

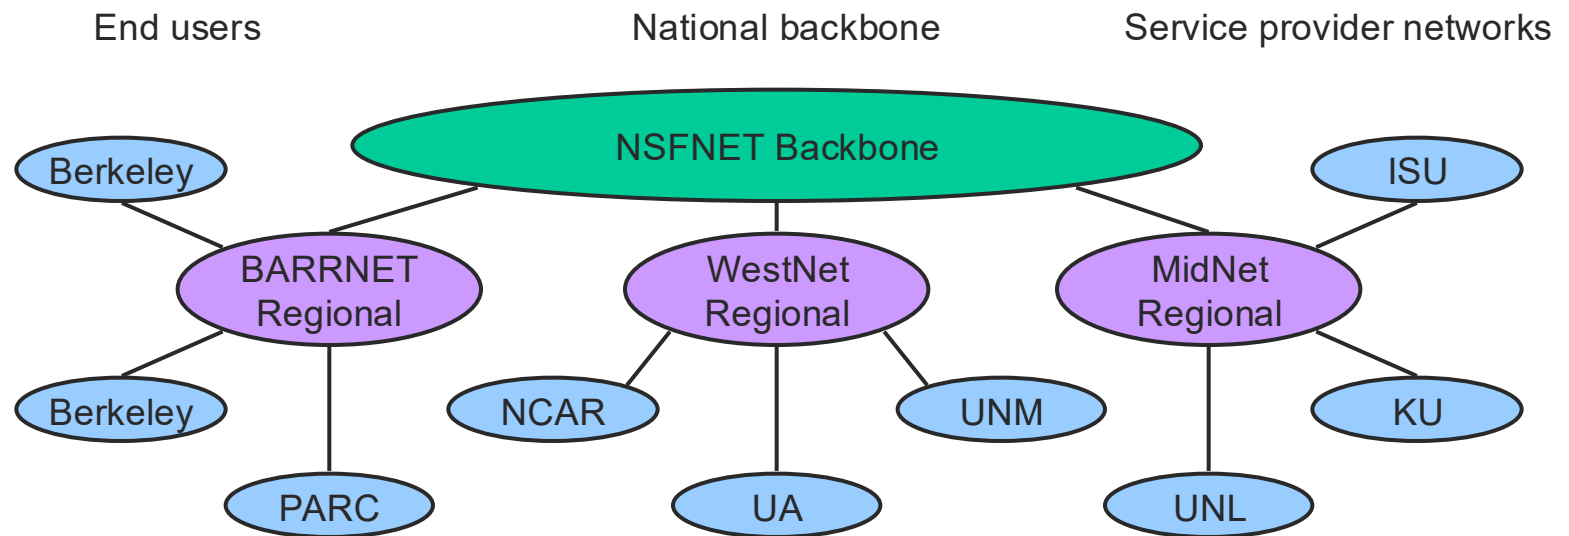


IP Addressing


Evolution of Internet Structure

■ Internet c. 1990


- Tree structure, centered around one backbone
- National Science Foundation (NSF) funded




[An Old Internet ISP Map]



[A New Internet Map]



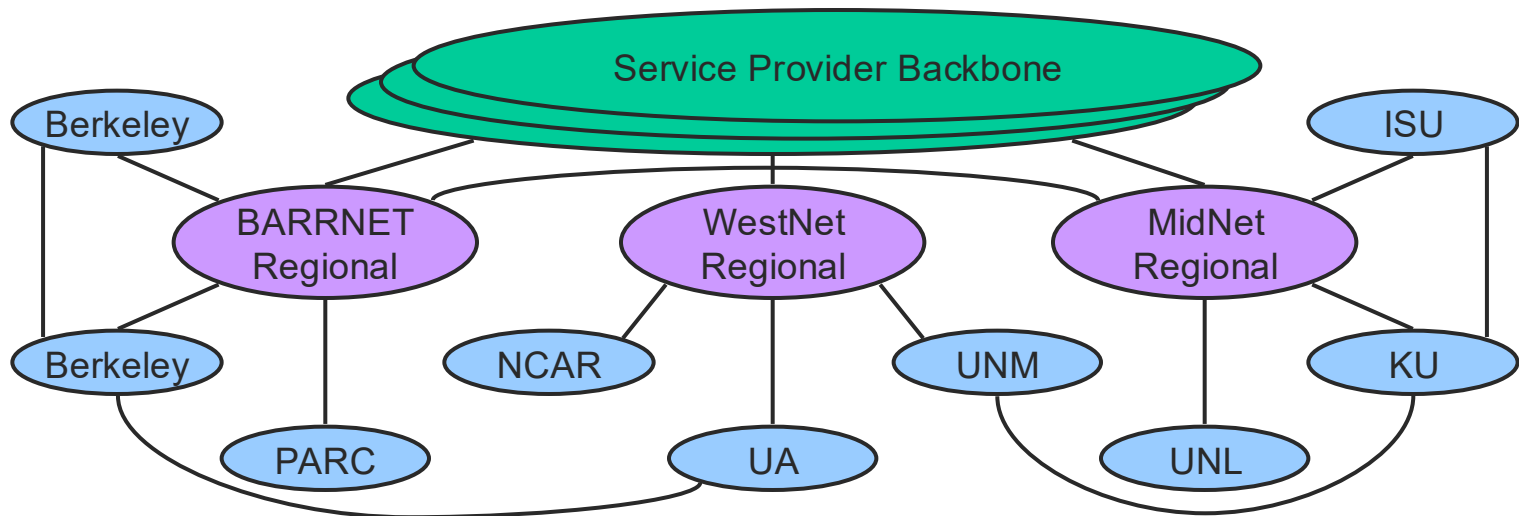
[Another Internet Map]



Evolution of Internet Structure

■ Today

- Multiple backbone service providers
- Arbitrary graph structure



[Problems of Scale]

- Main problems
 - Inefficient address allocation
 - Too many networks for routing

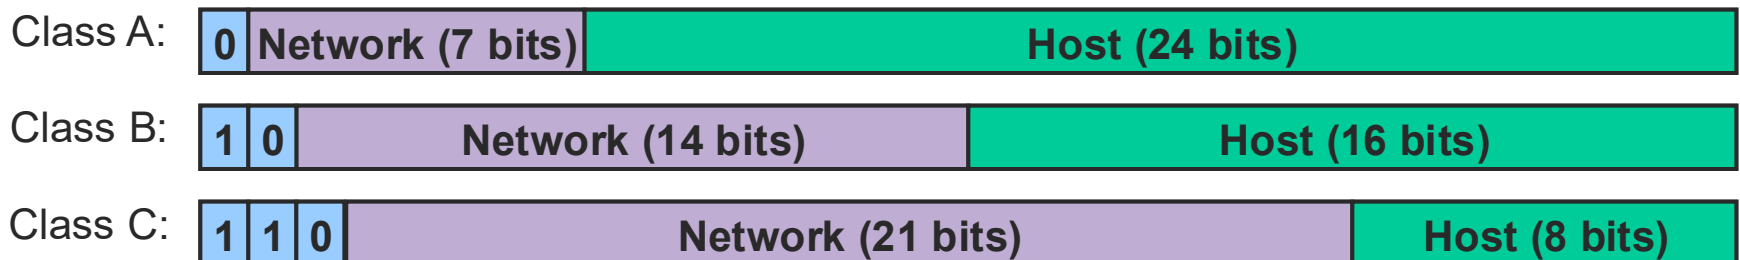


[IPv4 Address Model]

- Properties

- 32-bit address
- Hierarchical
 - Network, subnet, host hierarchy
- Maps to logically unique network adaptor
 - Exceptions: service request splitting for large web servers

- Three Class Model



[IPv4 Address Model]

- IP addresses
 - Host in class A network
 - 1.0.0.0 to 126.255.255.255
 - Host in class B network
 - 128.0.0.0 to 191.255.255.255
 - Host in class C network
 - 192.0.0.0 to 223.255.255.255



[IPv4 Address Model]

Class	Network ID	Host ID	# of Addresses	# of Networks
A	0 + 7 bit	24 bit	$2^{24}-2$	126
B	10 + 14 bit	16 bit	$65,536 - 2$	2^{14}
C	110 + 21 bit	8 bit	$256 - 2$	2^{21}
D	1110 + Multicast Address		IP Multicast	
E	Future Use			



Basic Datagram Forwarding with IP

- Hosts and routers maintain forwarding tables
 - List of **<prefix, next hop>** pairs
 - IP = 69.2.1.2 = 01000101 00000010 00000001 00000010
 - 24-bit prefix = 69.2.1.0/24
= 01000101 00000010 00000001 *****
 - Often contains a default route
 - Pass unknown destination to provider ISP
 - Simple and static on hosts, edge routers
 - Complex and dynamic on core routers



Basic Datagram Forwarding with IP

- Packet forwarding
 - Compare network portion of address with **<network/host, next hop>** pairs in table
 - Send directly to host on same network
 - Send to indirectly (via router on same network) to host on different network
 - Use ARP to get hardware address of host/router



[IPv4 Address Model]

- IP addresses
 - Number of hosts in class A network
 - Approximately 16.7 million
 - Number of hosts in class B network
 - Approximately 65 thousand
 - Number of hosts in class C network
 - 254
- Questions
 - What networks should be allocated to a company with 1000 machines?
 - What about a company with 100 machines?
 - What about a company with 2 machines that plans to grow rapidly?



[Problems of Scale]

- Pressure mostly on class B networks
 - Most companies plan to grow beyond 255 machines
 - Renumbering is time consuming and can interrupt service
 - Approximately 16,000 class B networks available
- Class B networks aren't very efficient
 - Few organizations have $O(10,000)$ machines
 - More likely use $O(1,000)$ of the 65,000 addresses
- Scaling problems with alternatives
 - Multiple class C networks
 - Routing tables don't scale
 - Protocols do not scale beyond $O(10,000)$ networks



IP Address Hierarchy Evolution

- Began with class-based system
 - Subnetting within an organization
 - Network can be broken into smaller networks
 - Recognized only within the organization
 - Implemented by packet switching
 - Smaller networks called subnets

Class A:



Class B:



Class C:



[Subnetting]

- Simple IP
 - All hosts on the same network must have the same **network** number
- Assumptions
 - Subnets are close together
 - Look like one network to distant routers
- Idea
 - Take a single IP network number
 - Allocate the IP addresses to several physical networks (subnets)
- Subnetting
 - All hosts on the same network must have the same **subnet** number



[Subnetting]

- Enables a domain to further partition address space into smaller networks
 - Subdivide host id into subnet ID + host ID
 - Subnet mask
- Only routers in the domain interpret subnet mask
 - Other routers treat IP address as normal class A, B or C address



[Subnet Example]

- Consider
 - A domain with a class B address
 - 135.104.*.*
- Without subnetting
 - Every router in the domain needs to know how to route to every host
- However
 - the domain itself is likely organized as a hierarchy of physical networks



[Subnet Example]

■ Solution

- Partition the 65K address in the class B network
 - 256 subnets each with 256 addresses
 - Subnet mask: 255.255.255.0
- If 135.104.5.{1,2,3} are all on the same physical network reachable from router 135.105.4.1
 - There only needs to be one routing entry for 135.104.5.* pointing to 135.105.4.1 as next hop



[Subnetting]

- Normal IP


Class B:



- Typical subnetting example

- Use first byte of host as subnet number

Class B:



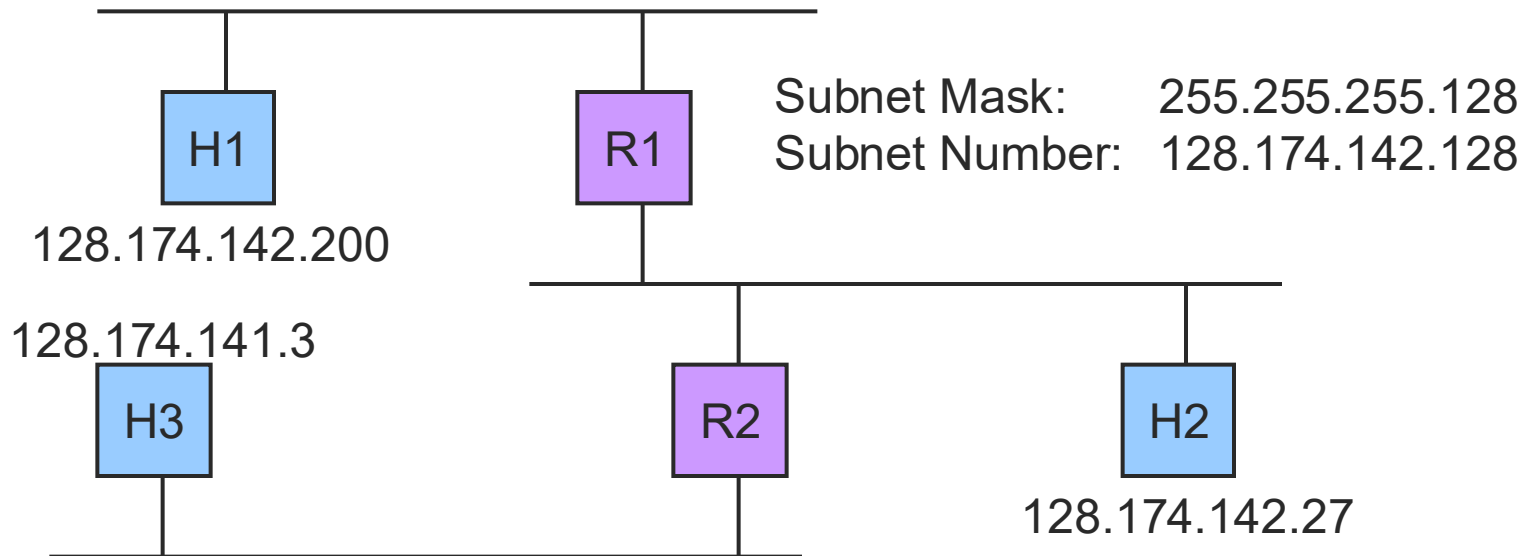
- Atypical example

- Non-contiguous 6-bit subnet number

Class B:



[Subnetting – Host 1]



Host 1: 128.174.142.200

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 1 0 0 1 0 0 0

Subnet Mask 255.255.255.128

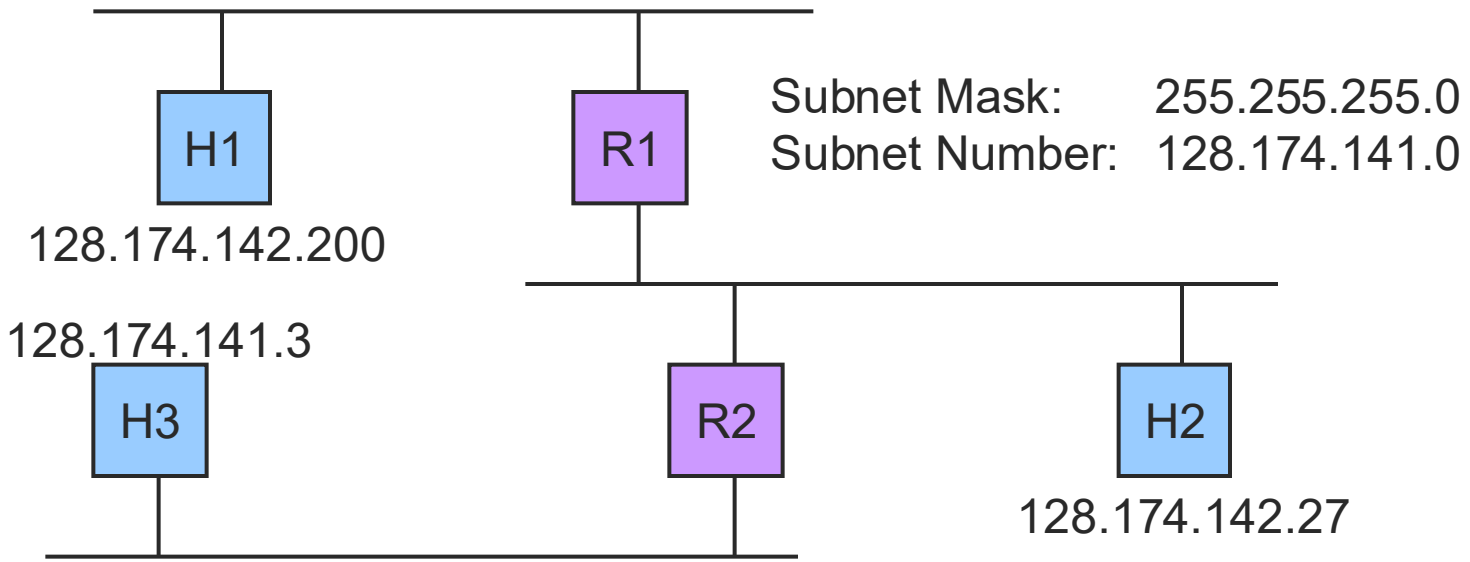
1 0 0 0 0 0 0 0

Subnet # 128.174.142.128

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 0 0 0 0 0 0 0



[Subnetting – Host 3]



Host 3: 128.174.141.3

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 0 1 0 0 0 0 0 0 1 1

Subnet Mask 255.255.255.0

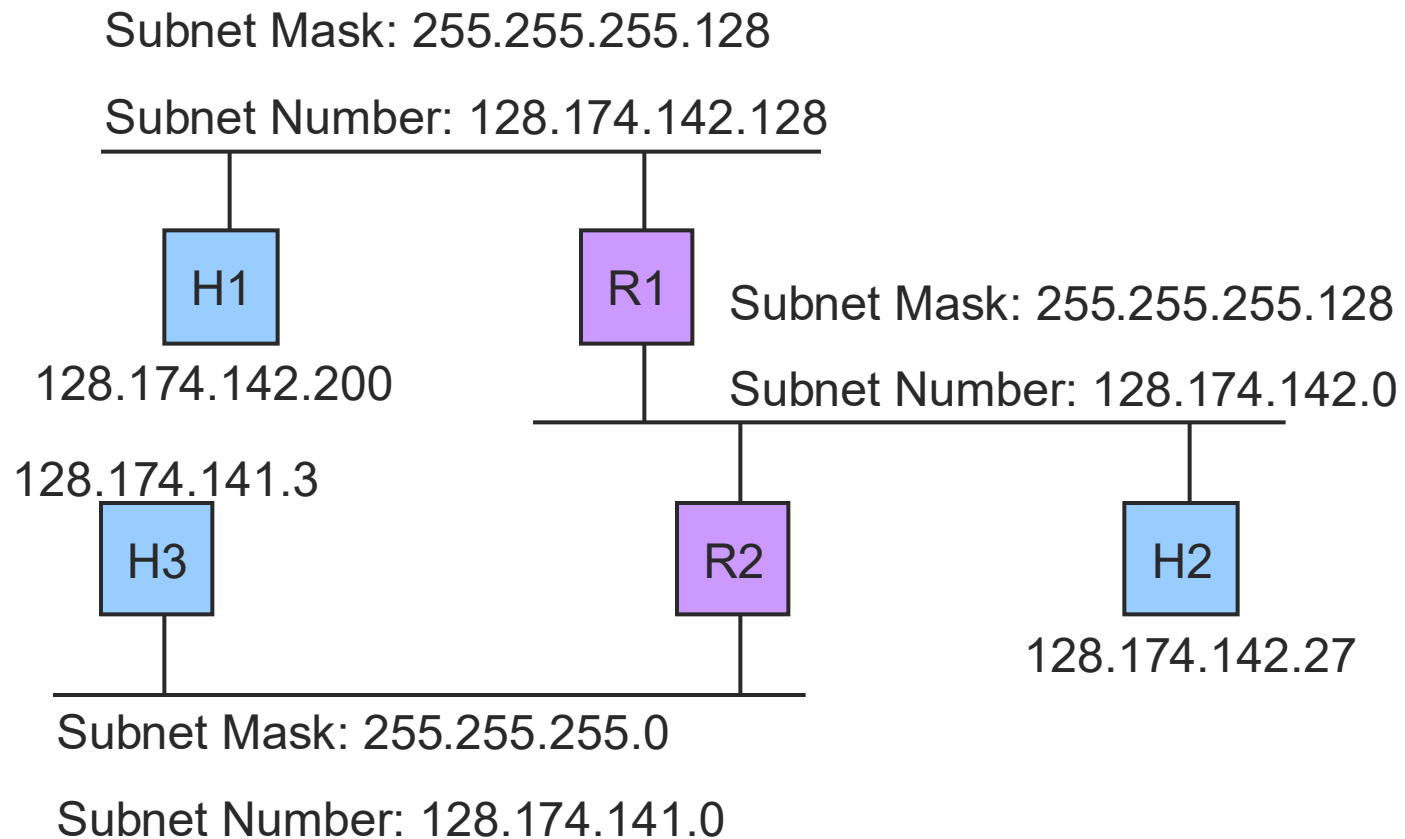
1 0 0 0 0 0 0 0 0

Subnet # 128.174.141.0

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 0 1 0 0 0 0 0 0 0 0

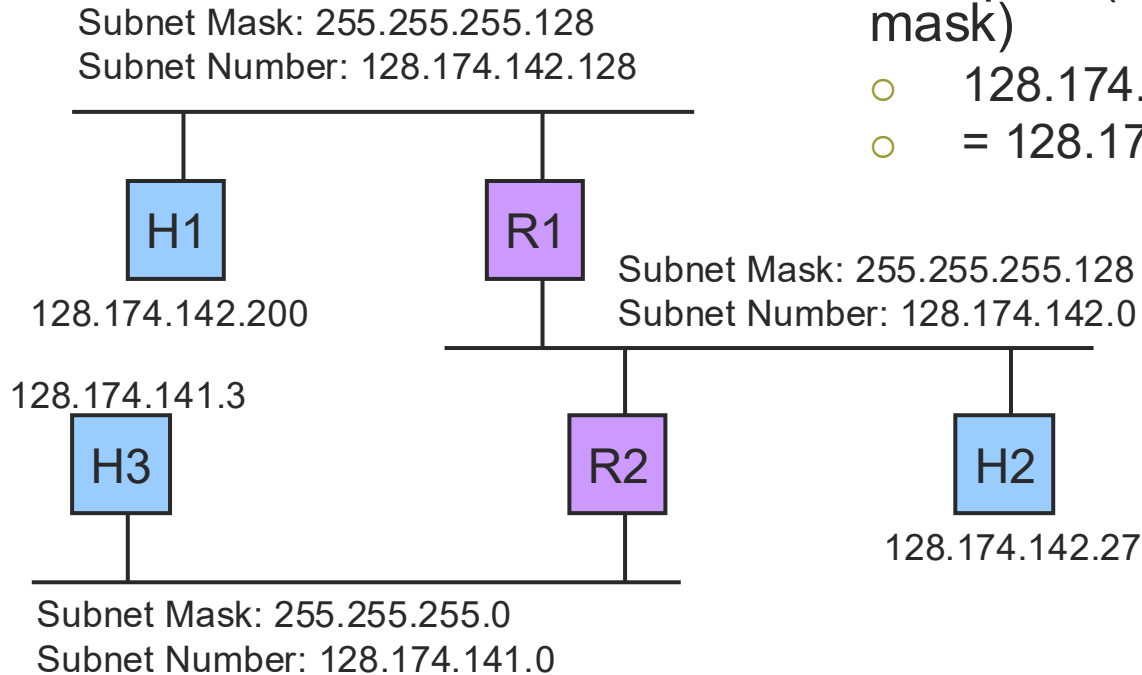


[Subnetting - Example]



[Subnetting]

Send from H1 to H3



- At H1:
- Compute (H3 AND H1's subnet mask)
 - 128.174.141.3 **AND** 255.255.255.128
 - = 128.174.141.0 (\neq 128.174.142.128)
- If result == H1's subnet number
 - H3 and H1 are on the same subnet
- else
 - route through appropriate router



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2
 - Next hop
 - 128.174.142.196
 - 128.174.142.95
 - 128.174.141.137
 - 128.174.145.18
 - 131.126.244.15

196 = 1100 0100 128 = 1000 0000
 141 = 1000 1101
 142 = 1000 1110
 145 = 1001 0001
 196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

$196 = 1100\ 0100$ $128 = 1000\ 0000$
 \neq
 $141 = 1000\ 1101$
 $142 = 1000\ 1110$
 $145 = 1001\ 0001$
 $196 = 1100\ 0100$



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

=

196 = 1100 0100 128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

196 = 1100 0100

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

196 = 1100 0100

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2
 - Next hop
 - 128.174.142.196 to R1
 - 128.174.142.95
 - 128.174.141.137
 - 128.174.145.18
 - 131.126.244.15

196 = 1100 0100 128 = 1000 0000
 141 = 1000 1101
 142 = 1000 1110
 145 = 1001 0001
 196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

?

- Example Table from R2
 - Next hop
 - 128.174.142.196 to R1
 - 128.174.142.95
 - 128.174.141.137
 - 128.174.145.18
 - 131.126.244.15

196 = 1100 0100 128 = 1000 0000
 141 = 1000 1101
 142 = 1000 1110
 145 = 1001 0001
 196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

95 = 0101 1111

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

95 = 0101 1111

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

≠

95 = 0101 1111
128 = 1000 0000

141 = 1000 1101
 142 = 1000 1110
 145 = 1001 0001
 196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

95 = 0101 1111

to Interface 1

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2
 - Next hop
 - 128.174.142.196
 - 128.174.142.95
 - 128.174.141.137
 - 128.174.145.18
 - 131.126.244.15

137 = 1000 1001
 128 = 1000 0000
 141 = 1000 1101
 142 = 1000 1110
 145 = 1001 0001
 196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

≠

137 = 1000 1001 128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

137 = 1000 1001

to Interface 0

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

■ Example Table from R2

○ Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

18 = 0001 0010 ≠

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

≠

18 = 0001 0010 128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

18 = 0001 0010

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

18 = 0001 0010

to R3

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

15 = 0000 1111

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



[Routing with Subnetting]

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2

- Next hop

- 128.174.142.196
- 128.174.142.95
- 128.174.141.137
- 128.174.145.18
- 131.126.244.15

15 = 0000 1111

to R3

128 = 1000 0000

141 = 1000 1101

142 = 1000 1110

145 = 1001 0001

196 = 1100 0100



Routing with Subnetting

Subnet #	Subnet Mask	Next Hop
128.174.141.0	255.255.255.0	Interface 0
128.174.142.0	255.255.255.128	Interface 1
128.174.142.128	255.255.255.128	R1
128.174.0.0	255.255.0.0	R3
Default	0.0.0.0	R3

- Example Table from R2
 - Next hop
 - 128.174.142.196 to R1
 - 128.174.142.95 to Interface 1
 - 128.174.141.137 to Interface 0
 - 128.174.145.18 to R3
 - 131.126.244.15 to R3



[Subnetting]

■ Notes

- Non-contiguous subnets are difficult to administer
- Multiple subnets on one physical network
 - Must be routed through router

■ Pros

- Helps address consumption
- Helps reduce routing table size



[The Crisis]

- Fixed 32-bit address space for IPv4
- Network allocation based on Classic A, B, C Model
- Central allocation authority
 - Randomly assigning addresses
- Problems
 - Router table explosion
 - Address space exhaustion



Classless Interdomain Routing (CIDR)

- CIDR/Supernetting
 - Problem with subnetting
 - Allows hierarchy within organizations
 - Does not reduce class B address space pressure
 - Solution
 - Aggregate routes in routing tables
 - Eliminate class notation
 - Generalize subnet notion
 - Allow only contiguous subnet masks
 - Specify network by <network #, # of bits in subnet mask>
 - Equivalent to <network #, # of hosts>
 - Blocks of class C networks can now be treated as one network



[CIDR]

- Route aggregation
 - Use contiguous blocks of Class C addresses
 - Example:
 - 192.4.16 – 192.4.31
 - 20 bit subnet mask
 - Block size must be a power of 2
 - Network number may be any length

192.4.16.0

1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 1 0 0 0 0 0 0 0 0 0 0 0 0

192.4.31.0

1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0 0 0 1 1 1 1 1 0 0 0 0 0 0 0 0

Subnet Mask

1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0



[CIDR]

Subnet # / length	Next Hop
128.174.141.0 / 24	Interface 0
128.174.142.192 / 27	Interface 1
128.174.142.128 / 25	R1
128.174.0.0 / 16	R3
Default	R3

- CIDR is similar to subnetting
 - Trend is for increasing amounts of overlap in routing table entries
 - Example: 128.174.142.200
 - Matches second, third and fourth lines
 - Route to entry with longest match



[CIDR]

Subnet: 128.174.141.0

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 0 1 0 0 0 0 0 0 0 0

Subnet Mask length = 24 (255.255.255.0)

1 0 0 0 0 0 0 0 0

Host: 128.174.142.200

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 1 0 0 1 0 0 0

Resulting Subnet Number: 128.174.142.0 (\neq 128.174.141.0)

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 0 0 0 0 0 0 0 0

Subnet: 128.174.142.192

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 1 0 0 0 0 0 0

Subnet Mask length = 27 (255.255.255.224)

1 0 0 0 0 0 0 0

Host: 128.174.142.200

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 1 0 0 1 0 0 0

Resulting Subnet Number: 128.174.142.192 (= 128.174.142.192)

1 0 0 0 0 0 0 0 1 0 1 0 1 1 1 0 1 0 0 0 1 1 1 0 1 1 0 0 0 0 0 0 0




[CIDR]

Subnet: 128.174.142.128



Subnet Mask length = 25 255.255.255.192




Host: 128.174.142.200



Resulting Subnet Number: 128.174.142.128 (= 128.174.142.128)



Subnet: 128.174.0.0




Subnet Mask length = 16 255.255.0.0



Host: 128.174.142.200



Resulting Subnet Number: 128.174.0.0 (= 128.174.0.0)



[CIDR]

- Subnetting
 - Share one address (network number) across multiple physical networks
- Supernetting
 - Aggregate multiple addresses (network numbers) for one physical network



[CIDR]

- Allows hierarchical development
 - Assign a block of addresses to a regional provider
 - Ex: 128.0.0.0/9 to provider
 - Regional provider subdivides address and hands out block to sub-regional providers
 - Ex: 128.132.0.0/16 to sub-provider
 - Sub-regional providers can divide further for smaller organizations
 - Ex: 128.132.32.0/1 to organization



[Pros and Cons]

- Provides a fast easy solution
- Was not intended to be permanent
- Multihomed sites cannot benefit from aggregation
- Not backward compatible



[IPv6]

- History
 - Next generation IP (AKA IPng)
 - Intended to extend address space and routing limitations of IPv4
 - Requires header change
 - Attempted to include everything new in one change
 - IETF moderated
 - Based on Simple Internet Protocol Plus (SIPP)



[IPv6]

- Wish list
 - 128-bit addresses
 - Multicast traffic
 - Mobility
 - Real-time traffic/quality of service guarantees
 - Authentication and security
 - Autoconfiguration for local IP addresses
 - End-to-end fragmentation
 - Protocol extensions
- Smooth transition!
- Note
 - Many of these functionalities have been retrofit into IPv4

