IPv6 Overview

CIS-3152, Spring 2013
Peter C. Chapin

What's the Problem?

- Current network protocol: IP version 4.
  - Limited addresses
    - 32 bit addresses only allows 4 billion possibilities
    - Sounds like a lot but they are wastefully allocated
      - Large ranges allocated but underutilized.
      - Example: VSC has 64K addresses, but maybe only 10,000 machines.
    - Also many devices coming on line.
      - India, China, Brazil
      - Many hand held devices, etc (cell phones, PDAs)
  - Poor routing structure
    - See next slide!

Routing Table Explosion

- Originally there was no correlation between IP address and physical location.
  - Routing tables expanded linearly with network size

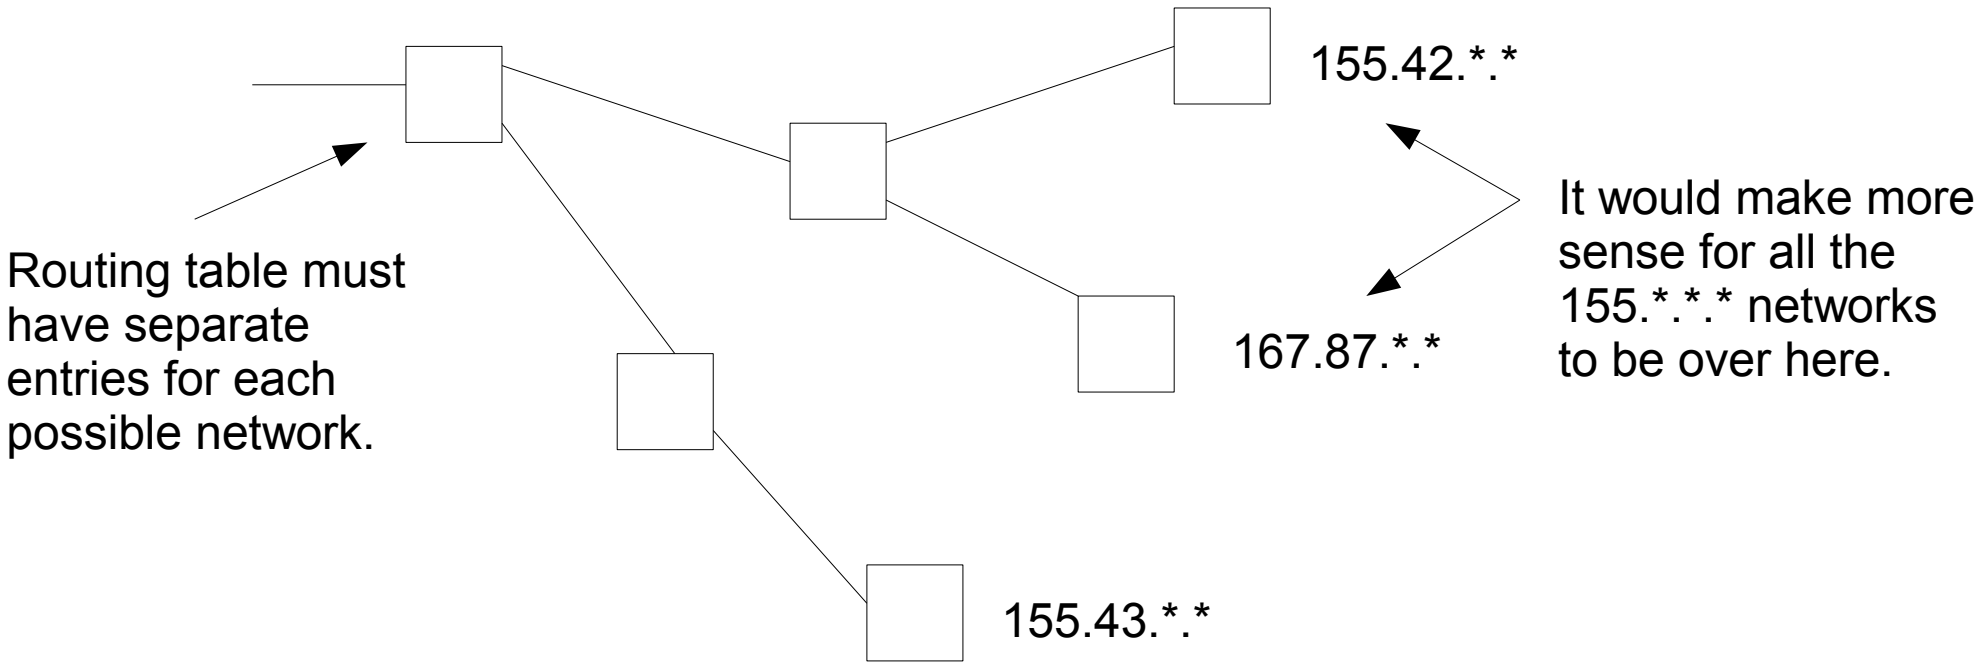

Network Collapse

- Growth toward destruction...
  - Network growth was exponential
  - Thus, routing table growth was exponential
  - Processor speed increased more slowly
  - Thus routers spent more time searching routing tables.
  - Eventually...
    - Routers wouldn't be able to keep up.
    - Packet lose would skyrocket.
    - The Internet would fail.
  - This was expected to happen by 2000.

CIDR

- Classless Interdomain Routing.
  - IDEA...
    - Allocate IP addresses according to network topology.
    - Similar addresses can be treated in one routing table entry.
  - Addresses with a common prefix all in the same “direction.”
    - 155.42.0.0/16
      - The trailing “16” means that the first 16 bits are the prefix.

CIDR In Action

- Top level router uses prefix to specify routes for many networks.
  - Reduces routing table size dramatically.
  - Tables now grow logarithmically with network size.

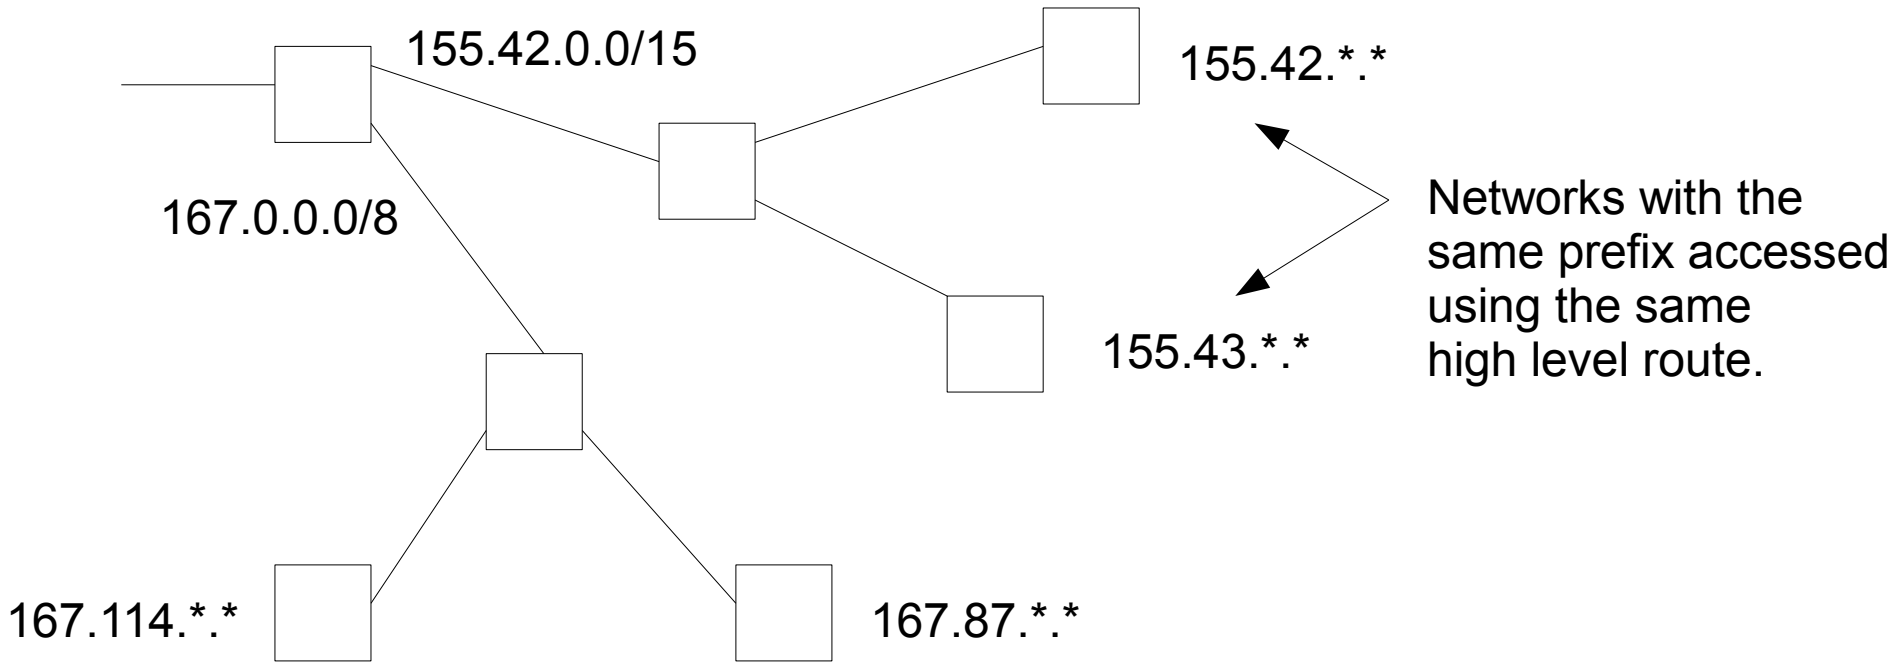

CIDR Deployed

- CIDR was deployed for all new allocations during the 1990s.
  - Collapse of the Internet was averted!
- However, old (class-style) address allocations remain.
  - A burden that can never be removed.

Network Address Translation

- NAT can be used to reduce pressure on IP addresses.
  - Maps single global address to entire network.

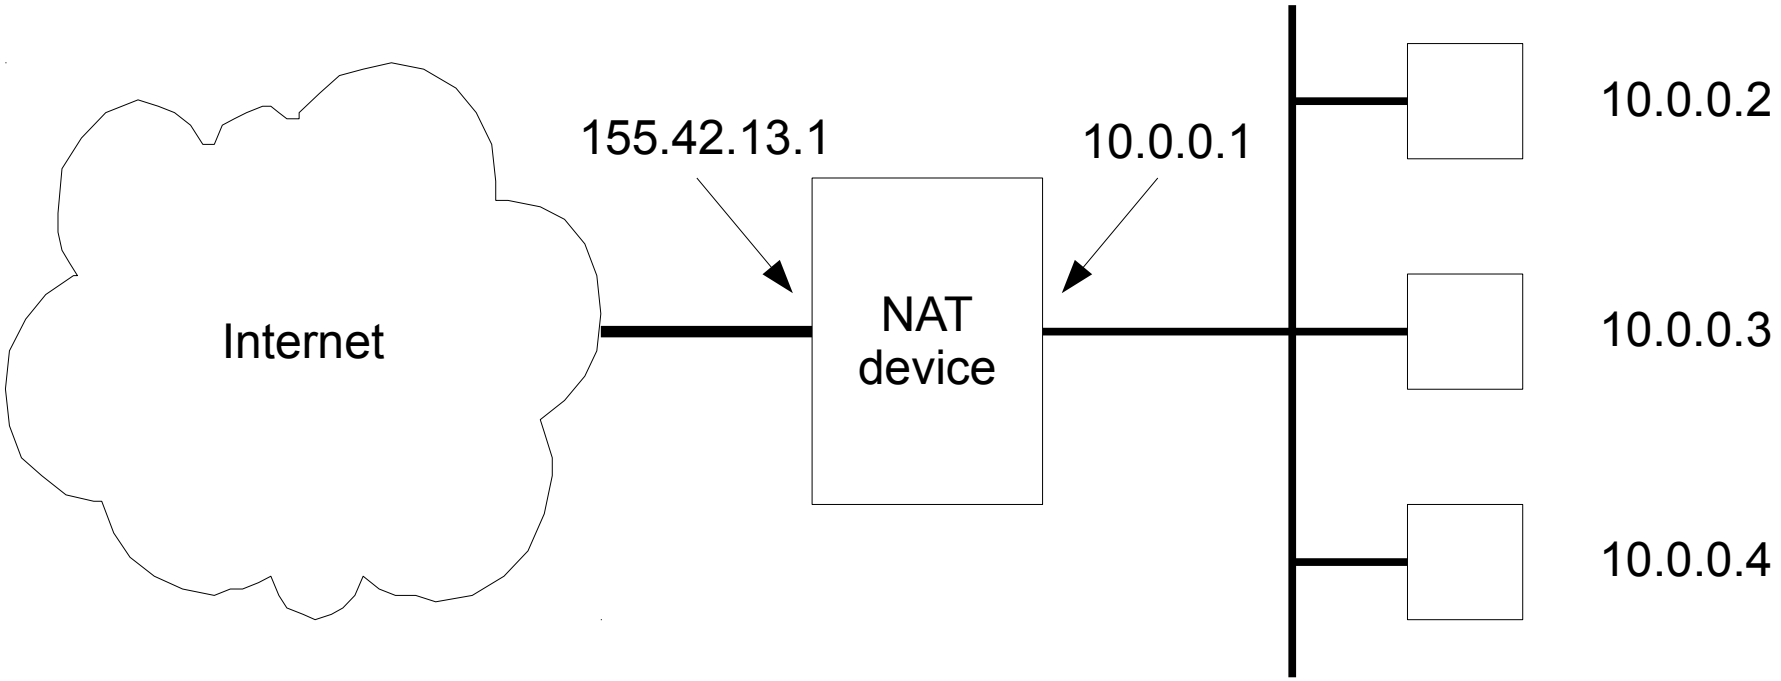

NAT Device

- Acts as a proxy.
  - Outgoing connections remapped.
    - NAT makes real connection using global interface.
    - Associates the port number it uses with internal client.
    - Forwards all data between client and Internet.
    - Allows a single global address to support many machines!
  - Incoming connections are harder.
    - NAT device must be configured to listen for each internal server.
      - And so... must know about all internal servers.
  - Note: NAT is a natural location for firewall functions.

Problems Remain

- NAT not suitable for all situations.
  - Certain protocols don't work well through NAT devices.
    - Security protocols especially (the NAT device is a kind of “man in the middle” and security protocols are specifically designed to prevent such an agent from tampering).
    - Also protocols that “leak” to higher layers (IP addresses in application headers, etc).
  - Limitation of 64K ports limits connection mappings.
- Still running out of addresses.
- Long term solution: Create a new protocol!

Introducing IPv6

- 128 bit addresses!
  - $2^{128} = 3.403 \times 10^{38}$ (big)
- Address format
  - 0123:4567:89AB:CDEF:0123:4567:89AB:CDEF
    - Displayed in hex, 16 bits at a time.
    - 8 sections separated by colon characters.
  - Zero sections can be dropped (zeros understood)
    - 0123:0000:0000:0000:0123:4567:89AB:CDEF
    - 123::123:4567:89AB:CDEF
    - Saves space; many IPv6 addresses have many zeros.
    - Can use :: only once in an address (to avoid ambiguity)

Special Addresses

- 0::0
  - The “unspecified” address. Used when an interface has no address (yet).
- ::1
  - The loop back address. Used to refer to yourself (like 127.0.0.1 in IPv4).
- ::FFFF:*IPv4-addr*
  - Example, ::FFFF:155.42.13.24
  - Used by IPv6 software to communicate with IPv4 nodes.
    - Emitted packets are actually IPv4.

Address Scopes

- An interface can have several addresses
  - Global unicast address
    - An address that is unique over the whole world
    - This is the “normal” case.
  - Site local address
    - An address only meaningful inside a site (such as VTC)
  - Link local address
    - An address only meaningful on a particular link.
- Link local addresses are generated automatically by IPv6 software.
  - Link communication possible with no external help.

Link Local Addresses

- Always have prefix FE80::/10
  - Notice: IPv6 uses CIDR from the start.
  - Prefix bits...
    - 1111,1110,10...
    - All addresses that start with FE8, FE9, FEA, and FEB are link local.
  - Format: FE80::*interface-id*
    - Where *interface-id* is a 64 bit number derived from MAC address of interface.
    - No need for user assignment; easy configuration.
  - Routers never forward link-local addresses.
    - Same address on different links is okay.

Site Local Addresses

- Always have prefix FEC0::/10
  - Prefix bits...
    - 1111,1110,11...
    - All addresses that start with FEC, FED, FEE, and FEF are site local.
  - Format: FEC0::*subnet-id*:*interface-id*
    - Where *subnet-id* is a 16 bit subnetwork number defined by the site administrators.
    - Where *interface-id* is a 64 bit number derived from MAC address of interface.
    - Allows an organization to use IPv6 internally.
  - Replaces IPv4 “unroutable” addresses.

Unique Local Addresses

- Problems with site local addresses:
  - Not unique.
    - Addresses sometimes “leak” into higher level protocols and are then exposed to the Internet.
- RFC-4193
  - Reserves FC00::/7 for “Unique Local IPv6 Unicast Addresses”

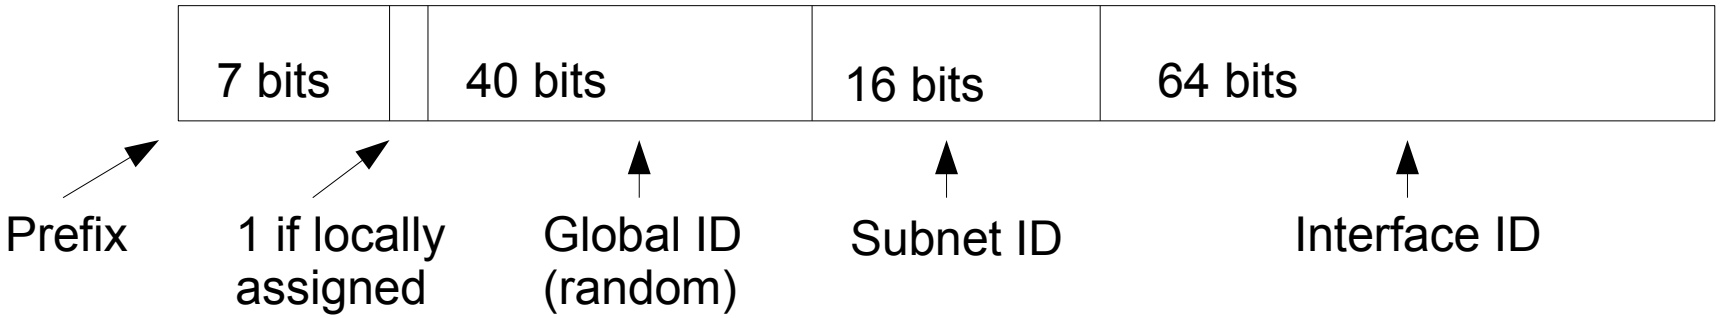

Computing Global ID

- Algorithm described in RFC-4193
  - http://www.goebel-consult.de/ipv6/createLULA
  - VTC's: FD25:F376:7B60::/48
- Random global ID unlikely to conflict.
  - ... But more likely than you might think.
  - Due to “birthday surprise” conflict probability is high after 2^{20} global IDs have been allocated.
    - That's only about 1,000,000 (base 10).

Global Addresses

- Addresses *can* appear on the open Internet.
  - Prefix bits 001 (2000::/3)
  - Must be assigned by a global authority.

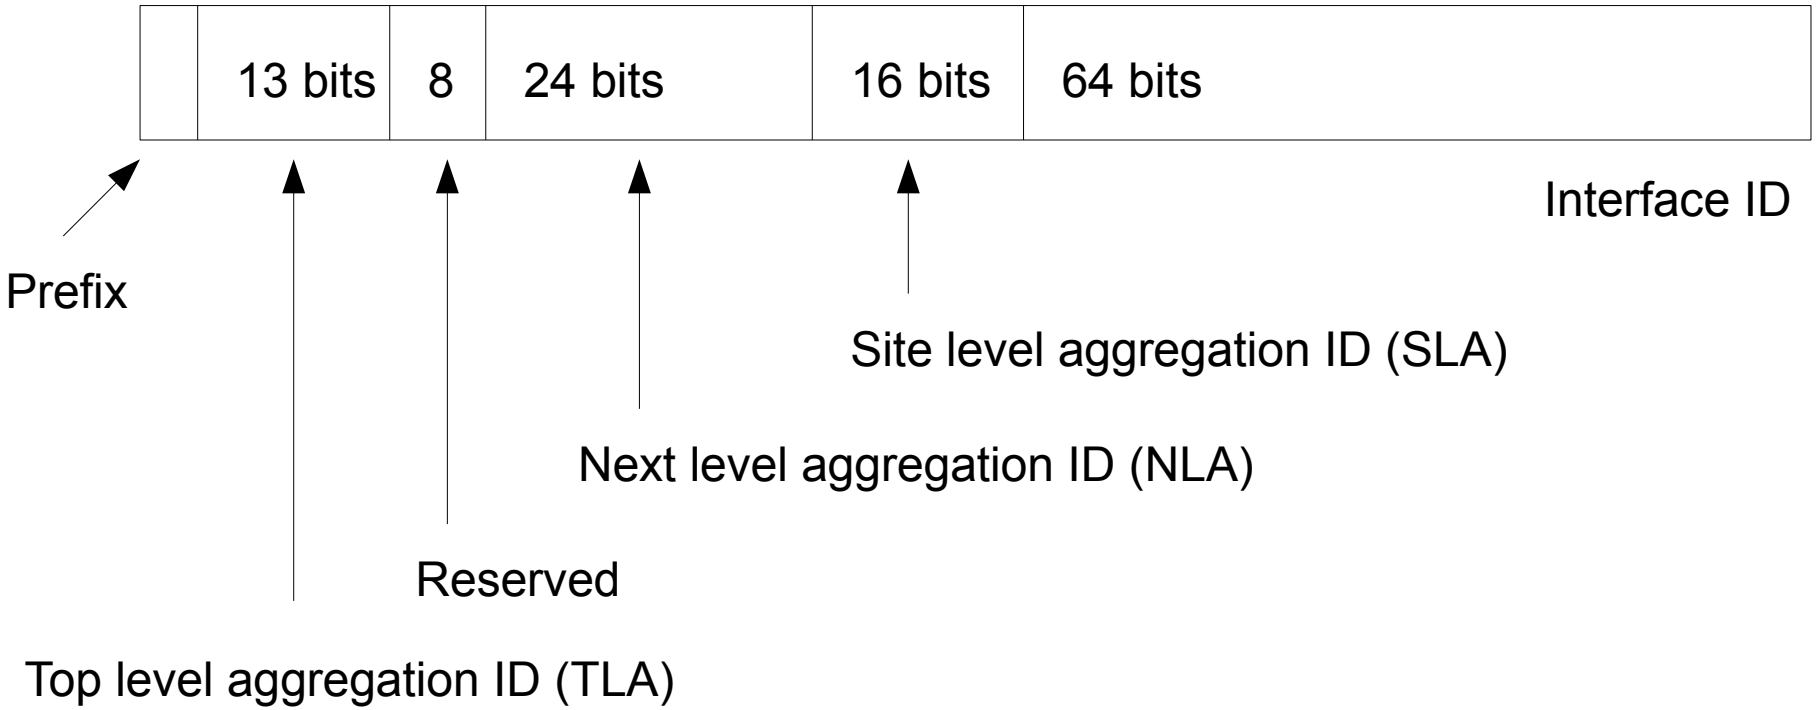

Tunneling

- How to do IPv6 over IPv4 infrastructure?
  - Obviously can't change entire network in one step
  - But... IPv6 can't be routed by IPv4 only routers.
- Solution...
  - Use IPv4 for on-the-wire communication.
    - Negates many of the advantages of IPv6.
  - Encapsulate IPv6 traffic in IPv4 packets.

Configured Tunnel

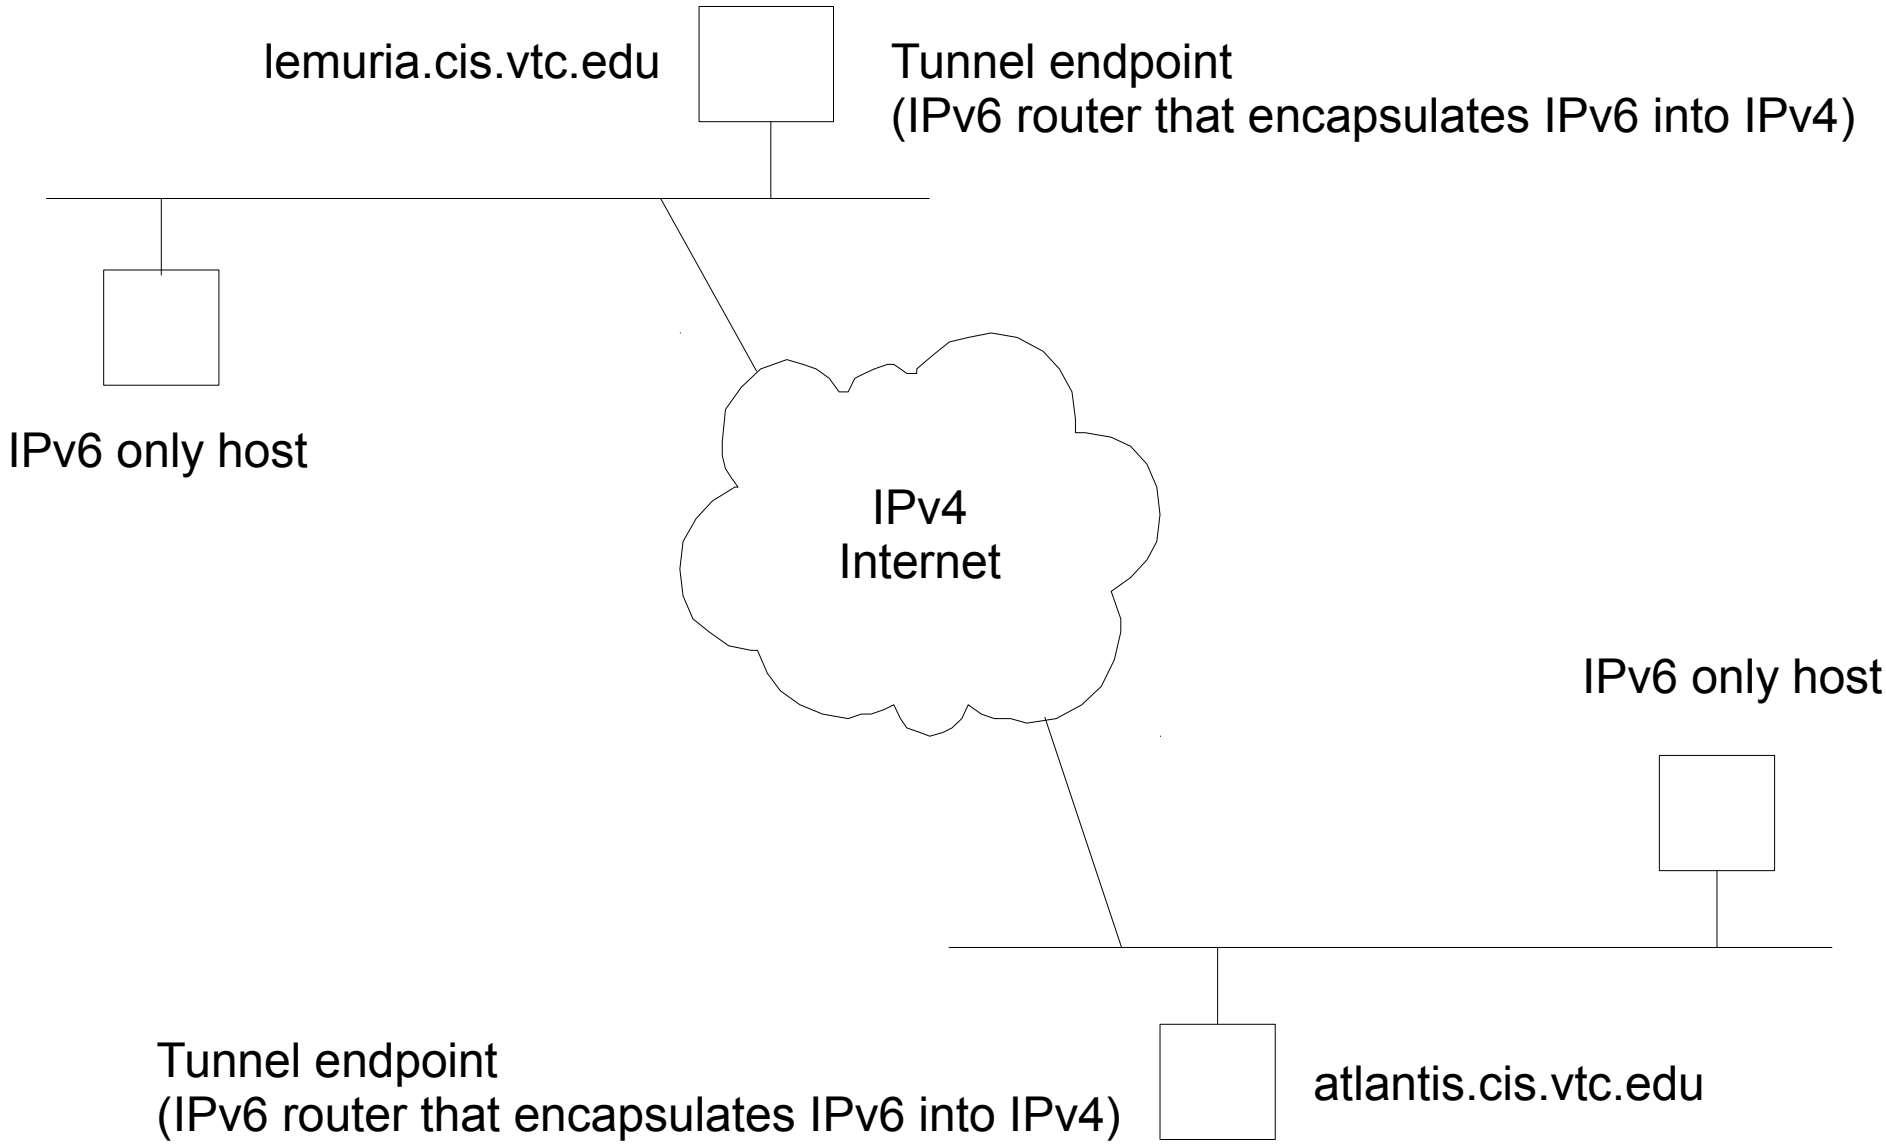

Automatic Tunnel

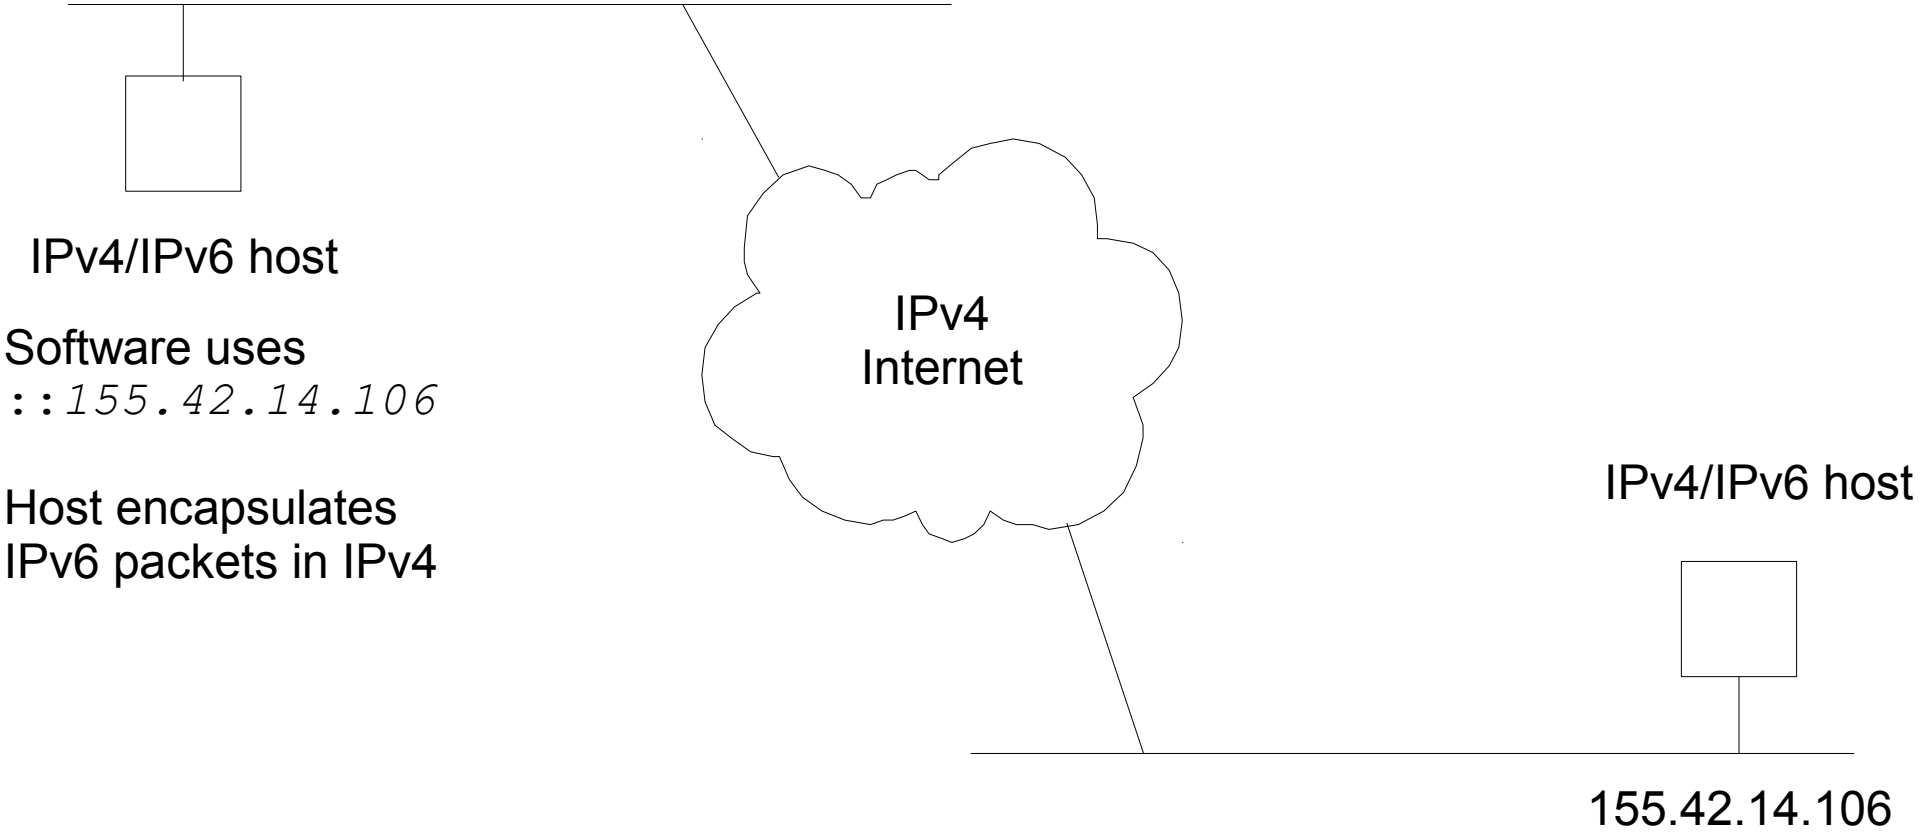

Automatic Protocol Conversion

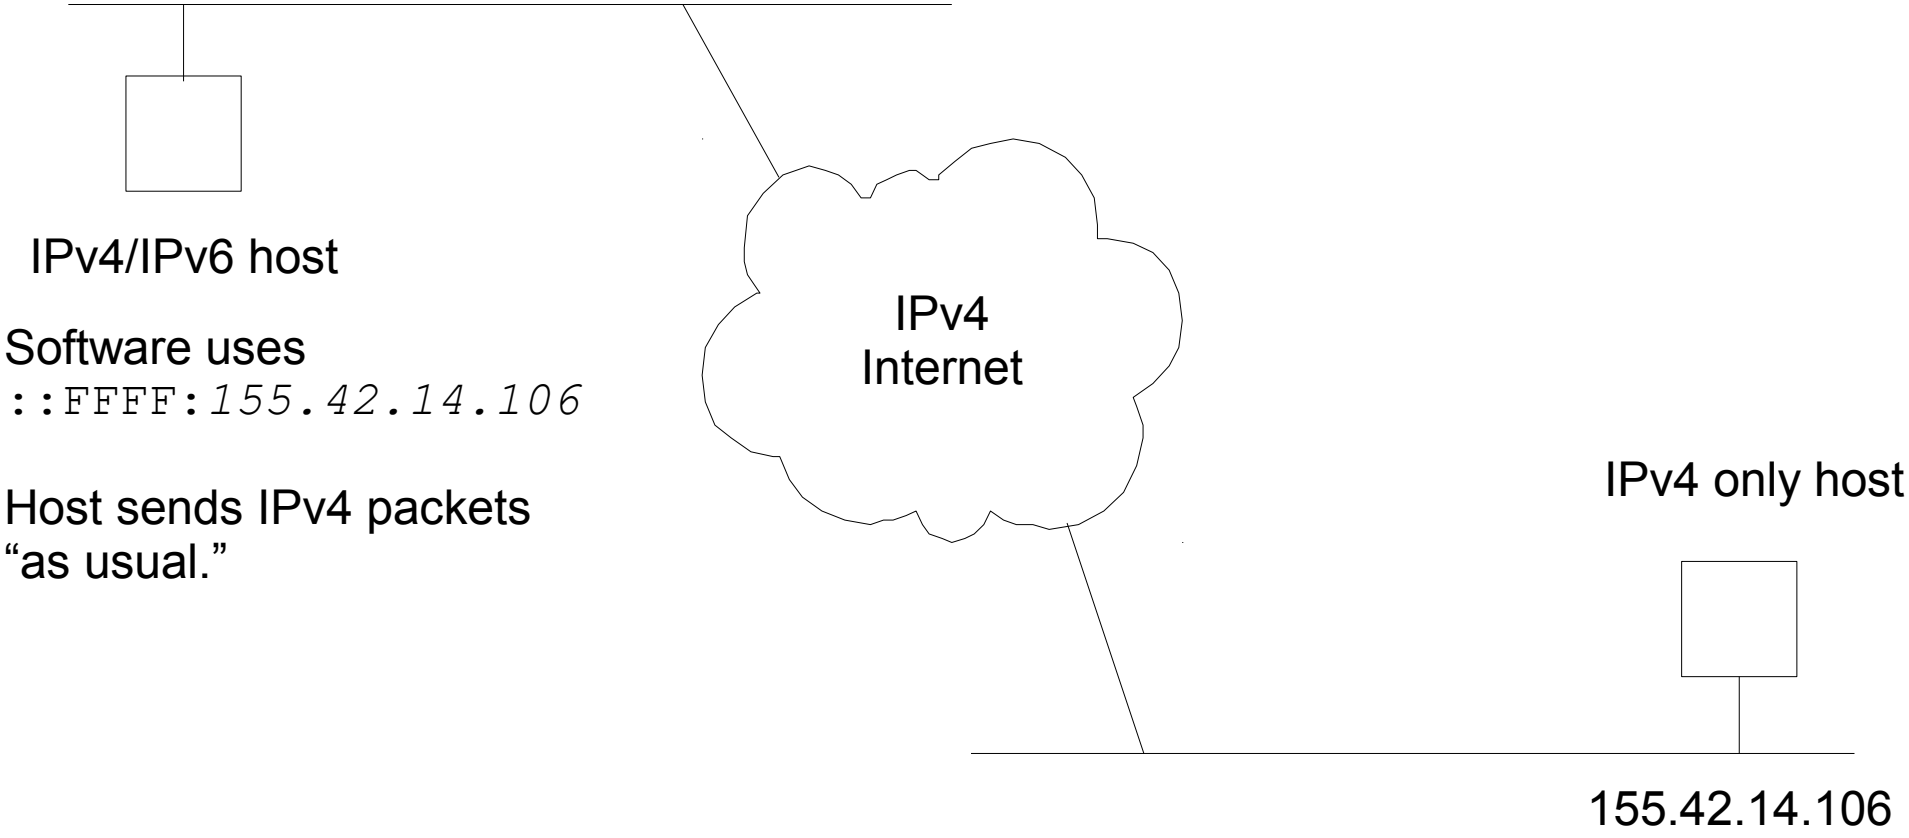

6-to-4 Tunneling

- RFC-3056
  - Allows IPv6 sites or hosts to communicate with each other over IPv4 infrastructure.
    - Does not require automatic or configured tunneling.
- Prefix 2002::/16
- A site with a valid IPv4 address gets prefix of
  - 2002:*IPv4-address*::/48
- Gateway machine does encapsulation, etc.

6-to-4 Example