Network Layer

Network Layer - Moving down the Stack

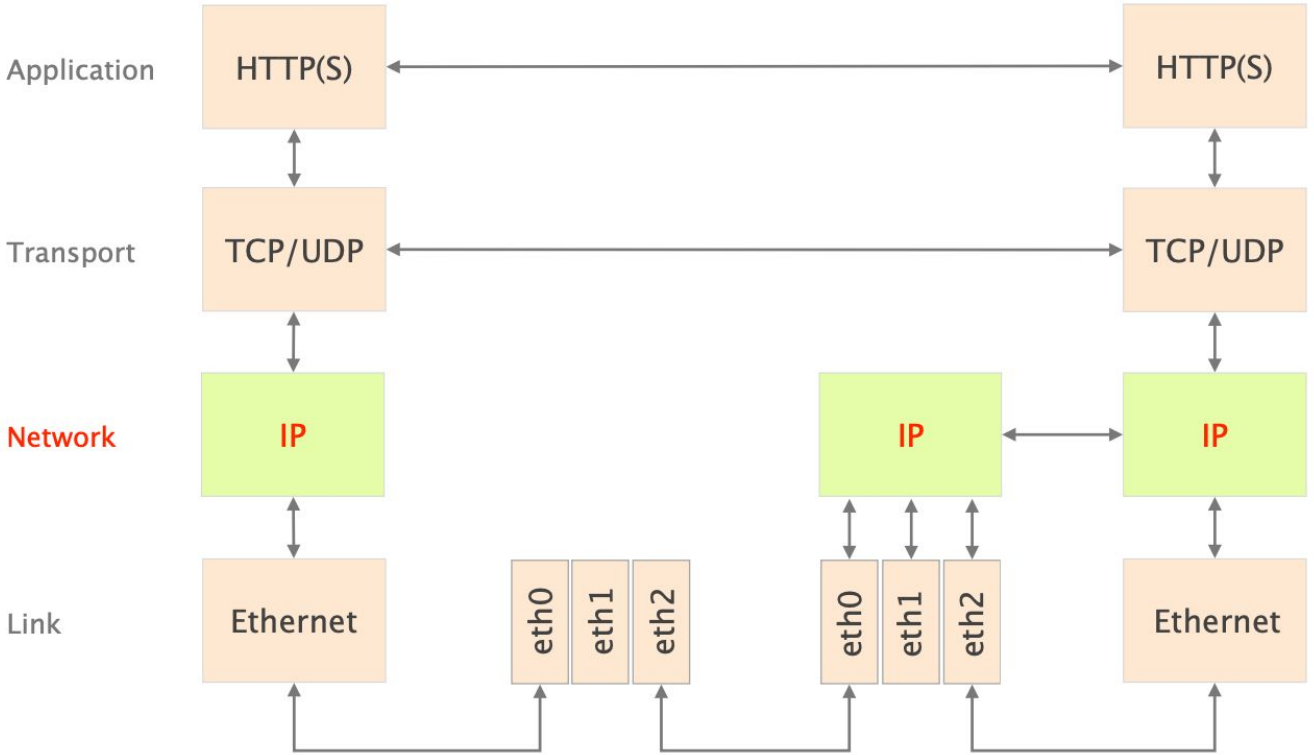

Network (IP) Layer

1. IP addresses
    - use, structure, allocation
2. IP forwarding
    - longest prefix match rule
3. IP header
    - IPv4 and IPv6, wire format

Network (IP) Layer

1. IP addresses
    - use, structure, allocation
2. IP forwarding
    - longest prefix match rule
3. IP header
    - IPv4 and IPv6, wire format

IP(v4) Addresses are Unique 32-bit Numbers Associated with a Host/Router

IP addresses are usually written

using dotted-quad notation

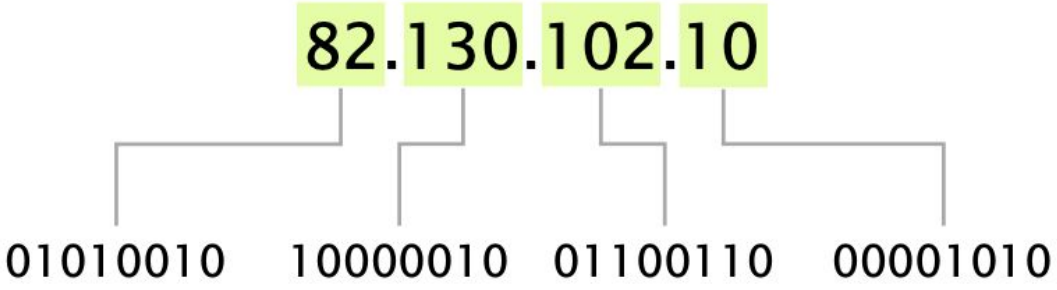

IP(v6) Addresses are Unique 128-bit Numbers Associated with a Host/Router

Notation	8 groups of 16 bits each separated by colons (:) Each group is written as four hexadecimal digits
Simplification	Leading zeros in any group are removed **One** section of zeros is replaced by a double colon (::) Normally the longest section
Examples	1080:0:0:0:8:800:200C:417A → 1080::8:800:200C:417A FF01:0:0:0:0:0:0:0101 → FF01::101 0:0:0:0:0:0:0:1 → ::1

Routers Forward Based on IP Destination

If IPs Were Individually Handed Out, Routers Would Require Massive Forwarding Tables

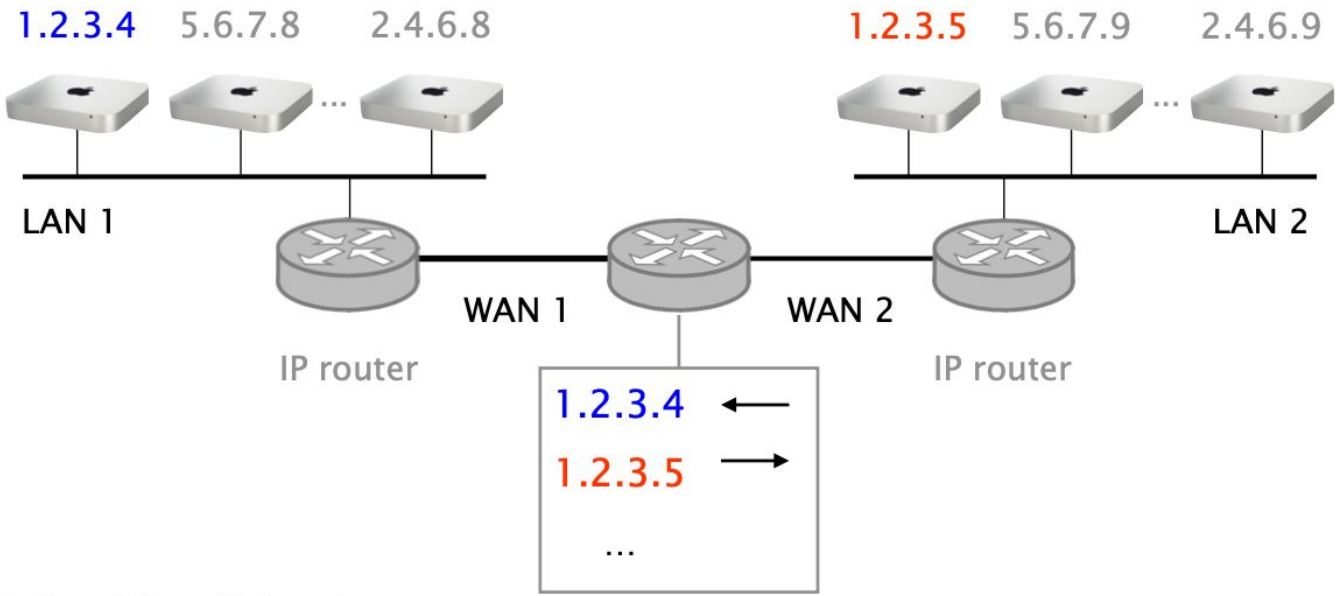

LAN Local Area Network

WAN Wide Area Network

Universal Answers to Such Problems

When you need… more flexibility,

you add… a layer of indirection

When you need… more scalability,

you add… a hierarchical structure

IPs are Hierarchically Allocated

Example: Mail service

- Deliver the letter to the post office responsible for the zip code
- Assign letter to the mail person covering the street
- Drop letter into the mailbox attached to the building
- Hand in the letter to the appropriate person

IPs are Hierarchically, Composed of Prefix (network address) and Suffix (host address)

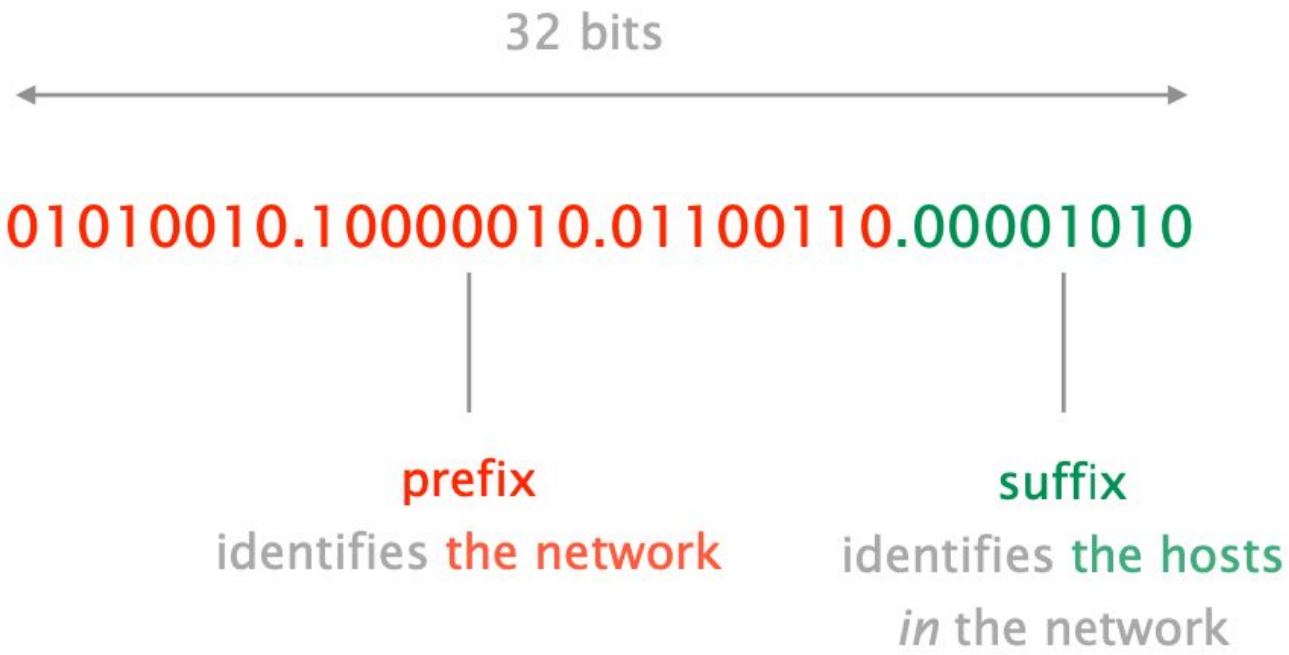

Prefixes Have Varying Lengths, Usually Written Using “slash notation”

/24 Means We Have 8 bits for Host Addresses, enough for 256 Hosts

82.130.102.0 /24

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0
01010010.10000010.01100110.	00000001	82.130.102.1
01010010.10000010.01100110.	00000010	82.130.102.2
01010010.10000010.01100110.	11111110	82.130.102.254
01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0
01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0

All 0s Identifies the Network Itself

01010010.10000010.01100110.	11111111	82.130.102.255

In Practice the First and Last IP Addresses of a Prefix are not Usable

prefix part	host part	IP address
01010010.10000010.01100110.	00000000	82.130.102.0

All 1s Identifies the Broadcast Address

01010010.10000010.01100110.	11111111	82.130.102.255

Prefixes Can Also Be Specified Using an Address and a Mask

Address	82.130.102.0
	01010010.10000010.01100110. 00000000
	11111111.11111111.11111111. 00000000
Mask	255.255.255.0

ANDing the Address and the Mask Gives you a Prefix

Address 82.130.102.0

01010010.10000010.01100110. 00000000

11111111.11111111.11111111. 00000000

Mask 255.255.255.0

Example

Given this IP prefix: 82.130.0.0/17:

# of addresses

the prefix mask

network address

1st host address

last host address

broadcast address

Example

Given this IP prefix: 82.130.0.0/17:

# of addresses
32,768 (32 bits - 17 = 15; 2^{15} = 32,768)

the prefix mask
255.255.128.0

network address
82.130.0.0

1st host address
82.130.0.1

last host address
82.130.127.254

broadcast address
82.130.127.255

Google “CIDR calculator” if you want to try on your own

Routers Forward Towards Destination Based on Network, NOT Host

Allows scalable forwarding tables

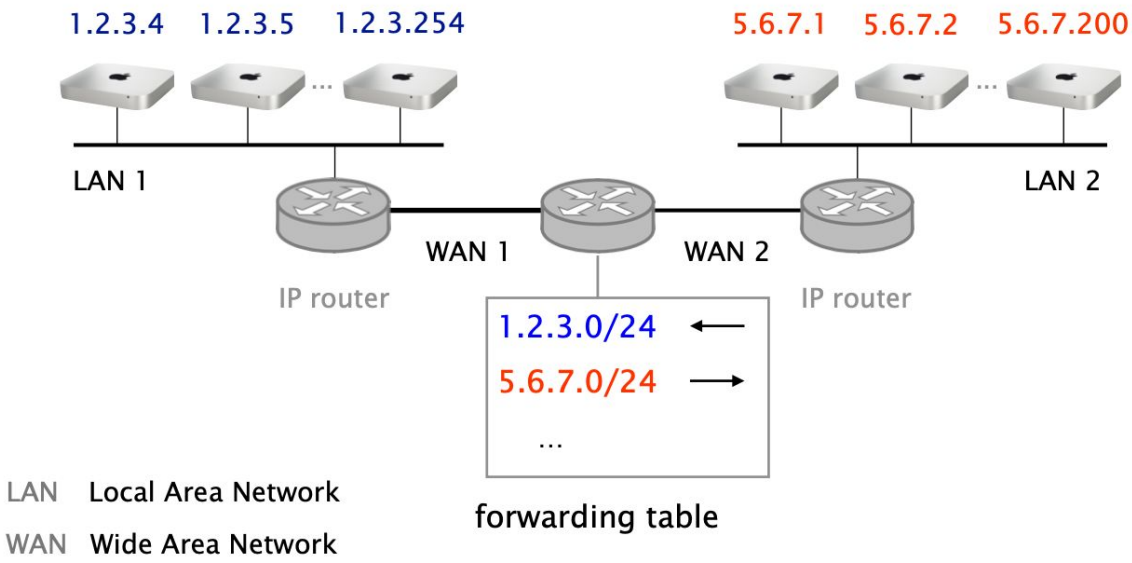

Hierarchical Addressing Allows Host Changes Without Forwarding Changes

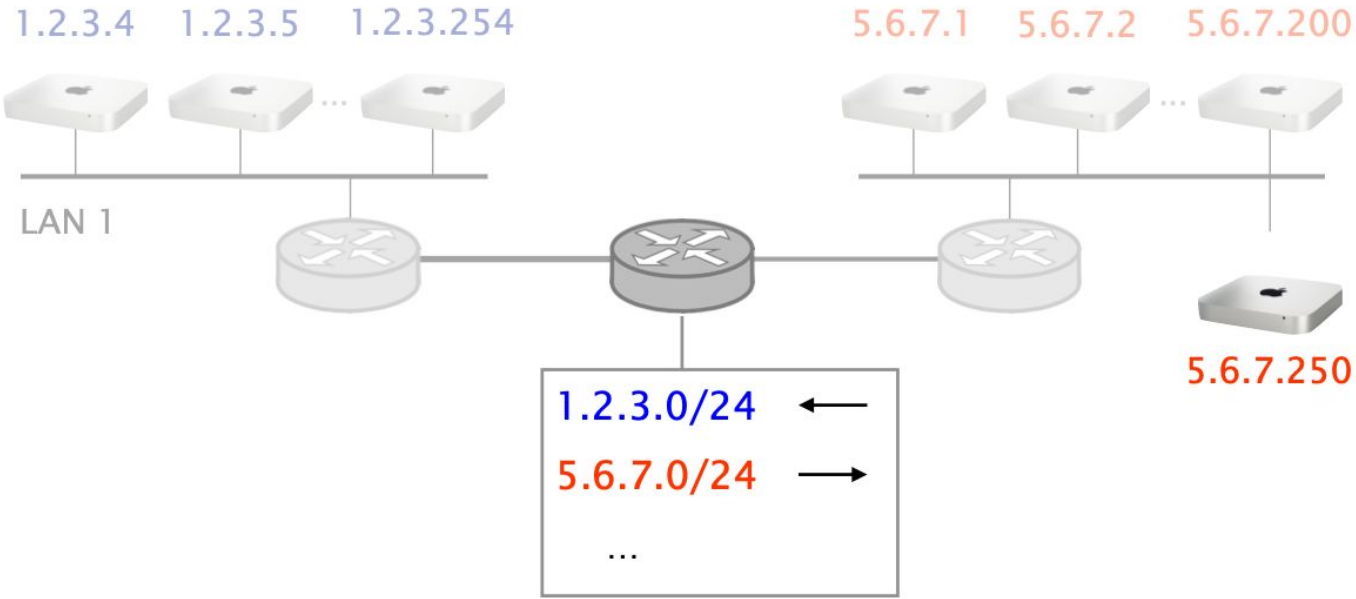

Originally There Were Fixed Allocation Sizes, Known as Classful Networking

	leading bits	prefix length	# hosts	start address	end address
class A	0	8	2^{24}	0.0.0.0	127.255.255.255
class B	10	16	2^{16}	128.0.0.0	191.255.255.255
class C	110	24	2^{8}	192.0.0.0	223.255.255.255
class D multicast	1110			224.0.0.0	239.255.255.255
class E reserved	1111			240.0.0.0	255.255.255.255

Classful Networking is Wasteful

problem

- Class C was too small, so everybody requested class B
  - but class Bs are too large, which led to wasted space

solution

- Classless Inter-Domain Routing (CIDR)
  - introduced in 1993

Classful Networking is Wasteful

Example (network needs 500 addresses)

with…	it gets a…	leading to a waste of…
Classful	class B (/16)	99%
CIDR	/23 (=2 class C’s)	2%

IP Address Allocation is Also Hierarchical

Root is held by the Internet Corporation for Assigned Names and Numbers, aka ICANN

ICANN Allocates Large Prefix Blocks to Regional Internet Registries

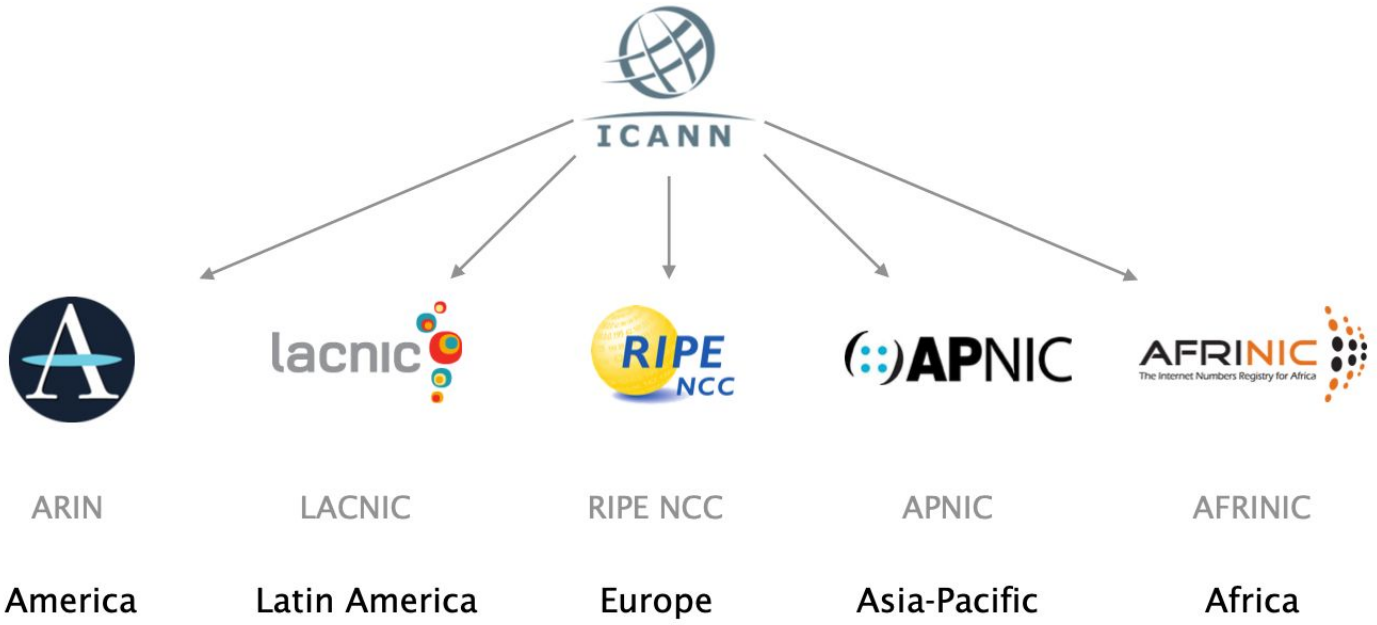

RIRs Allocate to ISPs and Large Organizations

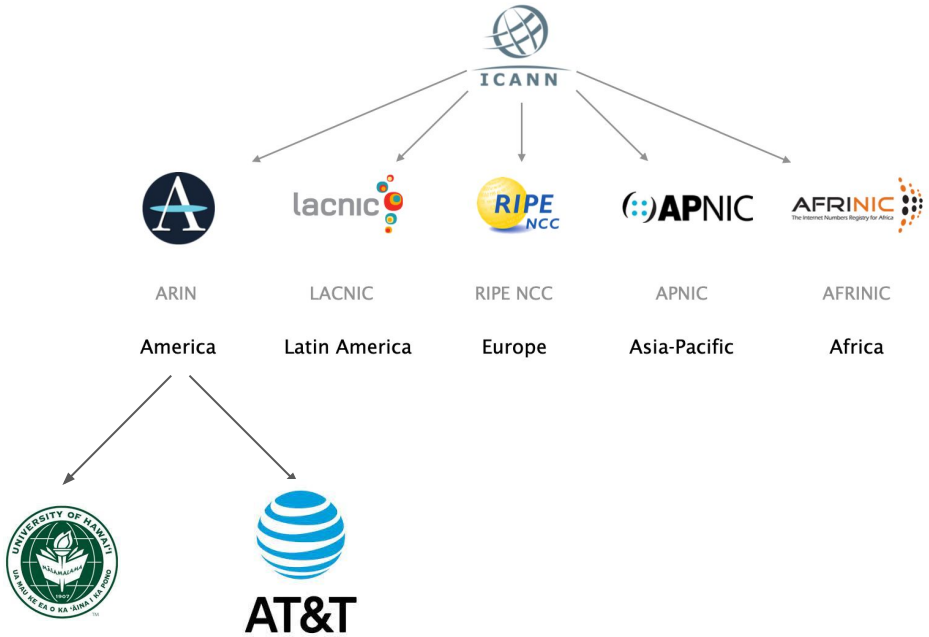

ISPs and Large Organizations Can Allocate Further

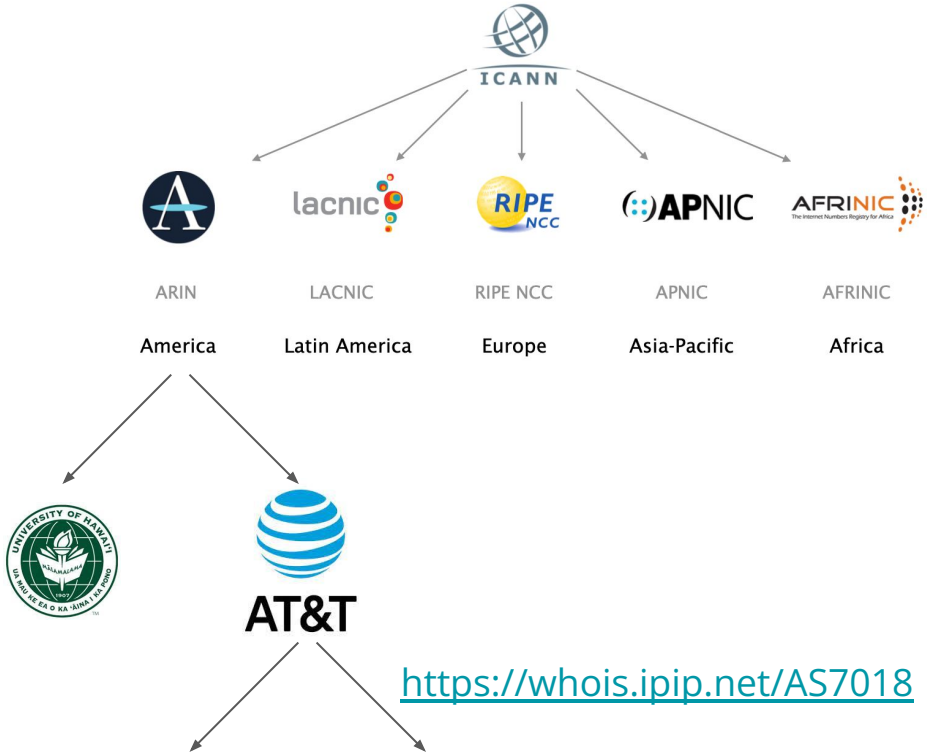

https://whois.ipip.net/AS7018

Network (IP) Layer

1. IP addresses
    - use, structure, allocation
2. IP forwarding
    - longest prefix match rule
3. IP header
    - IPv4 and IPv6, wire format

What's Inside a Router

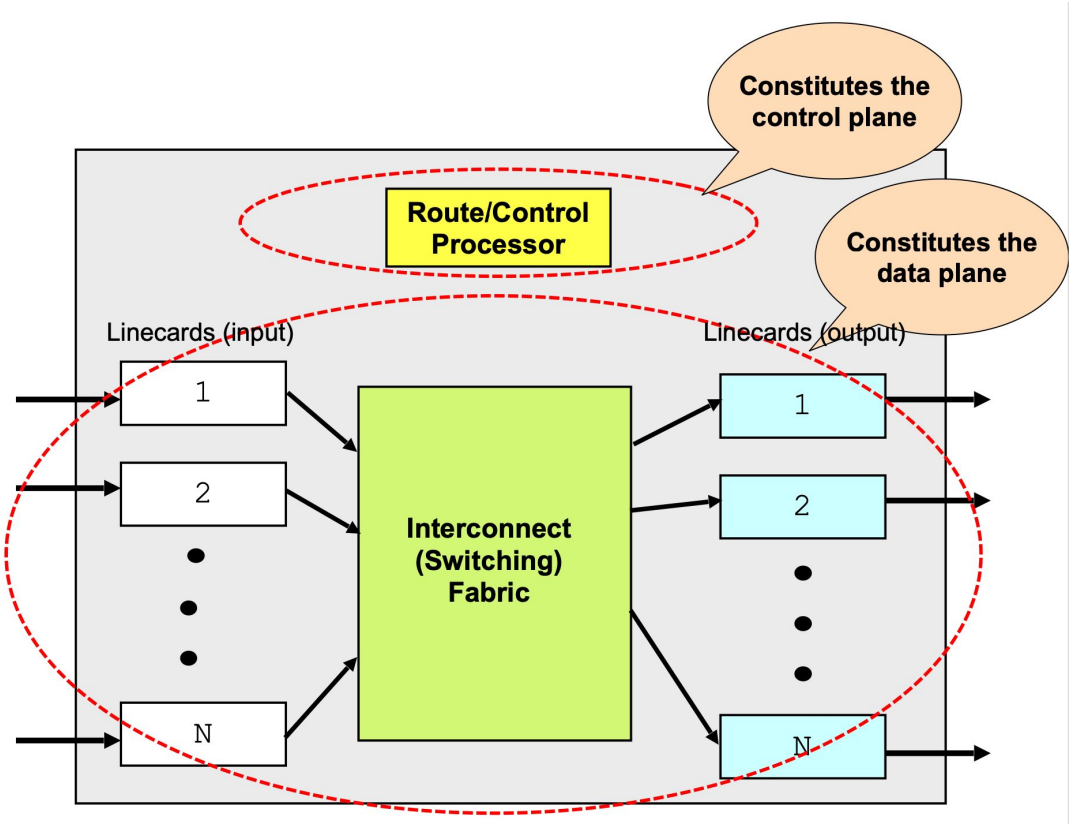

What's Inside a Router

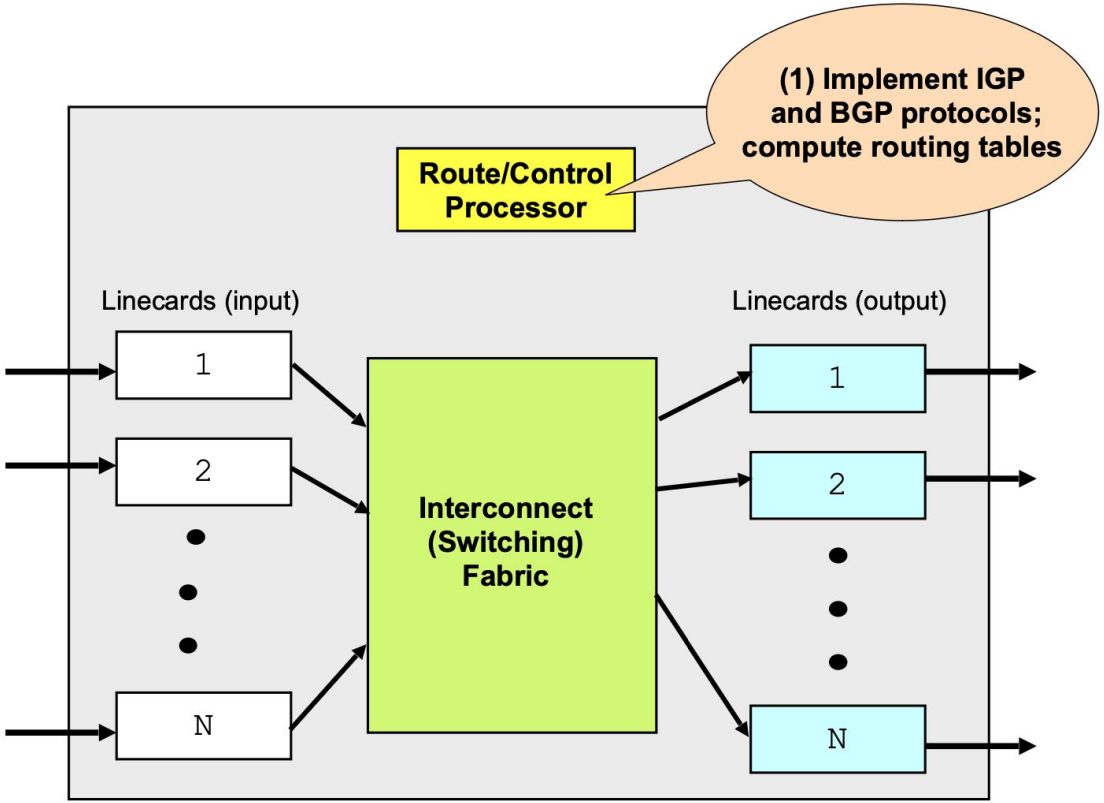

What's Inside a Router

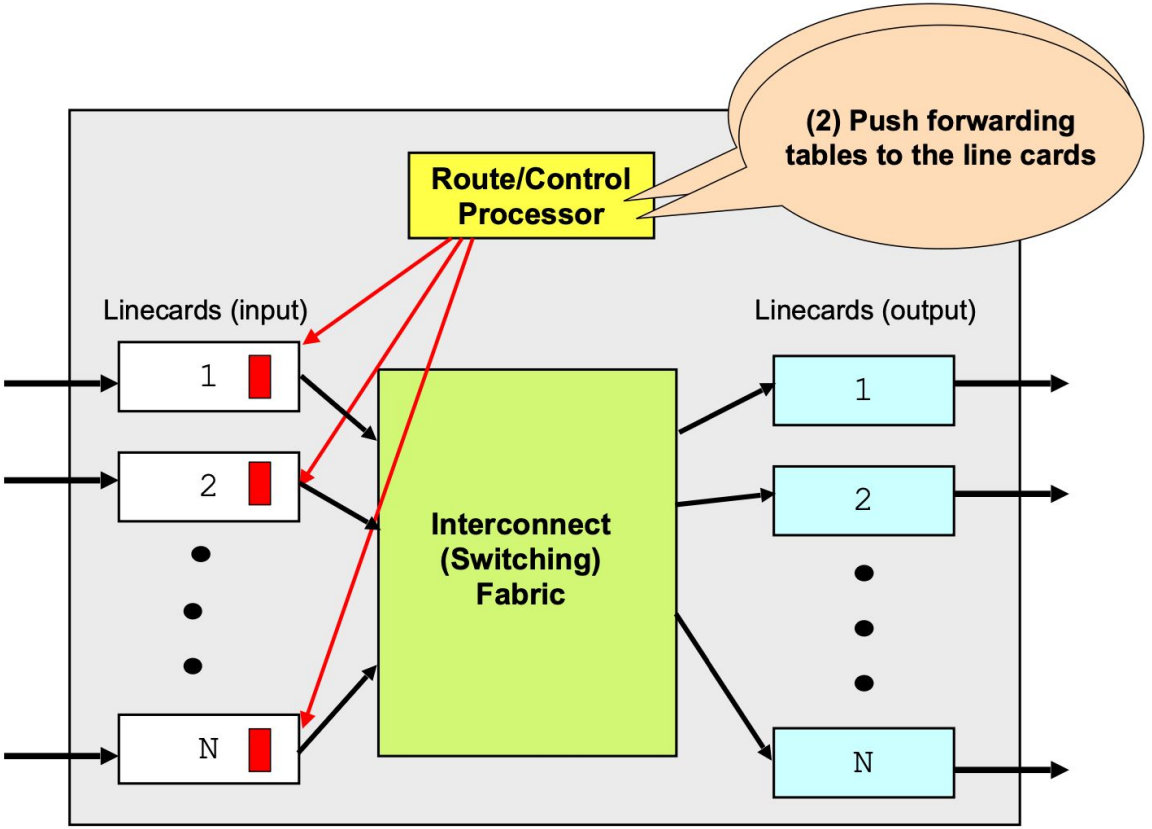

Routers Maintain Forwarding Entries for all Internet Prefixes

Provider 2's Forwarding table

IP prefix	Output
129.0.0.0/8	IF#2
129.132.1.0/24	IF#2
129.132.2.0/24	IF#2
129.133.0.0/16	IF#3

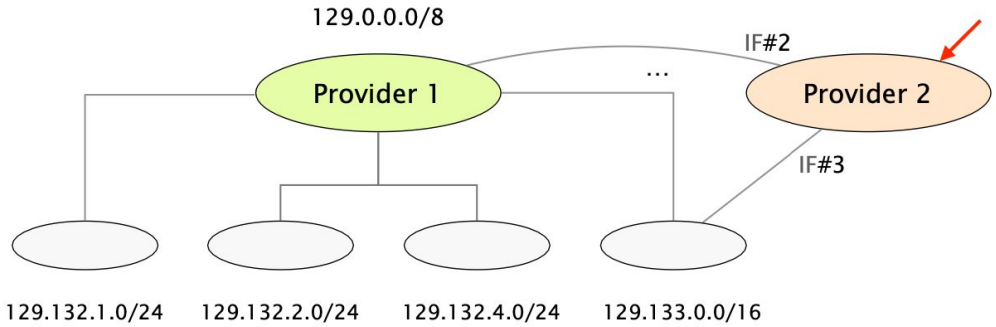

Routers Maintain Forwarding Entries for all Internet Prefixes

Provider 2's Forwarding table

IP prefix	Output
129.0.0.0/8	IF#2
129.132.1.0/24	IF#2
129.132.2.0/24	IF#2
129.133.0.0/16	IF#3

Packet arrives for 129.0.1.1

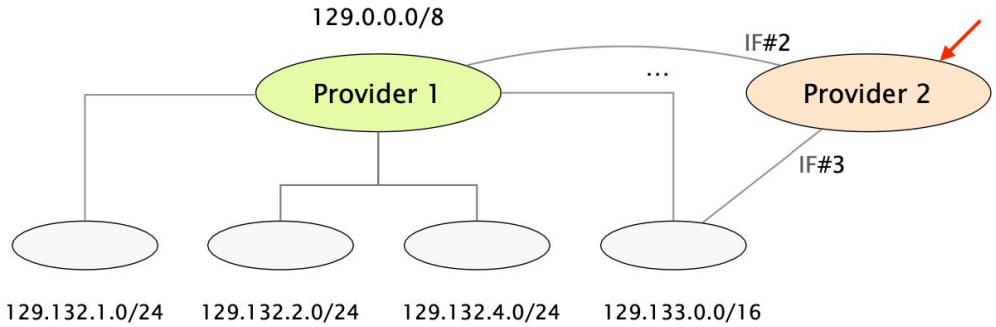

Routers Maintain Forwarding Entries for all Internet Prefixes

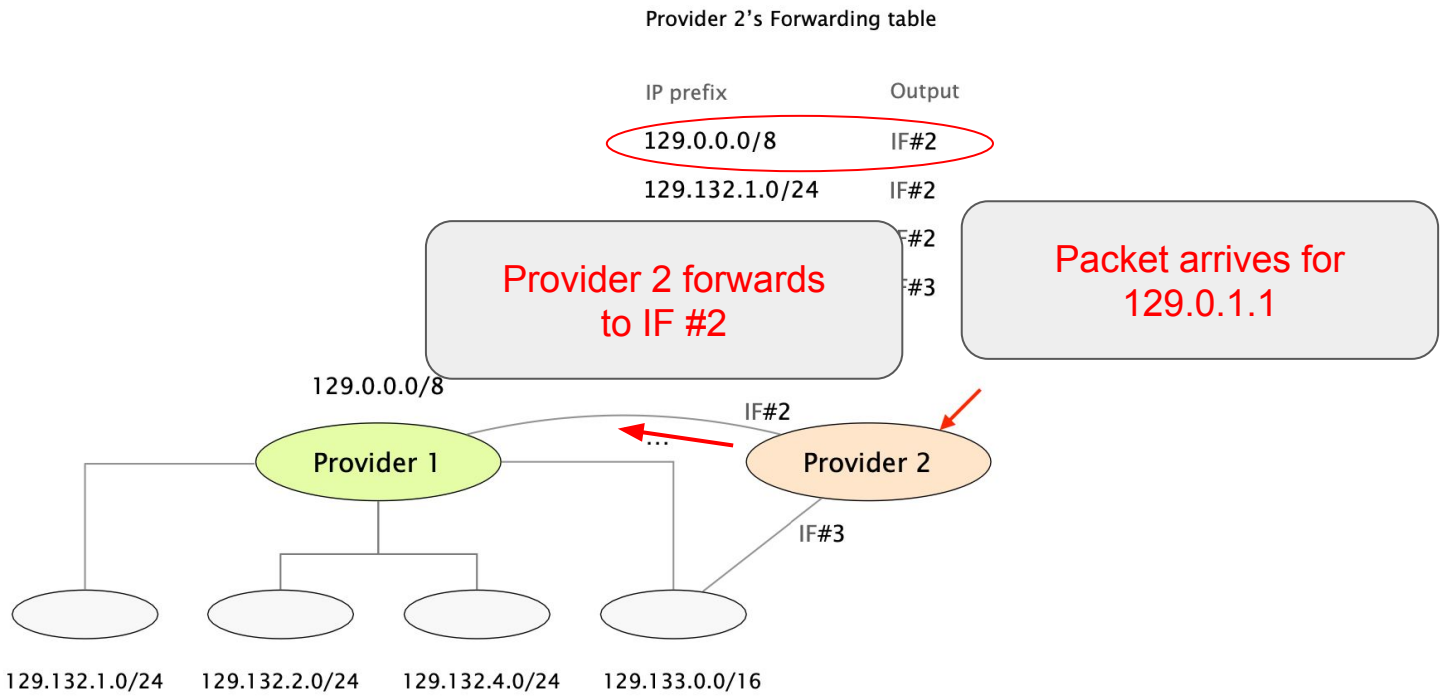

Does CIDR Make This Easier or Harder?

Provider 2's Forwarding table

IP prefix	Output
129.0.0.0/8	IF#2
129.132.1.0/24	IF#2
129.132.2.0/24	IF#2
129.133.0.0/16	IF#3

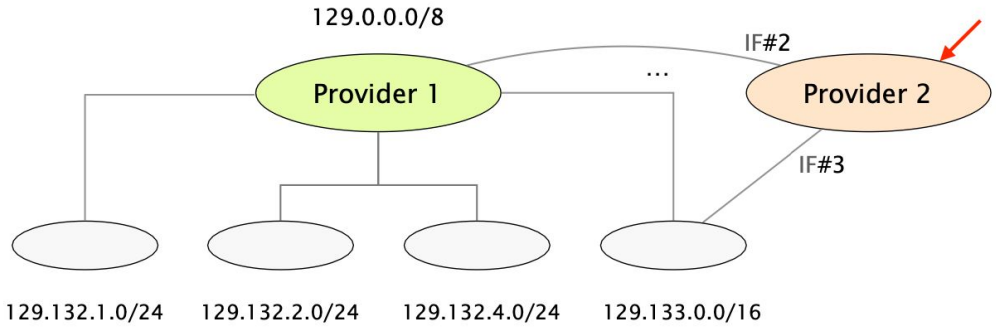

Does CIDR Make This Easier or Harder?

Multiple Matches

Provider 2's Forwarding table

IP prefix	Output
129.0.0.0/8	IF#2
129.132.1.0/24	IF#2
129.132.2.0/24	IF#2
129.133.0.0/16	IF#3

Packet arrives for 129.133.0.1

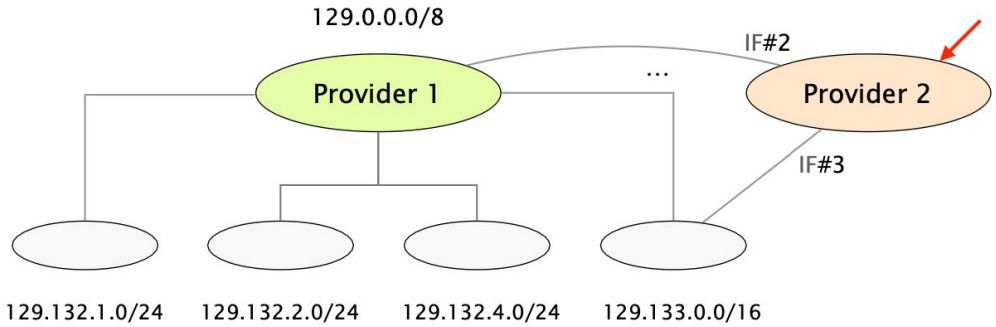

What Should We Do?

What Should We Do?

To resolve ambiguity, forwarding is done along the most specific prefix (i.e., the longer one)