

# Lecture 10.

## Subnetting & Supernetting

# Outline

## → Subnetting

⇒ Variable Length Subnet Mask (VLSM)

## → Supernetting

⇒ Classless Inter-Domain Routing (CIDR)

# medium org: N x class C? Class B?

→ **Class C addresses:**

⇒ *Undersized (254 hosts)*

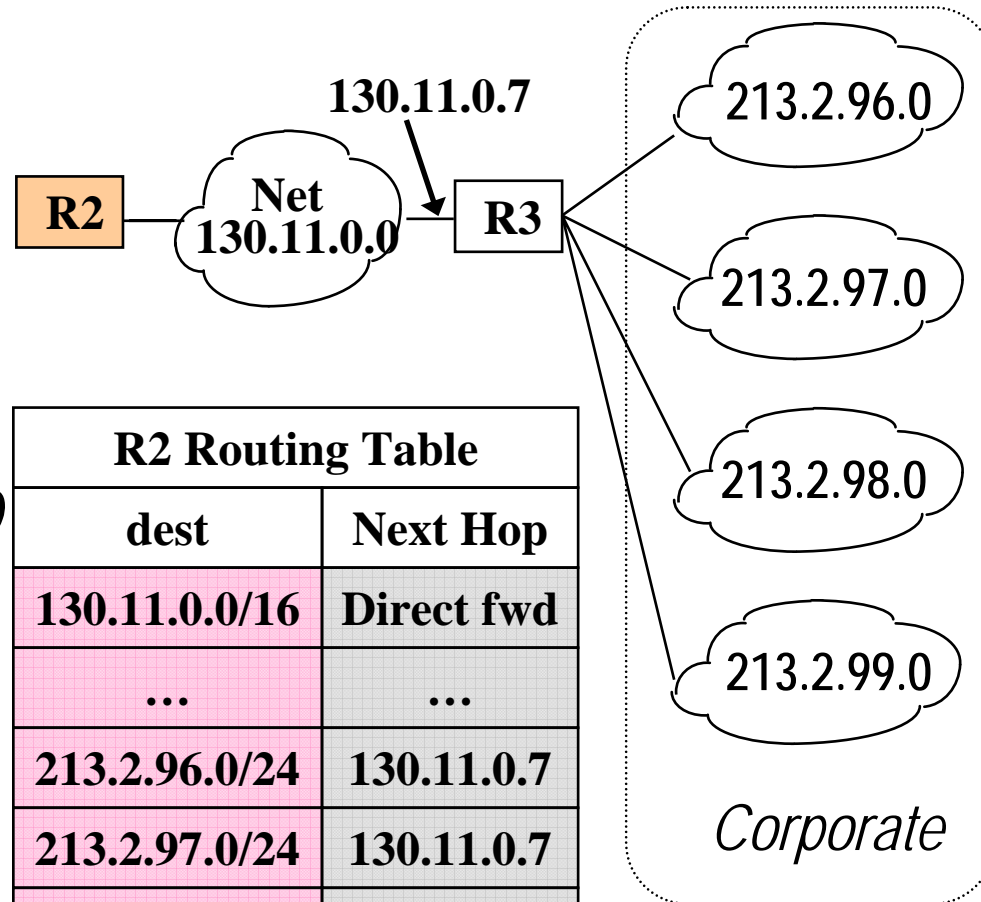
→ **Class B addresses:**

⇒ *Much more than enough (65534 hosts)*

→ **N x class C:**

⇒ *Unwise: exponential growth of routing tables*

→ **Result: Class B addresses were largely preferred**




R2 Routing Table	
dest	Next Hop
130.11.0.0/16	Direct fwd
...	...
213.2.96.0/24	130.11.0.7
213.2.97.0/24	130.11.0.7
213.2.98.0/24	130.11.0.7
213.2.99.0/24	130.11.0.7

*The aftermath: 10 bit class C design would have been much better...*

# Need for subnetting

## → Net\_id-Host\_id:

⇒ place host\_id on physical network net\_id



CLASS B:

From: 131.175.0.1

To: 131.175.255.254

65534 hosts on a same physical network????

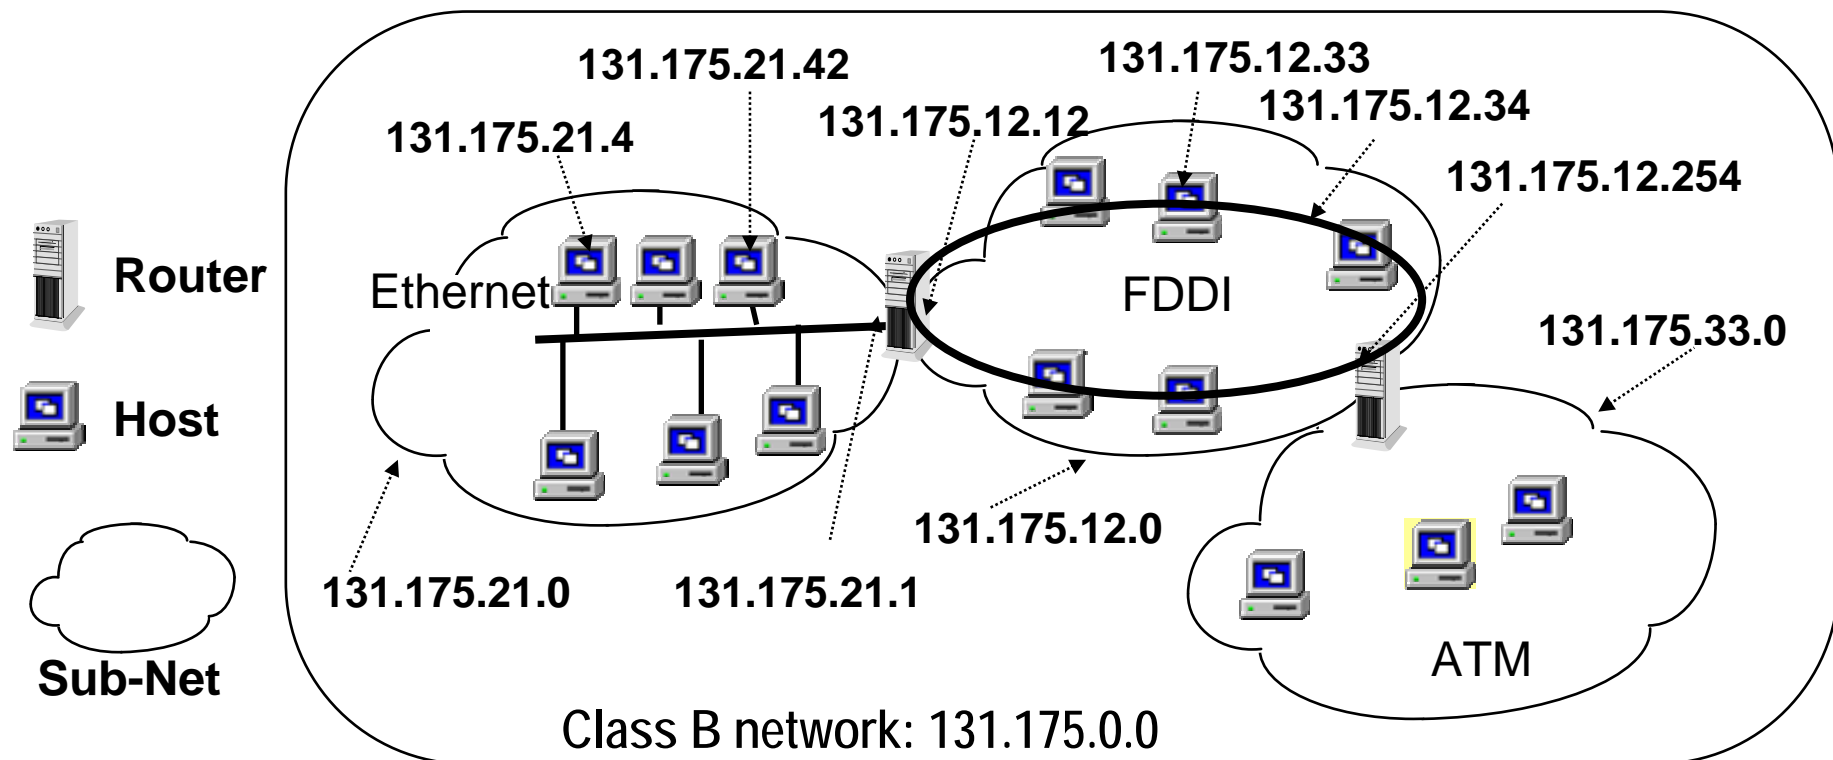
- performance?

- management?

# Idea: further hierarchy level

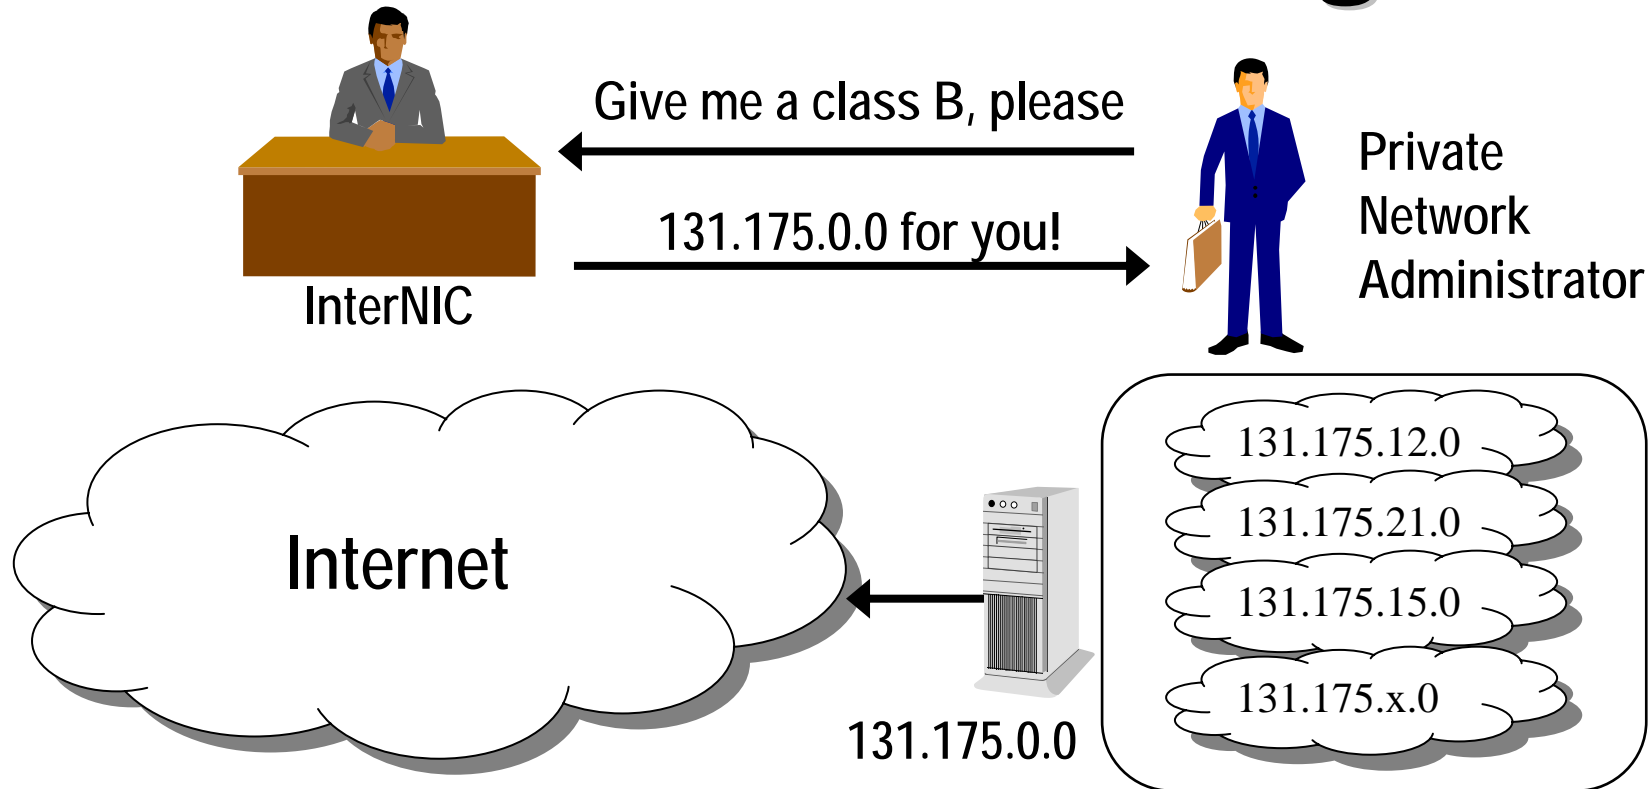
⇒ subdivide a network in several subnetworks

⇒ each subnet = a physical network (Ethernet, FDDI, X.25, ATM, Frame Relay, etc....)



May use third byte to identify subnet: 131.175.X.0 (or may not!)

# Subnet creation and management



## Best for local administrator:

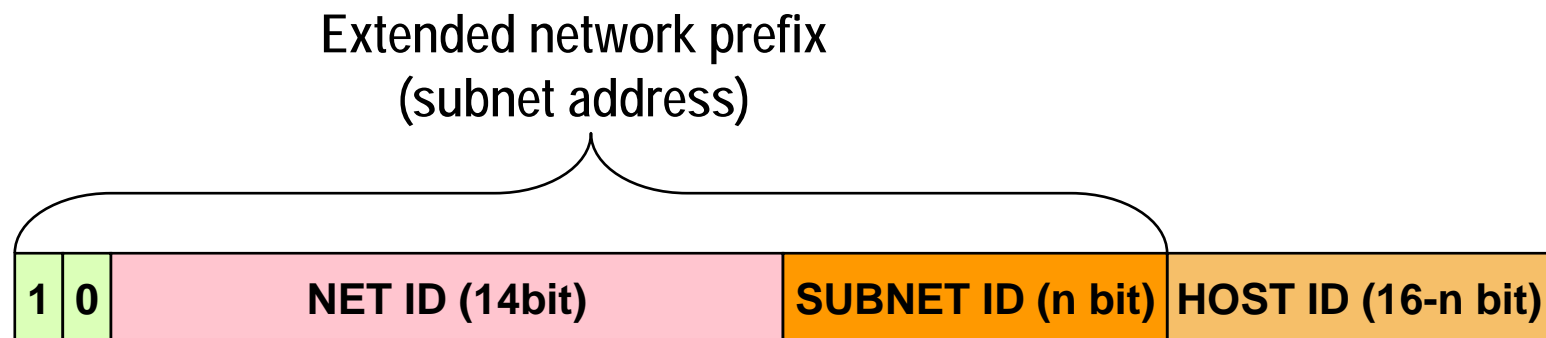

flexibility to create new networks without asking InterNIC new classful addresses.

## Best for Internet:

Route flapping in the private domain do not affect Internet  
One single entry in core router tables address all subnetworks

# Subnetting

## Class B address example



# Subnet Address & Mask

→ Host IP address:

159.100.9.18      10011111.01100100.00001001.00010010

→ Class B - network mask:

255.255.0.0      11111111.11111111.00000000.00000000

→ Subnet Mask

⇒ *Longer than natural class mask; Length set by administrator*

⇒ *Tells where the boundary network-host really is*

→ Example: class B address with 5 bits subnet\_id

⇒ *subnet mask = /21*

11111111.11111111.11111000.00000000

⇒ /prefix-length notation

⇒ *subnet mask = 255.255.248.0*

⇒ (dot decimal notation)

⇒ *159.100.0.0 = net\_id*

10011111.01100100.00001000.00000000

⇒ *159.100.8.0 = extended network address (net\_id+subnet\_id)*

⇒ *To avoid ambiguity: 159.100.8.0/21*



# Typical class B subnetting

→ **Class B address = /16 network prefix**

→ network address = 131.175.0.0

→ natural mask = 255.255.0.0

→ **Subnetted with /24 network prefix**



⇒ 255.255.255.0 subnet mask

⇒ subnet ID = third number in dotted notation

→ 131.175.**21**.0


No technical reasons to use /24 subnets, but convenient for humans  
(subnet boundary clearly visible in dotted notation)

# Remember: subnetting is arbitrary!


## Example: subnetting Class C 193.1.1.0 Address

Base net      11000001.00000001.00000001.00000000      193.1.1.0/24

Class C  
/24 prefix



Subnetted  
255.255.255.224  
/27prefix



Subnet # 0	11000001.00000001.00000001.00000000	<b>193.1.1.0/27</b>
Subnet # 1	11000001.00000001.00000001.00100000	<b>193.1.1.32/27</b>
Subnet # 2	11000001.00000001.00000001.01000000	<b>193.1.1.64/27</b>
Subnet # 3	11000001.00000001.00000001.01100000	<b>193.1.1.96/27</b>
Subnet # 4	11000001.00000001.00000001.10000000	<b>193.1.1.128/27</b>
Subnet # 5	11000001.00000001.00000001.10100000	<b>193.1.1.160/27</b>
Subnet # 6	11000001.00000001.00000001.11000000	<b>193.1.1.192/27</b>
Subnet # 7	11000001.00000001.00000001.11100000	<b>193.1.1.224/27</b>

Remember: maximum  $30(2^5-2)$  hosts attachable to each subnet

# Possible netmask values

128	64	32	16	8	4	2	1	
1	0	0	0	0	0	0	0	= 128
1	1	0	0	0	0	0	0	= 192
1	1	1	0	0	0	0	0	= 224
1	1	1	1	0	0	0	0	= 240
1	1	1	1	1	0	0	0	= 248
1	1	1	1	1	1	0	0	= 252
1	1	1	1	1	1	1	0	= 254
1	1	1	1	1	1	1	1	= 255

# Example: route 193.205.102.36

193								205								102								36							
1	1	0	0	0	0	0	1	1	1	0	0	1	1	0	1	0	1	1	0	0	1	1	0	0	0	1	0	0	1	0	0

Class C address;

Outside private domain routed with mask 255.255.255.0

network																								host							
193								205								102								36							
1	1	0	0	0	0	0	1	1	1	0	0	1	1	0	1	0	1	1	0	0	1	1	0	0	0	1	0	0	1	0	0

Inside private domain, administrator has set netmask 255.255.255.248

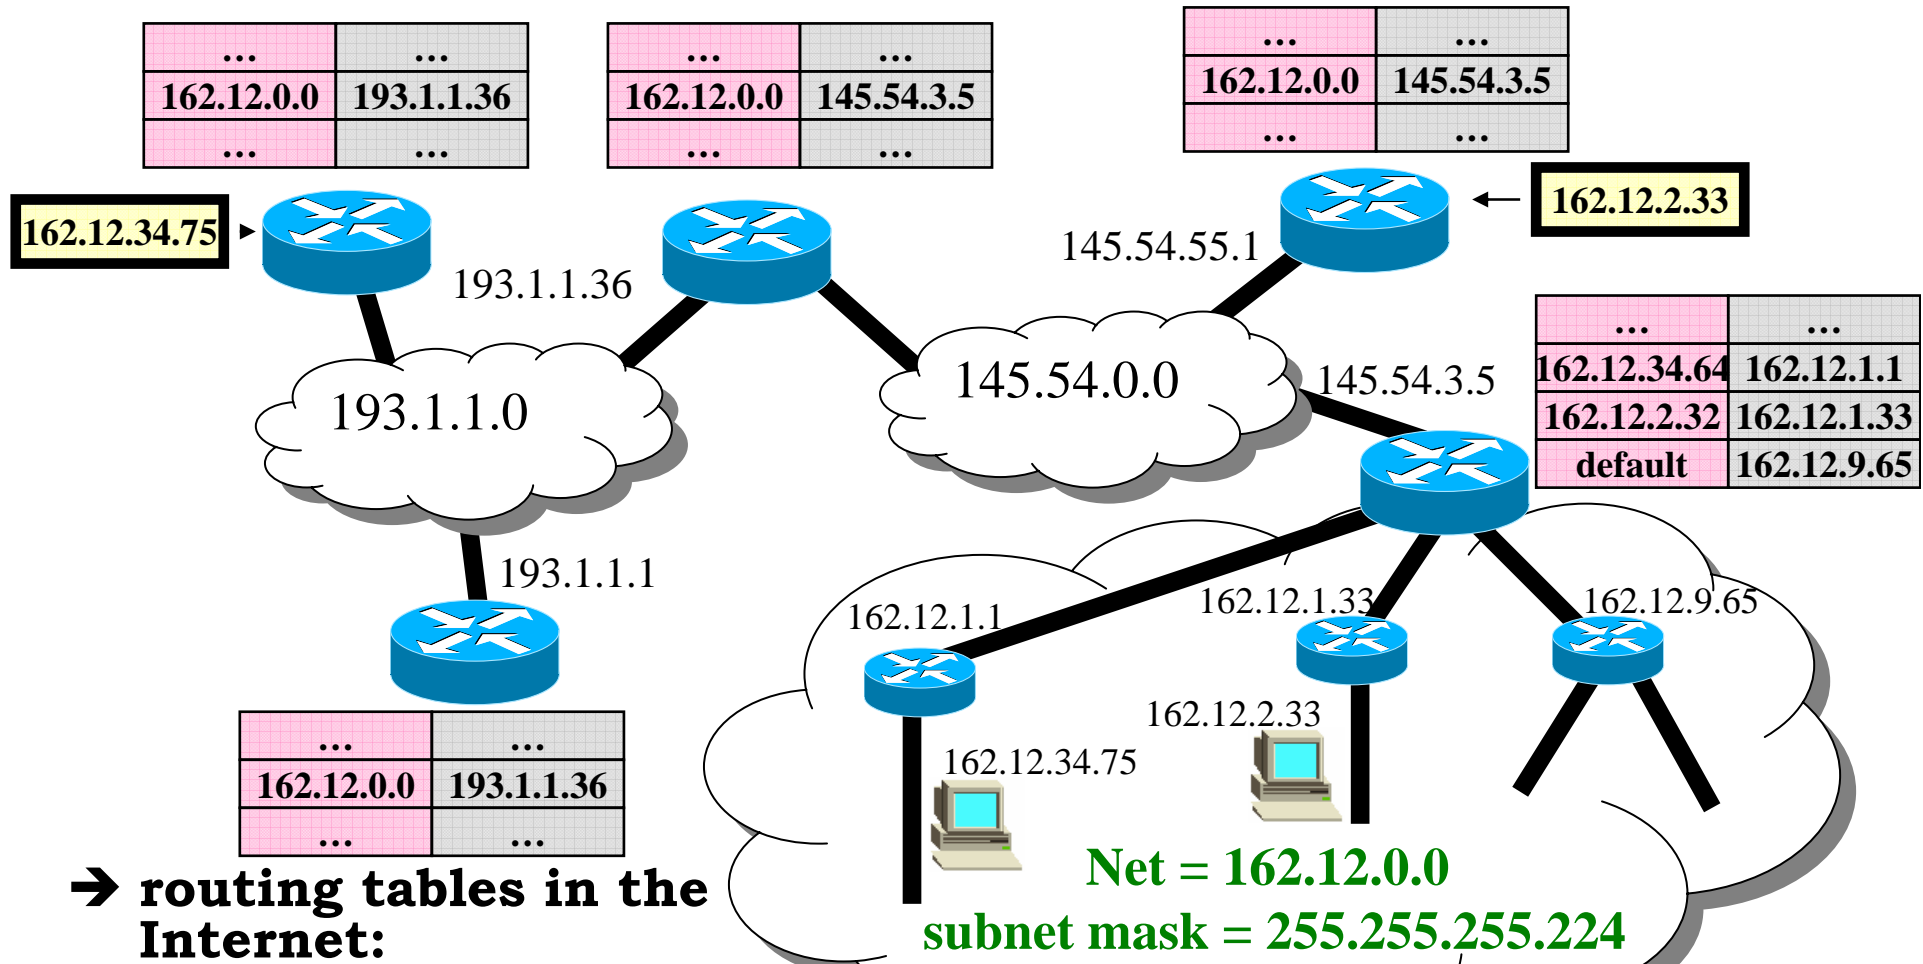
255								255								255								248							
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0

Hence, route to subnet address and then to host id, computed as:

network																								subnet				host			
1	1	0	0	0	0	0	1	1	1	0	0	1	1	0	1	0	1	1	0	0	1	1	0	0	0	1	0	0	1	0	0
193.205.102.32 /29																								4							

# Subnet routing – 2nd example

Core routers unaware of subnetting – route via class mask



→ **routing tables in the Internet:**

- ⇒ route according to net\_id
- ⇒ Use natural class mask

→ **Corporate routers & hosts:**

- ⇒ Route according to subnet\_id
- ⇒ Need to KNOW subnet mask

# Router configuration

## → Classful routing:

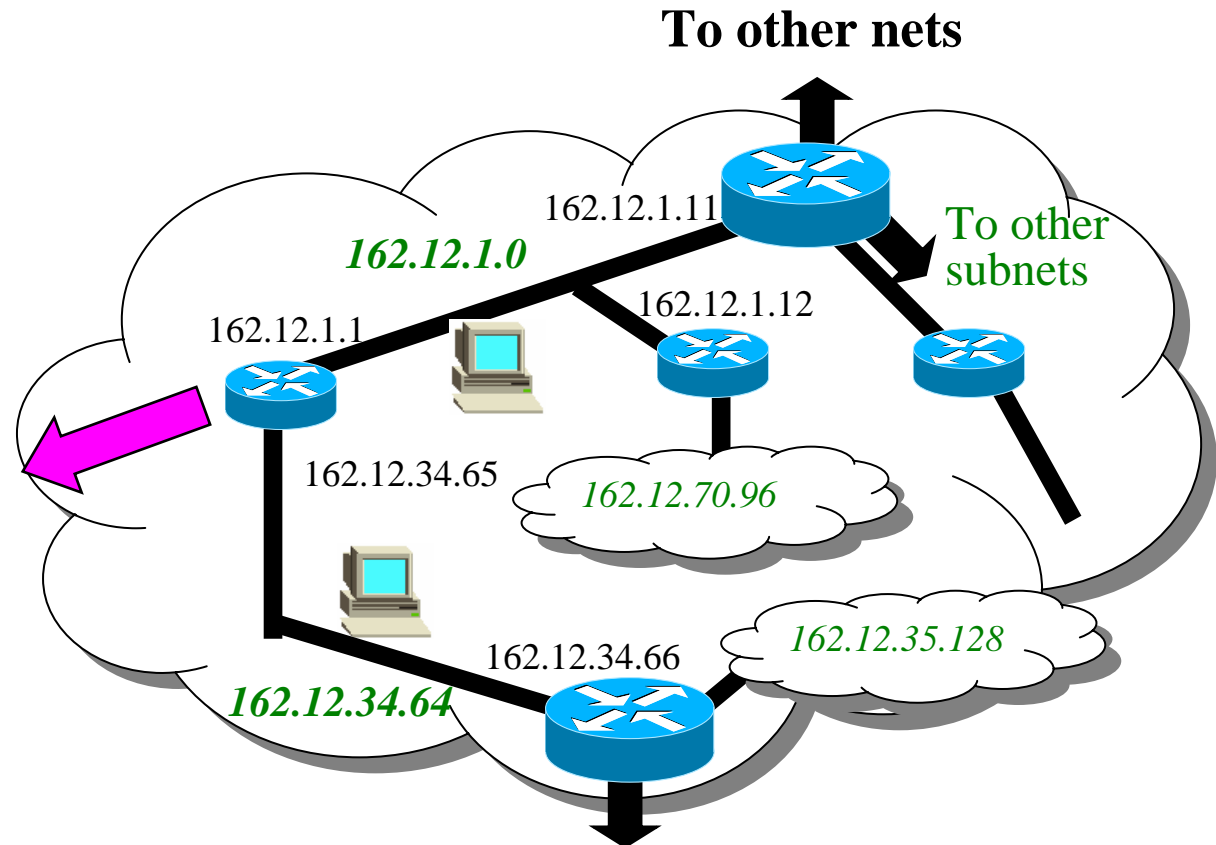
⇒ All necessary information included in Ipaddr

## → Subnet routing

⇒ Specific subnet mask (set by admin) required

Net = 162.12.0.0; subnet mask 255.255.255.224

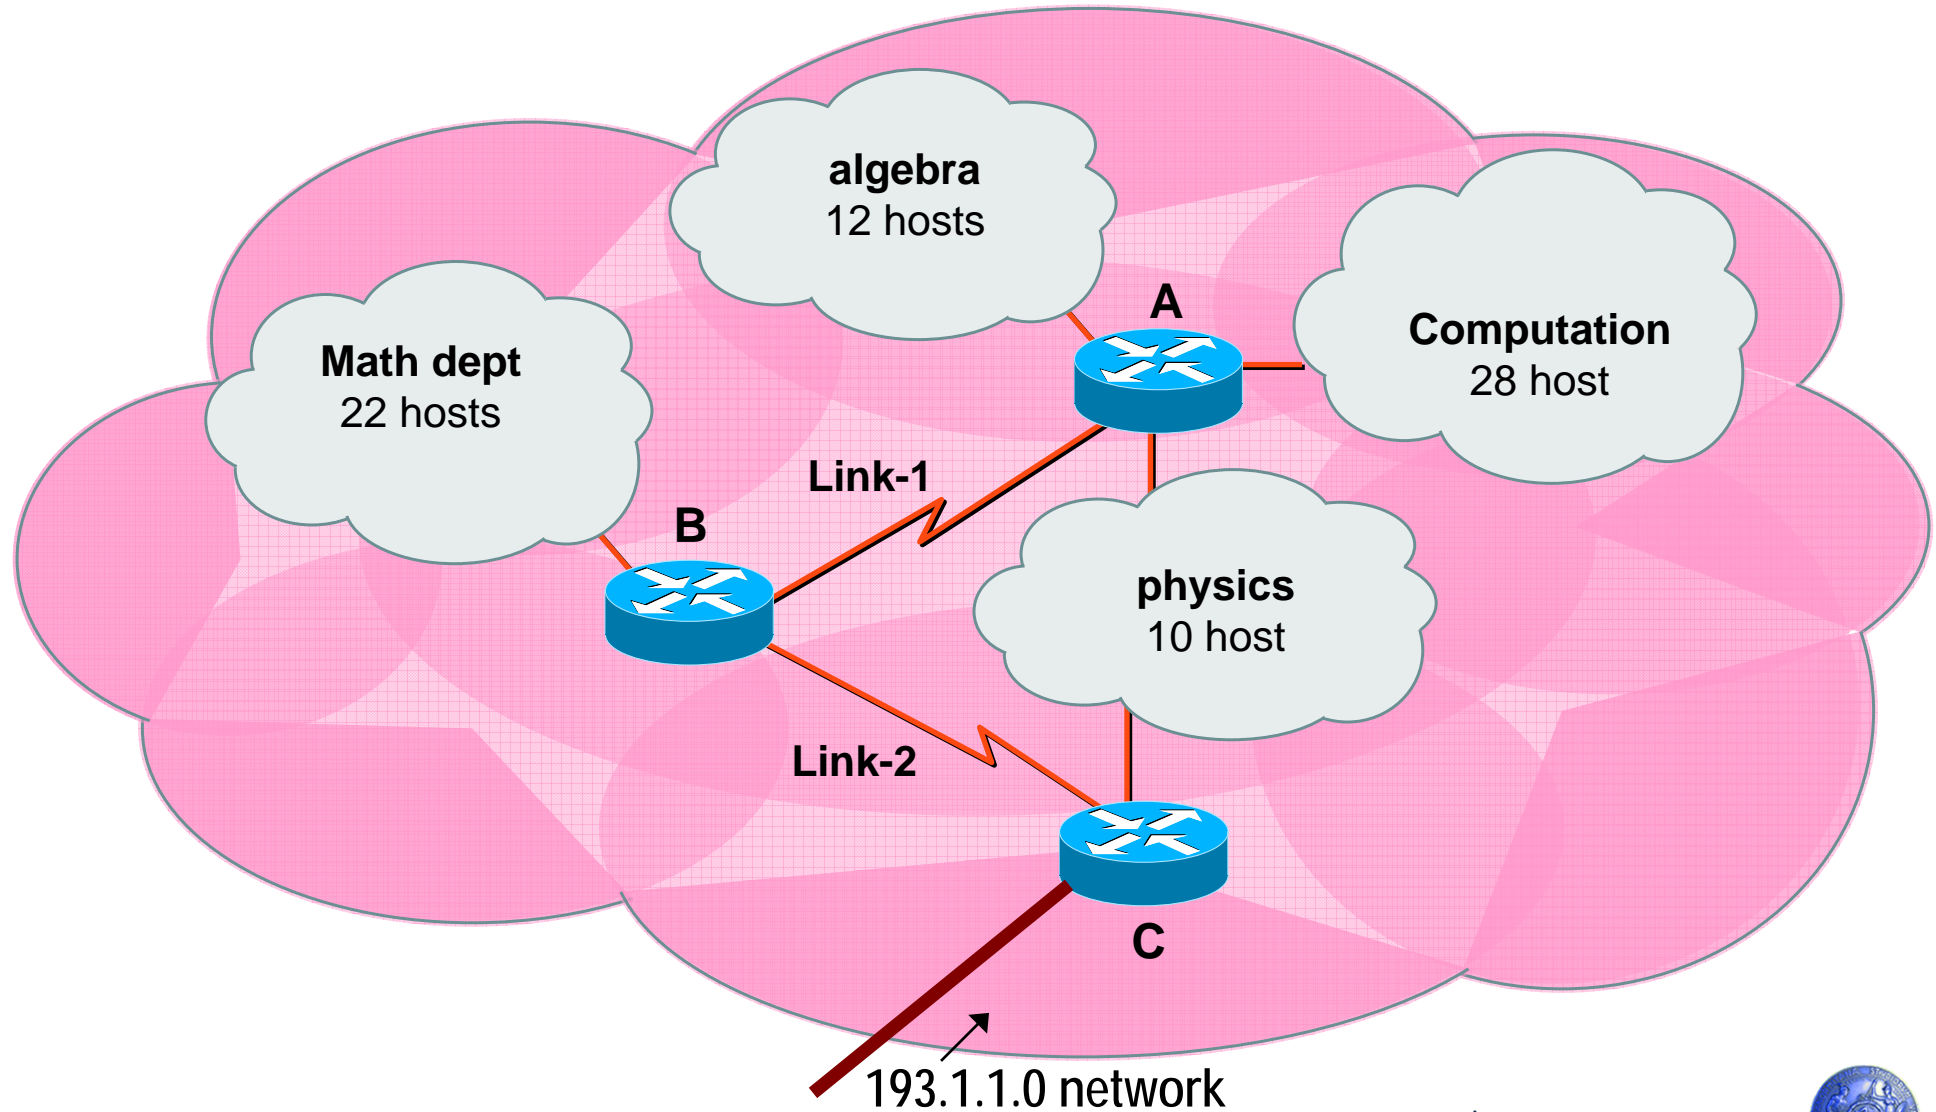
Routing Table	
Subnet mask: 255.255.255.224	
dest	Next Hop
162.12.1.0	Direct fwd
162.12.34.64	Direct fwd
162.12.35.128	162.12.34.66
162.12.70.96	162.12.1.12
131.175.0.0	162.12.34.66
131.176.0.0	162.12.34.66
default	162.12.1.11



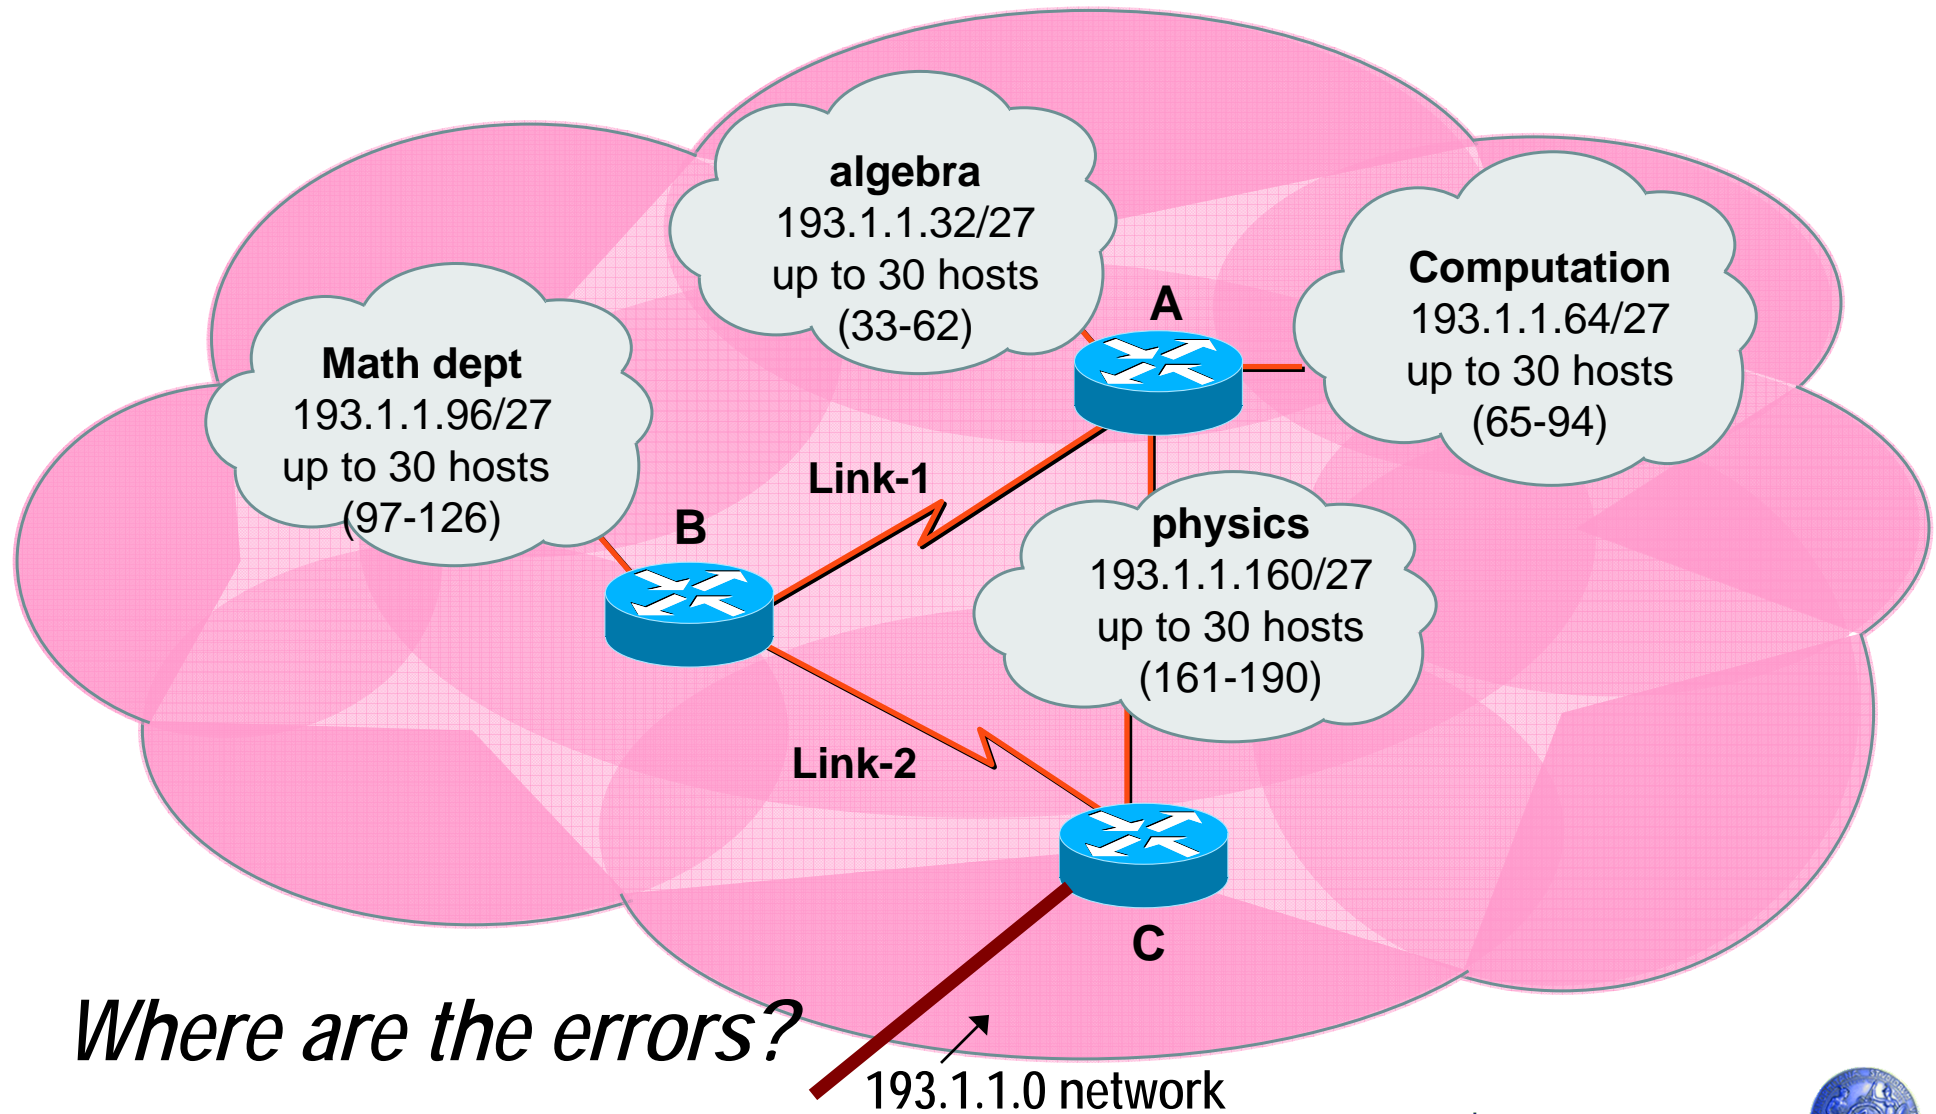
*May be quite a complex  
Routing table...  
VLSM will help (later)*

To  
131.175.0.0  
131.176.0.0

# Subnetting Example (problem)



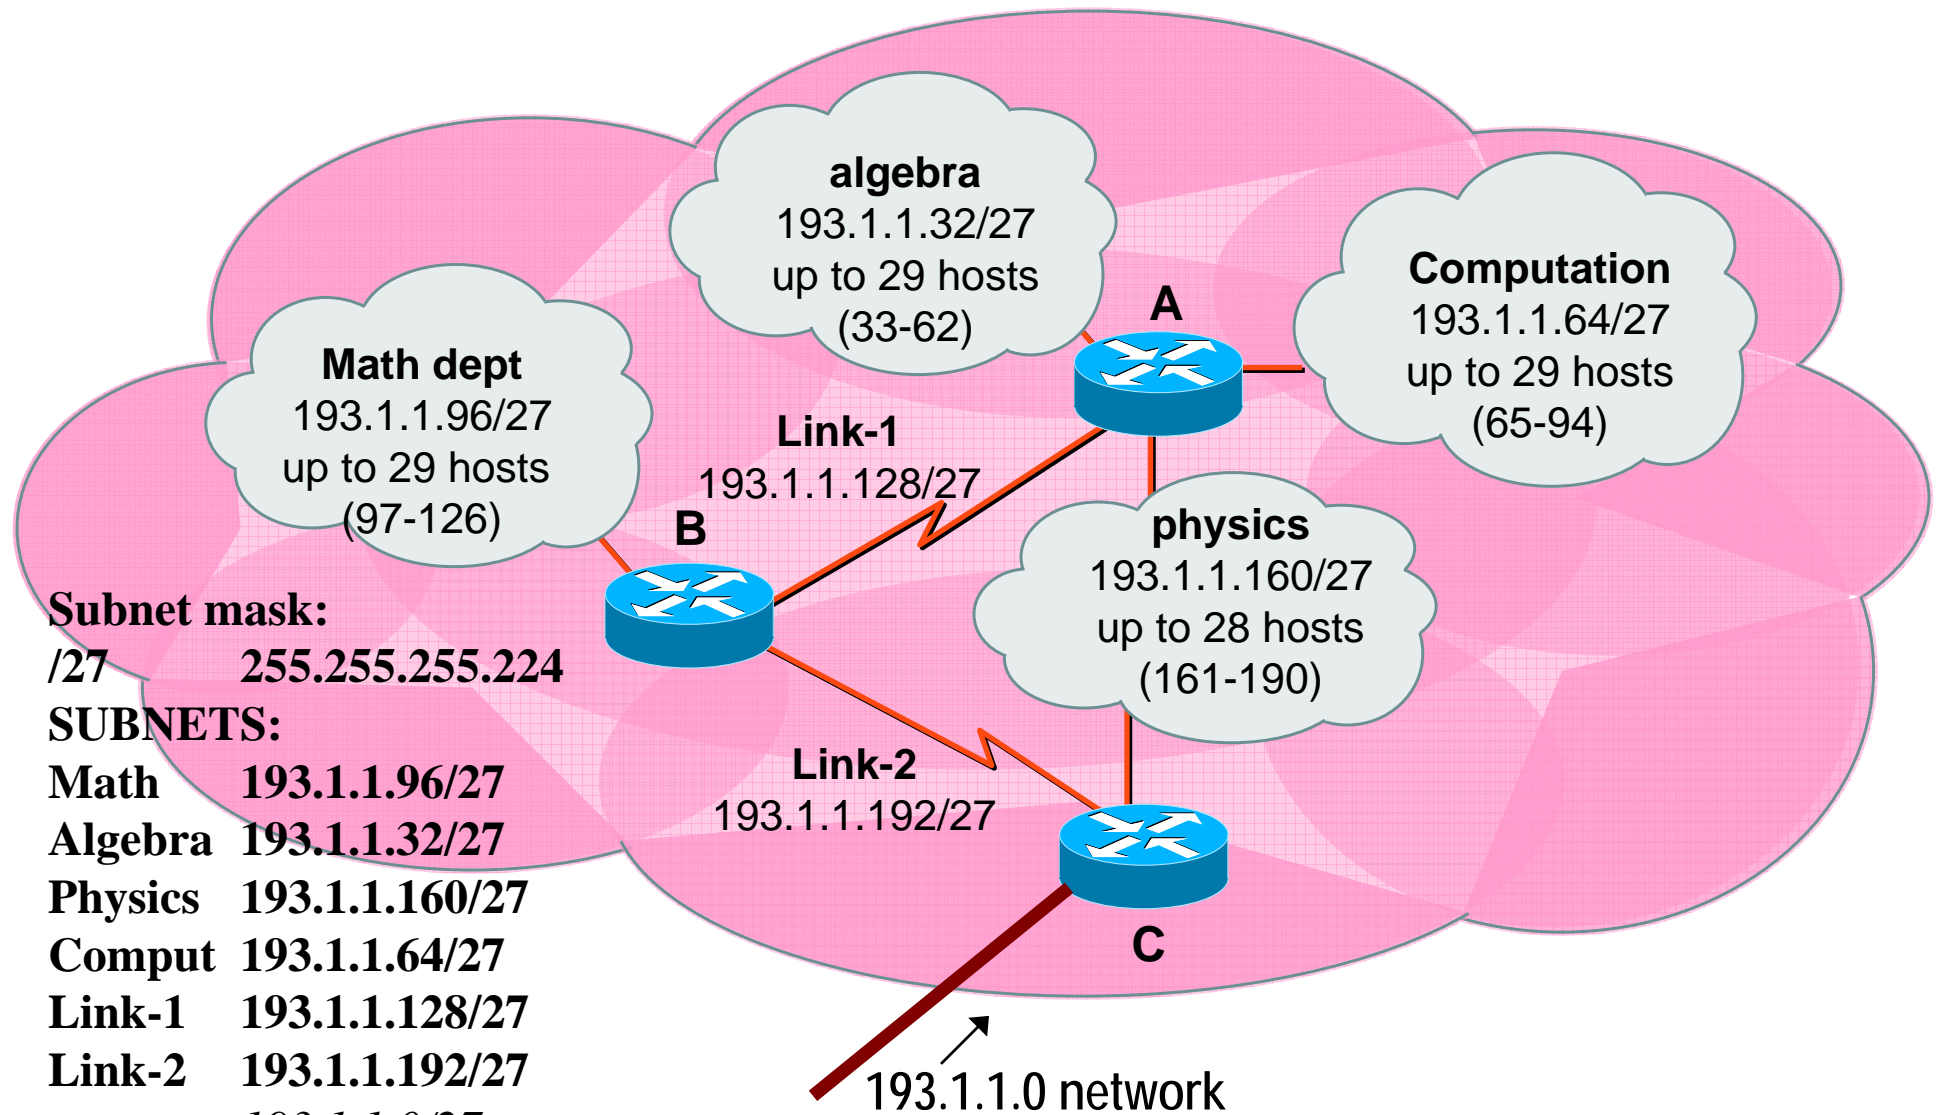
# Subnetting Example (solution?)



*Where are the errors?*



# Subnetting Example (solution!)



**Subnet mask:**

/27      255.255.255.224

**SUBNETS:**

**Math**      193.1.1.96/27

**Algebra** 193.1.1.32/27

**Physics** 193.1.1.160/27

**Comput** 193.1.1.64/27

**Link-1** 193.1.1.128/27

**Link-2** 193.1.1.192/27

---      193.1.1.0/27

---      193.1.1.224/27

**VLSM**  
**Variable Length Subnet Mask**  
**RFC 1009 (1987)**

# Variable Length Subnet Mask

**→ allows more than one subnet mask in the same network**

⇒ A) more efficient use of organization's IP address space

→ Subnets may significantly vary in relative size (computer room = 200 hosts, secretary = 4 hosts...)


→ consider a 4 host network with mask 255.255.255.0: wastes 250 IP addresses!

⇒ B) allows route aggregation, thus reducing routing information needed

**→ Needs further support by routing protocol**

⇒ e.g. RIP1 doesn't support VLSM

# A typical problem

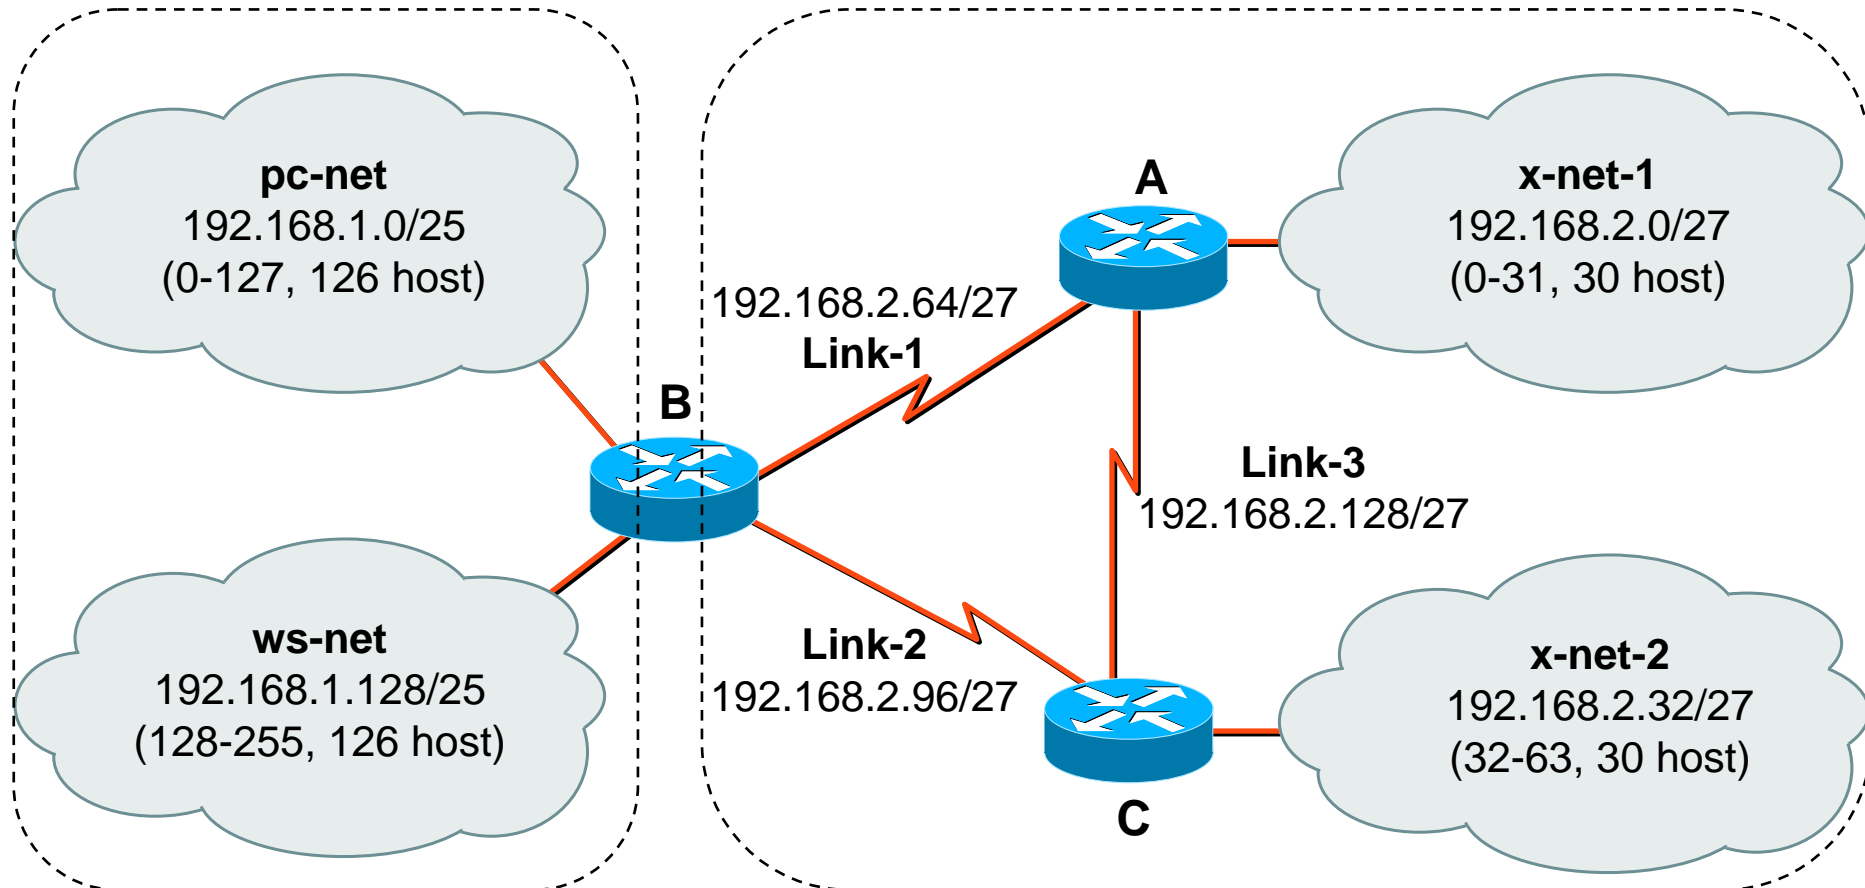


100+20+20+10 = 150 total hosts: 1 class C enough (including growth projections).  
7 subnets (4 LANS + 3 point to point links): 3 bit subnet ID (= up to 8 subnets)  
BUT then max 30 host per subnet: no way to accommodate pc-net!!



# Solution without VLSM

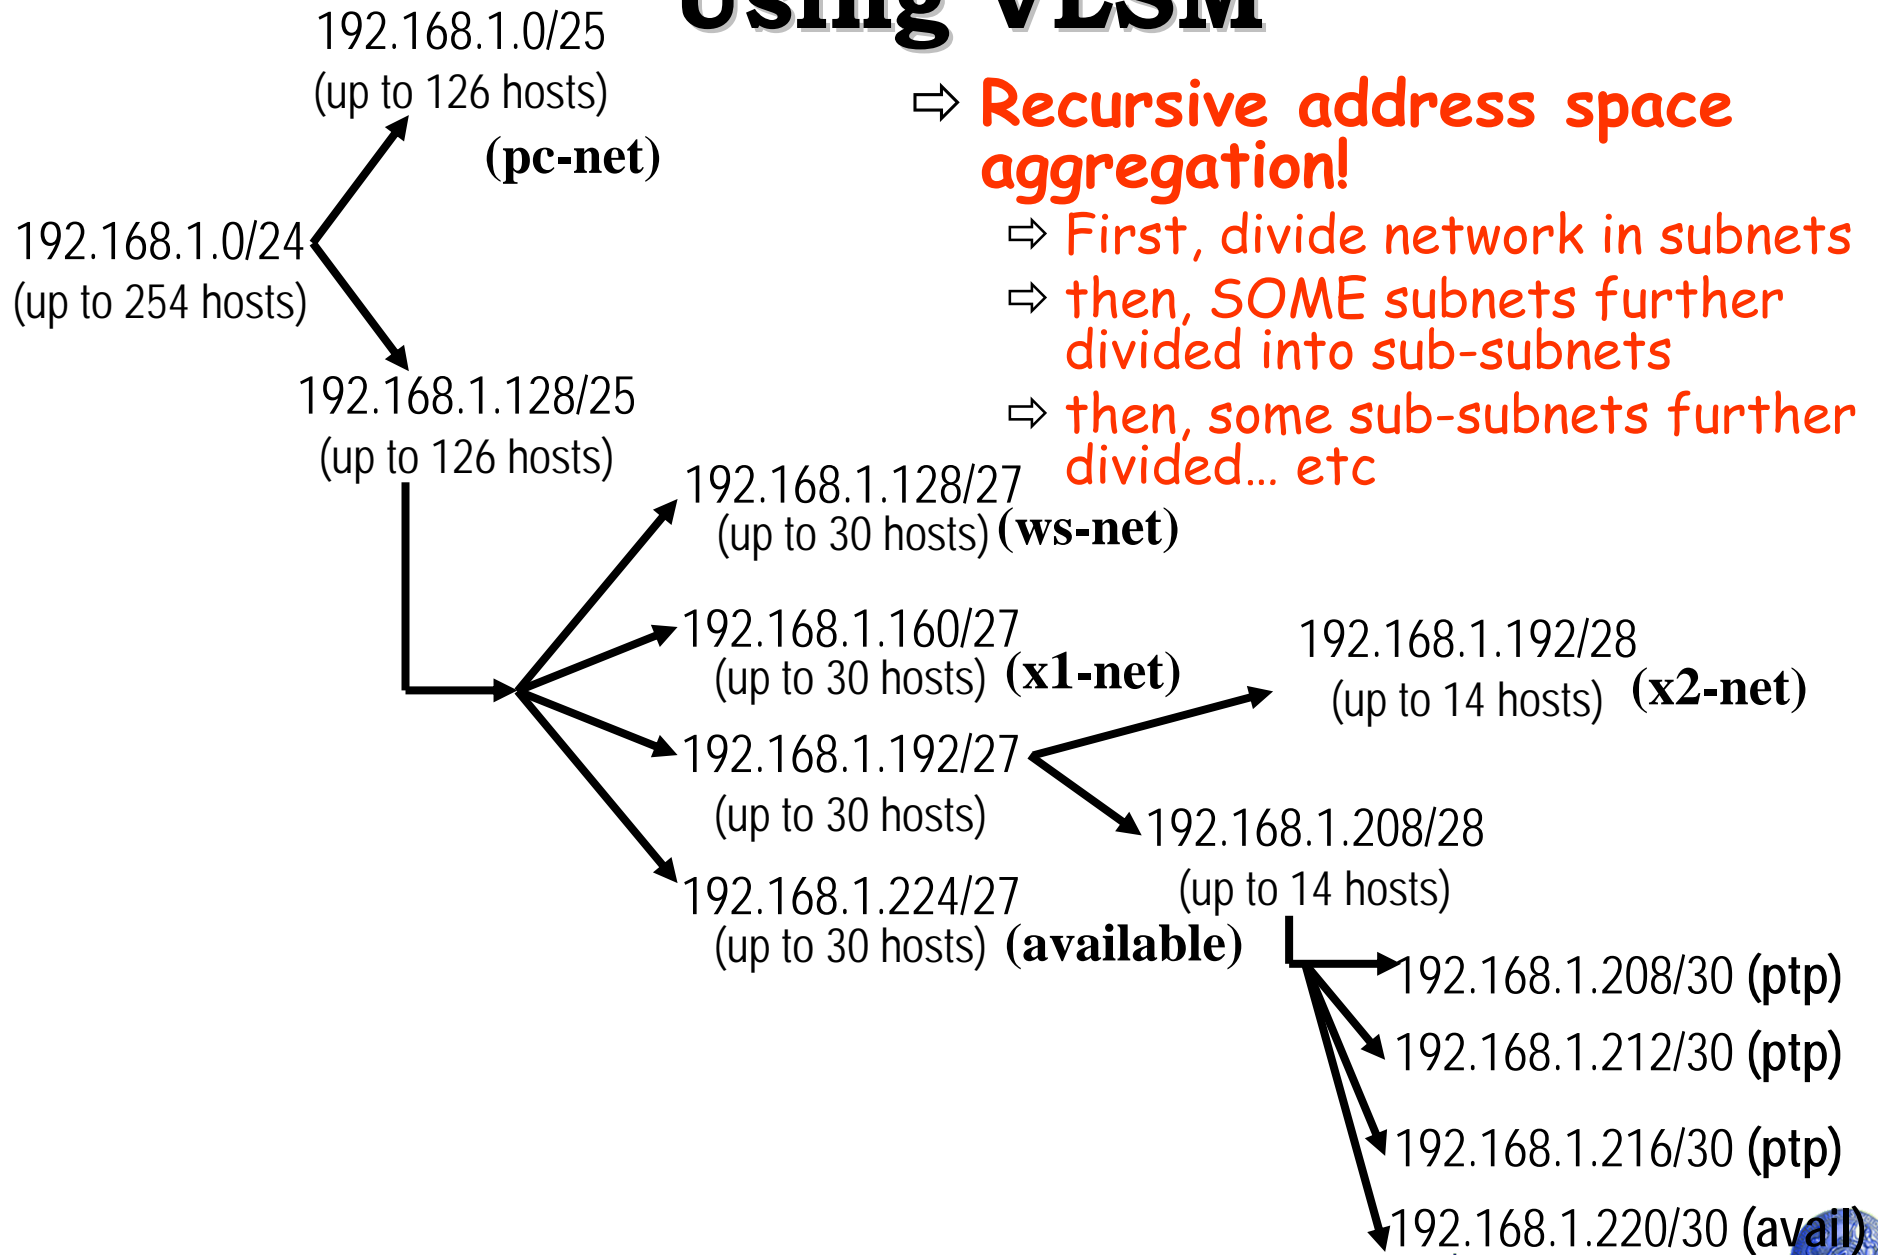
## need 2 class C address!



**192.168.1.0**  
mask 255.255.255.128

**192.168.2.0**  
mask 255.255.255.224

# Using VLSM



⇒ **Recursive address space aggregation!**

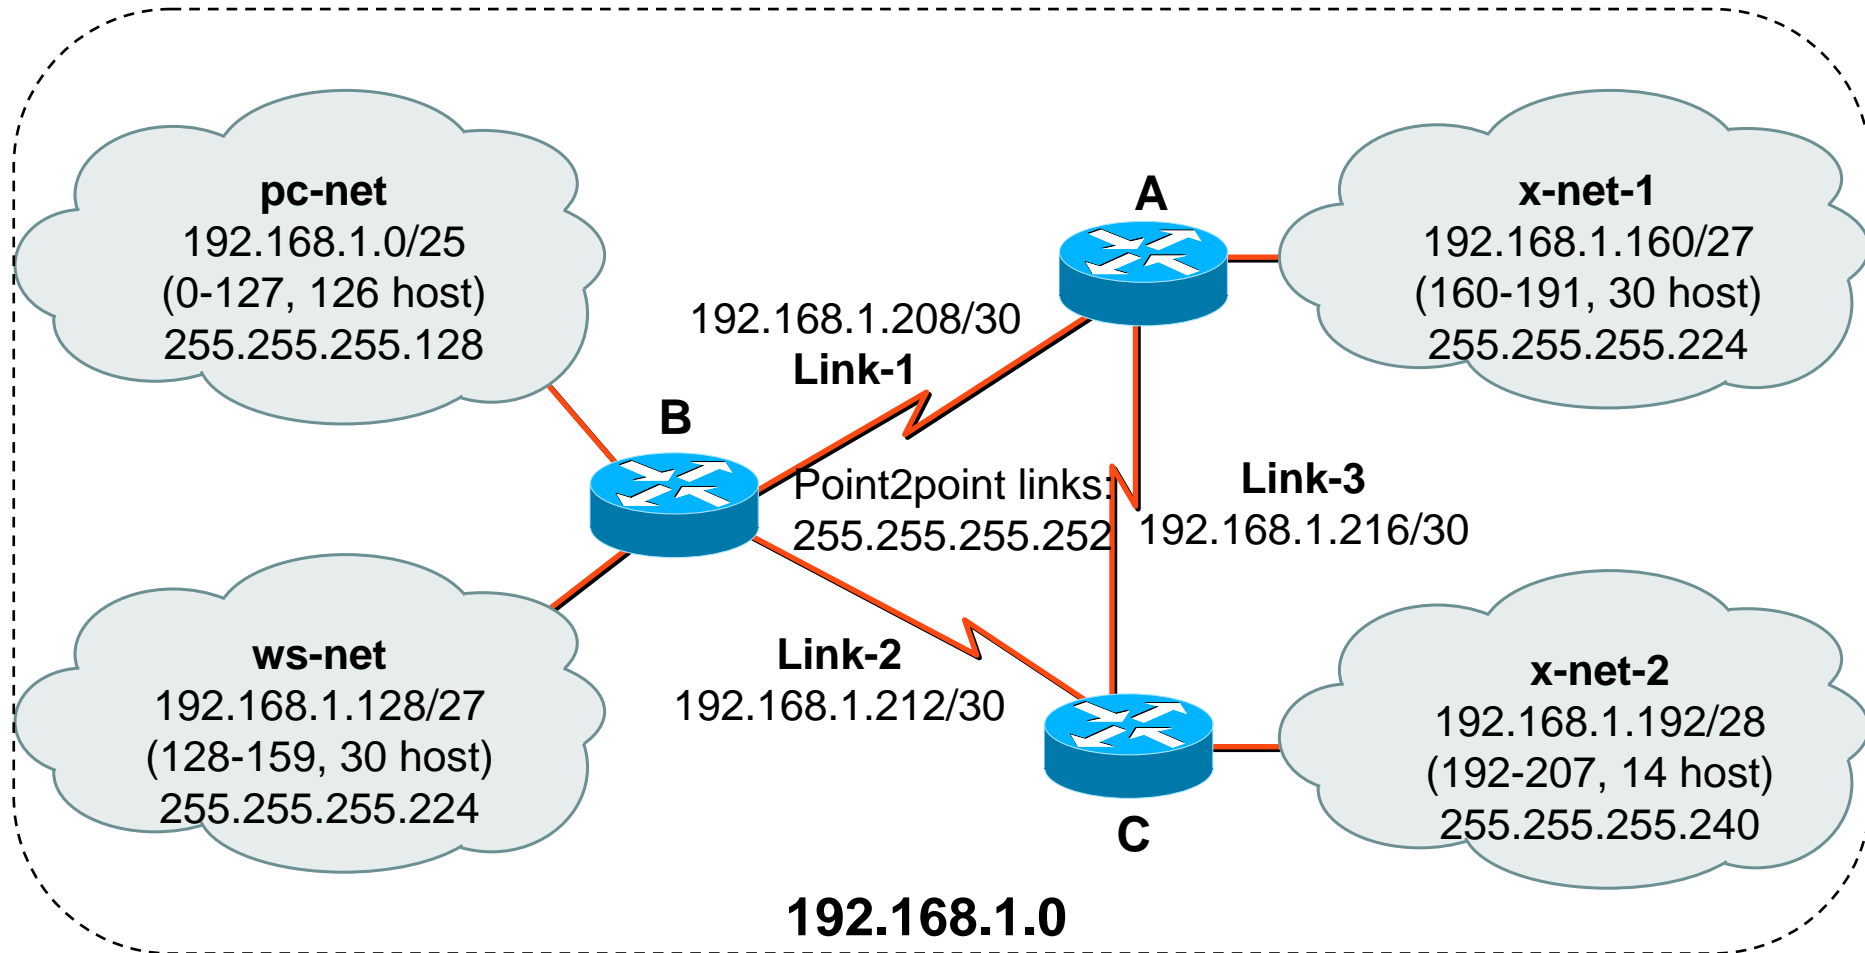
⇒ **First, divide network in subnets**

⇒ **then, SOME subnets further divided into sub-subnets**

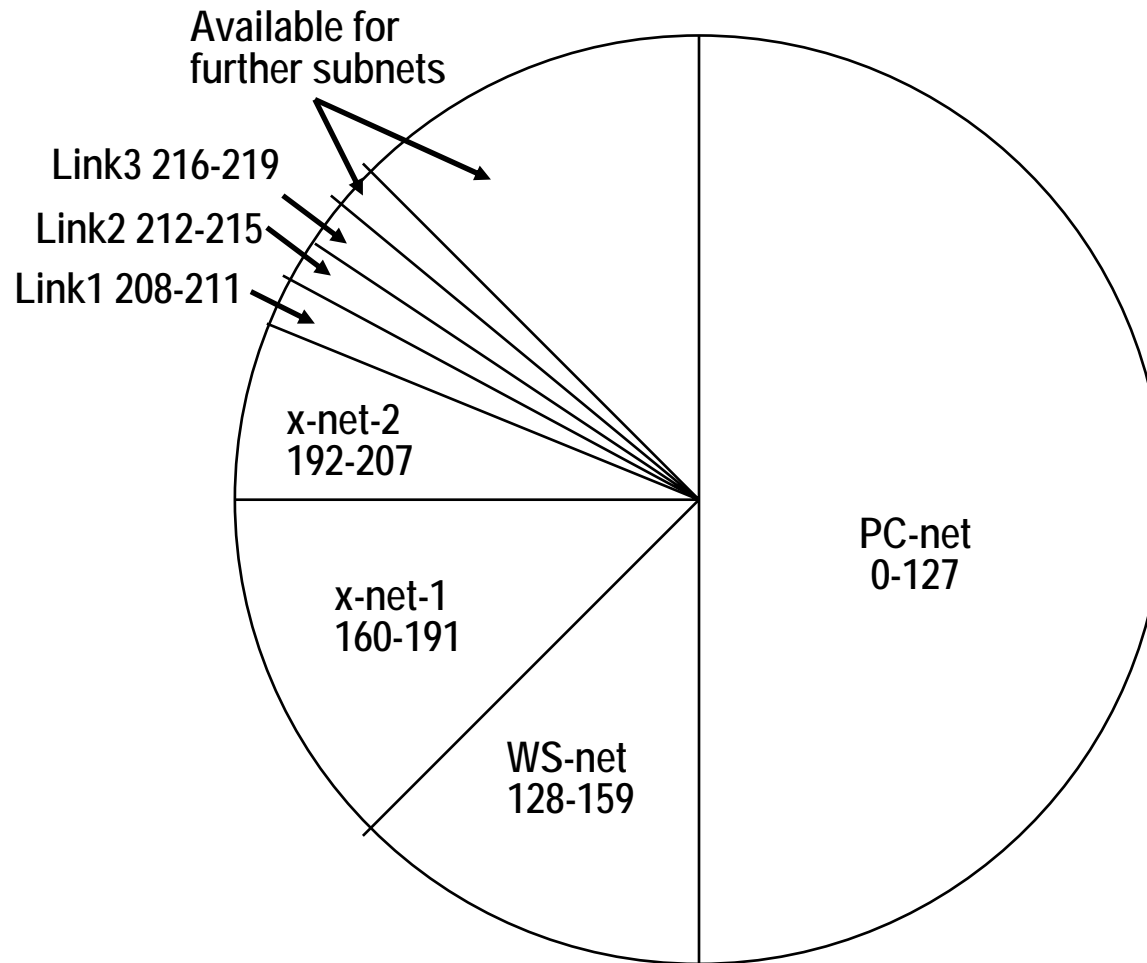
⇒ **then, some sub-subnets further divided... etc**

# Final solution with VLSM

## 1 C address is enough



# address pie for our sol.






# Requirements for VLSM support (1)

- Routing tables: need to specify extended network prefix information (subnet mask) per each entry
- Routing protocol: must carry extended network prefix information with each route advertisement

...	...	...
net	mask	route
...	...	...




New route advertise + *mask (or prefix len)*:

131.175.192.0      10000011.10101111.11000000.00000000  
255.255.240.0      11111111.11111111.11110000.00000000  
prefix /20

Without this feature: manually compiled tables (!!! Human error!!!)

*VLSM bottomline: need to use more complex routing protocols  
(e.g. OSPF) even for small org*

# Routing tables for previous example



Router C table

192.168.1.128	/27	192.168.1.213
192.168.1.0	/25	192.168.1.213
192.168.1.208	/30	192.168.1.213
192.168.1.192	/28	Direct fwd
192.168.1.192	/28	Direct fwd
192.168.1.212	/30	Direct fwd
192.168.1.216	/30	Direct fwd

192.168.1.0 network

# VLSM engineering

→ VLSM is a hierarchical subnet address assignment

⇒ BUT does not necessarily implies, by itself, a hierarchical routing!!

→ Effective designs combine:

⇒ address space reduction

⇒ with topologically significant address assignment

→ Substantial reduction of routing table sizes

→ Multiple route aggregation

# VLSM engineering

→ VLSM is a hierarchical subnet address assignment

⇒ BUT does not necessarily implies, by itself, a hierarchical routing!!

→ Effective designs combine:

⇒ address space reduction

⇒ with topologically significant address assignment

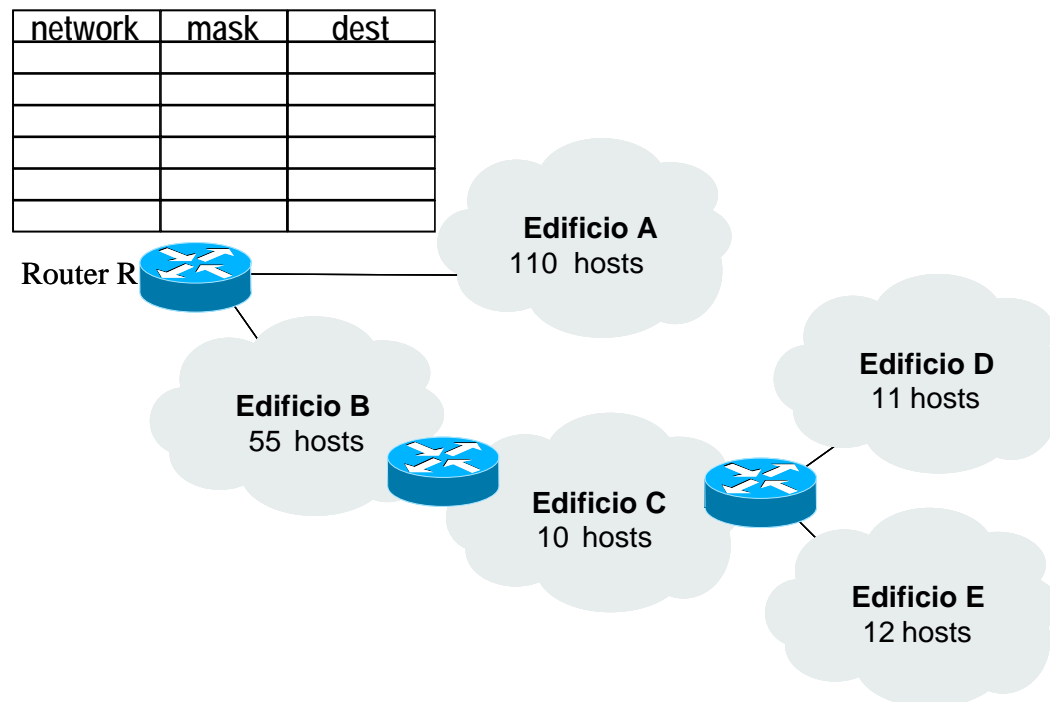
→ Substantial reduction of routing table sizes

→ Multiple route aggregation

# Complete example 1

Acquistando uno spazio di indirizzi il più piccolo possibile, da un provider che gestisce lo spazio 64.2.0.0 /16,

- Si divida in sottoreti la rete illustrata in figura in modo da soddisfare alle capacità richieste
- Si assegnino indirizzi IP alle interfacce dei router
- Si mostri la routing table del router R

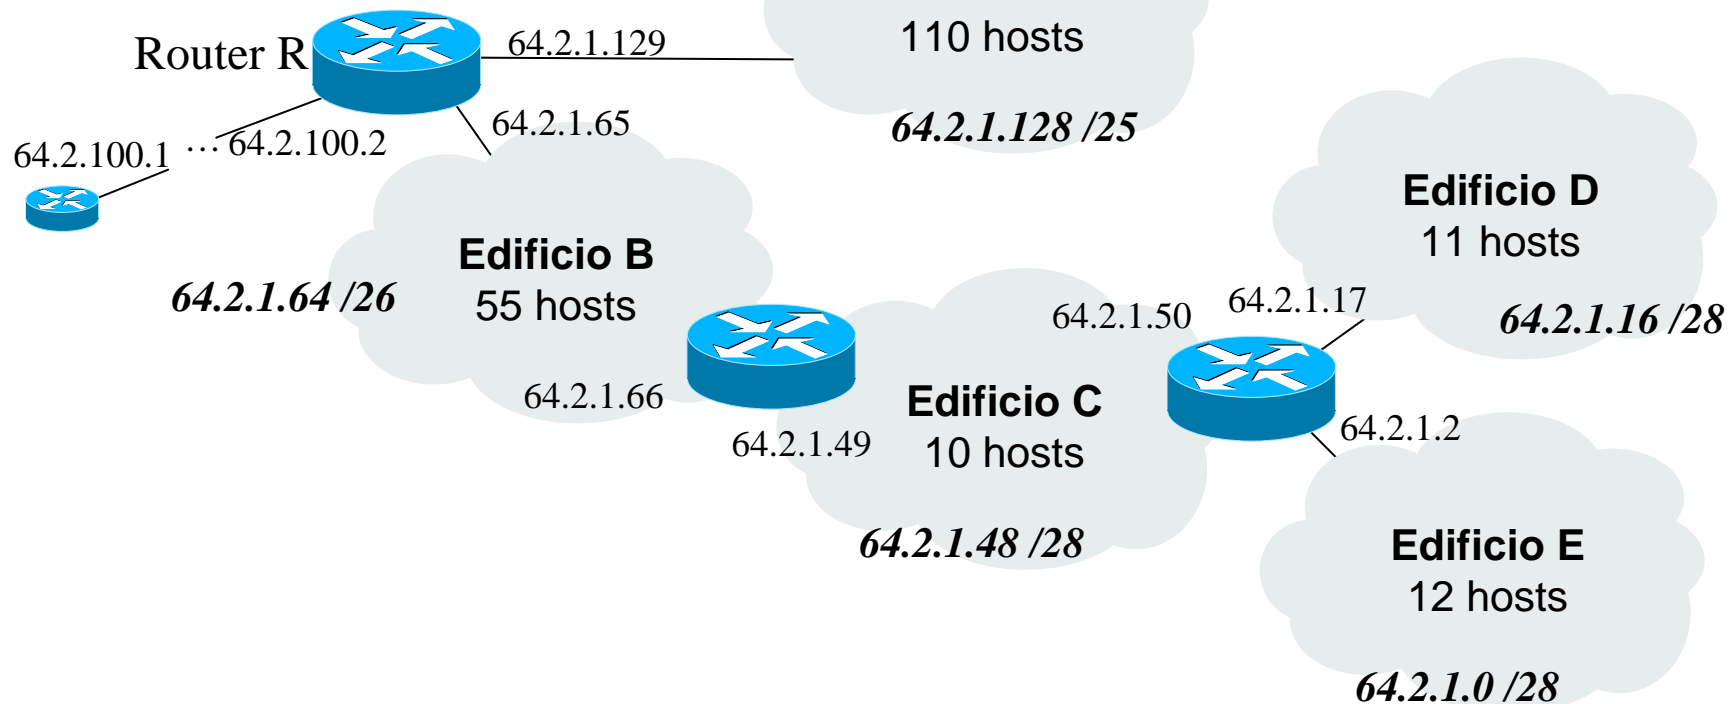


# Solution – no route aggregation

È sufficiente uno /24, es: 64.2.1.0 /24

Una soluzione possibile, con massima aggregazione dei route, è illustrata in figura (si assume che il routing esterno alla rete avvenga tramite l'interfaccia remota 64.2.100.1)

network	mask	next hop	interface
64.2.1.128	/25	64.2.1.129	64.2.1.129
64.2.1.64	/26	64.2.1.65	64.2.1.65
64.2.1.48	/28	64.2.1.66	64.2.1.65
64.2.1.0	/28	64.2.1.66	64.2.1.65
64.2.1.16	/28	64.2.1.66	64.2.1.65
0.0.0.0	/0	64.2.100.1	64.2.100.2

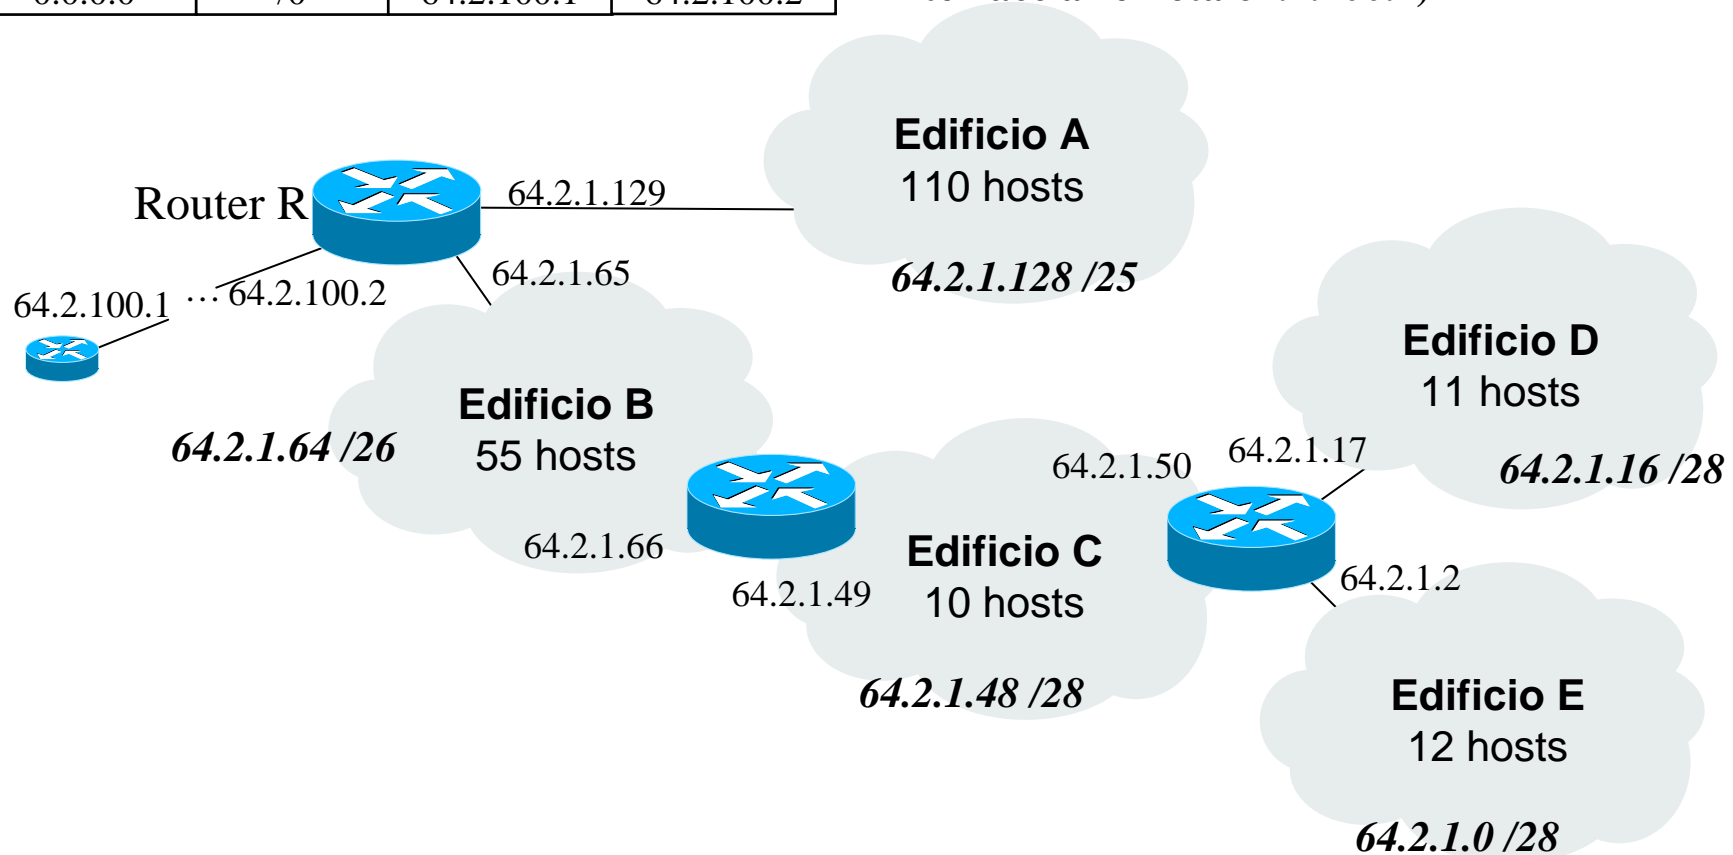


# Solution – final

network	mask	next hop	inteface
64.2.1.128	/25	64.2.1.129	64.2.1.129
64.2.1.64	/26	64.2.1.65	64.2.1.65
64.2.1.0	/26	64.2.1.66	64.2.1.65
0.0.0.0	/0	64.2.100.1	64.2.100.2

È sufficiente uno /24, es: 64.2.1.0 /24

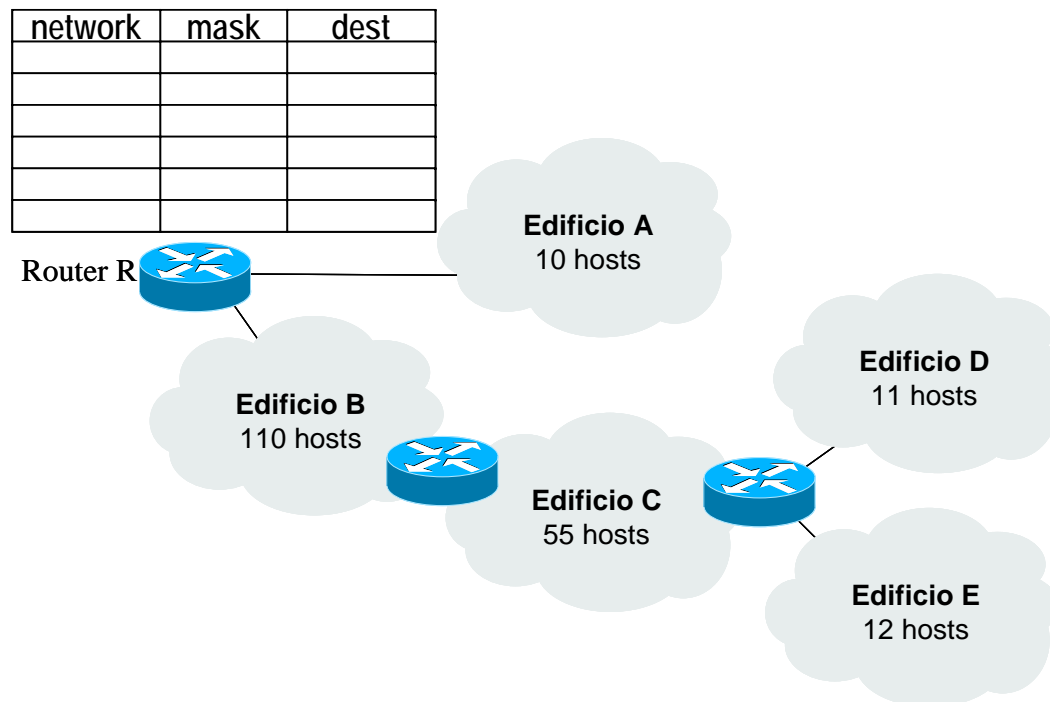
Una soluzione possibile, con massima aggregazione dei route, è illustrata in figura (si assume che il routing esterno alla rete avvenga tramite l'interfaccia remota 64.2.100.1)



# Complete example 2

Acquistando uno spazio di indirizzi il piu' piccolo possibile, da un provider che gestisce lo spazio 64.2.0.0 /16,

- Si subnetti la rete illustrata in figura in modo da soddisfare alle capacità richieste
- Si assegnino indirizzi IP alle interfacce dei router
- Si mostri la routing table del router R



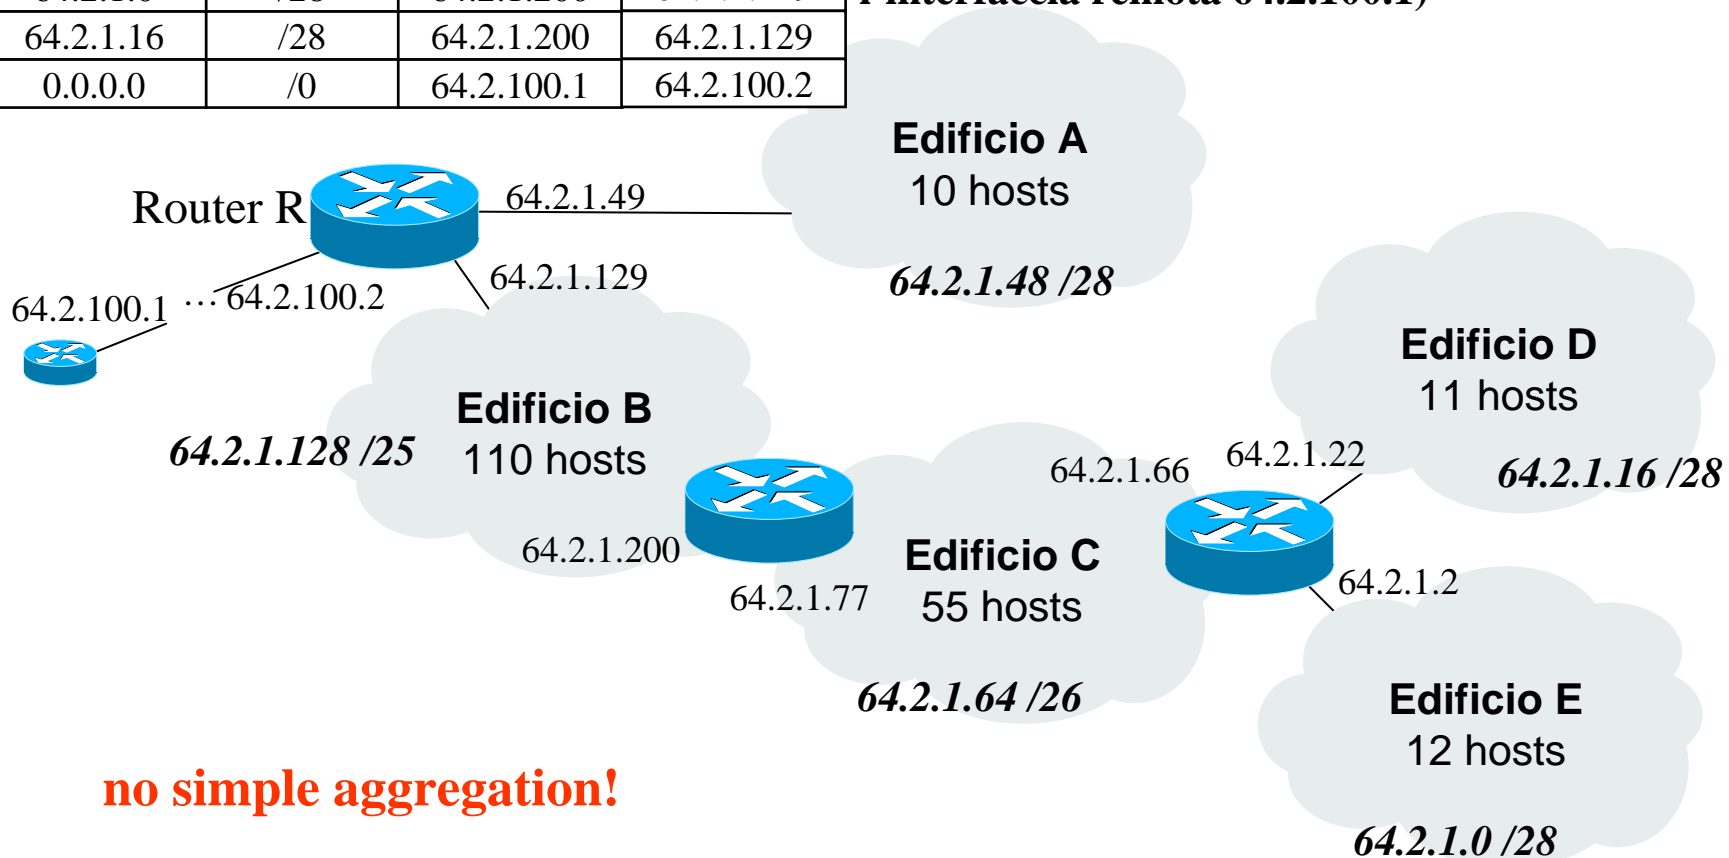


# Solution – no route aggregation

network	mask	next hop	interface
64.2.1.128	/25	64.2.1.129	64.2.1.129
64.2.1.64	/26	64.2.1.200	64.2.1.129
64.2.1.48	/28	64.2.1.49	64.2.1.49
64.2.1.0	/28	64.2.1.200	64.2.1.129
64.2.1.16	/28	64.2.1.200	64.2.1.129
0.0.0.0	/0	64.2.100.1	64.2.100.2


È sufficiente uno /24, es: 64.2.1.0 /24

Una soluzione possibile, con massima aggregazione dei route, è illustrata in figura (si assume che il routing esterno alla rete avvenga tramite l'interfaccia remota 64.2.100.1)



# Requirements for VLSM support (2)

## → “Longest Match” Forwarding Algorithm

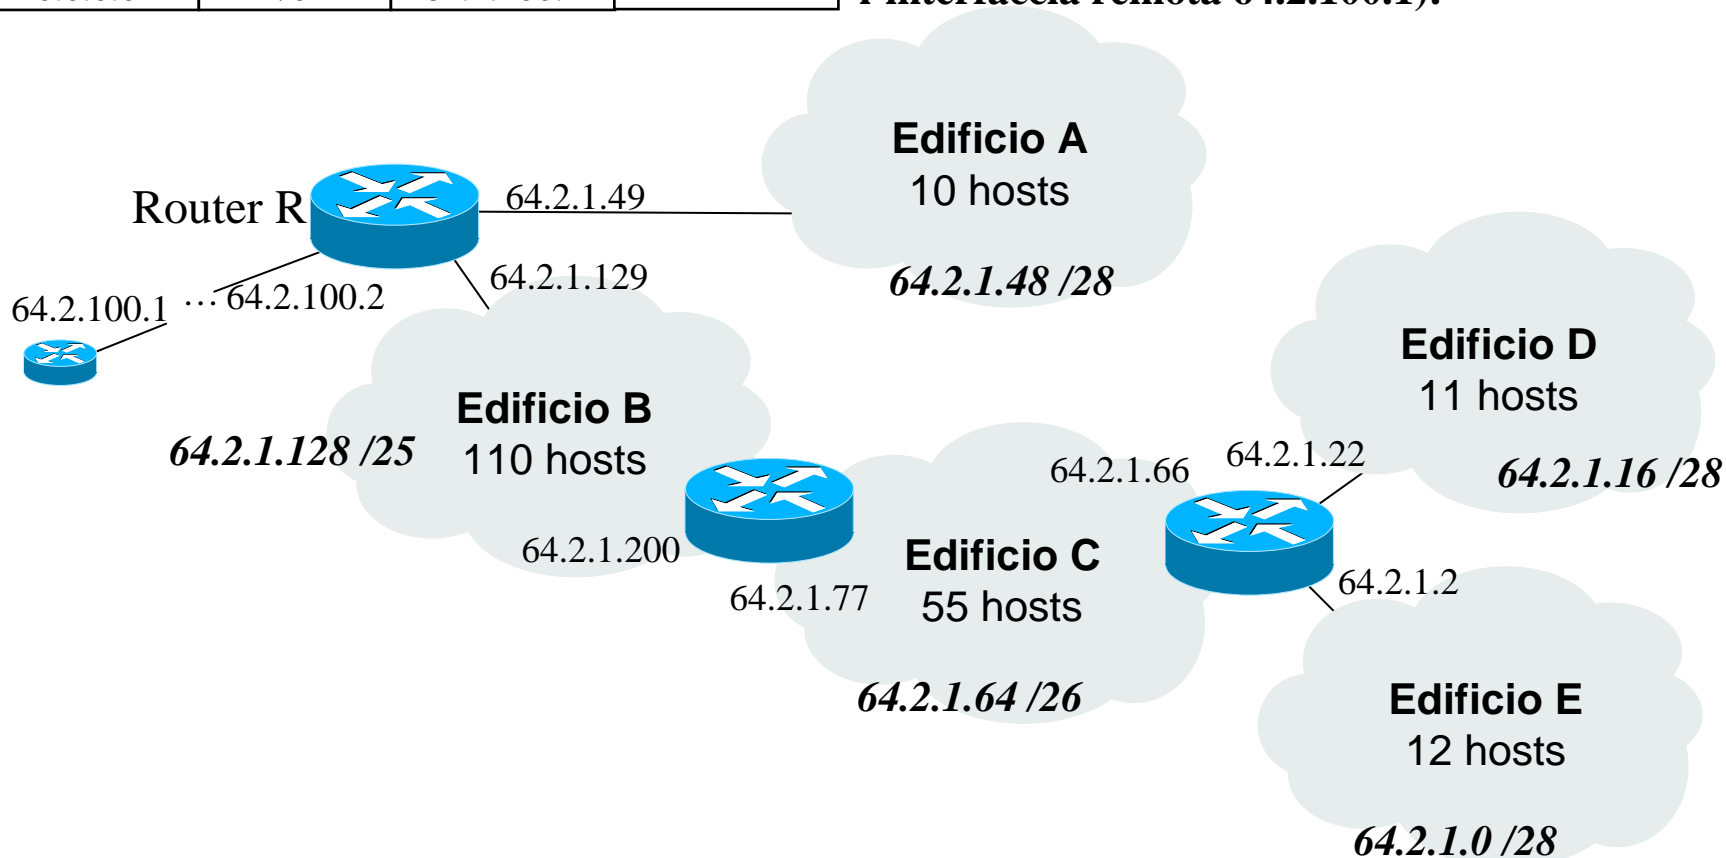


# Solution - final

network	mask	next hop	interface
64.2.1.128	/25	64.2.1.129	64.2.1.129
64.2.1.48	/28	64.2.1.49	64.2.1.49
64.2.1.0	/25	64.2.1.200	64.2.1.129
0.0.0.0	/0	64.2.100.1	64.2.100.2

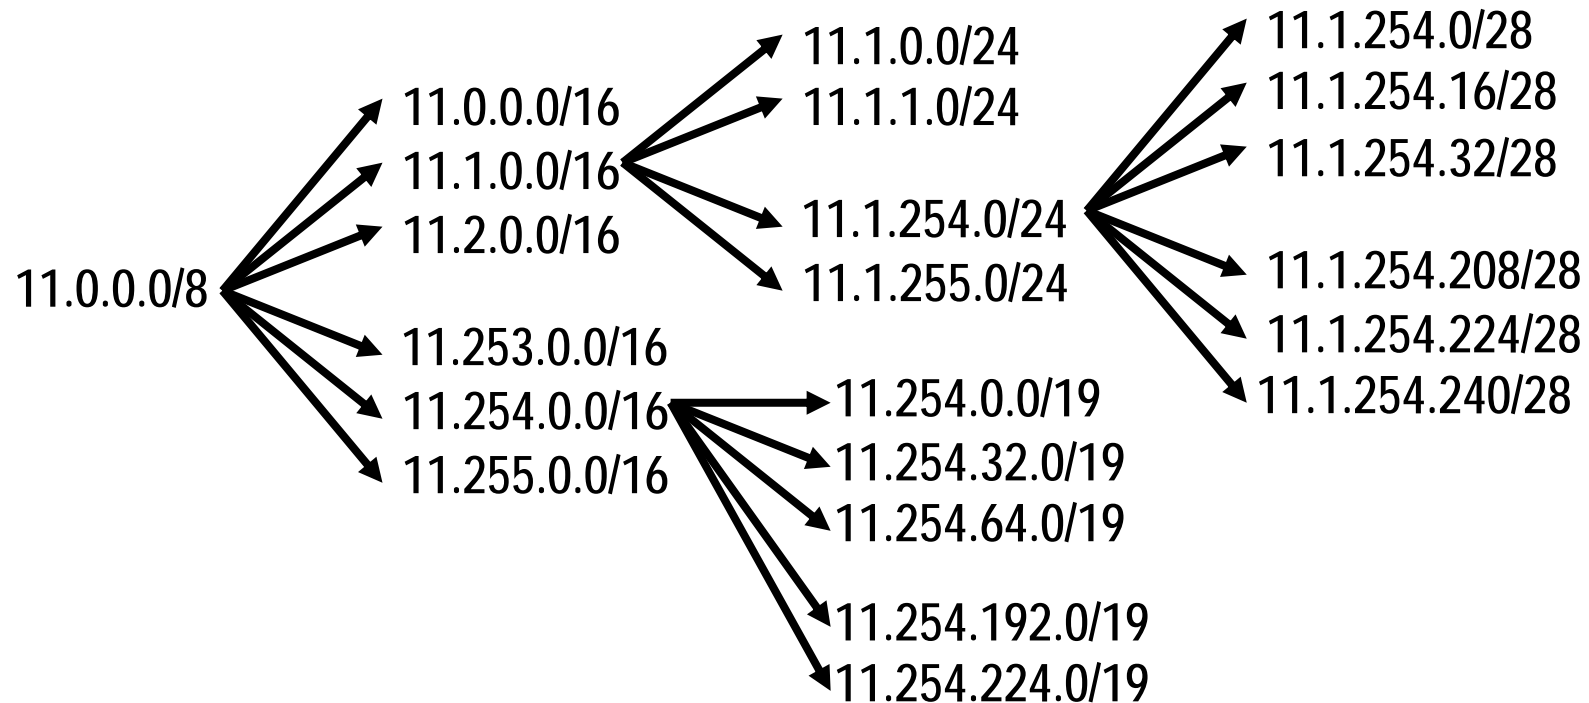
È sufficiente uno /24, es: 64.2.1.0 /24

Una soluzione possibile, con massima aggregazione dei route, è illustrata in figura (si assume che Il routing esterno alla rete avvenga tramite l'interfaccia remota 64.2.100.1):



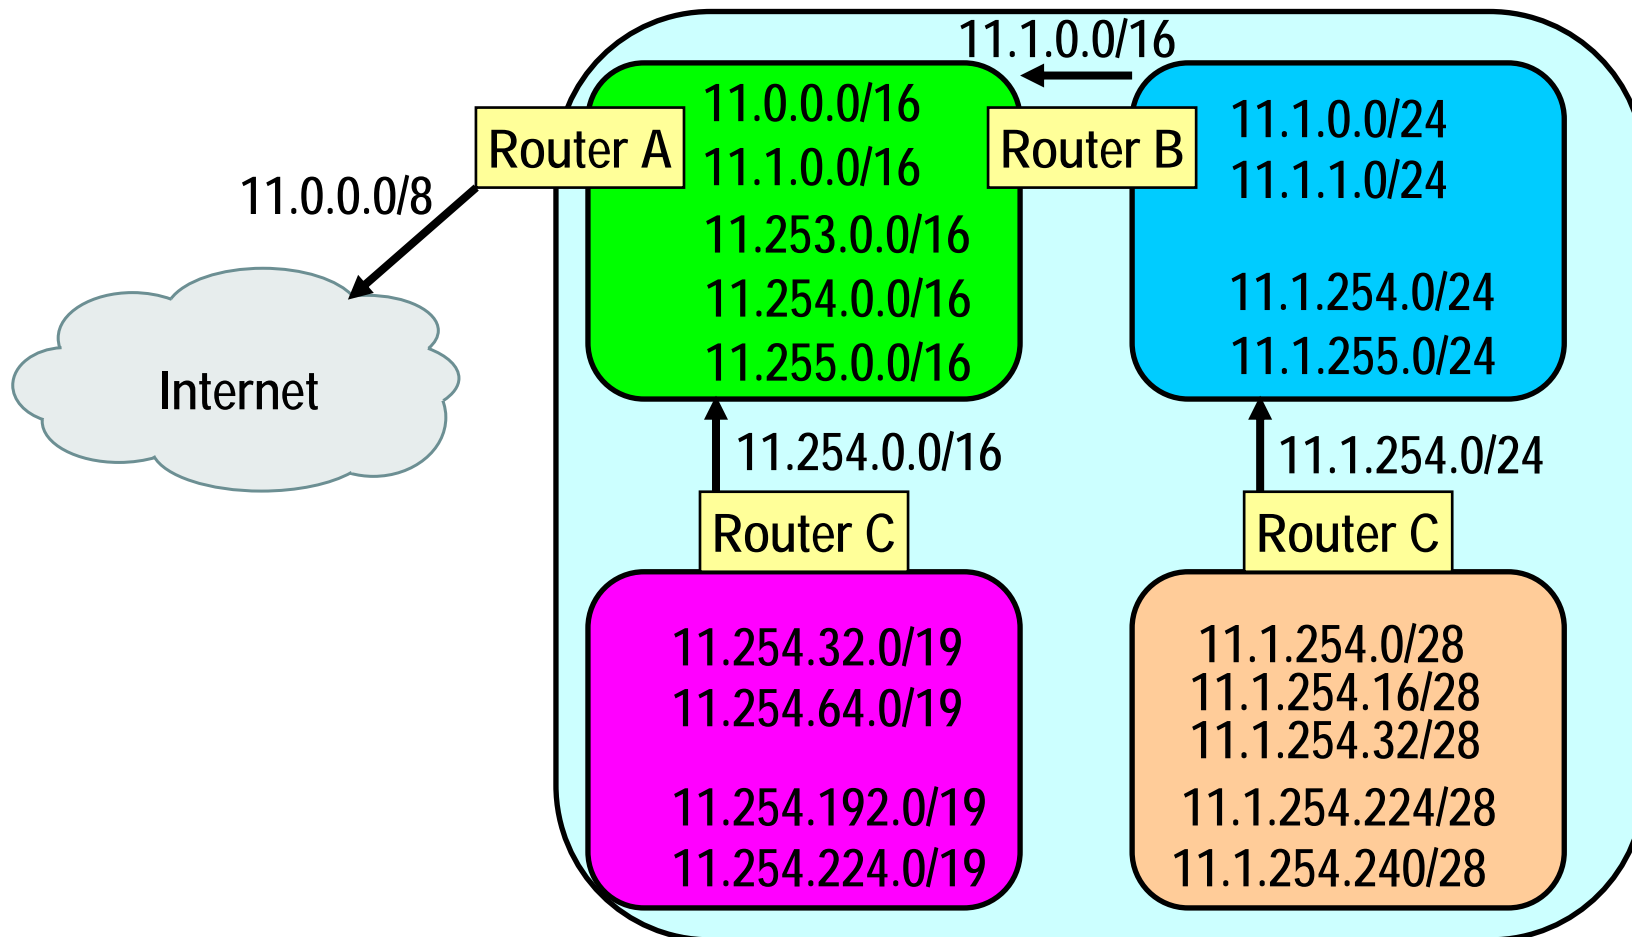
# Example: VLSM engineering

## VLSM subnetting of class A 11.0.0.0



# Route aggregation with VLSM

- VLSM allows to hide detailed structure of routing information for one subnet group from other routers - reducing routing table Size



**CIDR**  
**Classless Inter-Domain Routing**  
**RFC 1517 to 1520 (1993)**

# An historical perspective

## N x class C? Class B?

### → Class C addresses:

⇒ *Undersized (254 hosts)*

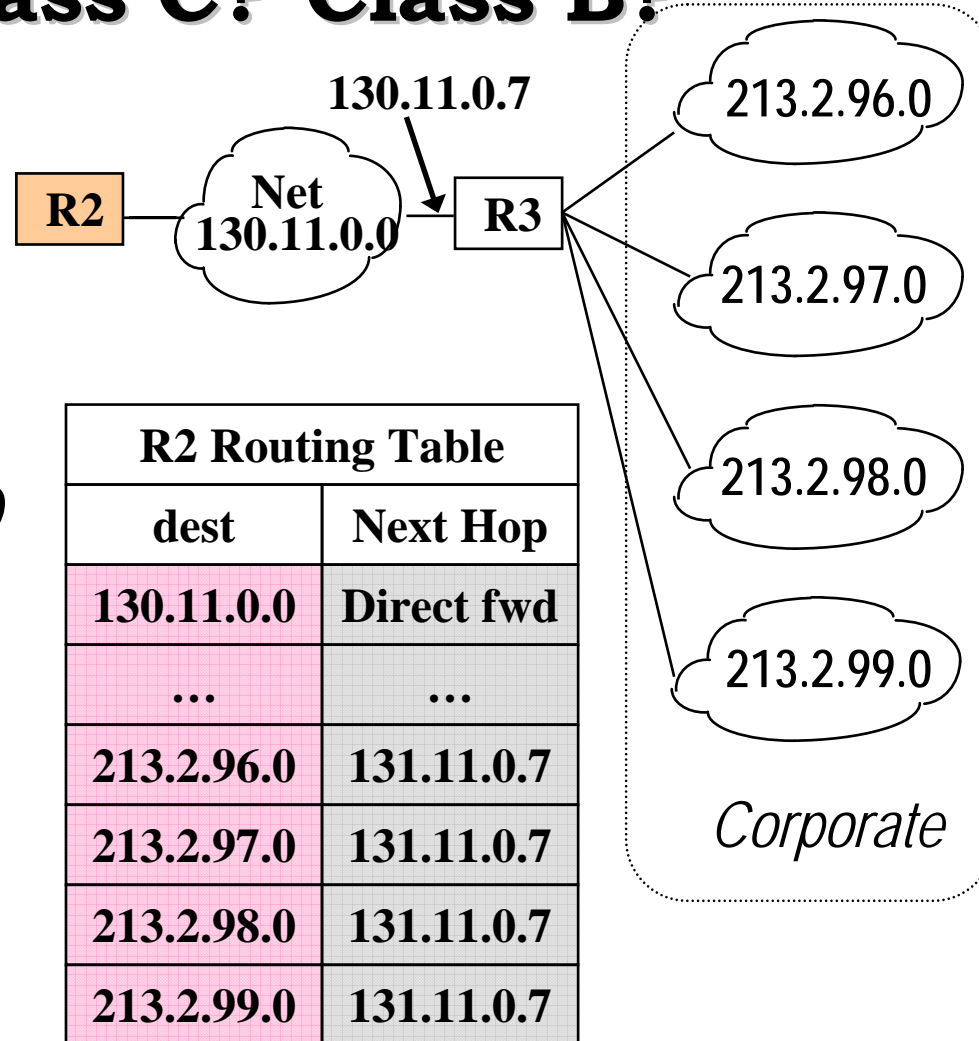
### → Class B addresses:

⇒ *Much more than enough (65534 hosts)*

### → N x class C:

⇒ *Unwise: exponential growth of routing tables*

→ **Result: Class B addresses were largely preferred**



# The 1992 Internet scenario

## → Near-term exhaustion of class B address space

- ⇒ In early years, Class B addresses given away!
- ⇒ Inefficient division into A, B, C classes
  - byte-word: unwise choice (class C too little, class B too big)
  - The aftermath: much better, e.g. C=10 bits, B=14 bits
- ⇒ Projections at the time: class B exhaustion by 1994/95




Figure 1: Assigned and Allocated Network Numbers

*Not a real problem: there are in principle 2M class C addresses!  
... what are we missing??*




# The problem

→ **Corporate has to build 4 physical networks (e.g. buildings)**

⇒ Example: networks up to 254 hosts

→ **Must “buy” 4 IP network addresses**

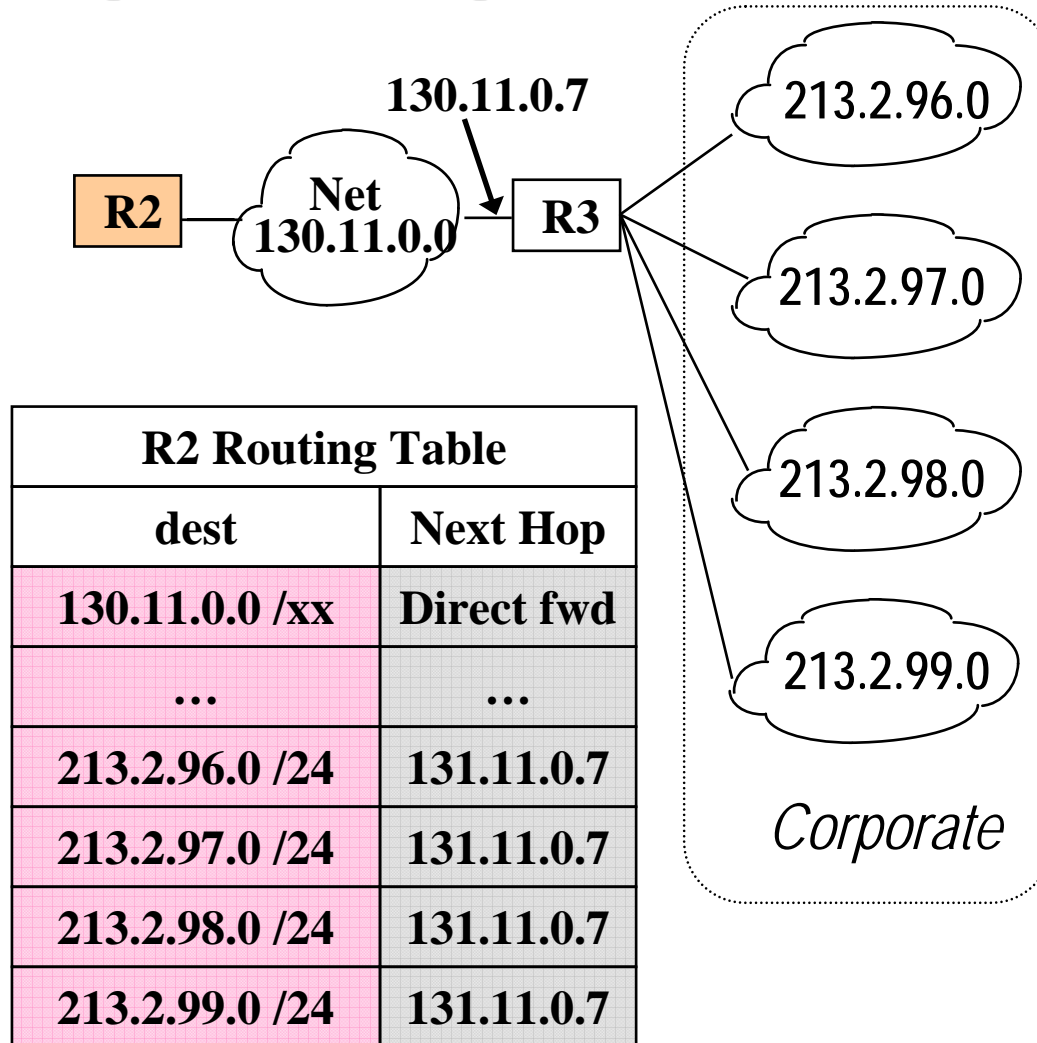
→ **Why this is bad?**



# Routing table growth

→ 4 x networks

⇒ *Unwise:*  
*exponential*  
*growth of*  
*routing tables*



# The 1992 Internet scenario


## → Exponential growth of routing tables

⇒ Multiple class C allocation dramatic for routing tables

→ necessary because of Class B exhaustion

→ 100.000 entries highly critical for performance

» *2M class C: WAY OUT of the capabilities of routing sw & hw*



## ⇒ Projections at the time

→ End 1990: 2190 routes; end 1992: 8500 routes;

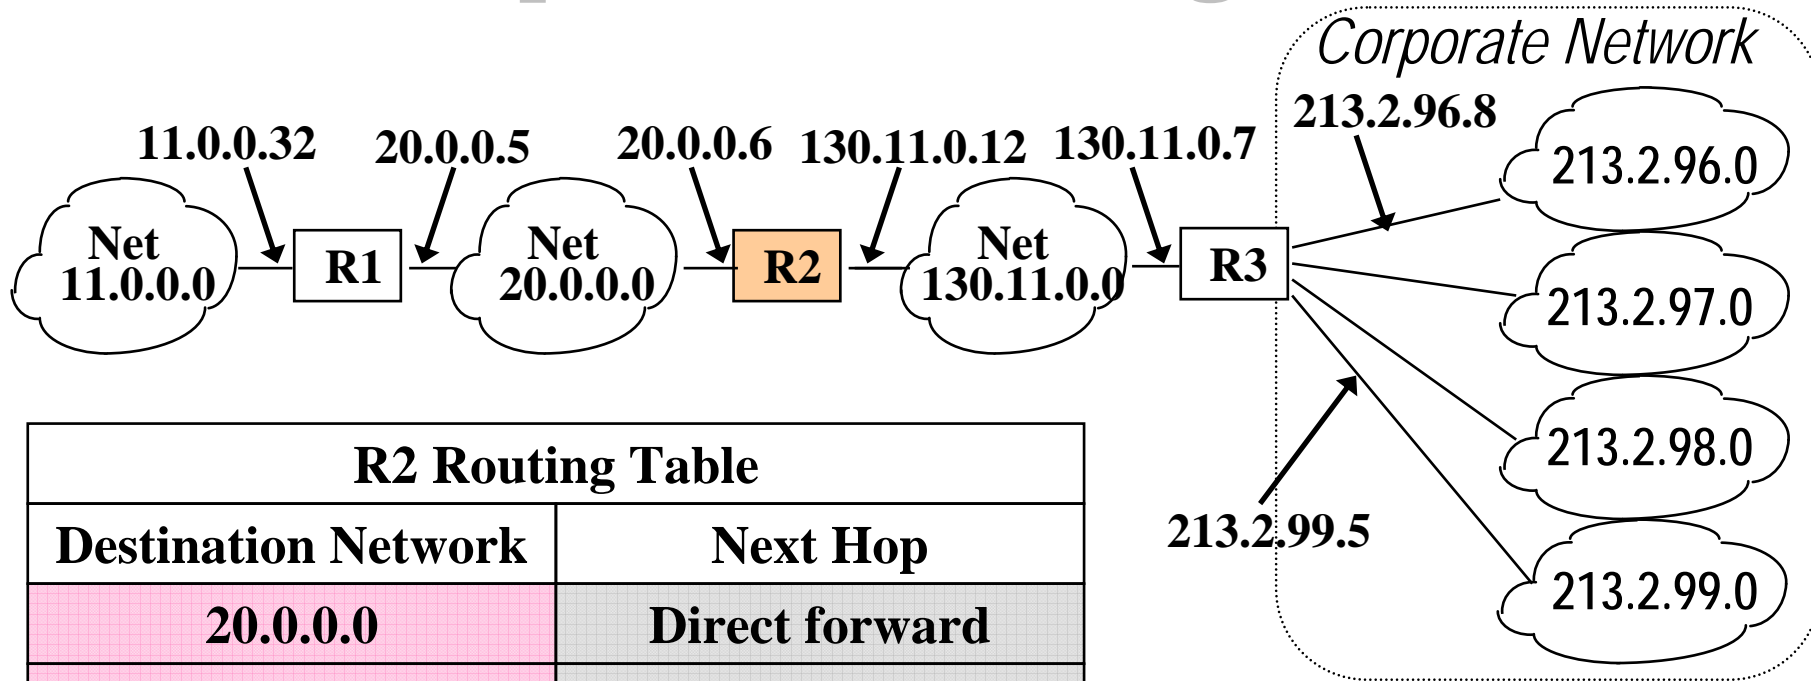
→ End 1995 projection: 70000 routes (critical);

→ End 1995 factual: 30000 routes thanks to classless routing

→ Mid 1999: 50000 routes

Figure 2: Growth of Internet Routing Tables

# Multiple class C assignment



R2 Routing Table	
Destination Network	Next Hop
20.0.0.0	Direct forward
130.11.0.0	Direct forward
11.0.0.0	20.0.0.5
213.2.96.0	130.11.0.7
213.2.97.0	130.11.0.7
213.2.98.0	130.11.0.7
213.2.99.0	130.11.0.7

- ➔ **Default routes: suboptimal traffic balancing**
- ➔ **Core routers: cannot have default routes (large tables)**
- ➔ **HW and SW limits on routing table lookup time**
- ➔ **Routing table updates are critical** (large tables traveling among routers for updates)



# Classless Inter-Domain Routing

## CIDR

- Officially developed in september 1993
  - ⇒ RFC 1517,1518,1519,1520
- CIDR also known as Supernetting
- Fundamental solution for Routing table problem
- Temporary solution to Internet address space depletion
  - ⇒ 32 bits: unwise choice
    - nobody could expect such an Internet growth
    - and Internet appliances will have a terrific impact
  - ⇒ unwise address assignment in early days
    - class B addresses with less than 100 hosts are common!!
  - ⇒ Projections (RFC 1752): address depletion between 2005 and 2011
  - ⇒ Ultimate solution: IPv6 (128 bits address!)

# CIDR model

## → Classless

⇒ Completely eliminates traditional concepts of Class A, B and C addresses

## → network prefix based

⇒ routers do not make any assumption on the basis of the three leading bits

⇒ they require an explicit network prefix to determine dividing point between net\_id and host\_id

⇒ clearly, capability of advertise prefix must be supported by routing protocol (e.g. BGP4)

**→ In essence: CIDR = VLSM applied to the WHOLE Internet!!**

# CIDR addresses

10.23.64.0/20      00001010.00010111.01000000.00000000

130.5.0.0/20      10000010.00000101.00000000.00000000

200.7.128.0/20      11001000.00000111.10000000.00000000

*Regardless the traditional class, all these addresses are similar!  
All address a network composed of as much as 4094 hosts*

Interpreting 200.7.128.0/20: a SINGLE NETWORK, contiguous block of 16 class C addr

200.7.128.0	200.7.132.0	200.7.136.0	200.7.140.0
200.7.129.0	200.7.133.0	200.7.137.0	200.7.141.0
200.7.130.0	200.7.134.0	200.7.138.0	200.7.142.0
200.7.131.0	200.7.135.0	200.7.139.0	200.7.143.0

# CIDR = supernetting


→ Organization assigned  $2^n$  class C addresses

⇒ with contiguous address space

→ addressing: use network bits with host\_id meaning

⇒ the opposite of subnetting!

*Example: 4 class C addresses appear to networks outside as a single network*





# Supernet Address

## → 4 address-contiguous networks:

⇒ 213.2.96.0	11010101.00000010.011000 <b>00</b> .00000000
⇒ 213.2.97.0	11010101.00000010.011000 <b>01</b> .00000000
⇒ 213.2.98.0	11010101.00000010.011000 <b>10</b> .00000000
⇒ 213.2.99.0	11010101.00000010.011000 <b>11</b> .00000000

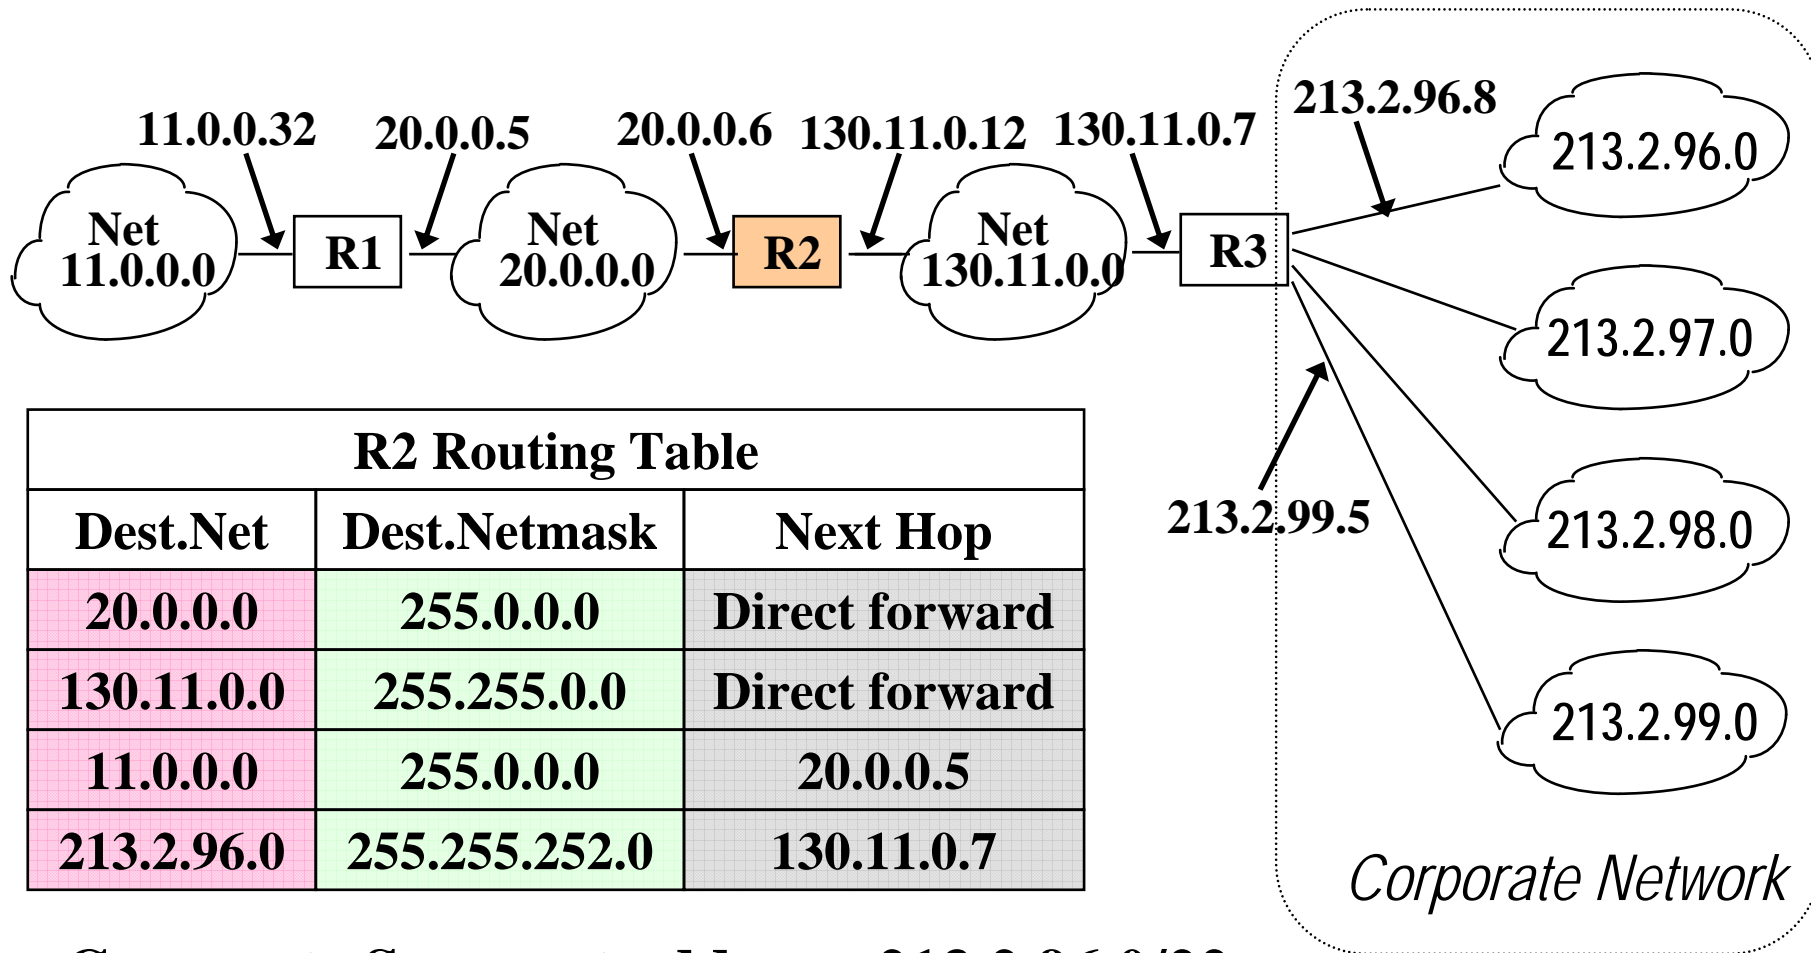
## → supernet mask:

⇒ 255.255.252.0

## → supernet address: 213.2.96.0/22

⇒ 11010101 . 00000010 . 011000 **00 . 00000000**

# Routing with CIDR



R2 Routing Table		
Dest.Net	Dest.Netmask	Next Hop
20.0.0.0	255.0.0.0	Direct forward
130.11.0.0	255.255.0.0	Direct forward
11.0.0.0	255.0.0.0	20.0.0.5
213.2.96.0	255.255.252.0	130.11.0.7

**Corporate Supernet address: 213.2.96.0/22**  
**11010101 . 00000010 . 011000 00 . 00000000**

# Large networks deployment

- Organization assigned 2<sup>n</sup> class C addresses
- may arbitrarily deploy subnetworks with more than 254 hosts!

⇒ This was impossible with class C, as natural netmask was /24

- BUT Software running on all the subnet hosts need to accept larger masks than natural one

⇒ e.g. setting netmask = 255.255.252.0 for host IP address 193.21.34.54 may be forbidden by sw

# **Requirements for CIDR support**

**→ Same of VLSM (but on a worldwide scale)**

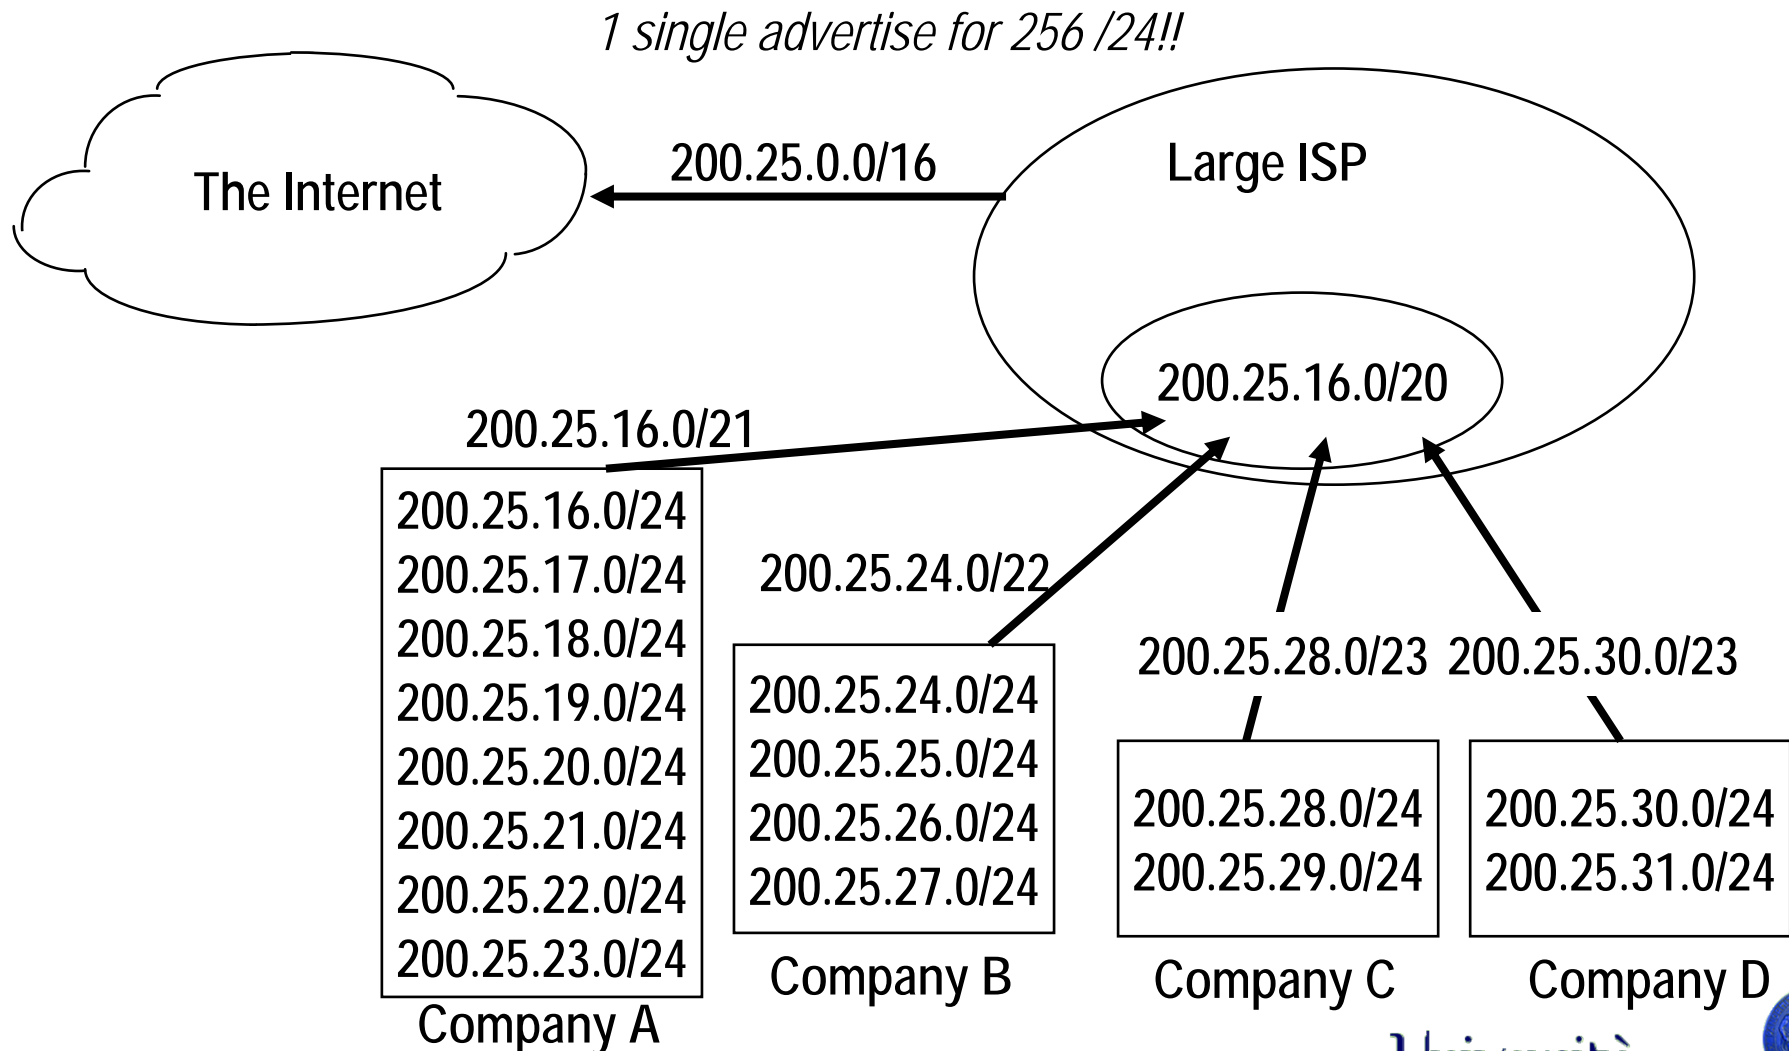
**⇒ Routing protocol must carry network prefix information with each route advertising**

**⇒ all routers must implement a consistent forwarding algorithm based on the “longest match”**

**⇒ for route aggregation to occur, addresses must be assigned to be topologically significant**

# Route aggregation

## control of internet tables growth



# CIDR allocation

## topological allocation of ex class-C addresses


Multi regional	192.0.0.0 - 193.255.255.255
Europe	194.0.0.0 - 195.255.255.255
Others	196.0.0.0 - 197.255.255.255
North America	198.0.0.0 - 199.255.255.255
Central-South America	200.0.0.0 - 201.255.255.255
Pacific Rim	202.0.0.0 - 203.255.255.255
Others	204.0.0.0 - 205.255.255.255
Others	206.0.0.0 - 207.255.255.255
IANA reserved	208.0.0.0 - 223.255.255.255

All are class C blocks, since class B blocks are no more allocated...

Recent trends: "attack" unused class A addresses  
(address space 64.0.0.0/2: from 64.0.0.0 to 126.0.0.0)

# Longest match forwarding

IP packet  
Destination: 203.22.66.5  
11001011 . 00010110 . 01000010 . 00000101



Routing table	
203.0.0.0 /11	Route 1
203.20.0.0 /14	Route 2
203.22.64.0 /20	Route 3

Three matches

Best (longest) match

R1: 11001011 . 00010110 . 01000010 . 00000101

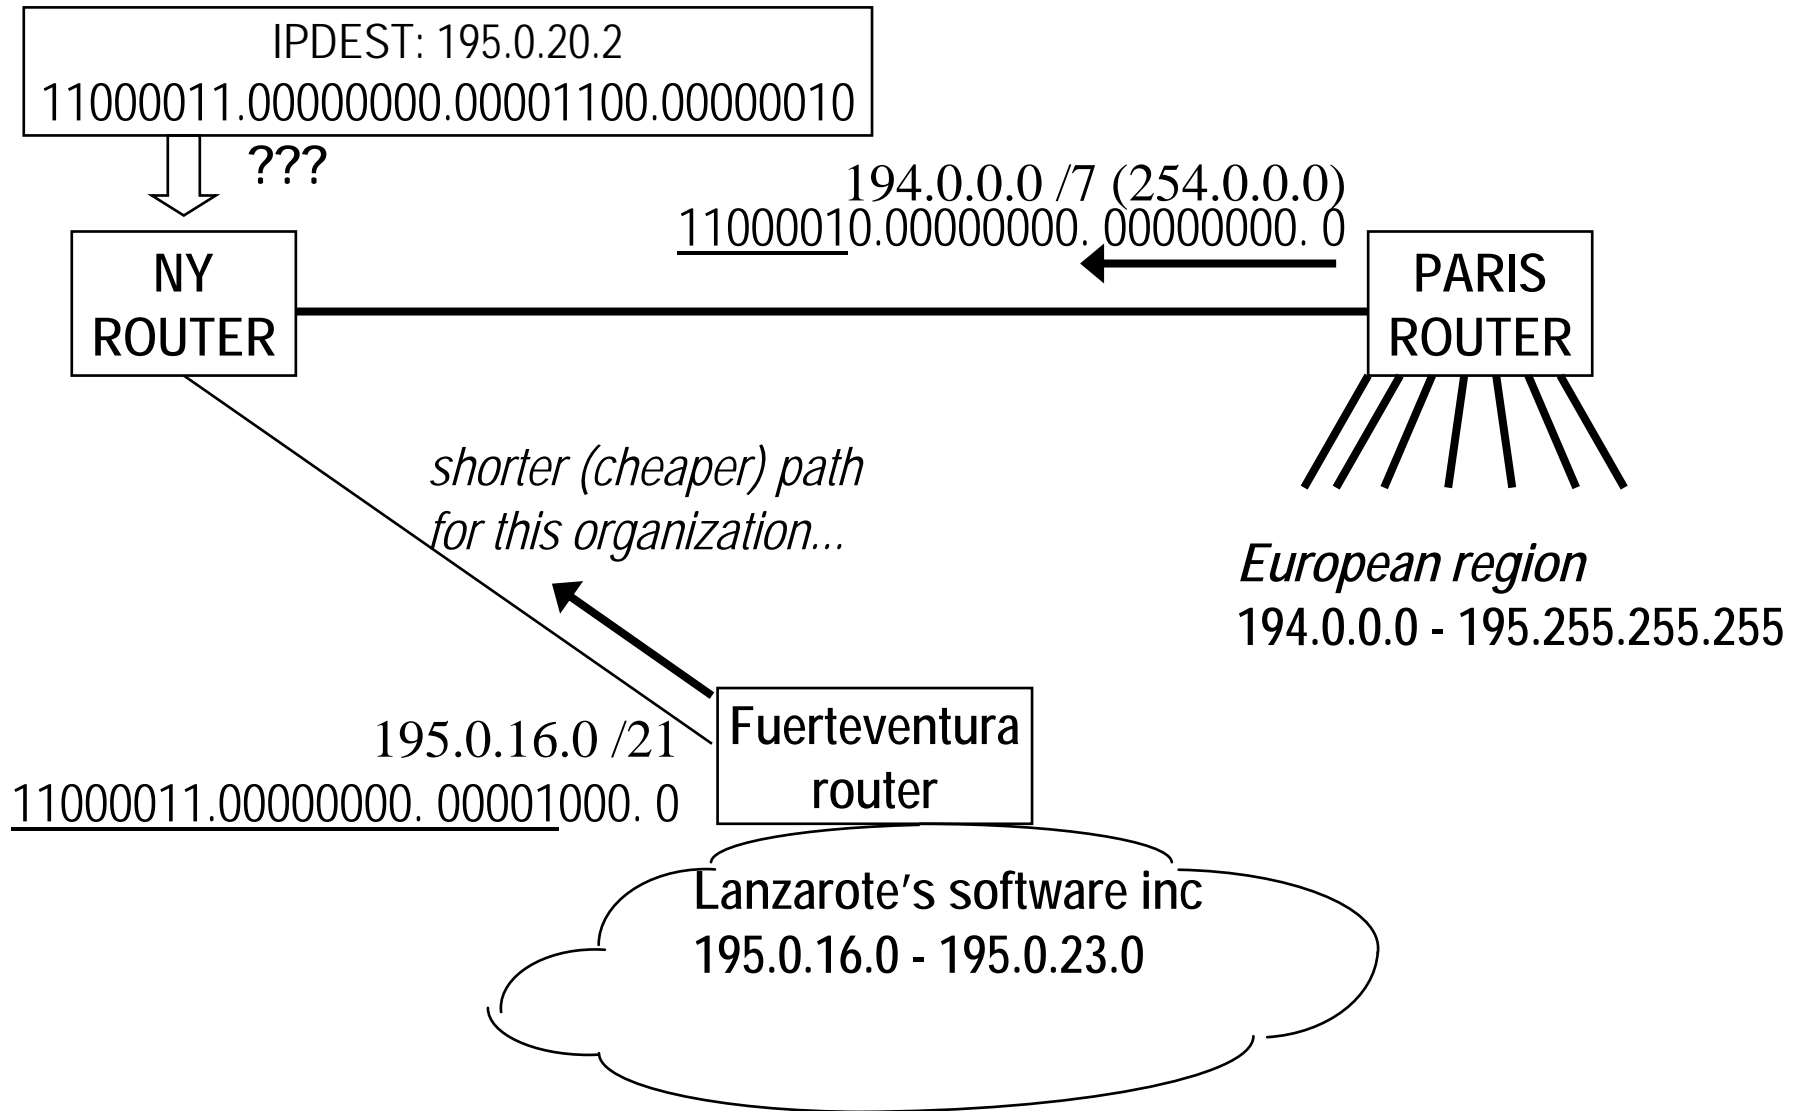
R2: 11001011 . 00010110 . 01000010 . 00000101

R3: 11001011 . 00010110 . 01000010 . 00000101

Longest match(R3) = smaller network

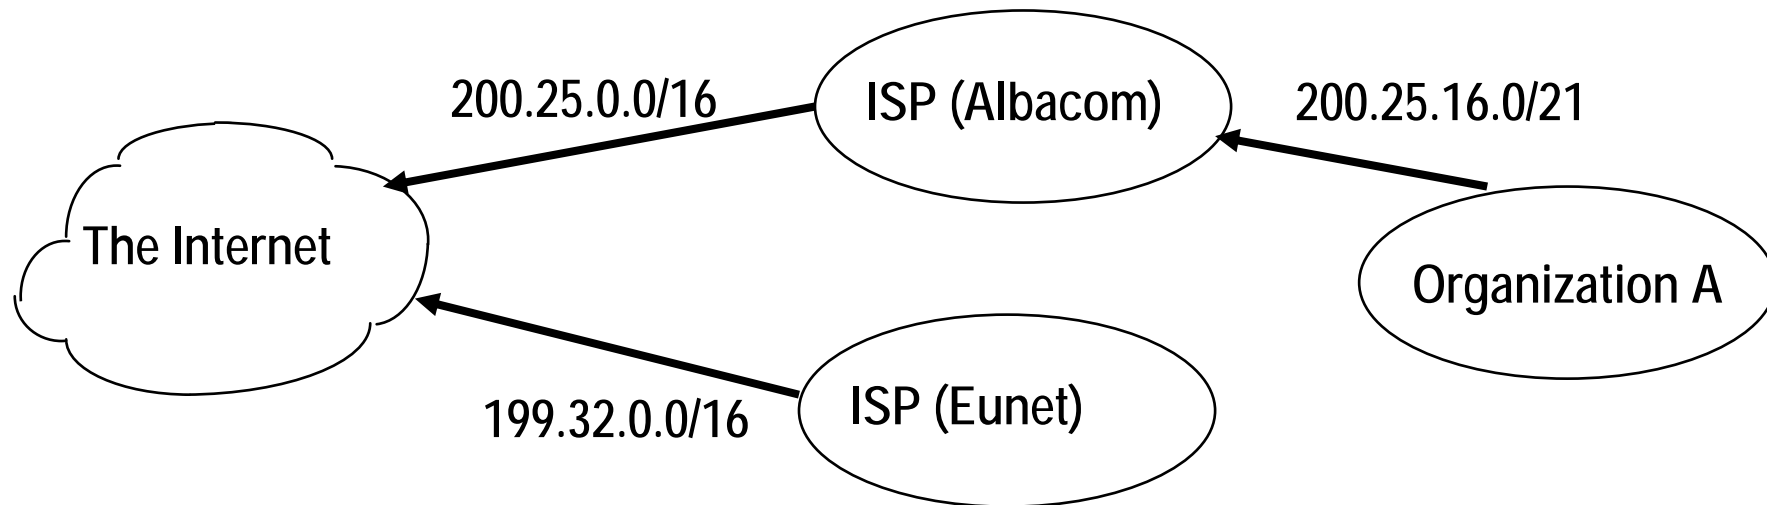
*But why longest match is ever needed???*

# Exception route





# Common exception route case



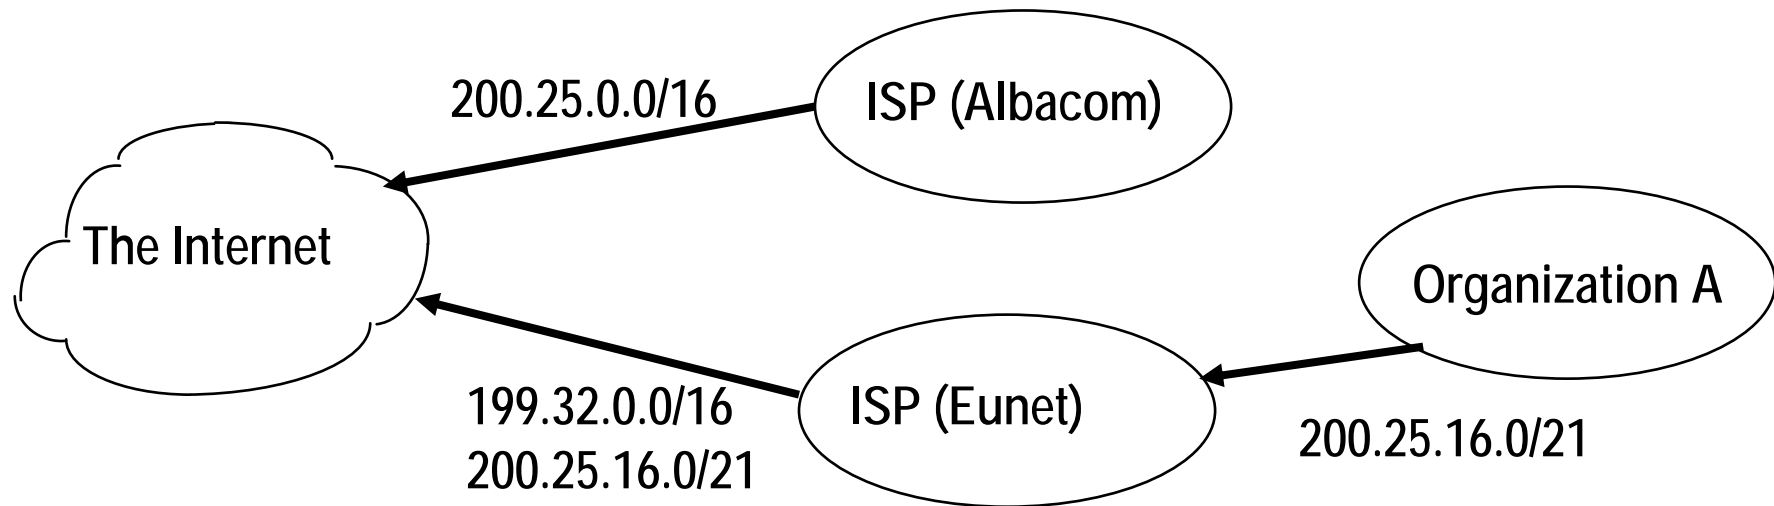
**→ At a point in time, organization A selects Eunet as new ISP!**

⇒ Best thing to do (for the Internet): obtain a new block of addresses and renumber

⇒ virtually impossible for a reasonably complex organization...

→ and even think to organizations that re-sells subnets...

# Common exception route case



→ **Then organization A keeps the same address block**

⇒ Eunet is in charge to advertise the new block, too, by injecting in the internet more specific route infos

⇒ This has created a new entry in routing tables, to be solved with longest match

# The open problems of CIDR

## 1. Still exist pre-CIDR routers

- ⇒ Non CIDR routers: Need to rely on “default” routes to keep reasonable routing table sizes
- ⇒ Consequence: not optimal routing (longer paths)

## 2. The number of exceptions is raising

- ⇒ recent trends indicate a return to exponential routing tables growth!
  - *Address ownership (portable blocks): dramatic*
    - » *Proposals (not accepted) to allows ownership only up to /9 ISPs*
    - » *Current “rule”: ownership starts from 8192 host networks (/19)*
  - *Address lending*
    - » *Renumbering necessary when changing ISP*

## 3. Shortage of IP addresses remains a hot problem

- ⇒ Appeals to return unused IP addresses (RFC 1917)
  - *unlikely, as they are viewed as assets!!*

# Address blocks for private Internets (RFC 1918)

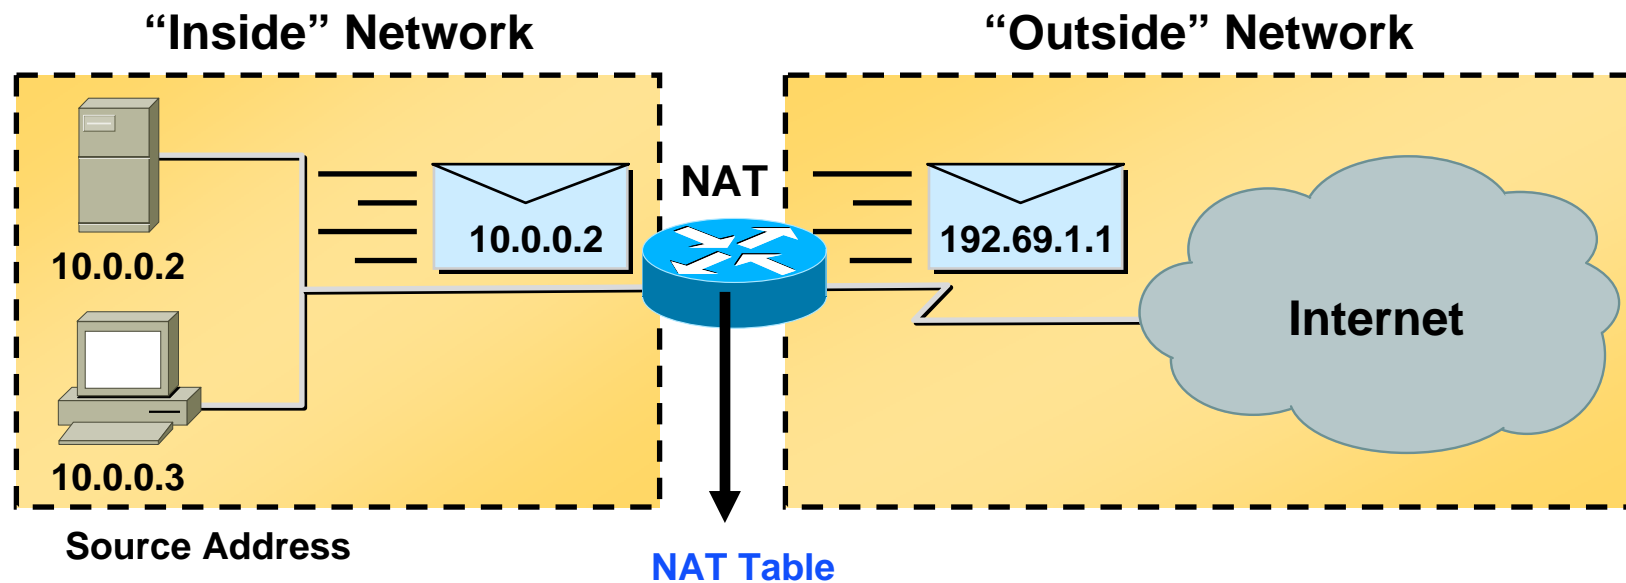
## IANA-Allocated, Non-Internet Routable, IP Address Schemes

Class	Network Address Range
A	10.0.0.0-10.255.255.255
B	172.16.0.0-172.31.255.255
C	192.168.0.0-192.168.255.255

To be used by private organizations not connected to the Internet  
No need to ask to IANA or InterNIC for these addresses.

Use Network Address Translator when external connectivity needed

# Network Address Translator



Inside Local IP Address	Inside Global IP Address
10.0.0.2	192.69.1.1
10.0.0.3	192.69.1.2

→ Map external address with Internal ones (may be a subset)

# IPv6

## (IP next generation - IPng)

### → The ultimate address space solution

⇒ 128 bit addresses

⇒ some other very important corrections and improvements to IPv4

→ although mostly designed to be as close as possible to IPv4

### → Prices to pay:

⇒ Double IP header size (40 bytes versus 20)

⇒ Difficult and slow transitory from IPv4 to IPv6