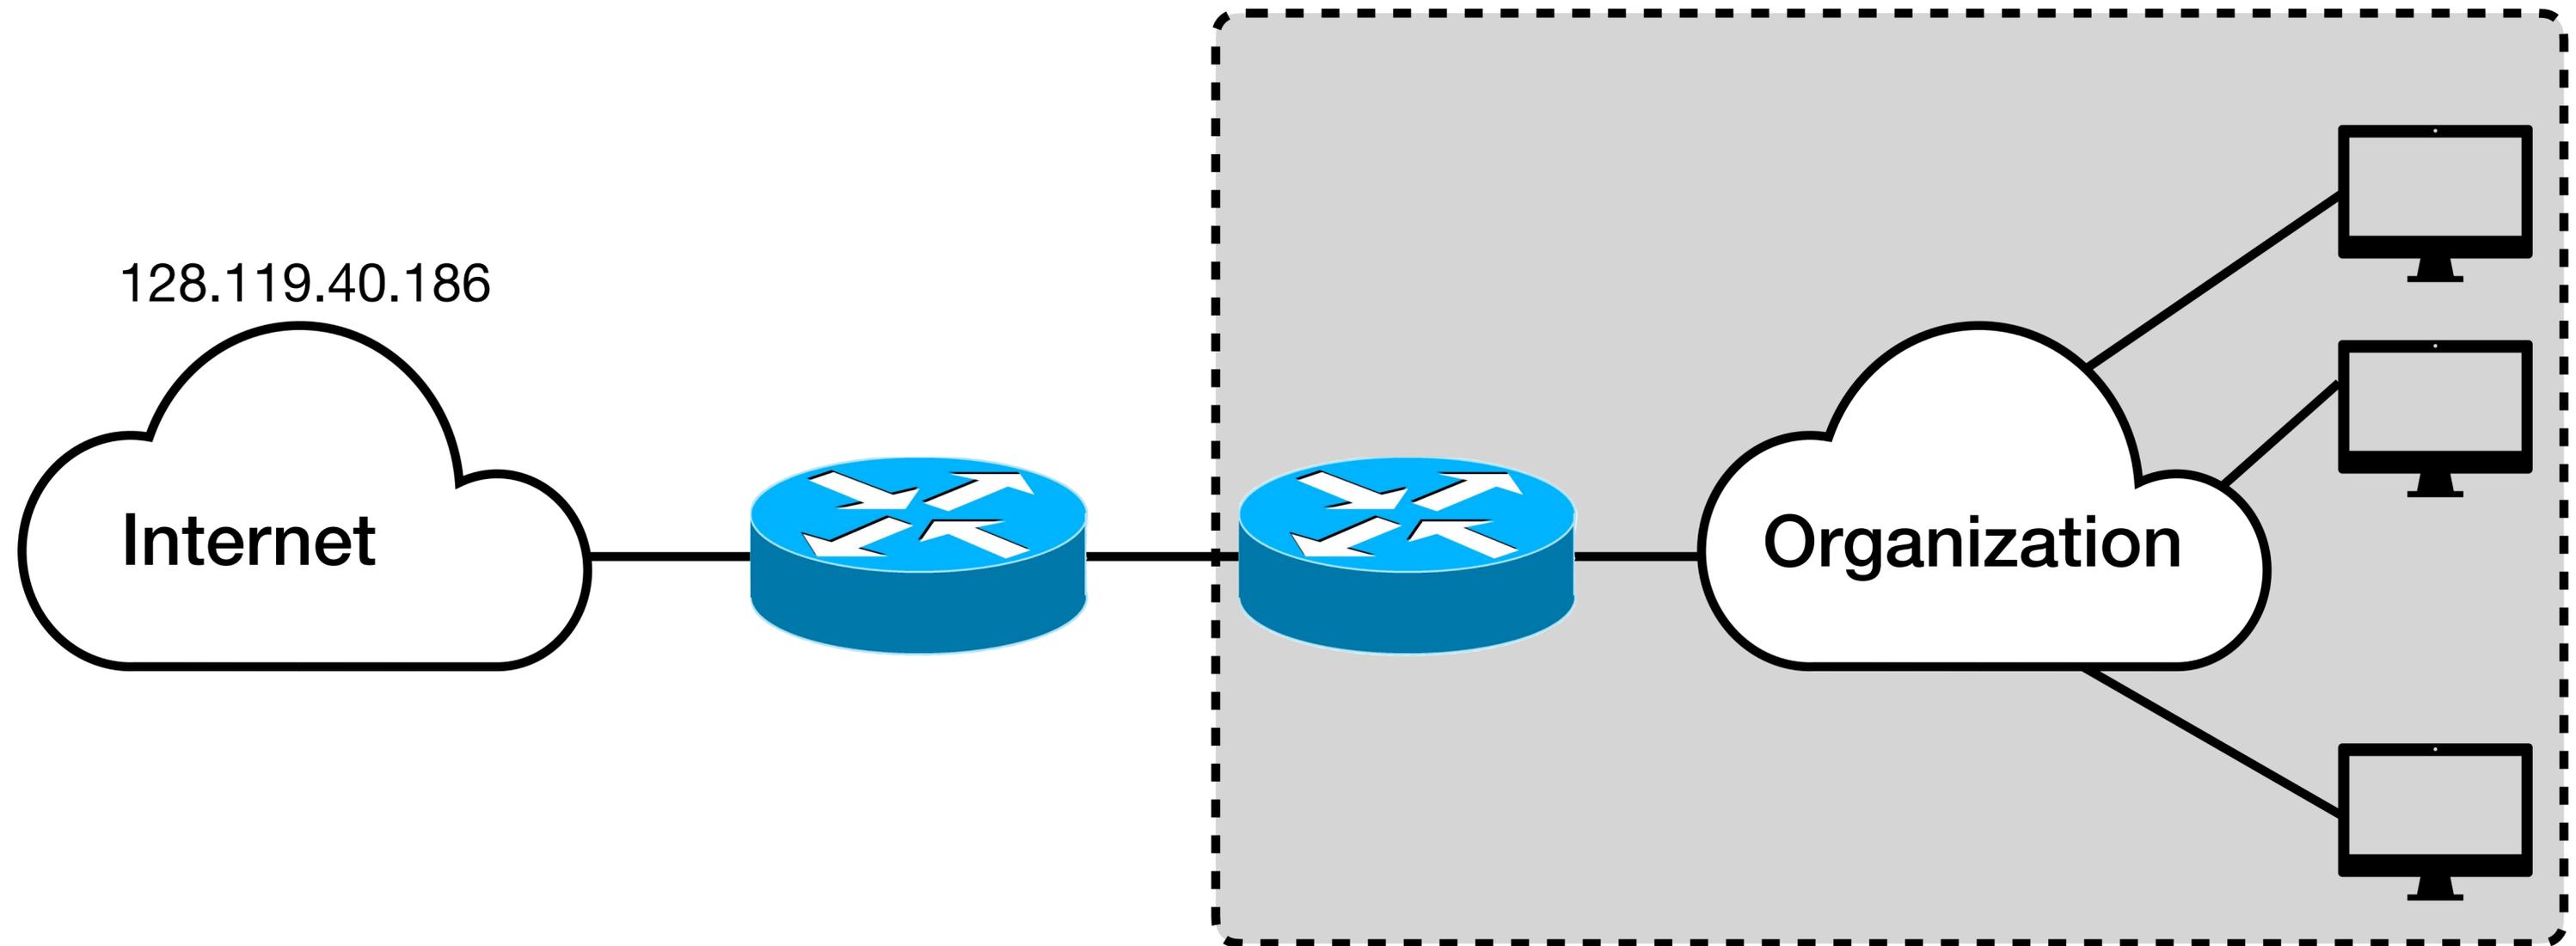
- How can we handle the host increase?
- How can we handle host replacement?
- How can we reduce the cost?
- Can we manage the IP address by ourselves, but still make the host connect to the Internet?

# Network Address Translation (NAT)

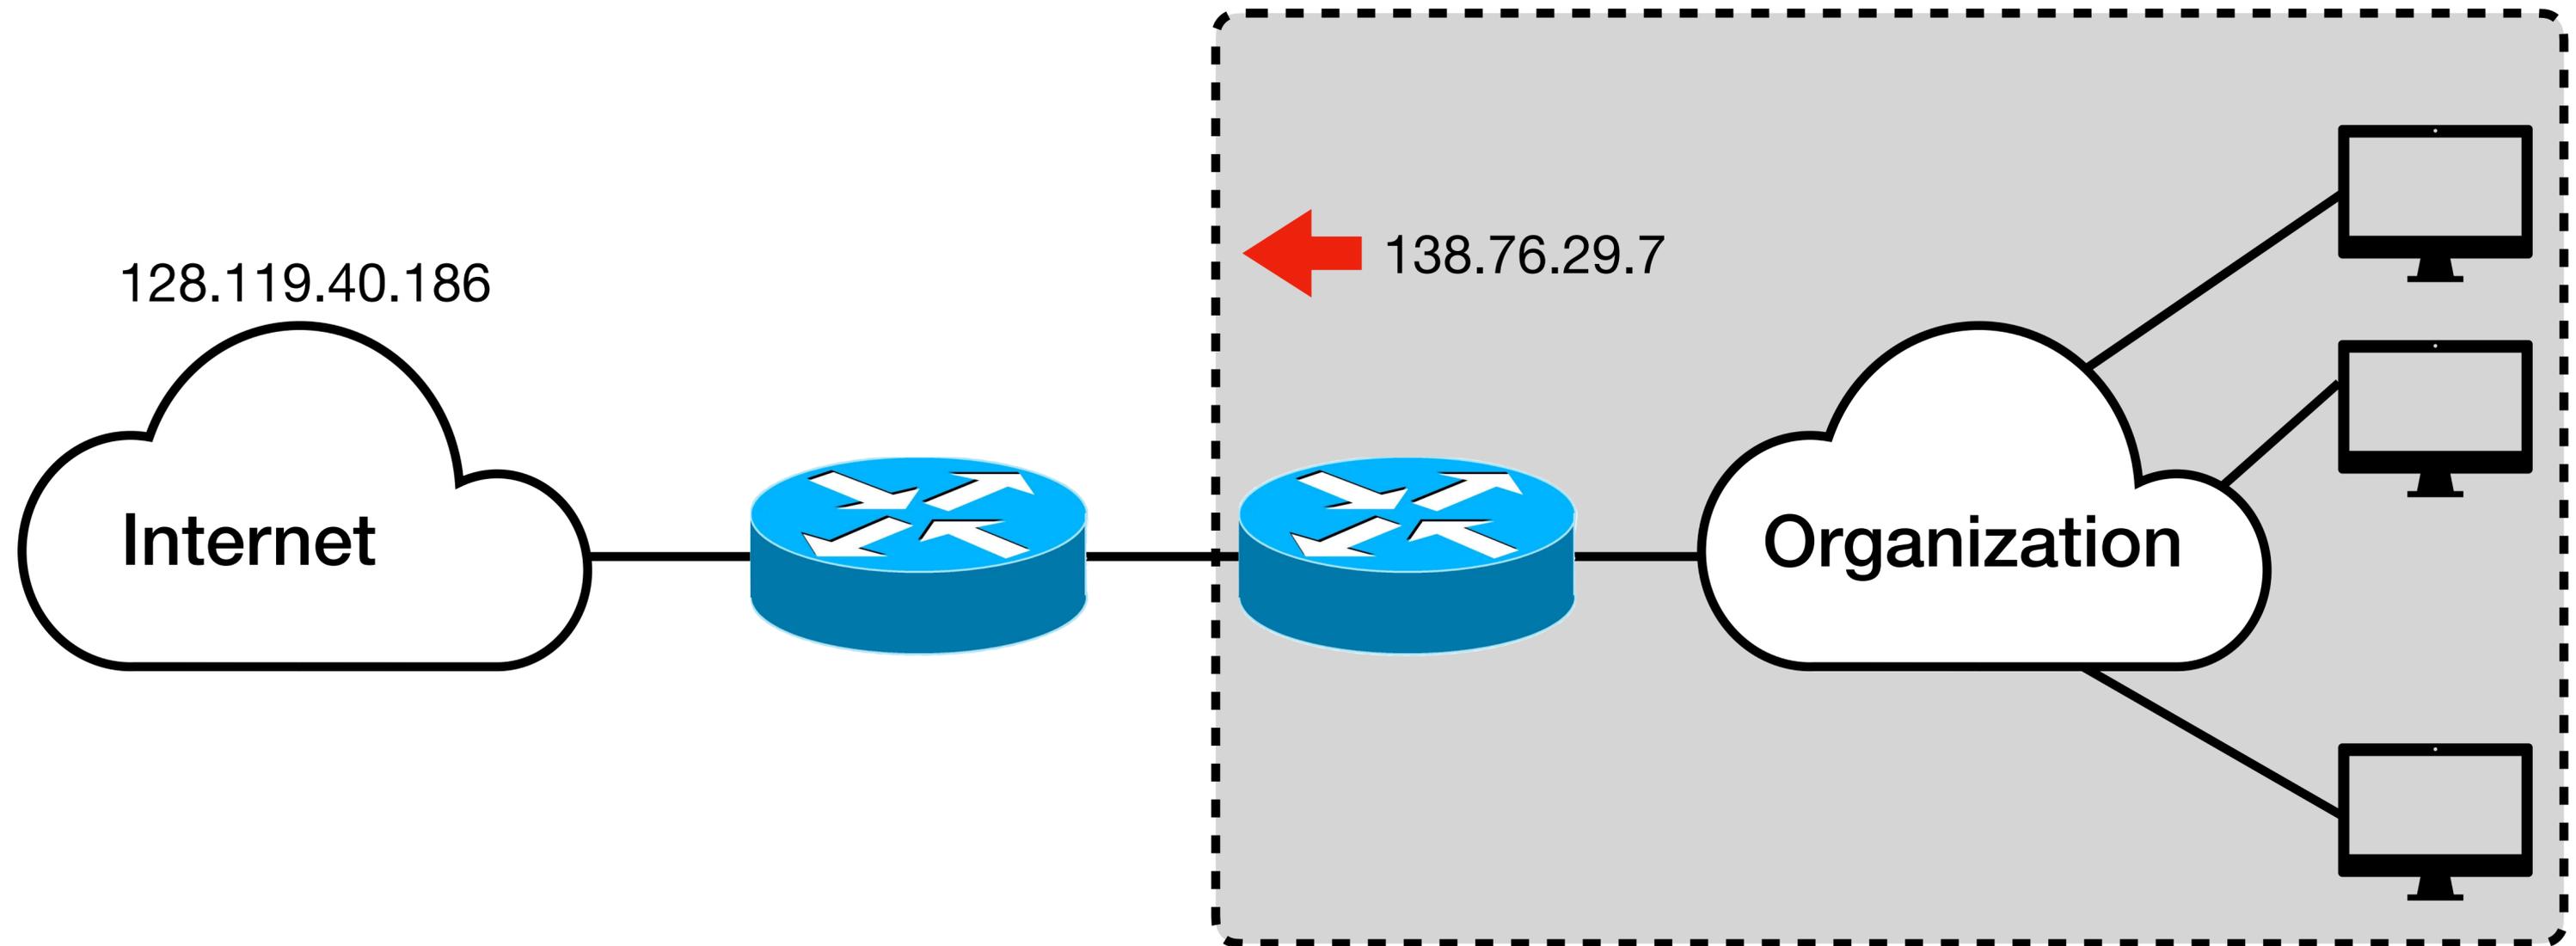
- Abstract the private network as a single device to the Internet
- RFC 2663; RFC 3022



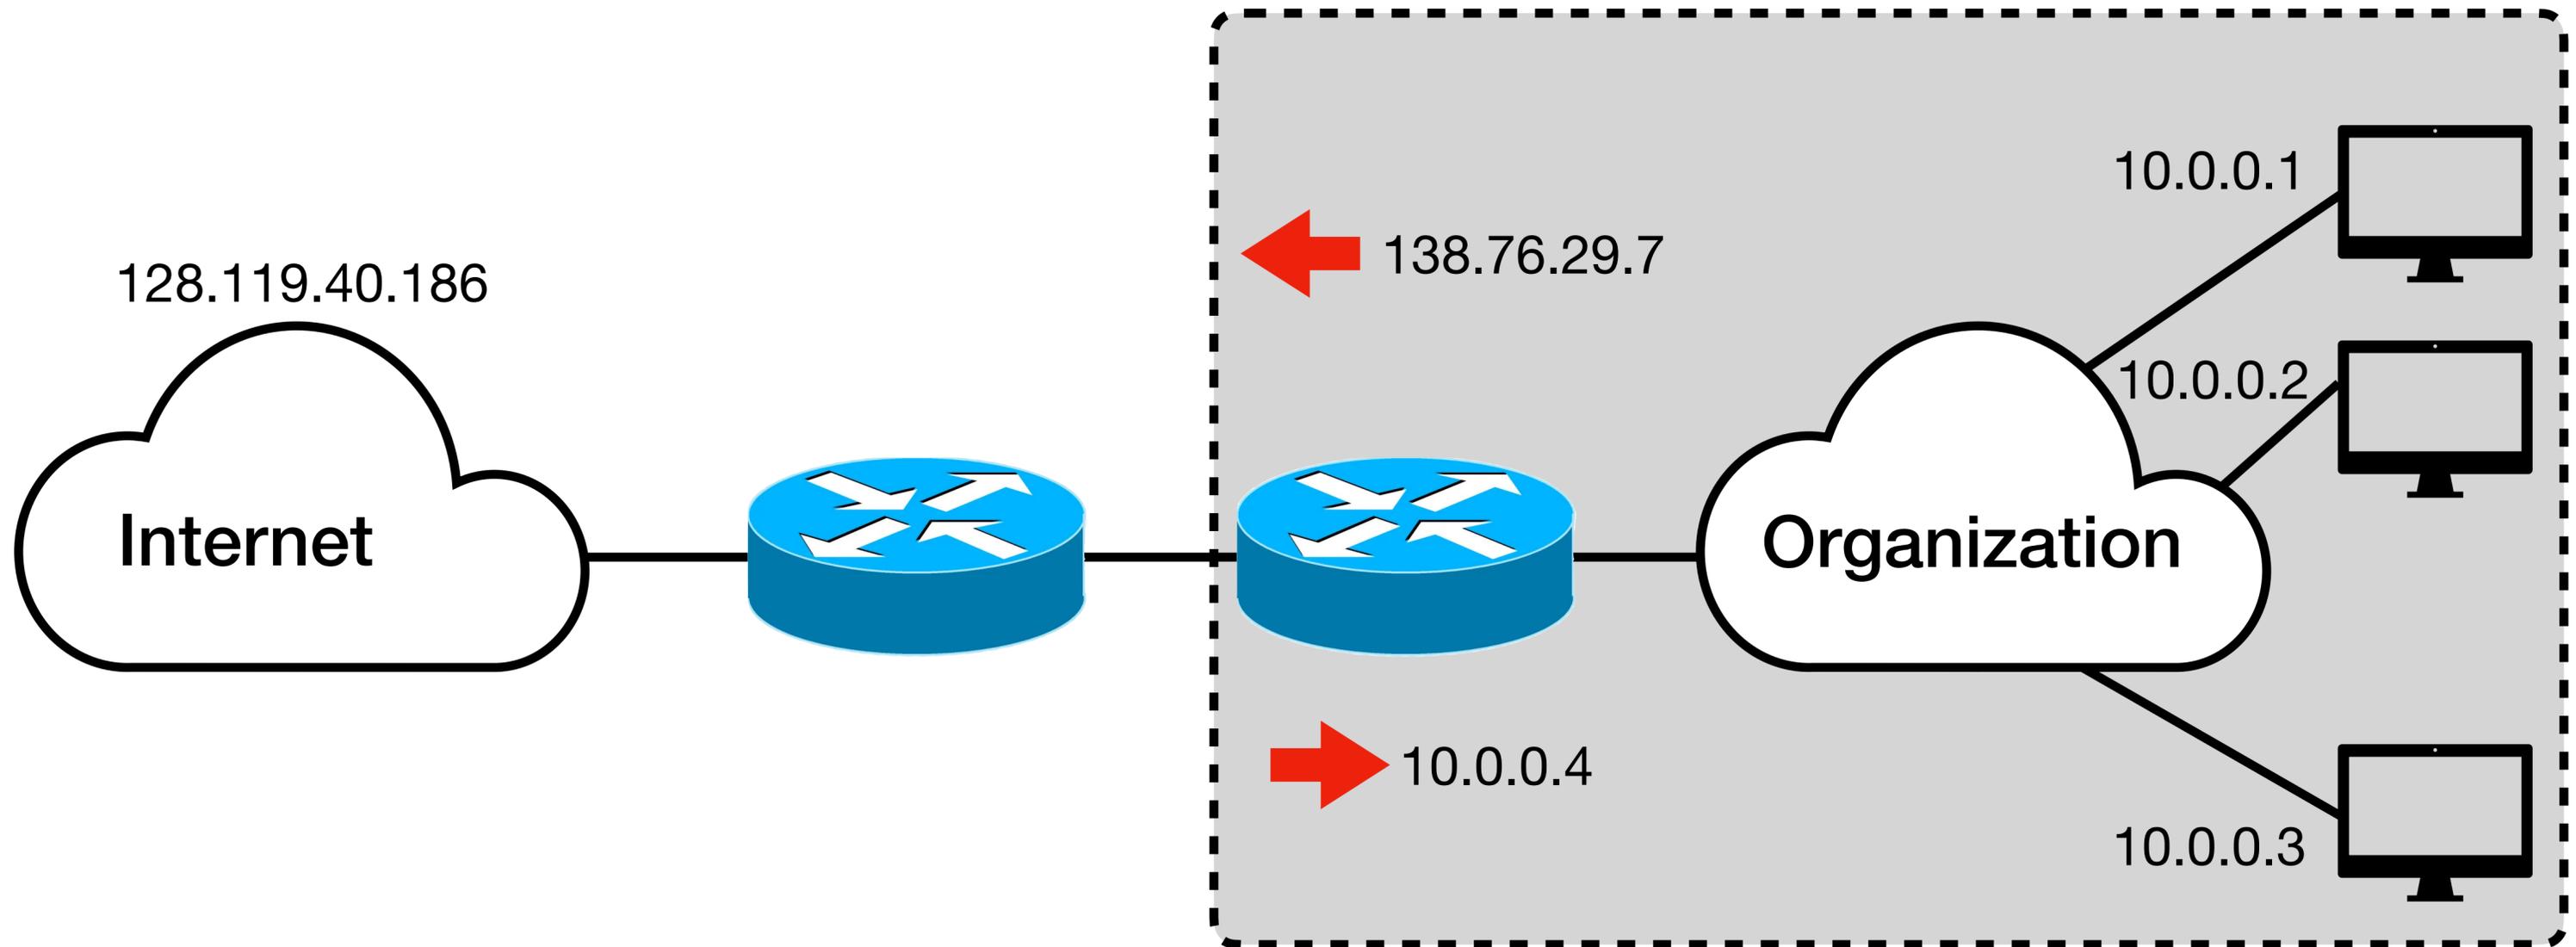
# How Does NAT Work?



# How Does NAT Work?

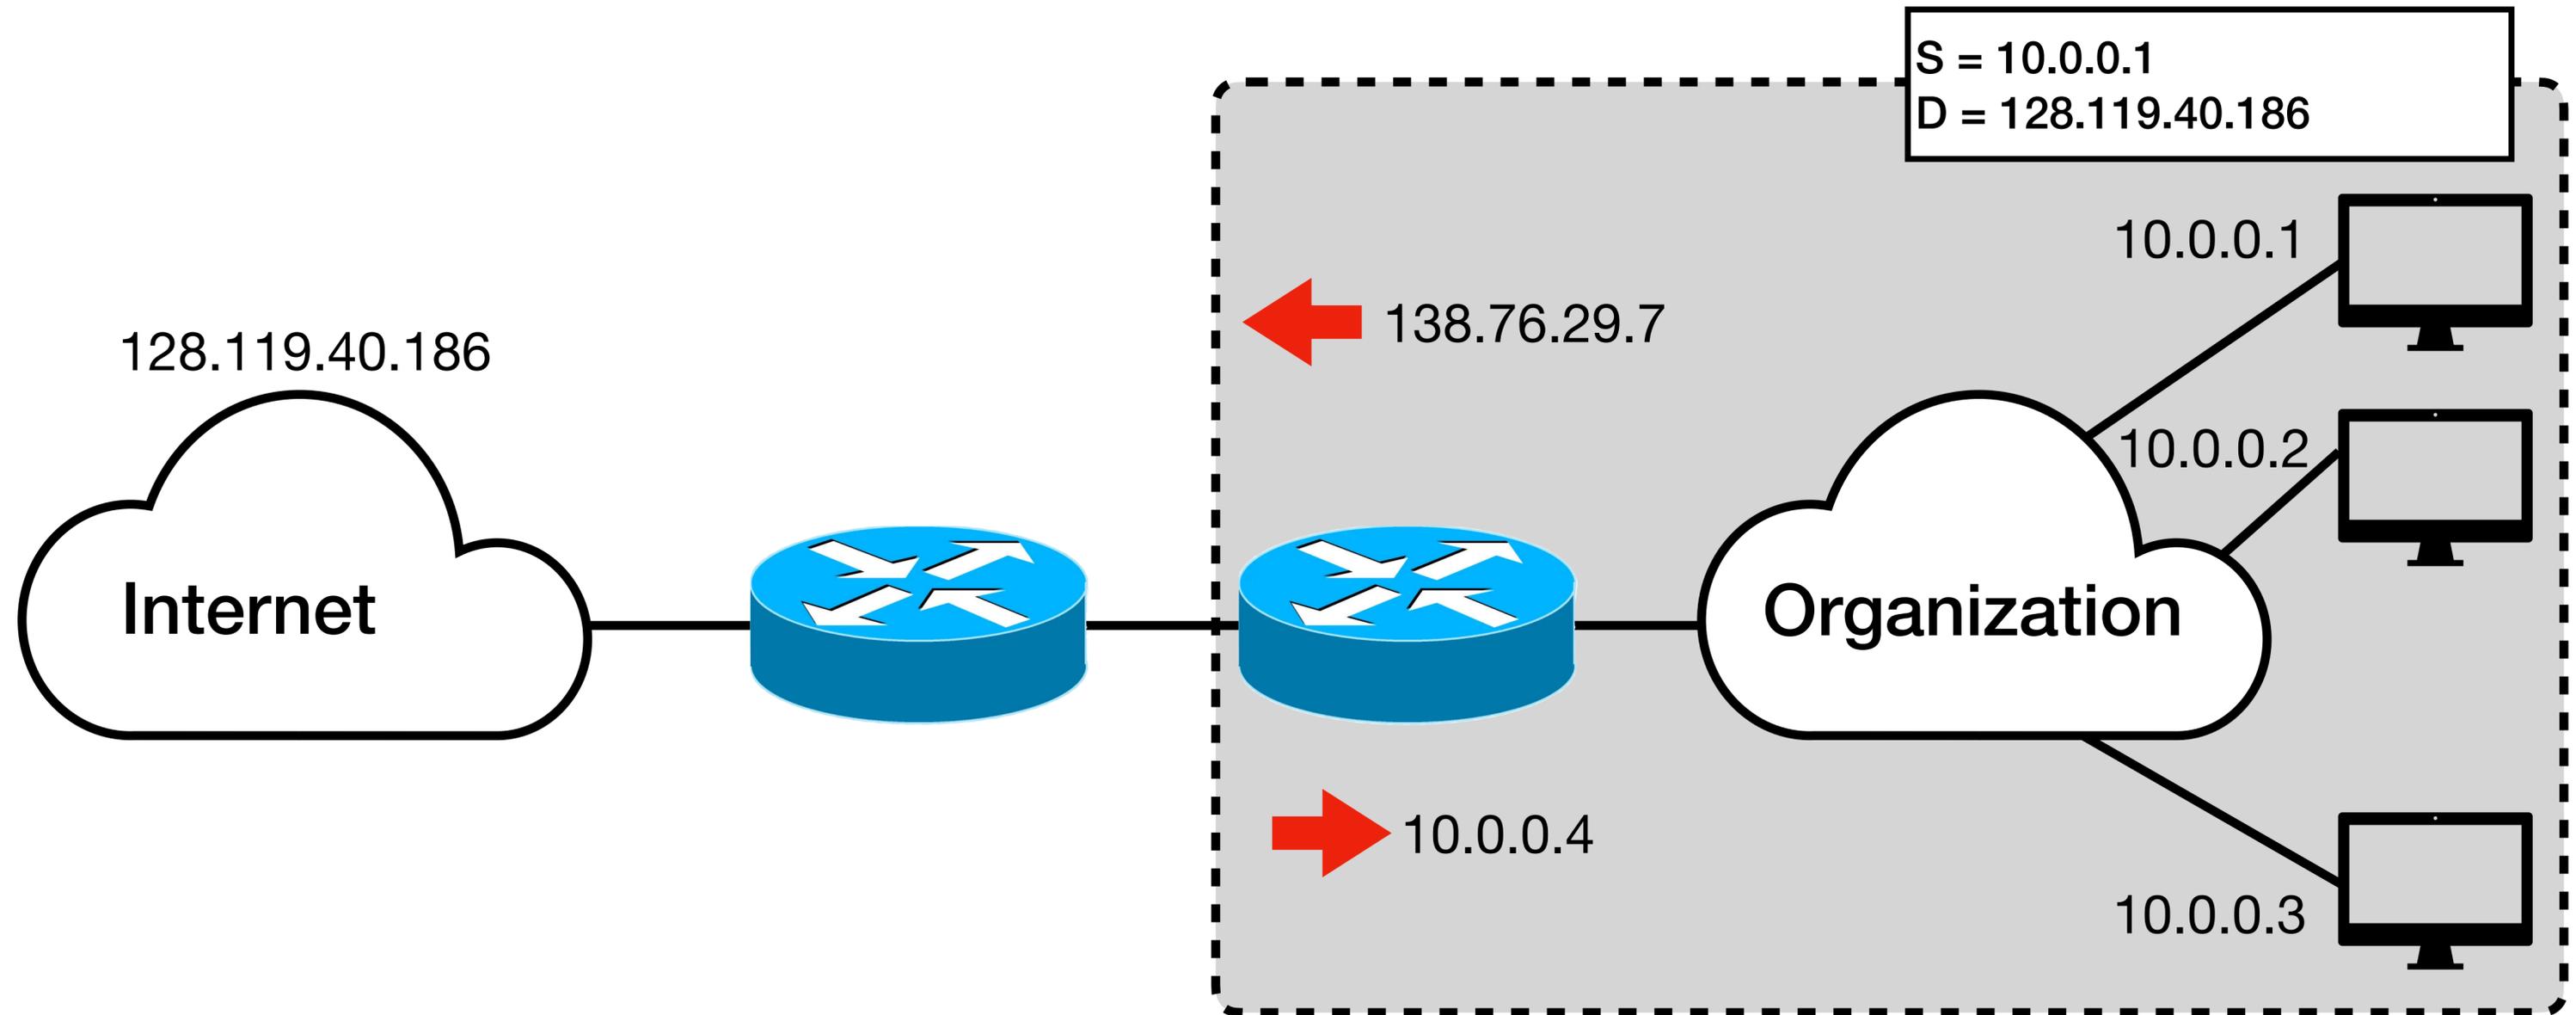


# How Does NAT Work?



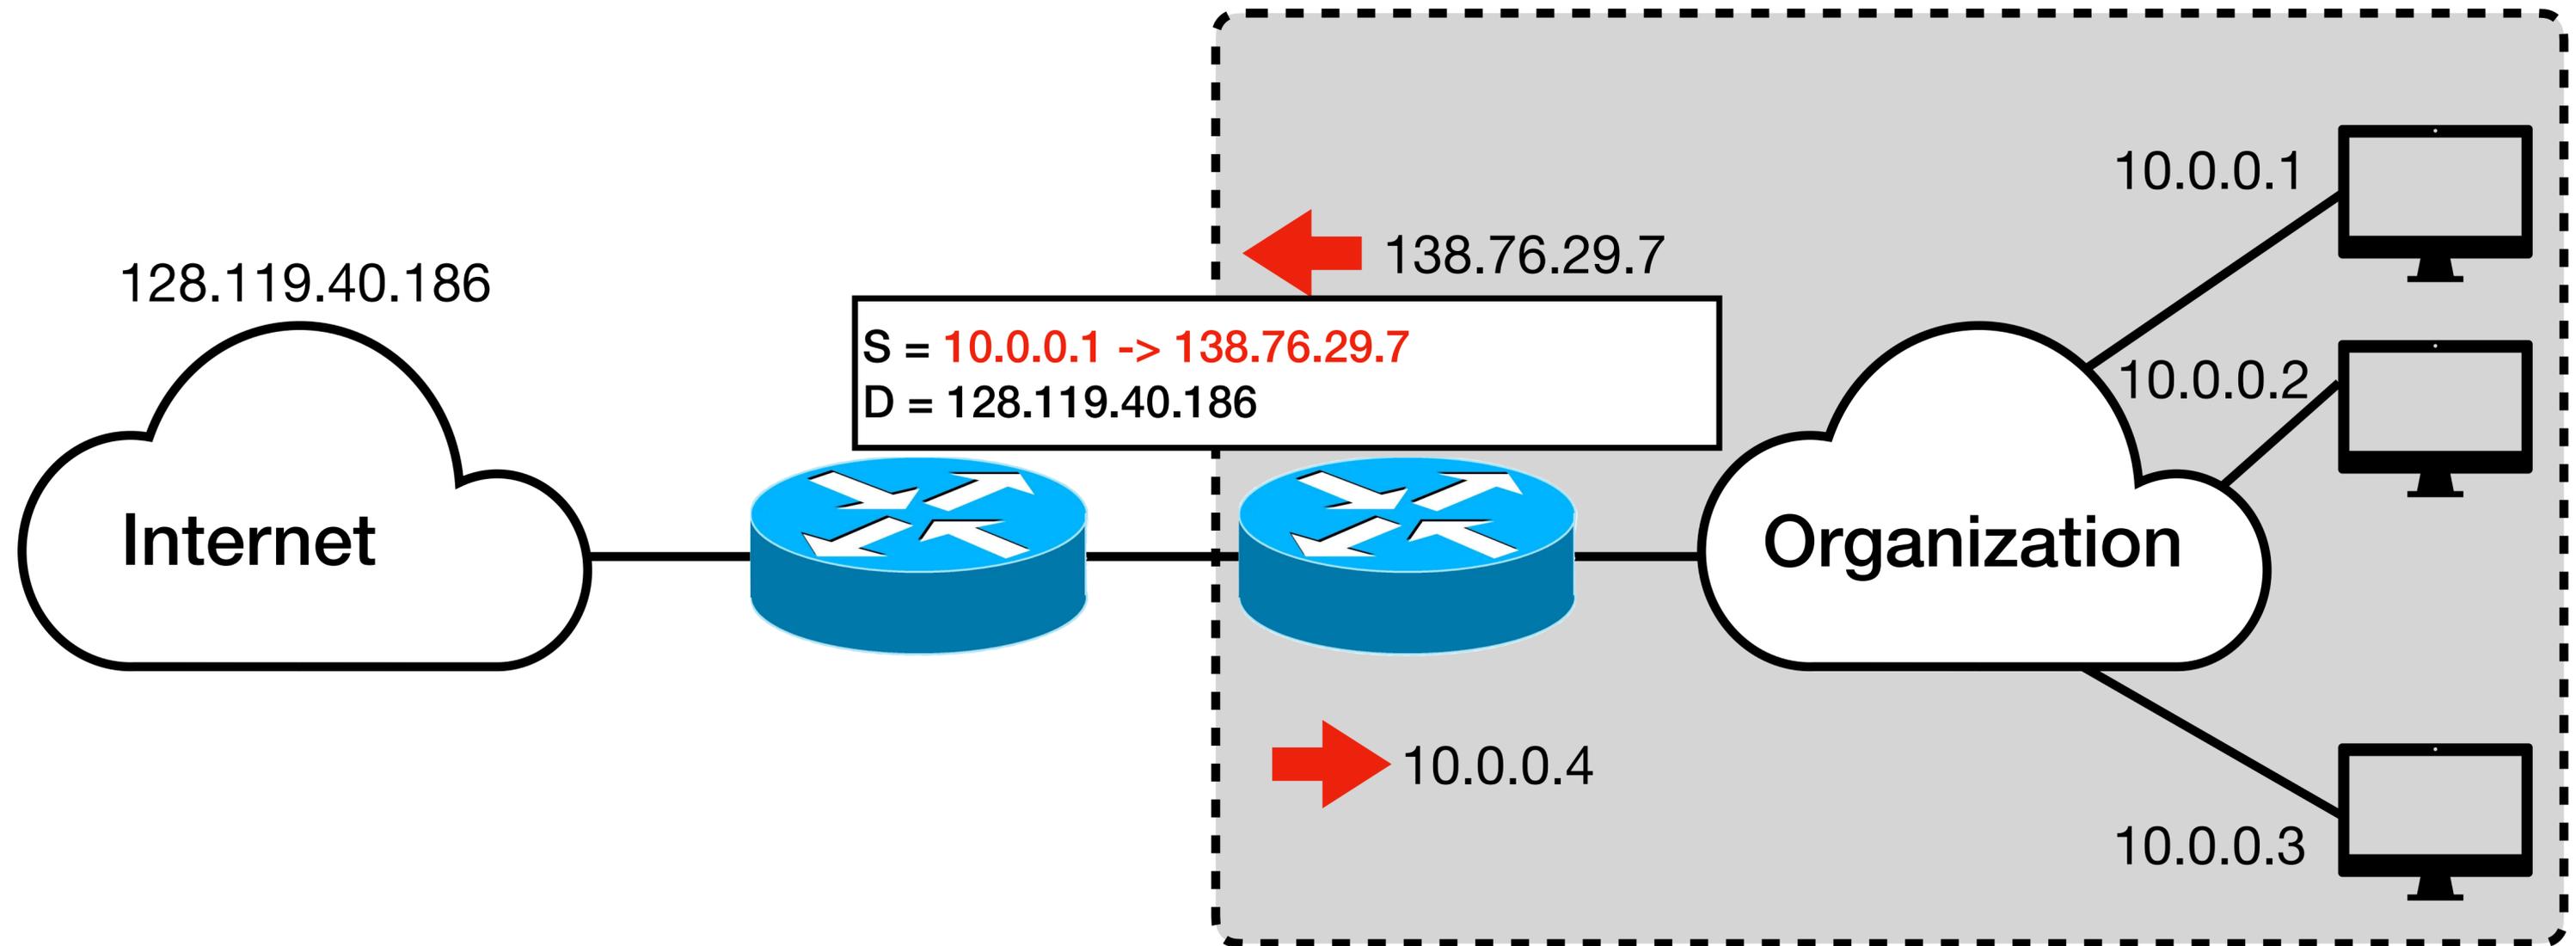
# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186



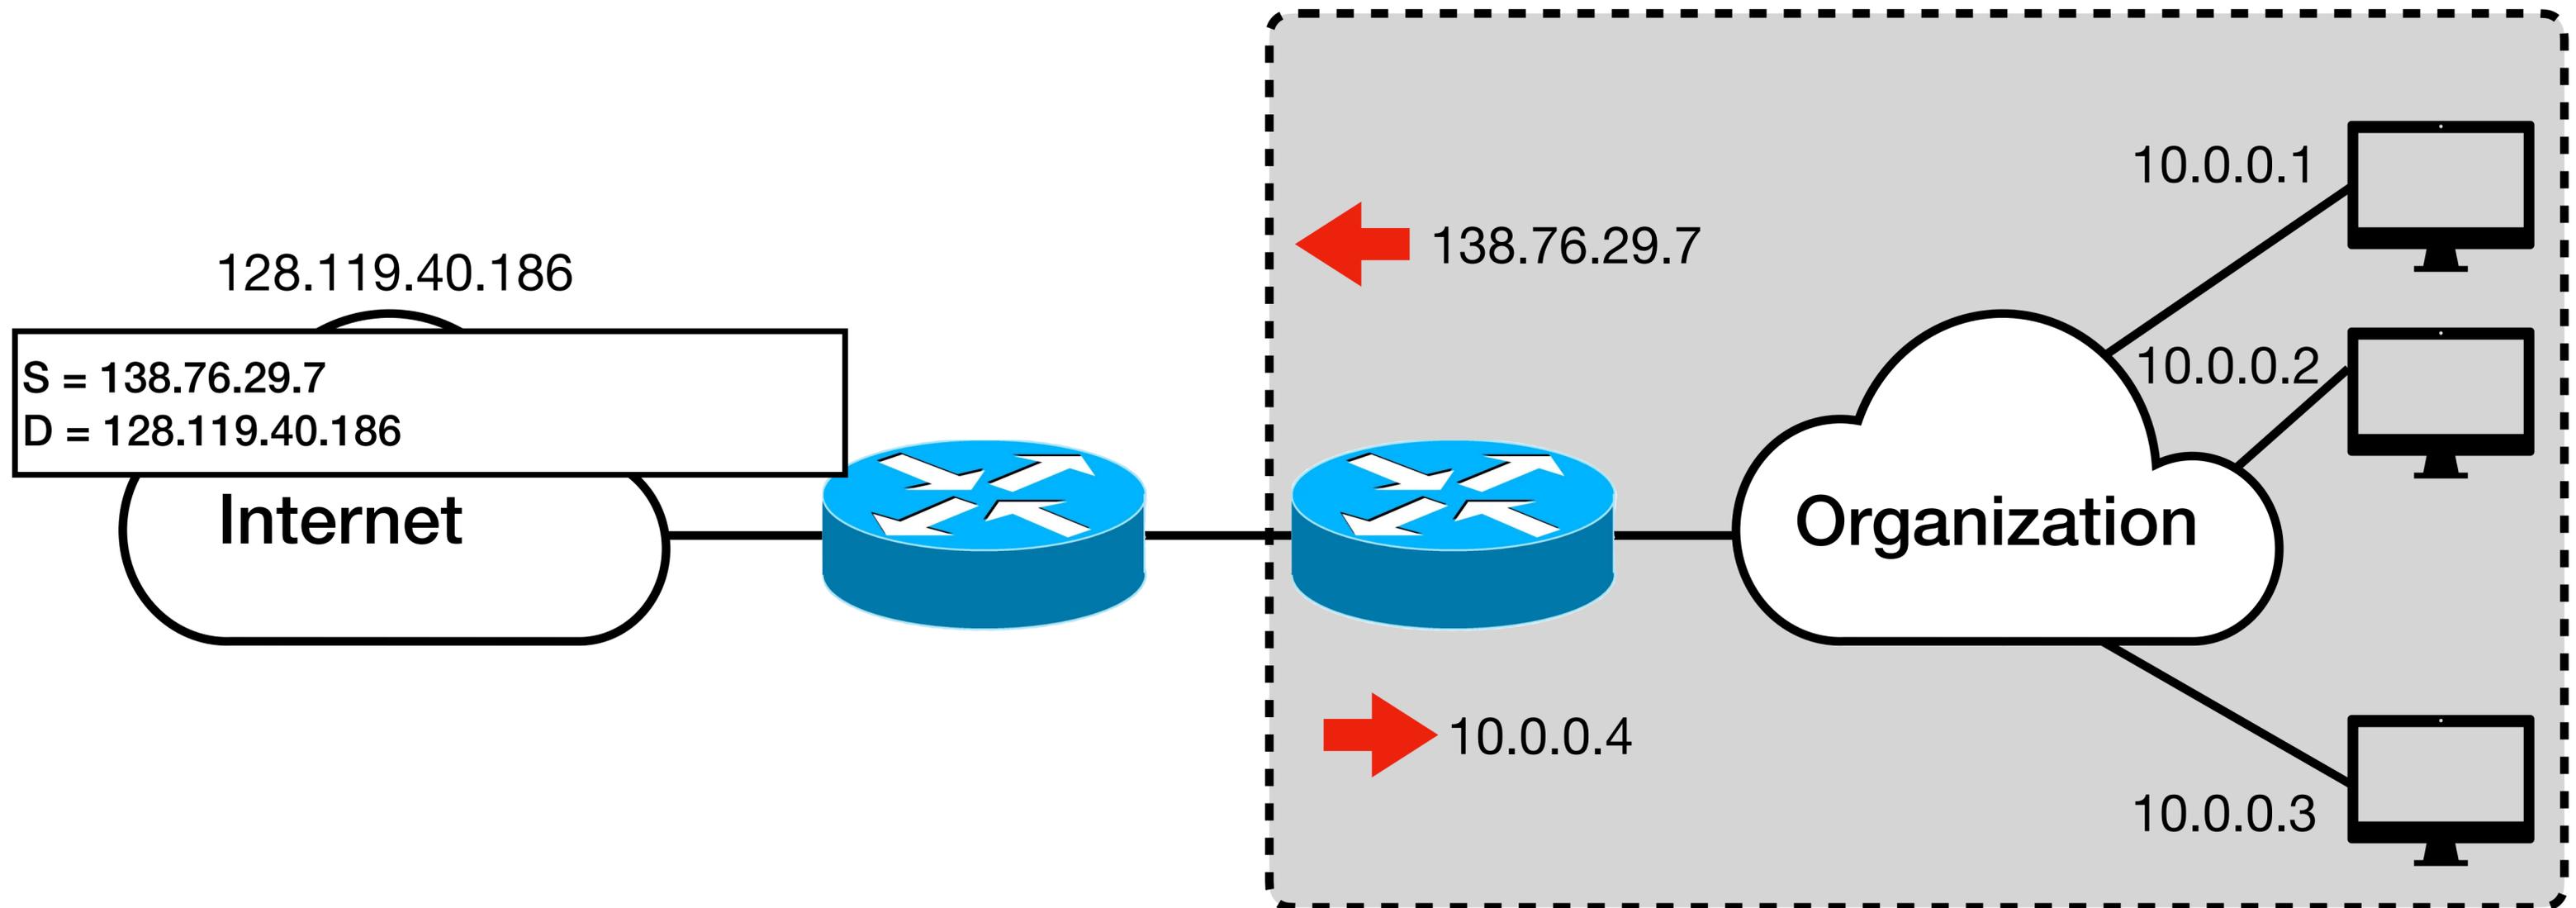
# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186



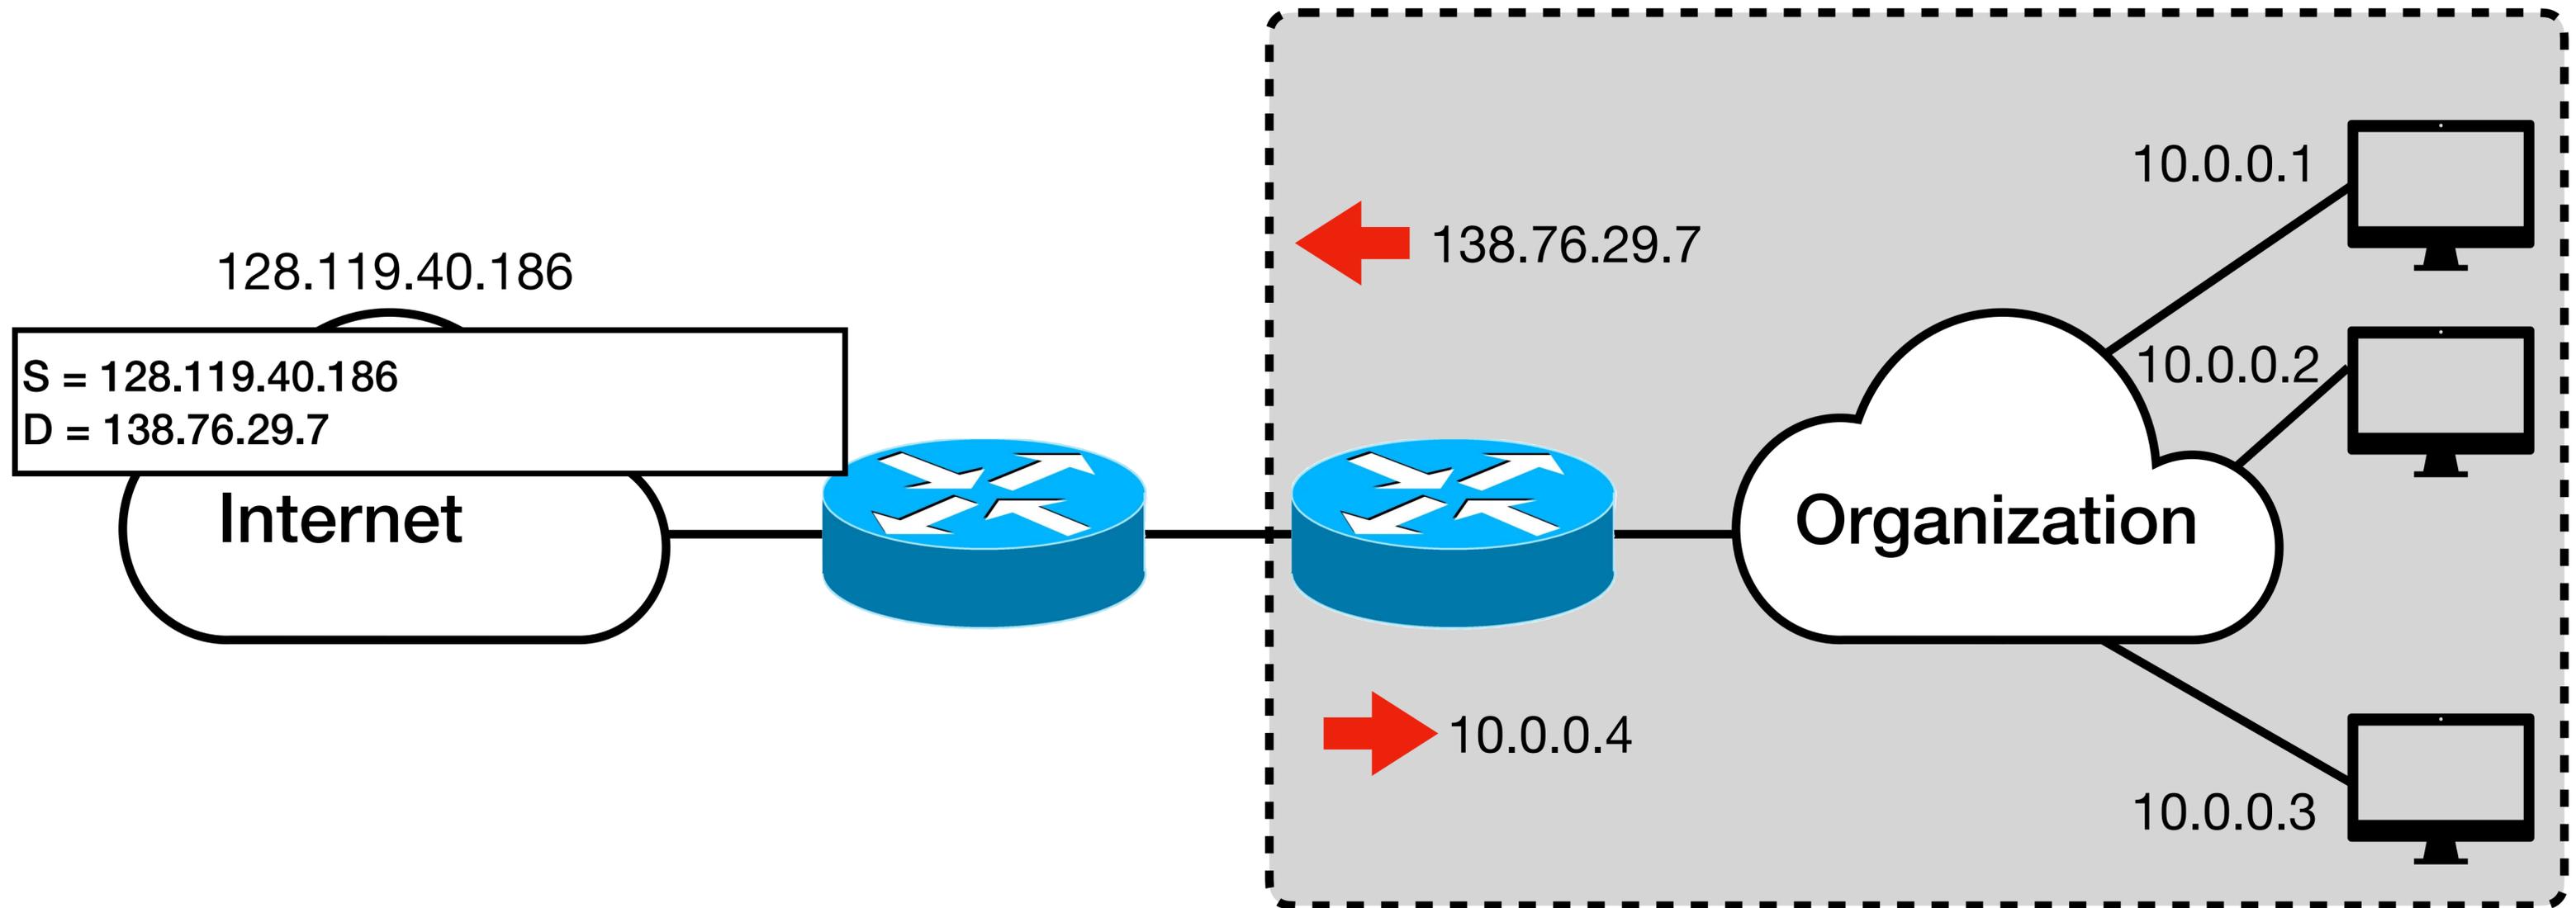
# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186



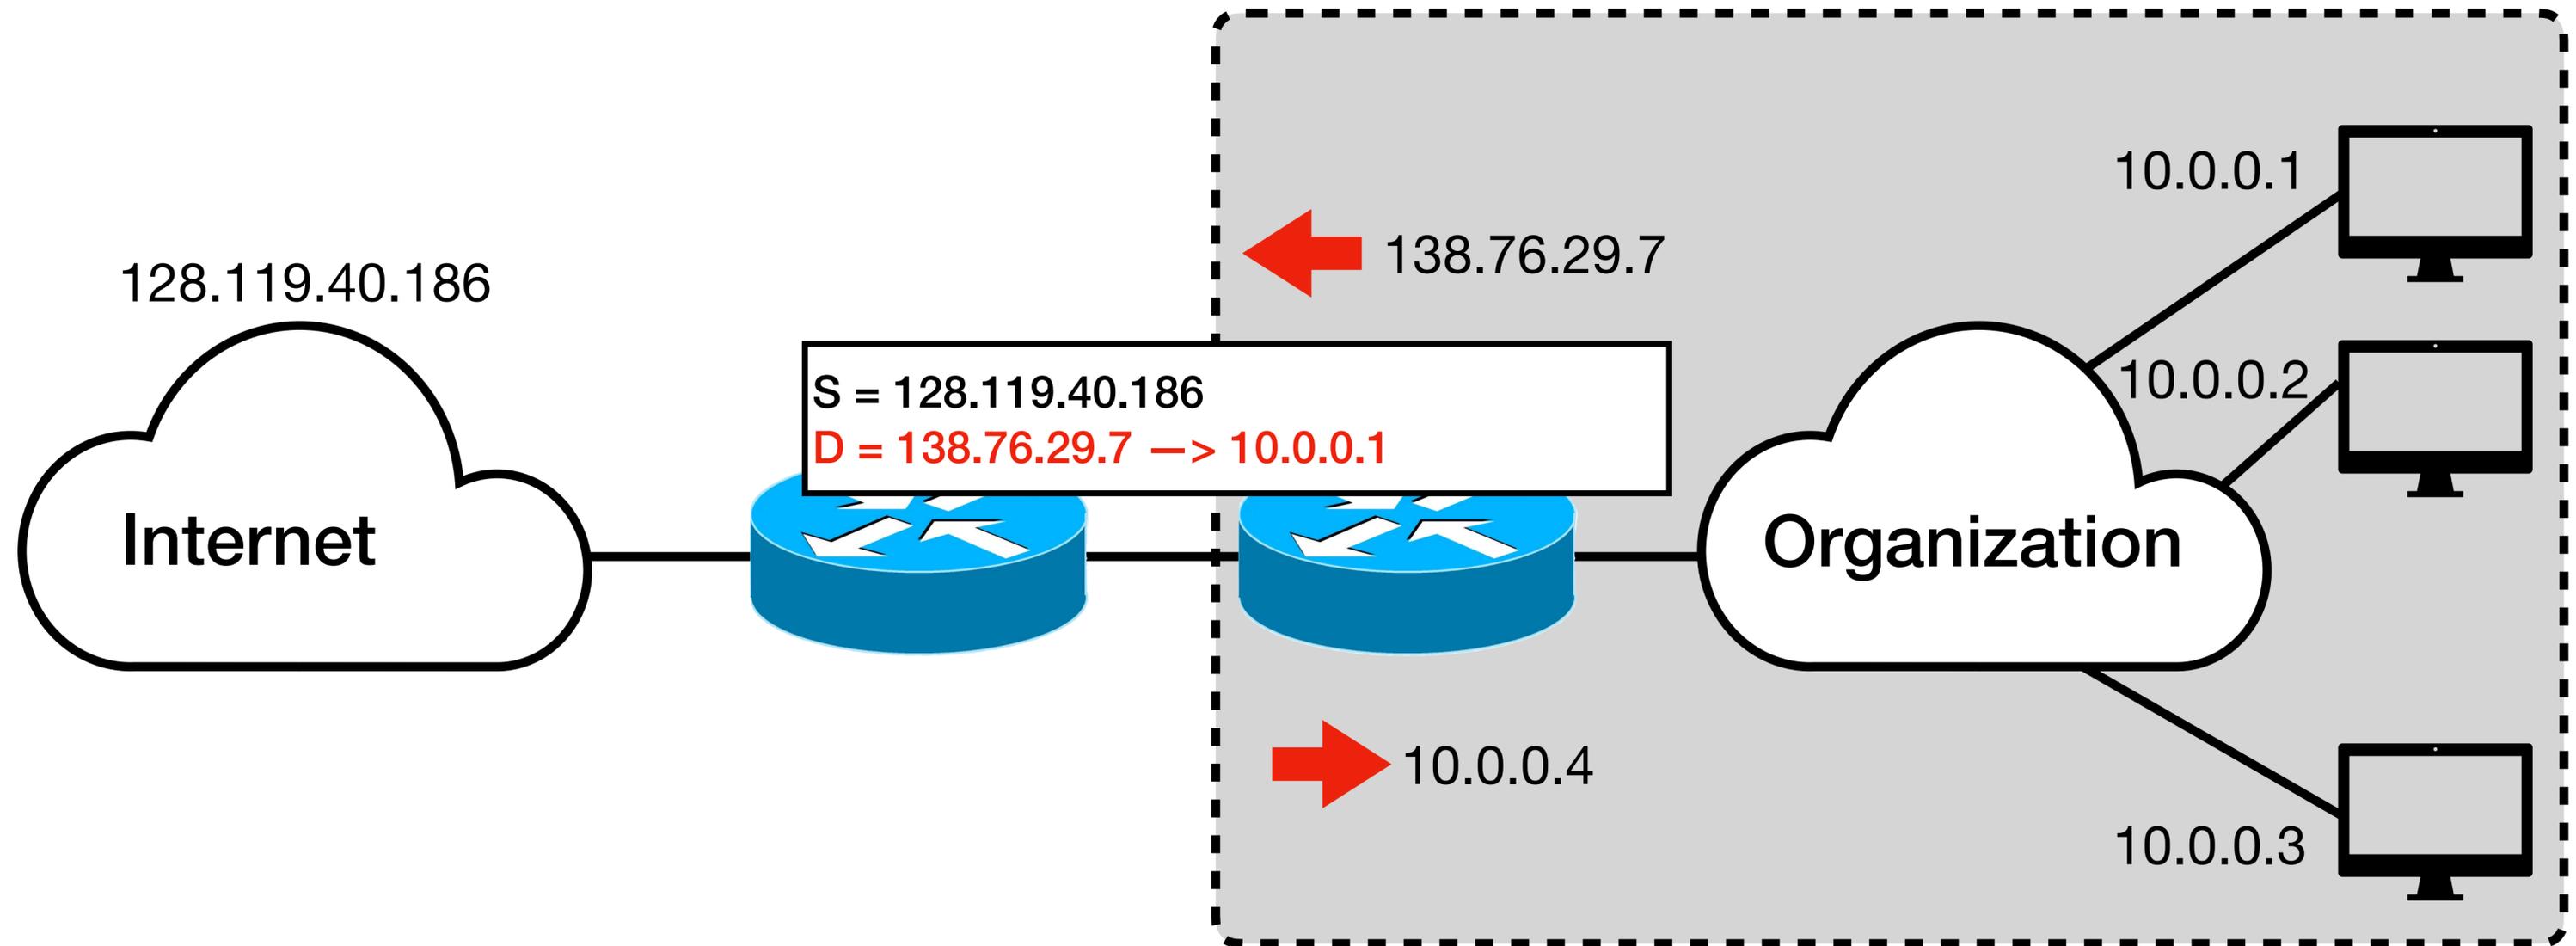
# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186
- Internet to Host: 128.119.40.186 → 10.0.0.1



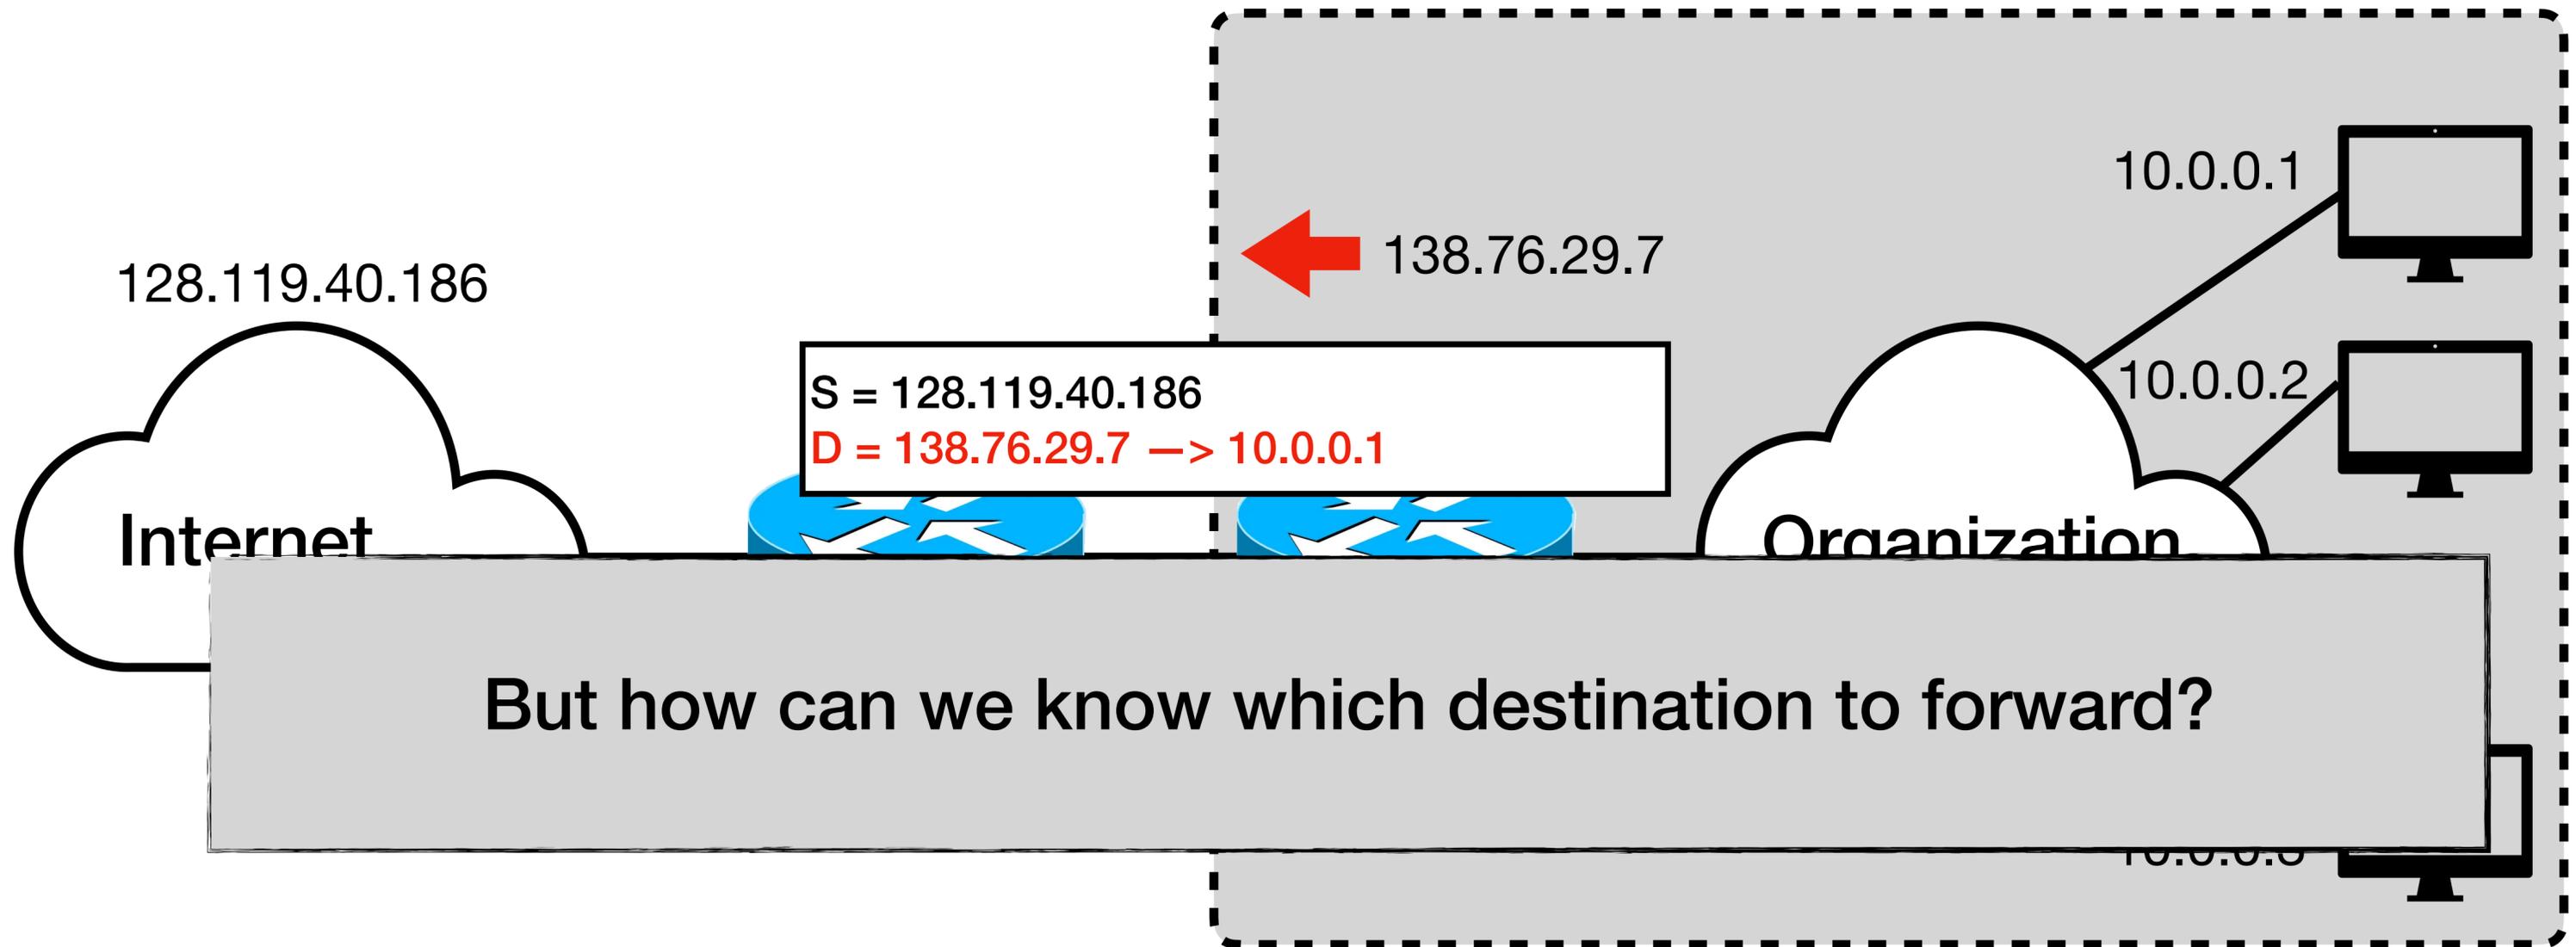
# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186
- Internet to Host: 128.119.40.186 → 10.0.0.1



# How Does NAT Work?

- Host to Internet: 10.0.0.1 → 128.119.40.186
- Internet to Host: 128.119.40.186 → 10.0.0.1



# NAT Table

NAT Translation Table	
Wide-Area Network (WAN)	Local-Area Network (LAN)
138.76.29.7	10.0.0.1
138.76.29.7	10.0.0.2

# NAT Table

- Use the port number (transport layer) to differentiate hosts
  - The NAT table is maintained by the NAT router
  - The NAT table is continuously updated

NAT Translation Table	
Wide-Area Network (WAN)	Local-Area Network (LAN)
138.76.29.7, 1234 (service port)	10.0.0.1, 4567 (service port)
138.76.29.7, 4321 (service port)	10.0.0.2, 7654 (service port)

# Private IPv4 Addresses

- Three ranges are reserved for private use
  - Class A: 10.0.0.0 to 10.255.255.255
  - Class B: 172.16.0.0 to 172.31.255.255
  - Class C: 192.168.0.0 to 192.168.255.255
- Private network: a network uses private addresses
  - Not allocated to any specific organization
  - IP packets whose source or destination address is a private IP address cannot be routed through the public Internet

# IPv6 Introduction

- IETF started to design a new version of IP in 1991
- Solicitation of suggestions from the community
  - The first version was completed in 1994

# IPv6 Planned Features

- #1: 128-bit address space
- #2: Support diverse Quality of Service (QoS) apps
- #3: Support security and authentication
- #4: Auto-configuration
  - Hosts are auto-configured with an IP address and a domain name
  - Try to make systems plug-and-play

## IPv6 Planned Features (cont'd)

- #5: Enhanced routing functionality (e.g., mobile hosts)
- #6: Support efficient multicast
- #7: Rely on simple protocol extensions
- #8: Enable a smooth transition path from IPv4

# IPv6 Address Space

- Allocation is classless
- Prefixes specify the unicast, multicast, and anycast cases
- Prefixes can be used to map between v4 and v6
- Lots of addresses with 128 bits!
  - ~1500 addresses per square foot of the earth's surface

# IPv6 Address Notation

- Set of eight 16-bit values separated by colons
  - E.g., 47CD:1234:3200:0000:0000:4325:B792:0428
- A large number of zeros omitted with a series of colons
  - E.g., 47CD:1234:3200::4325:B792:0428
- Address prefixes (slash notation) are the same as IPv4
  - E.g., FEDC:BA98:7600::/40 describes a 40-bit prefix

# IPv6 Address Prefix Assignments

0000 0000	Reserved
0000 0001	Unassigned
0000 001	Reserved for NSAP (non-IP addresses used by ISO)
0000 010	Reserved for IPX (non-IP addresses used by IPX)
0000 011	Unassigned
0000 1	Unassigned
1	Unassigned
001	<b>Unicast Address Space</b>
10	Unassigned
11	Unassigned
100	Unassigned
101	Unassigned
110	Unassigned
1110	Unassigned
1111 0	Unassigned
1111 10	Unassigned
1111 110	Unassigned
1111 1110 0	Unassigned
1111 1110 10	Link Local Use addresses
1111 1110 11	Site Local Use addresses
1111 1111	<b>Multicast addresses</b>

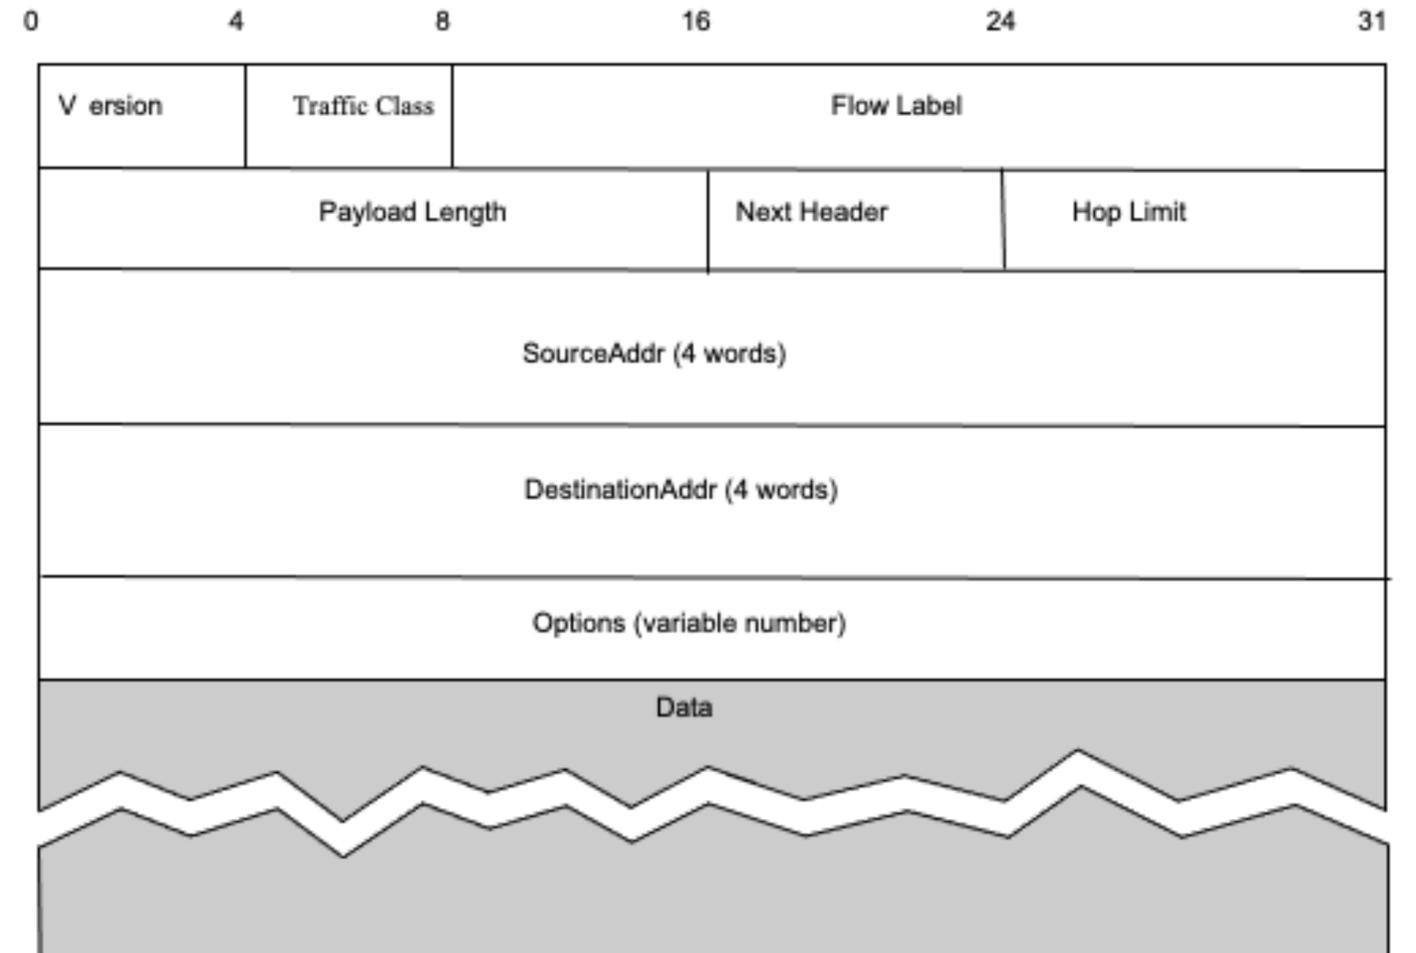
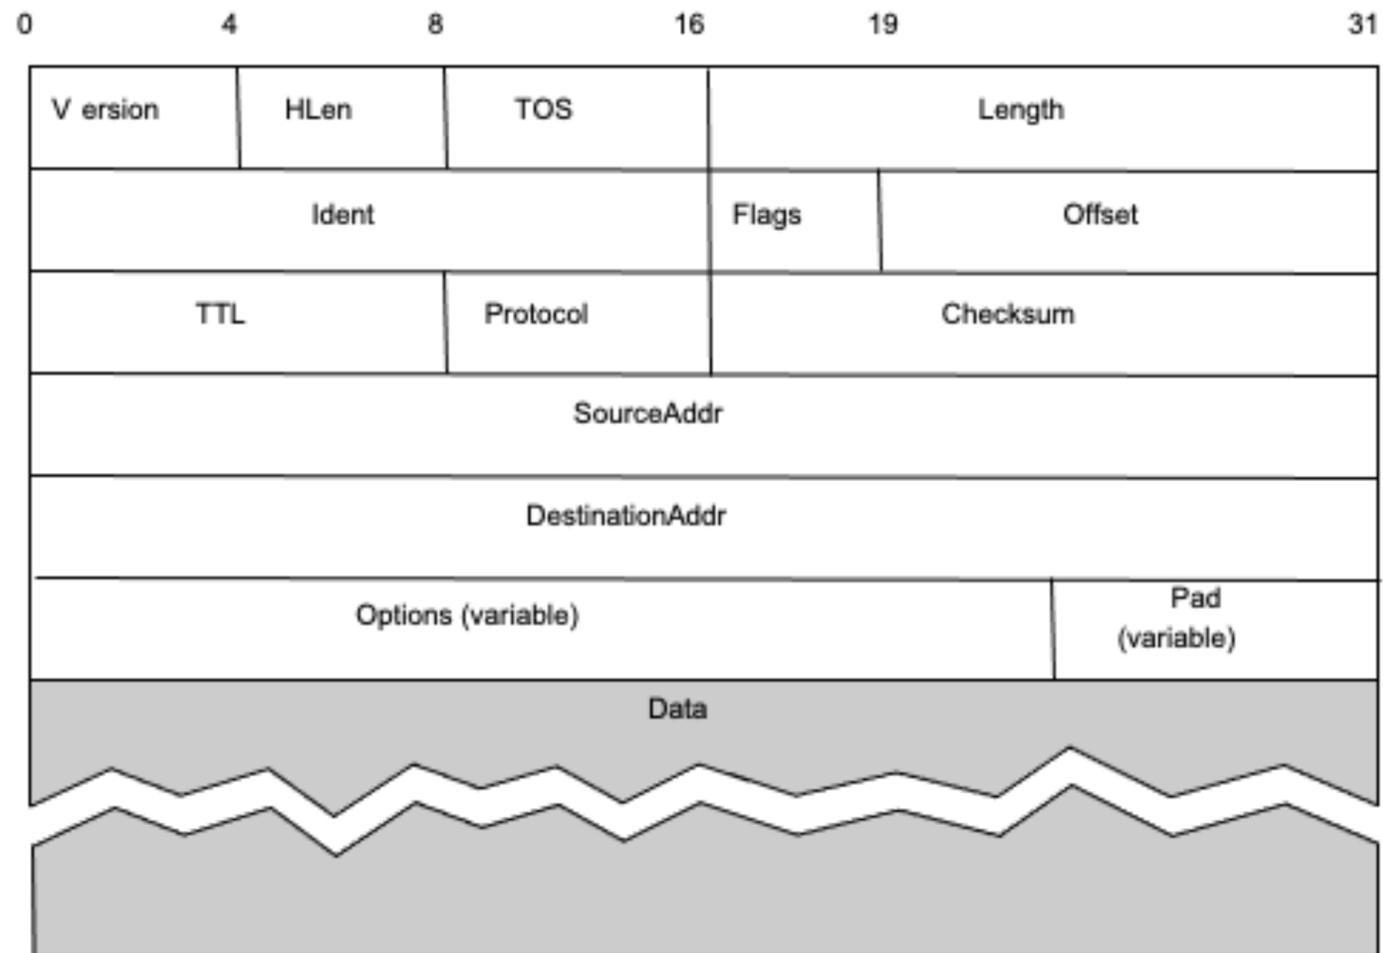
# IPv6 Unicast Assignment

- The unicast address assignment
  - Unicast addresses start with 001
  - Host interfaces belong to subnets
  - Addresses are composed of a subnet prefix and a host identifier
  - Prefix aggregation is also possible

# One Strategy

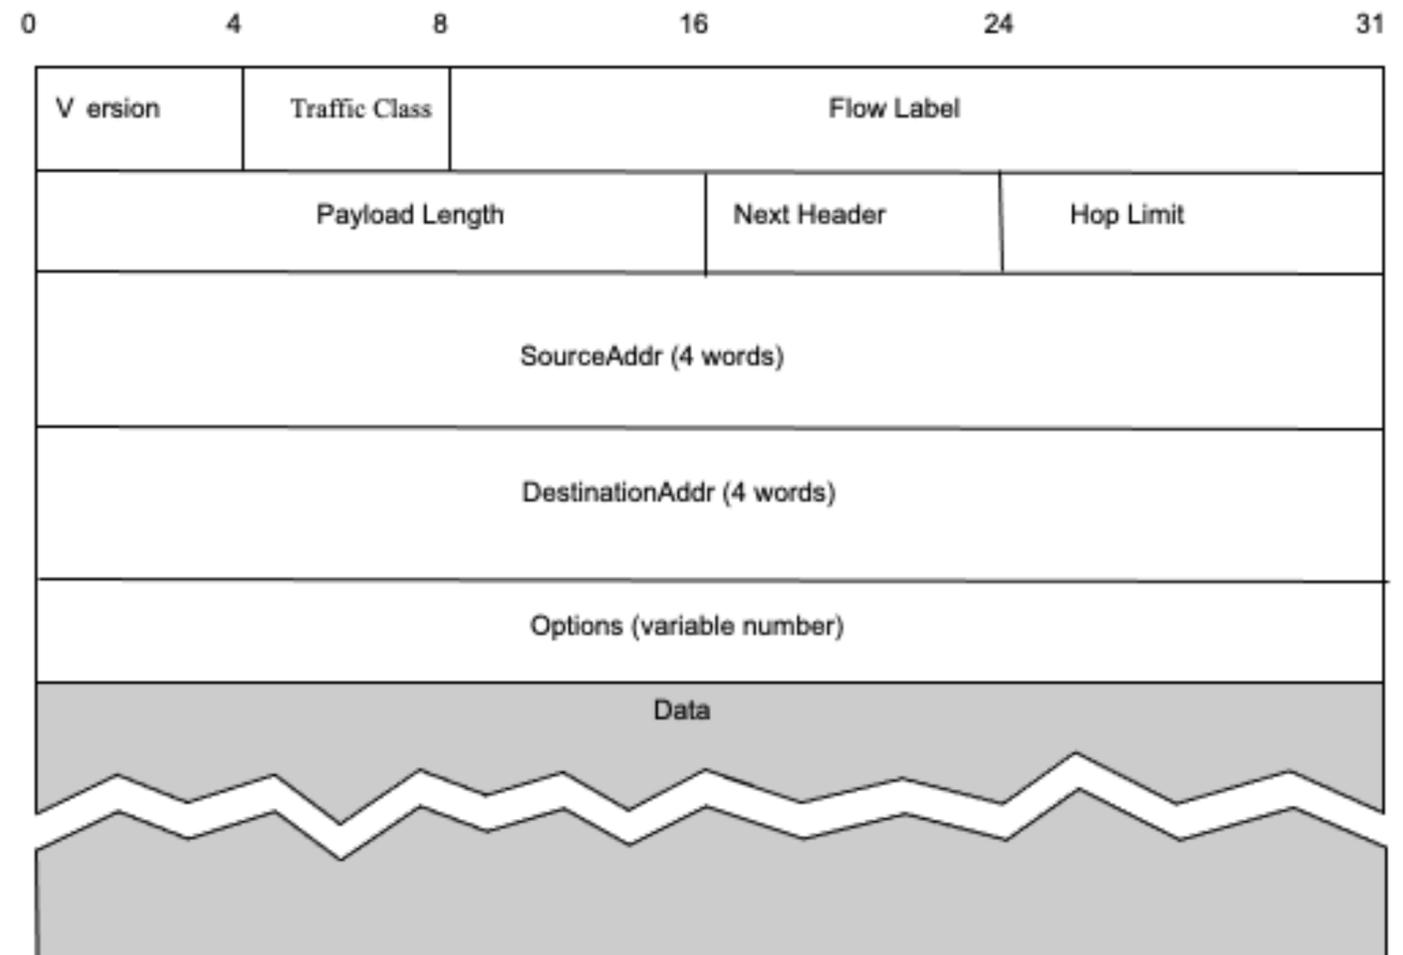
- Provider-based plan
  - The Internet is organized into a hierarchy of networks
  - 3 levels: region, provider, and subscriber
- Goal: aggregate routes to reduce the BGP overhead
  - The provider can advertise a single prefix for all its subscribers
  - Region = 13 bits, Provider = 24 bits, Subscriber = 16 bits, Subnet+Host
  - E.g., 001, region ID, provider ID, subscriber ID, subnet ID, host ID

# IPv4 v.s. IPv6 Packet Format



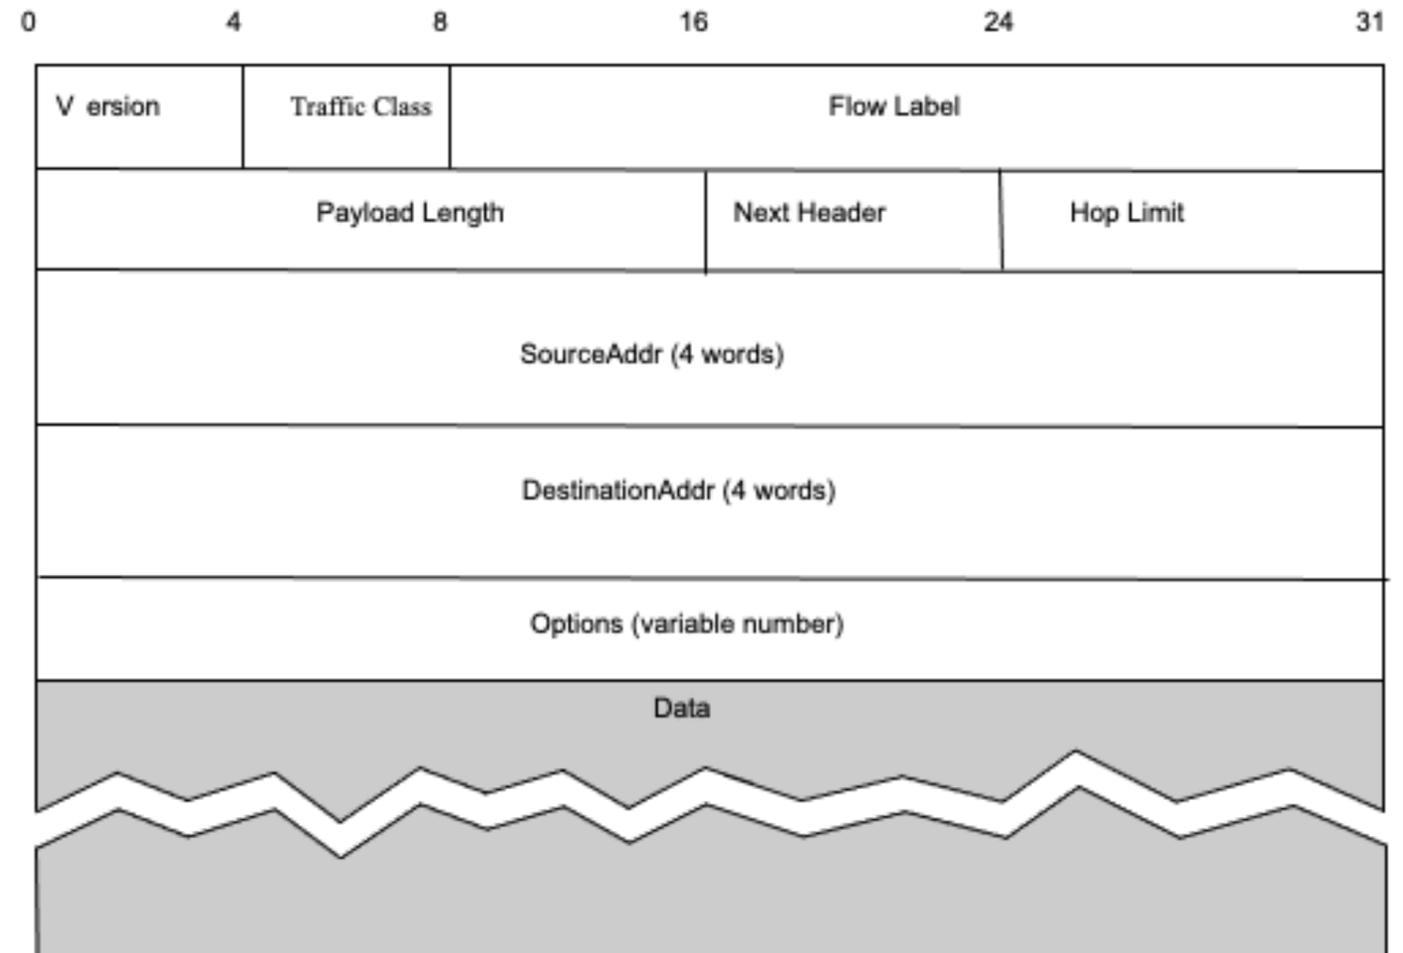
# IPv6 Packet Format Details

- #1: Simpler format than IPv4
- #2: Version = 6
- #3: Traffic class = IPv6 ToS
- #4: Treat all packets with the same flow label equally
  - Support QoS and fair bandwidth allocation



# IPv6 Packet Format Details (cont'd)

- #5: Payload length (16-bit)
- #6: Hop limit = IPv4 TTL field
- #7: Next header combines options and protocol
  - If no options, just the protocol field
- #8: Options are “extension header”
  - E.g., routing, fragment, authentication, encryption...



# Key Differences in Header

- No checksum
  - Bit-level errors are checked for all over the place
- No length variability in the header
  - Fixed format speeds processing
- No more fragmentation and reassembly in the header
  - Incorrectly sized packets are dropped
  - Hosts should do path MTU discovery

# Transition from IPv4 to IPv6

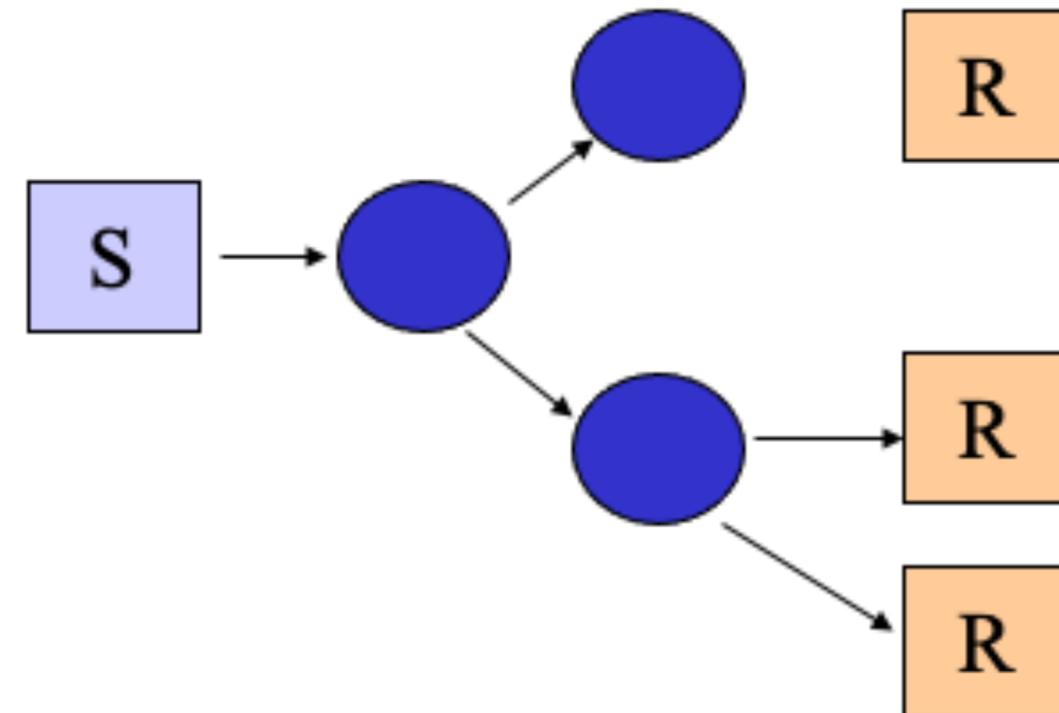
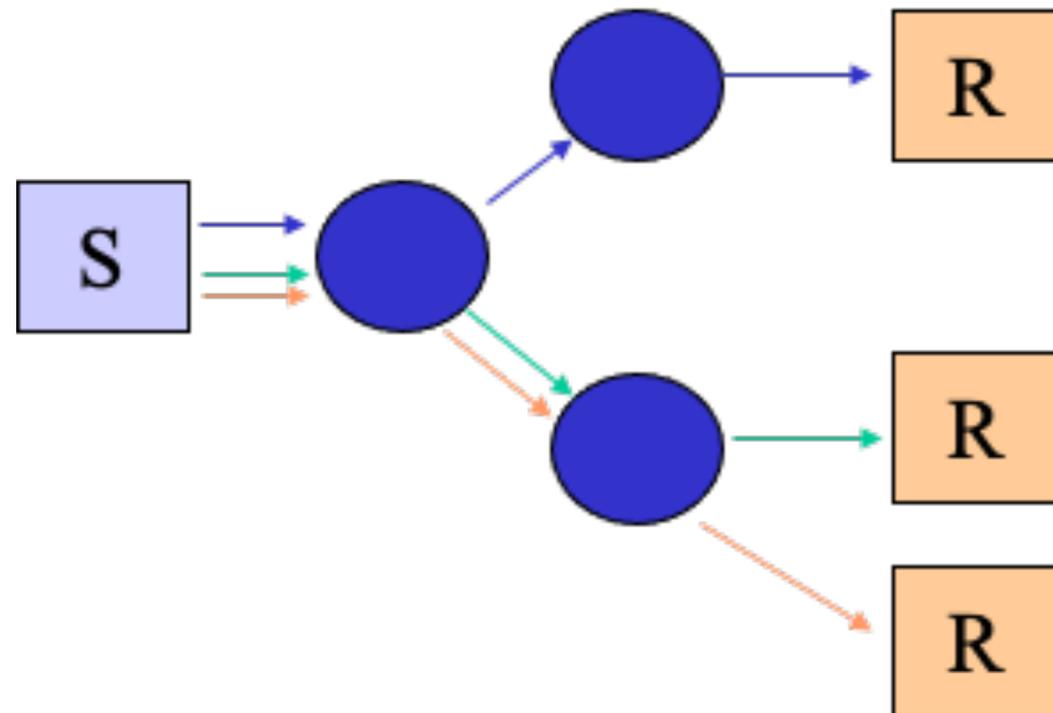
- Flag day is not feasible
- Dual stack operation
  - IPv6 nodes run in both IPv4 and IPv6 modes and use the version field to decide which stack to use
- Tunneling is used to deal with networks where IPv4 routers sit between two IPv6 routers
  - Encapsulate IPv6 packets in IPv4 packets until hitting the next IPv6 router

# The Scalability Challenge

- Q1: How can we support a rising number of hosts?
- Q2: How can we design the routing scheme to reduce traffic?
- *Q3: How can we build the underlying infrastructure to provide enough routing capacity?*

# One to Many Communication

- Application-level one-to-many communication
  - Multiple unicasts



# Why Multicast

- When sending the same data to multiple receivers
  - Better bandwidth utilization
  - Less host/router processing
  - Fast participation

# Why Multicast

- When sending the same data to multiple receivers
  - Better bandwidth utilization
  - Less host/router processing
  - Fast participation
- Benefit applications
  - Video/Audio broadcast (one sender)
  - Video conferencing (many senders)
  - Real-time news distribution
  - Interactive gaming

# IP Multicast Service Model

- Invented by Steve Deering (Ph.D. 1991)
  - It's a different way of routing datagrams
  
- RFC 1112: Host Extensions for IP Multicasting (1989)

# IP Multicast Workflow

- **#1: Configuration**
  - Members join and leave the group and indicate this to the routers
  - The “host group” is identified by a class D IP address
- **#2: Execution**
  - Senders transmit IP datagrams to a “host group”.
  - Routers listen to all multicast addresses managed by the multicast routing protocol
  - Members of the host group can be present anywhere on the Internet

# IP Multicast Group Address

- Class D address space
  - High-order three-bits are set
  - 224.0.0.0 ~ 239.255.255.255
- Allocation is essentially random — any class D can be used
  - Senders transmit IP datagrams to a “host group”.
  - Routers listen to all multicast addresses managed by the multicast routing protocol
  - Members of the host group can be present anywhere on the Internet

# Multicast Packets —> Endhosts

- Packets from remote sources will only be forwarded by IP routers onto a local network only if they know there is at least one recipient for that group on that network

# Multicast Packets —> Endhosts

- Packets from remote sources will only be forwarded by IP routers onto a local network only if they know there is at least one recipient for that group on that network
- Internet Group Management Protocol (IGMP, RFC 2236)
  - Used by hosts to signal that they want to join a specific multicast group
  - Used by routers to discover the mapping between hosts and multicast groups
  - Implemented directly over IP

# IGMP: Joining a Group

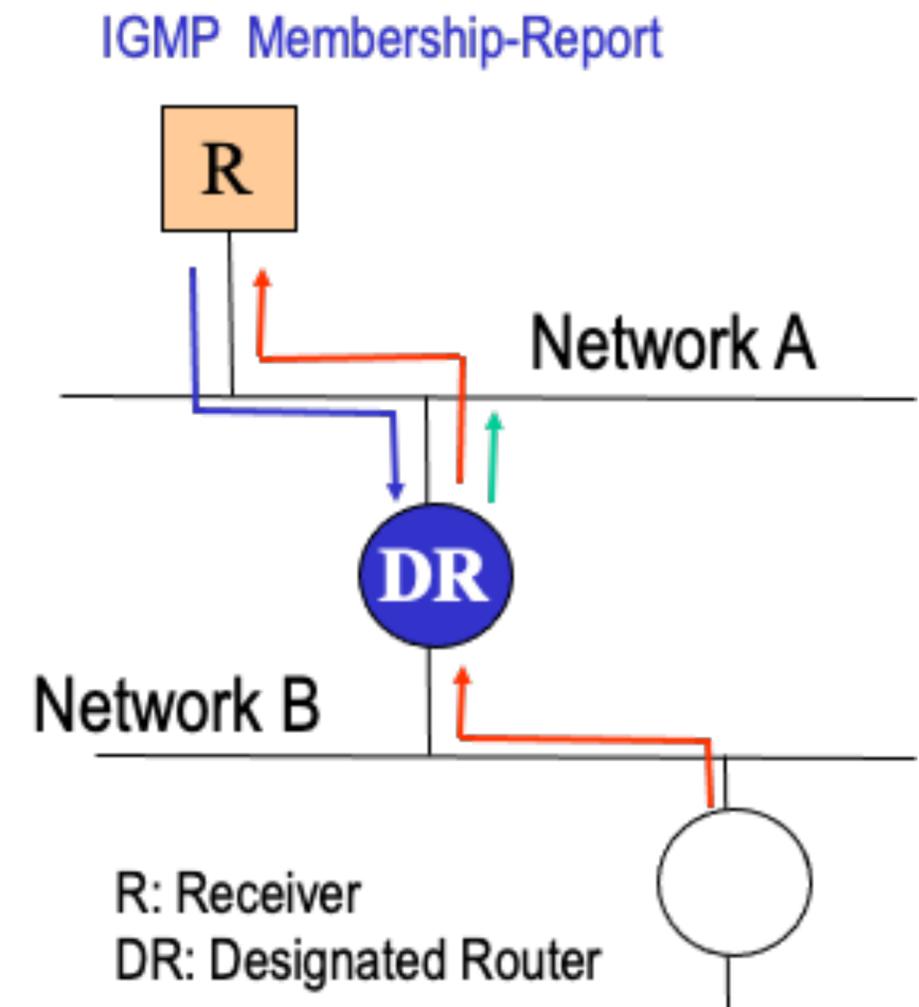
- Example: R joins to the group 224.2.0.1

#1: R sends an **IGMP Membership-Report** to **224.2.0.1**

#2: DR receives it. DR will start forwarding **packets** for **224.2.0.1** to Network A

#3: DR periodically sends **IGMP Membership-Query** to **224.0.0.1** (ALL-SYSTEM.MCAST.NET)

#4: R answers **IGMP Membership-Report** to **224.2.0.1**



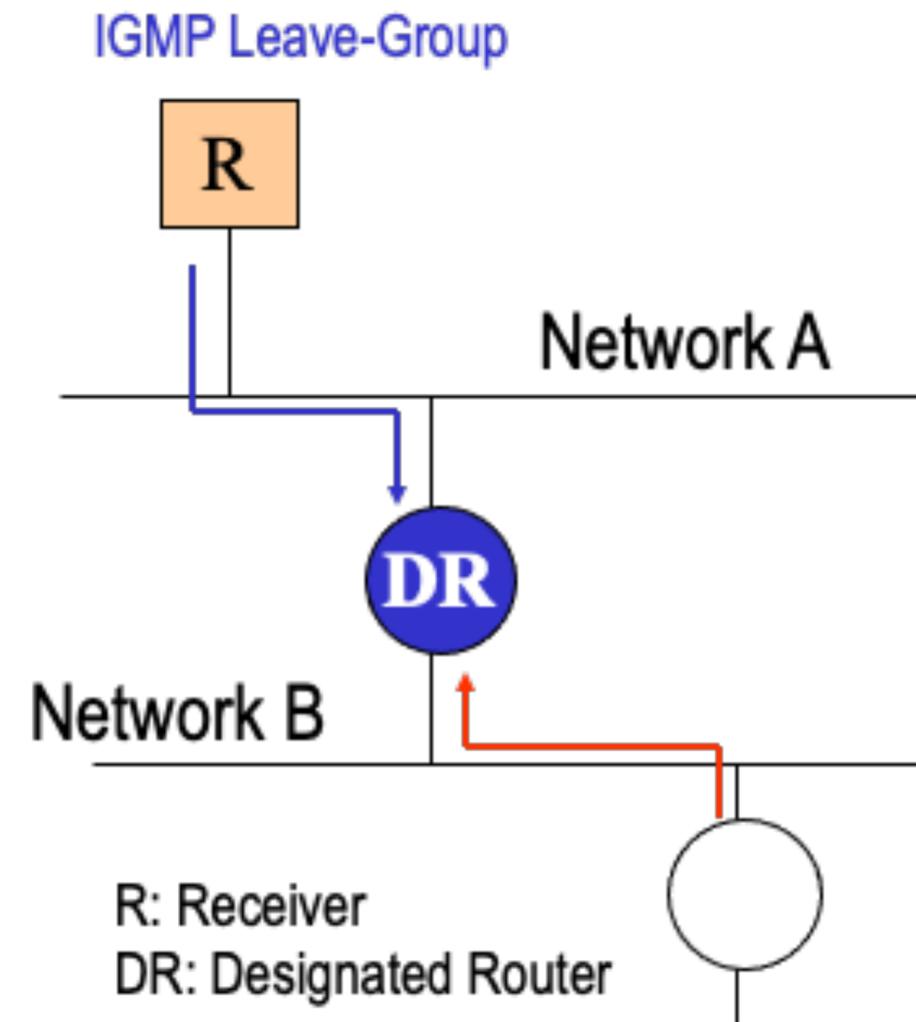
# IGMP: Leaving a Group

- Example: R leaves from the group 224.2.0.1

#1: R sends **IGMP Leave-Group** to 224.0.0.2  
(ALL-ROUTES.MCAST.NET)

#2: DR receives it

#3: DR stops forwarding **packets for 224.2.0.1** if  
no more 224.2.0.1 group members on Network A



# Challenges in the Multicast Model

- How can a sender restrict who can receive a packet?
  - Need authentication
  - Encryption of data
  - Key distribution
  - Still an active research area

# Multicast Requires Router Support

- Share the group information among routers to implement better routing for data distribution
- Distribution tree structure
- Routing protocols are used in conjunction with IGMP

# IP (Network) Layer Summary

Physical layer

A reliable (and efficient) bit delivery channel over a link

# IP (Network) Layer Summary

Link layer

A frame delivery channel between directly connected or switched hosts

Physical layer

A reliable (and efficient) bit delivery channel over a link

# IP (Network) Layer Summary

**Q1: How can we address any hosts in a network?**  
**Q2: How can we deliver data efficiently for an arbitrary communication path?**

**IP (Network)  
layer**

**A datagram delivery channel between hosts in any network  
under the best-effort service model**

**Link layer**

**A frame delivery channel between directly connected or  
switched hosts**

**Physical layer**

**A reliable (and efficient) bit delivery channel over a link**

# IP (Network) Layer Summary

**Q1: How can we address any hosts in a network?**

**=> IPv4, CIDR, IPv6, and NAT**

**Q2: How can we deliver data efficiently for an arbitrary communication path?**

**=> Intra-domain/inter-domain routing, SDN, and Multicast**

**IP (Network)  
layer**

**A datagram delivery channel between hosts in any network under the best-effort service model**

**Link layer**

**A frame delivery channel between directly connected or switched hosts**

**Physical layer**

**A reliable (and efficient) bit delivery channel over a link**

# Summary

- Today
  - NAT, IPv6, and Multicast
  
- Next lecture
  - Transport Introduction