Networking Technologies and Applications

Rolland Vida
BME TMIT

October 14, 2016

Problems with Classful IP Addresses

Problem 3. Inflexible. Assume a company requires 2,000 addresses

- Class A and B addresses are overkill
- Class C address is insufficient (requires 8 Class C addresses)

Problem 4: Exploding Routing Tables: Routing on the backbone Internet needs to have an entry for each network address. In 1993, the size of the routing tables started to outgrow the capacity of routers.

Fix #2 (to both of these problems): Classless Interdomain Routing (CIDR)

CIDR - Classless Interdomain Routing

- **Goals:**
  - Restructure IP address assignments to increase efficiency
  - Hierarchical routing aggregation to minimize route table entries

Key Concept: The length of the network id (prefix) in IP addresses is arbitrary/flexible and is defined by the network hierarchy.

- Consequence:
  - Routers use the IP address **and** the length of the prefix for forwarding.
  - All advertised IP addresses must include a prefix

CIDR Example

- CIDR notation of a network address:

  192.0.2.0/18

  - "18" says that the first 18 bits are the network part of the address
- The network part is called the network prefix
- Example:
  - Assume that a site requires an IP network domain that can support 1000 IP host addresses
  - With CIDR, the network is assigned a continuous block of 1024 = 2^{10} (>1000) addresses with a 32-10 = 22-bit long prefix

CIDR: Prefix Size vs. Host Space

CIDR Block Prefix	# of Host Addresses
/27	32 hosts
/26	64 hosts
/25	128 hosts
/24	256 hosts
/23	512 hosts
/22	1,024 hosts
/21	2,048 hosts
/20	4,096 hosts
/19	8,192 hosts
/18	16,384 hosts
/17	32,768 hosts
/16	65,536 hosts
/15	131,072 hosts
/14	262,144 hosts
/13	524,288 hosts

CIDR and Address assignments

- IANA – Internet Assigned Numbers Authority
  - The RIRs get short prefix CIDR blocks
    - Regional Internet Registries
  - E.g., 62.0.0.0/8 assigned to RIPE NCC
    - Réseaux IP Européens Network Coordination Centre
- RIRs fragment and redistribute parts of the address space
  - Backbone ISPs obtain large blocks of IP address space and then reallocate portions of their address blocks to their customers.

Example:

- Assume that an ISP owns the address block 206.0.64.0/18, which represents 16,384 ($2^{32-18}=2^{14}$) IP host addresses
- Suppose a client requires 800 host addresses

01000100

> ➢ $512=2^9<800<1024=2^{10}$ -> 32-10 = 22,
>
> ➢ Assigning a /22 block, i.e., 206.0.68.0/22 -> gives a block of 1,024 (2^{10}) IP addresses to client.

CIDR example

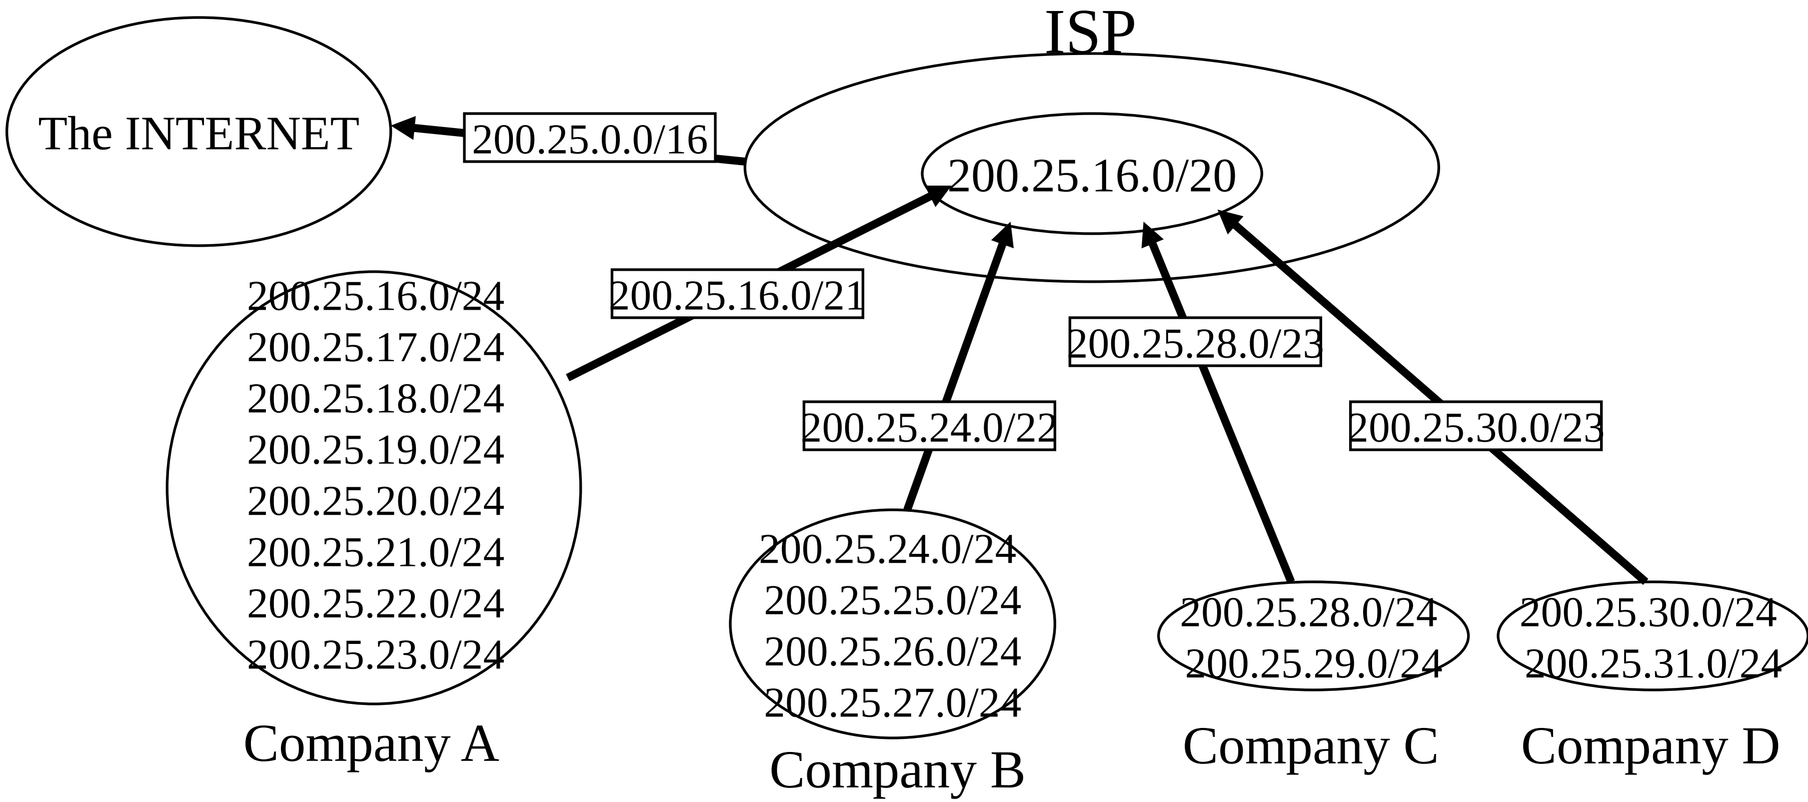

CIDR and Routing

- **Aggregation** of routing table entries:
  - 128.143.0.0/16 and 128.142.0.0/16 can be represented as 128.142.0.0/15 at a router.
    - 143 = 128.10001111.0.0 142 = 128.10001110.0.0
- **Longest prefix match**: Routing table lookup finds the routing entry that matches the longest prefix
  - Why?

E.g., What is the outgoing interface for destination IP address: 128.143.137.0?

Prefix	Interface/outgoing link
128.143.128.0/17	interface #1
128.128.0.0/9	interface #2
128.0.0.0/4	interface #5

Routing table

Problems with Classful IP Addresses

Problem 5. The Internet is going to outgrow the 32-bit addresses

- **Fix #3: IP Version 6**

Exhaustion of IPv4 addresses

- No problem in the US
  - „Internet Heaven”
- Serious problem everywhere else
  - In China they have asked IP addresses for connecting 60.000 schools, they received a Class B address (65.534 addresses)
  - Many European or African countries received just a Class C address (254 addresses)
- Fast development of the Internet outside Northern America
  - Asia (2.5 billion people), Eastern Europe (250 million), Africa (800 million), South and Latin America (500 million)
- New communication devices need IP addresses
  - Mobile phones, PDAs, sensors, cars, etc.
- The exhaustion of IPv4 addresses was always projected for the next month/years (for more than 10 years)

Is the Internet „full”?

- On February 1st, 2011, the last /8 IPv4 address block allocated to the RIRs

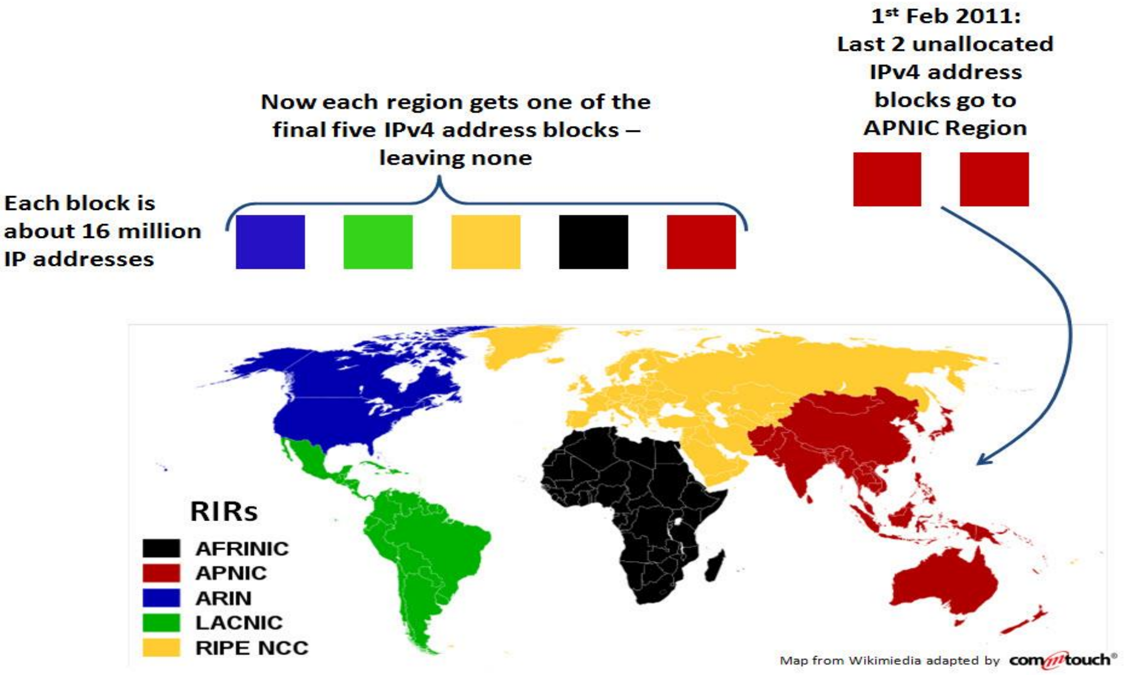

Reuse of private addresses

- Private addresses are not „visible” from the Internet
- Network address translation

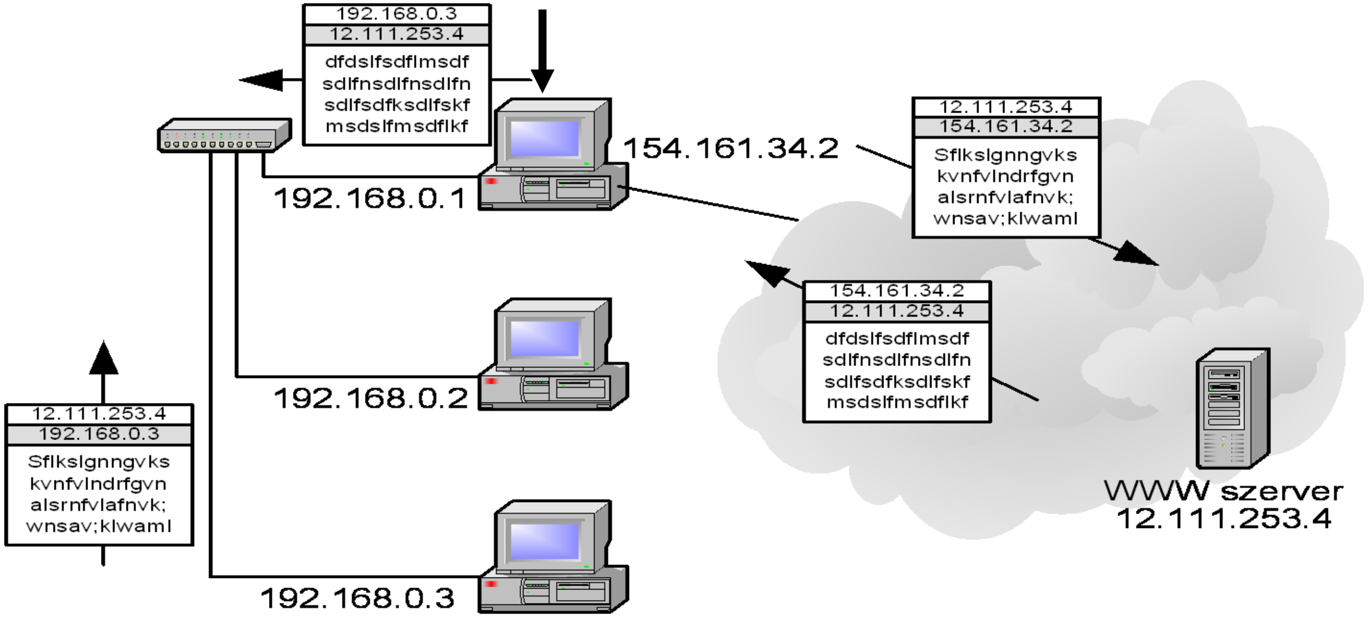

NAT problems

- Just an intermediate solution
  - Cannot establish a connection from outside with a host behind a NAT
  - More and more applications require public IP addresses
    - VoIP, videoconferencing, network games
  - Many protocols do not work on a network with NAT

IPv6 chronology

- **TUBA** (1992)
  - TCP and UDP over Bigger Addresses
  - Based on the OSI CLNP (Connection-Less Network Protocol)
  - abandoned
- **SIPP** (1993)
  - Simple IP Plus
  - 64 bit addresses
- **IPng**, based on an extended SIPP version (1994)
  - 128 bit addresses
  - From December 1995 officially called IPv6

IPv6 addressing scheme

- The IPv6 address pool is huge
  - 2^{128} = 340.282.366.920.938.463.463.374.607.431.768.211.456
  - 67 billion billion addresses for each cm^2 on the surface of the Earth
  - 10^{30} address for each person on the Earth
  - The address distribution and routing requires a hierarchical structure

IPv6 addressing scheme

- IPv6 addressing quite similar to IPv4
  - 128 bit long addresses, instead of 32 bits
- Three address types:
  - Unicast addresses
    - Identify a unique interface
  - Multicast addresses
    - Identify a group of interfaces, each of these will receive the message
    - Replace the broadcast addresses as well
  - Anycast addresses
    - Identify a group of interfaces, message will be delivered to one of these interfaces

Writing IPv6 addresses

- 128 bits = 16 bytes = 32 x 4bits = 32 hexadecimal digits grouped in 8 segments

  FECD:BA98:0000:0000:00CD:BA98:0000:3200

- The opening 0 digits in each segment can be neglected

  Instead of FECD:BA98:0000:0000:00CD:BA98:0000:3200

  we write FECD:BA98:0:0:CD:BA98:0:3200

- Adjacent 0 segments can be neglected, if there is only one siuch case in an address

  FECD:BA98::CD:BA98:0:3200

- Netwrok prefix is encoded as in case of IPv4 CIDR

  entire IPv6 address/prefix length in bits

  12AB:0000:0000:CD30:FFFF:DEC8:0000:0000/60

  12AB:0:0:CD30:0:0:0:0/60

  12AB:0:0:CD30::/60