



CIDR – VLISM – AS

รศ.ดร. อนันต์ พลเพิ่ม

Asso. Prof. Anan Phonphoem, Ph.D.

anan.p@ku.ac.th

<http://www.cpe.ku.ac.th/~anan>

Computer Engineering Department

Kasetsart University, Bangkok, Thailand



Outline

- Classless Interdomain Routing (CIDR)
- Variable Length Subnet Mask (VLSM)
- Autonomous System (AS)



IP Addresses Revisited

- Potential exhaustion of IPv4 address space (due to inefficiency)
 - Class B is too big
 - Class C is too small (many are available)
- Growth of back bone routing tables
 - Lots of small networks causes large routing tables
 - Route calculation and management requires high computational overhead




Classless InterDomain Routing (CIDR)

- Try to balance two competing effects
 - Address utilization
 - Router complexity
- CIDR allows routers to break the rigid interpretation of IP address structures
- Also called “Supernet”
 - Opposite of “Subnet”

Classful & Classless addressing

Classful



Obsolete

- *inefficient*
- *depletion of B space*
- *too many routes from C space*

Classless

Hosts	Prefix	Classful
2	/31	
4	/30	
8	/29	
16	/28	
32	/27	
64	/26	
128	/25	
256	/24	1 C
...
4096	/20	16 C
8192	/19	32 C
16384	/18	64 C
32768	/17	128 C
65536	/16	1 B
...

Best Current Practice



Prefix Length

<i>/n</i>	<i>Mask</i>	<i>/n</i>	<i>Mask</i>	<i>/n</i>	<i>Mask</i>	<i>/n</i>	<i>Mask</i>
/1	128.0.0.0	/9	255.128.0.0	/17	255.255.128.0	/25	255.255.255.128
/2	192.0.0.0	/10	255.192.0.0	/18	255.255.192.0	/26	255.255.255.192
/3	224.0.0.0	/11	255.224.0.0	/19	255.255.224.0	/27	255.255.255.224
/4	240.0.0.0	/12	255.240.0.0	/20	255.255.240.0	/28	255.255.255.240
/5	248.0.0.0	/13	255.248.0.0	/21	255.255.248.0	/29	255.255.255.248
/6	252.0.0.0	/14	255.252.0.0	/22	255.255.252.0	/30	255.255.255.252
/7	254.0.0.0	/15	255.254.0.0	/23	255.255.254.0	/31	255.255.255.254
/8	255.0.0.0	/16	255.255.0.0	/24	255.255.255.0	/32	255.255.255.255



CIDR Example

What is the first address in the block if one of the addresses is **167.199.170.82/27**?

Solution

Address in binary: 10100111 11000111 10101010 01010010
Keep the left 27 bits: **10100111 11000111 10101010 01000000**

Result in CIDR notation: 167.199.170.64/27




Supernetting: CIDR

- Enable network number to be any length (No Class)
- Collapse multiple addresses assigned to a single AS to one address
- All routers must understand CIDR addressing
 - Need both Address and Mask (prefix and suffix)
 - Slash notation (123.10.16.0 /20)
- Some prefixes are reserved for private add.
 - 10/8, 172.16/12, 192.168/16, 169.254/16
 - These are not routable in the Internet

Example of CIDR

- Consider an ISP providing IP connection to a number of private companies
- If the IP addresses for the companies are carefully selected
 - a border router need only advertise one “aggregated” route for all the companies



Example of CIDR (Supernetting)

- If ISP needs 16 class C addresses
 - make them **contiguous**
- Eg. 199.23.16.0 to 199.23.31.0
 - enables a 20-bit network number

199.23.0001	0000.0	→	199.23.16.0
199.23.0001	0001.0	→	199.23.17.0
199.23.0001	0010.0	→	199.23.18.0
199.23.0001	0011.0	→	199.23.19.0
...			
199.23.0001	1111.0	→	199.23.31.0

Example of CIDR

Without CIDR

199.23.16.0

199.23.17.0


...

199.23.31.0

With CIDR

199.23.16.0/20

Border Gateway



Aggregation

- Some pairs of consecutive prefixes
- Example: routes within the same AS:
AS has 2 address blocks:

1.2.2.0/24 =	0000001.00000010.00000010.00000000/24
1.2.3.0/24 =	0000001.00000010.00000011.00000000/24

Can announce 1.2.2.0/23



CIDR: Longest prefix match

- Because prefixes of arbitrary length allowed, **overlapping prefixes** can exist.
- Example:
router hears 124.39.0.0/16 from one neighbor
and 124.39.11.0/24 from another neighbor
- Router forwards packet according to most specific forwarding information, called longest prefix match
 - Packet with destination 124.39.11.32 will be forwarded using /24 entry.
 - Packet with destination 124.39.22.45 will be forwarded using /16 entry



CIDR: Longest prefix match

- Implicit ordering in the routing table
 - longer prefixes higher up the table
 - So, the first match is the right one
- Explicit route to directly attached host
 - a netmask of 0.0.0.0



VLSM




VLSM


- Variable-length subnet mask
- Classful allows only one subnet in a network
 - > one subnet in an autonomous system
- Maximizing the use of address (Subnet Zero)
- “Subnetting a Subnet”
- Routing that support VLSM
 - OSPF, Integrated IS-IS, EIGRP, RIPv2, and static routing

Subnet with VLSM

192.168.10.0/24



Regular Subnet




192.168.10.0/24

7 subnets; The largest subnet needs **60** hosts

If **3** bits for subnet (8 subnets) → **5** bits for host (32 hosts)

If **6** bits for host (64 hosts) → **2** bits for subnet (4 subnets)

Subnet with VLSM



- Select the biggest first

- 192.168.10.0/24

- 192.168.10.0/26

← Perth

- 192.168.10.64/26

- 192.168.10.128/26

- 192.168.10.192/26


192.168.10.64/26

- 192.168.10.64/27

← KL

- 192.168.10.96/27

Subnet with VLSM



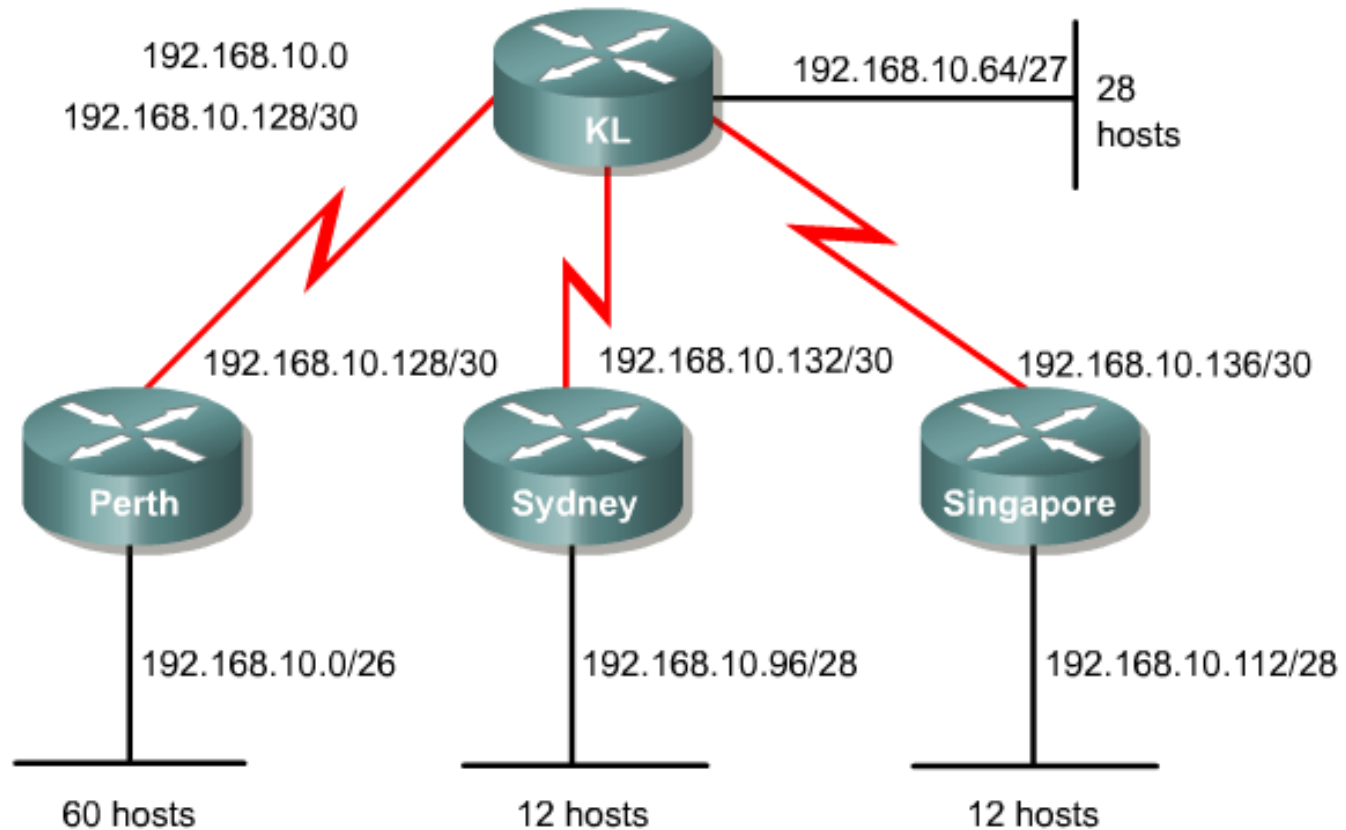
- 192.168.10.96/27

- 192.168.10.96/28 ← Sydney
- 192.168.10.112/28 ← Singapore

- 192.168.10.128/26


- 192.168.10.128/30 ← Perth – KL
- 192.168.10.132/30 ← Sydney – KL
- 192.168.10.136/30 ← Singapore – KL
- 192.168.10.140/30
- ...

Final: Subnet with VLSM





Assignment

250 Stations





CPE



700 Stations



165.23.208.0/20

IE



500 Stations


EE



100 Stations

Solution

250 Stations



700 Stations 500 Stations 100 Stations

- 165.23.208.0/20
- 165.23.208.0/22** 1022 Hosts
- 165.23.212.0/22
- 165.23.216.0/22
- 165.23.220.0/22

The largest subnet needs **700** hosts
→ **10** bits for host (1024 hosts)

- 165.23.212.0/22
- 165.23.212.0/23** 510 Hosts
- 165.23.214.0/23

- 165.23.214.0/23
- 165.23.214.0/24** 254 Hosts
- 165.23.215.0/24


- 165.23.215.0/24
- 165.23.215.0/25** 126 Hosts
- 165.23.215.128/25

- 165.23.215.128/25 2 Hosts
- 165.23.215.128/30**
- 165.23.215.132/30**
- 165.23.215.136/30
- ...



Final Solution

165.23.214.0/24

250 Stations



CPE



700 Stations

165.23.208.0/22


165.23.208.0/20

165.23.215.128/30

IE


165.23.215.132/30

EE



500 Stations

165.23.212.0/23



100 Stations


165.23.215.0/25



Notes for CIDR


- CIDR was actually intended as a **quick fix**
 - Solve addressing crisis until IPv6 was deployed
- Unfortunately, CIDR has been **widely adopted**
 - IPv6 deployment has proven to be very, very slow
- CIDR is **currently** deployed
 - However, IPv6 is not compatible with IPv4
 - Generates a big migration problem

Growth in Routing Table Size



Pre-CIDR (1988-1994): Steep Growth Rate

Growth in Routing Table Size



CIDR Deployment (1994-1996): Much Flatter



Autonomous System (AS)

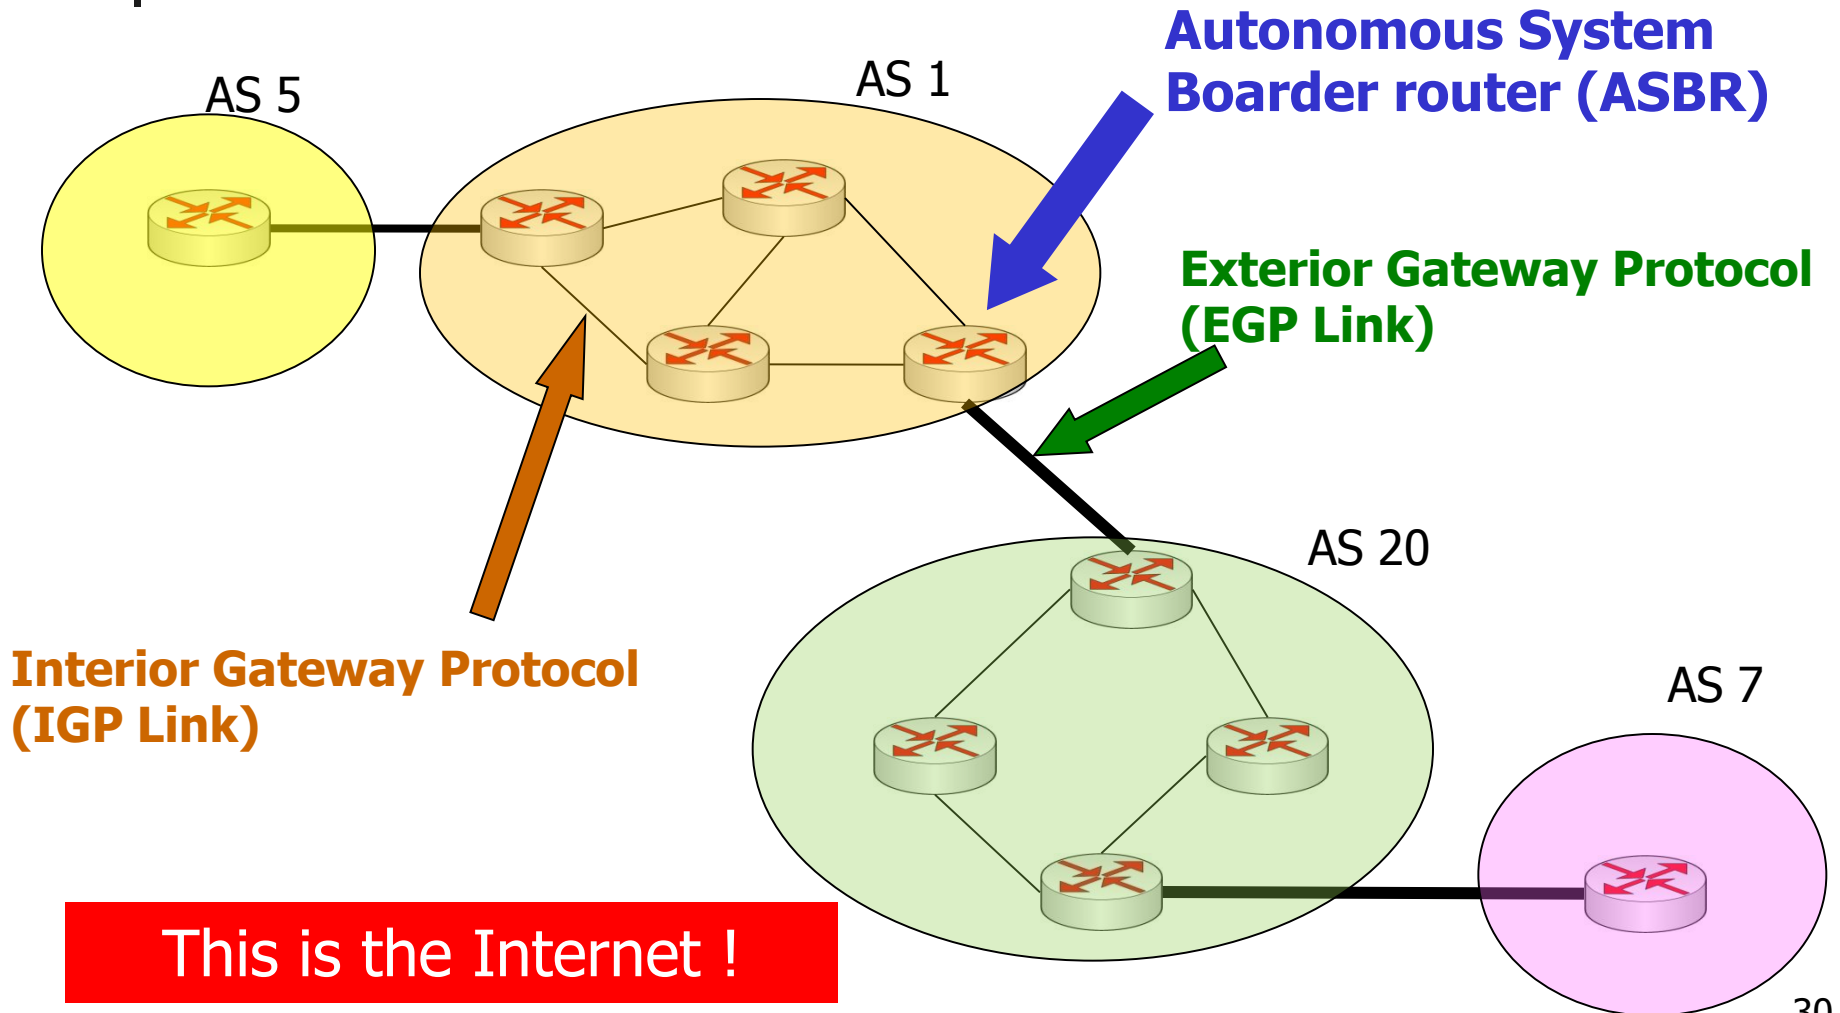
- A single network domain
- Grouping of computers/routers
- Operate in isolation from other groups
- A single network administrative entity



Autonomous System (AS)

- Need protocols for distribute routing information in the AS
 - Interior Gateway Protocols (IGPs)
 - Intradomain routing algorithms
- Between AS
 - Need interdomain routing algorithms
 - Exterior Gateway Protocols (EGPs)
 - More complex task

Autonomous System (AS)





Types of AS


- **Stub AS**
 - Only has a single connection to one other AS
 - only carries *local traffic*
- **Multihomed AS**
 - Connect to more than one other AS
 - But will not carry *transit traffic*
- **Transit AS**
 - Connect to more than one other AS
 - Can carry both *local* and *transit traffic*

Transit vs. Nontransit AS

Transit traffic = traffic whose **source** and **destination** are **outside the AS**

Nontransit AS: does not carry transit traffic

- Advertise own routes only
- Do not propagate routes learned from other AS's



Transit AS: does carry transit traffic

- Advertises its own routes PLUS routes learned from other AS's

