Internet Protocol (version 4)

CIS-2151

Peter Chapin

Vermont Technical College

IP Addresses

- IPv4 addresses are 32-bit binary values.
  - Traditionally written as four octet values in decimal separated by dots.
  - 155.42.13.4 → 0x9B,2A,0D,04 → 1001,1011,0010,1010,0000,1101,0000,0100
  - The network software deals with these 32 bit addresses. The traditional appearance is just for human convenience.
- Addresses assigned to network interfaces (hardware). Not machines.
  - A machine can have multiple IP addresses if there are multiple interfaces.
  - An interface can also have multiple IP addresses.

Network Addresses

- All interfaces connected to the same link have IP addresses with a common *prefix*.
  - e. g. 155.42.13.27, 155.42.13.35, 155.42.13.178 all share a prefix of 155.42.13.xx
  - When an interface is configured you must provide both the IP address for that interface and who many bits form the prefix shared with other interfaces on that link.
  - 155.42.13.27/24. Here the “24” means the prefix is 24 bits long (i. e., the first three octets of the address).
  - Thus 155.42.13.00 is the “network address” (the address of the network link).
  - The count of prefix bits need not be a multiple of 8.

Netmask

- An alternative way to specify the prefix is with a *netmask*.
  - addr = 155.42.13.27, netmask = 255.255.255.0
  - Let's look at that netmask in binary:
    1111 1111 . 1111 1111 . 1111 1111 . 0000 0000
  - The 1 bits correspond to the part of the IP address that is the prefix (network address). The remaining bits for the *host address* on that network. With only 8 bits remaining there can be at most 256 hosts.
  - Another example: netmask = 255.255.192.0
    1111 1111 . 1111 1111 . 1100 0000 . 0000 0000
  - Now the first 18 bits form the network address and the last 14 bits form the host address. There can be at most 2^{14} = 16K hosts.

Vermont State University

- The entire VTSU system has been assigned a prefix of 155.42.0.0/16.
  - Every system on all the campuses thus has an IP address that starts with "155.42"
  - 2^{16} = 64K = 65536 hosts maximum in the entire VSC.
- The various campuses have been given more specific prefixes to work with… a process called *subnetting*.
- The IT department on each campus assigns appropriate, specific prefixes to the various links.
- The CIS server room and CIS lab in Williston is network 155.42.107.0/24. At most 2^8 = 256 hosts can exist on that network.

Broadcast?

- A special address with host ID all 1s is the *broadcast address*.
  - A packet transmitted to that address goes to all hosts on the same network (the same link).
  - 155.42.107.0/24 implies a broadcast address of 155.42.107.255.
  - Let's see how it works:

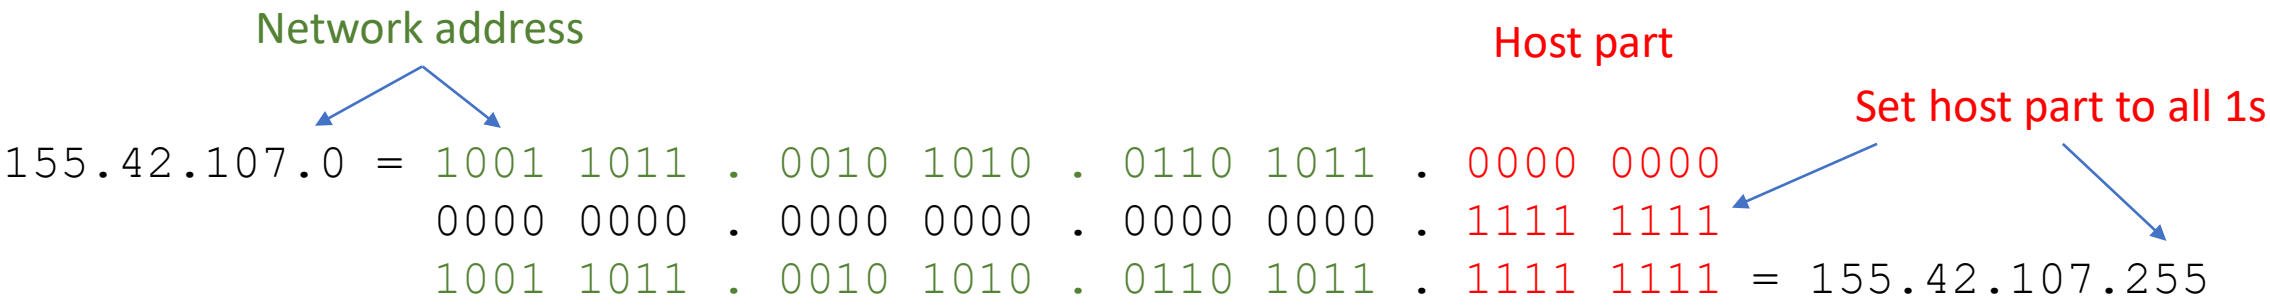

Lemuria's Interfaces

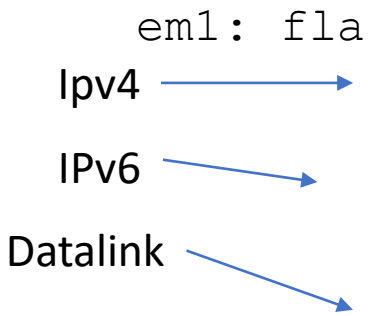


```
em1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 155.42.107.97  netmask 255.255.255.0  broadcast 155.42.107.255
        inet6 fe80::862b:2bff:fe65:afc4  prefixlen 64  scopeid 0x20<link>
        inet6 fd25:f376:7b60:1::1  prefixlen 64  scopeid 0x0<global>
        inet6 fd25:f376:7b60:1:862b:2bff:fe65:afc4  prefixlen 64  scopeid ...
        ether 84:2b:2b:65:af:c4  txqueuelen 1000  (Ethernet)

em2: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 10.0.0.254  netmask 255.255.255.0  broadcast 10.0.0.255
        inet6 fe80::862b:2bff:fe65:afc6  prefixlen 64  scopeid 0x20<link>
        inet6 fd25:f376:7b60:10::1  prefixlen 64  scopeid 0x0<global>
        inet6 fd25:f376:7b60:10:862b:2bff:fe65:afc6  prefixlen 64  scopeid ...
        ether 84:2b:2b:65:af:c6  txqueuelen 1000  (Ethernet)
```

Network Layer

- IP is a network layer protocol. What does that mean?
  - It's main concern is getting *protocol data units* (PDUs), usually called *packets* at the network layer to machines across multiple links.
- Data from the transport layer (one layer up) is stuffed into packets. Each packet is launched to its destination separately.
  - Packets traverse many links on the way to the destination.
  - Routers are specialized hosts with multiple interfaces that copy packets from one link to another.
  - Packets might take different paths across the network (but usually not).
- "Packet Switched"

Simple Network

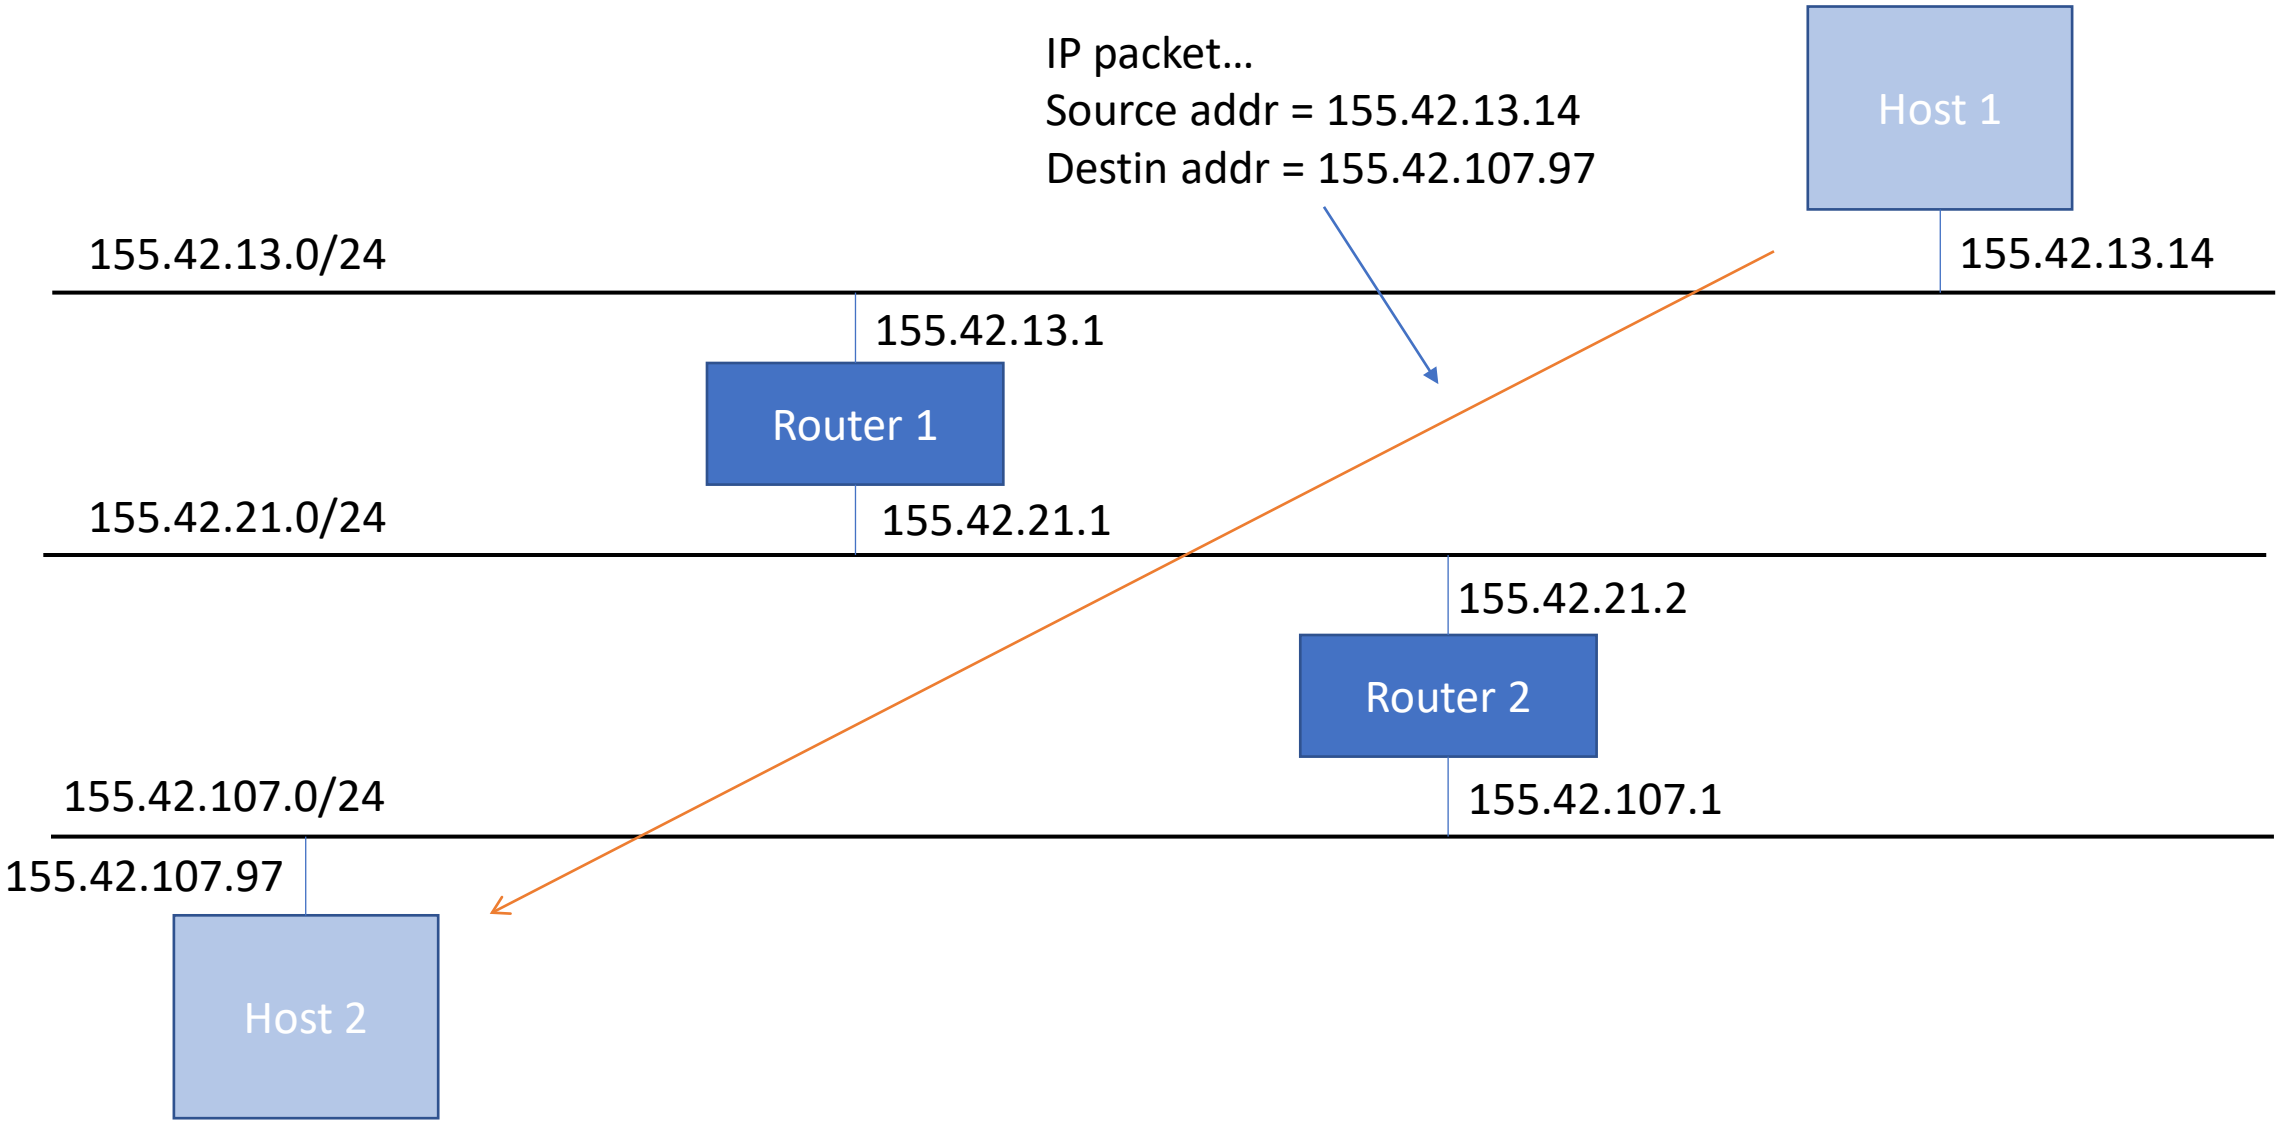

IP is Unreliable

- IP may not deliver packets properly
  - Some packets might never get delivered
  - Some packets might get delivered multiple times
  - Packets might get delivered out of order
- Isn’t this bad?
  - It simplifies router implementation
  - If a router is overloaded, it can simply throw packets away
- Reliability is implemented at a higher level (if needed)
  - TCP uses acknowledgements and sequence numbers to give reliability

Address Classes

- A largely obsolete feature, but still important to know about.
  - https://www.meridianoutpost.com/resources/articles/IP-classes.php
- Certain ranges of addresses are “private” and used for internal networks only. They are not routed on the Internet.
  - Packets that use internal addresses are dropped by Internet routers.
- Address classes have been replaced by Classless Interdomain Routing (CIDR).
  - IPv6 does not use address classes in the same sense.
  - Will discuss this more when we talk about routing.

IETF and RFCs

- The Internet Engineering Task Force (IETF) creates network standards for the Internet.
- The IETF produces Request For Comments (RFC) documents.
  - Some RFCs are informational
  - Some are experimental
  - Some are (or become) standards
- The RFCs never change once published (“immutable”).
  - If a standard evolves, a new RFC (with a new number) is issued instead of revising the original RFC. The RFC index links them all together.
- https://www.rfc-editor.org/

RFC-791: Internet Protocol

- RFC-791 describes IPv4 (published in 1981!)
  - In theory it is the most authoritative source.
  - However…
    - There are “update” RFCs that augment what RFC-791 says to account for changes since the original RFC was written.
    - Actual practice might not always follow the RFC exactly (for various reasons).
- https://www.rfc-editor.org/rfc/inline-errata/rfc791.html

RFC-792: Internet Control Message Protocol

- ICMP is used for “control messages”
  - Reporting errors
  - Controlling data flow
- ICMP has many message *types* and each type may have several *codes*
  - The type reflects the nature of the control message
  - The code further refines the message
- Many ICMP messages are not used or rarely used
  - They introduce their own problems
  - They can be security hazards
- Some ICMP messages are commonly used

What are the ICMP Message Types?

- Well… let's look at RFC-792!

Common ICMP Messages

- Echo Request/Reply
  - Used by ping to check if a remote host is “alive”
  - Sometimes administratively disabled (no replies sent) for security reasons
- Time limit Exceeded
  - Used by traceroute
    - Sender sends an Echo Request with a small TTL
    - When the TTL “expires” an ICMP Time Limit Exceeded message is returned
    - Sender notes the IP address that sent the Time Limit Exceeded message
    - Sender sends an Echo Request with the TTL set to one larger to probe one step farther

Traceroute Example

```
pchapin@lemuria:~$ traceroute www.microsoft.com
traceroute to e13678.dscb.akamaiedge.net (23.200.197.152), 64 hops max
  1   155.42.107.1  0.717ms  0.683ms  0.684ms
  2   10.254.234.1  0.155ms  0.148ms  0.079ms
  3   10.254.4.1  0.444ms  0.374ms  0.362ms
  4   10.42.9.1  2.466ms  1.755ms  5.215ms
  5   10.254.0.2  0.573ms  0.490ms  0.422ms
  6   155.42.255.3  0.641ms  0.552ms  0.606ms
  7   216.238.175.93  1.315ms  1.120ms  1.095ms
  8   66.152.98.137  2.577ms  1.227ms  1.235ms
  9   66.152.98.5  5.072ms  4.956ms  5.169ms
 10   66.152.97.41  5.048ms  4.745ms  4.496ms
 11   198.32.118.222  9.061ms  9.361ms  9.382ms
 12   192.168.225.181  8.545ms  8.390ms  8.358ms
 13   192.168.230.143  8.265ms  8.164ms  8.284ms
 14   192.168.245.17  8.704ms  8.659ms  8.643ms
 15   23.200.197.152  8.706ms  8.412ms  8.316ms
```

Private-use addresses?
These are internal to VTSU

More private-use addresses?
Perhaps these are internal to MS

There’s a problem…

- The underlying layer (data link) doesn’t know anything about IP
  - Doesn’t know how to deal with IP addresses at all. Never heard of them.
- So, how does the data link layer know where to send frames?

Ethernet

- Let’s look briefly at Ethernet, **a** common data link technology
- Ethernet has its own addressing scheme
  - 48-bit addresses called, variously:
    - Ethernet addresses
    - MAC (Media Access Control) addresses
    - Hardware addresses
    - Link layer addresses
    - etc!
- Lemuria’s em1 interface has a MAC address of:
  - 84:2b:2b:65:af:c4

Ethernet Address Format

- Ethernet addresses have two parts
  - Upper 24 bits (3 octets) is a vendor identifier
  - Lower 24 bits (3 octets) is a vendor generated
- Let's look at Lemuria's em1 MAC address again:

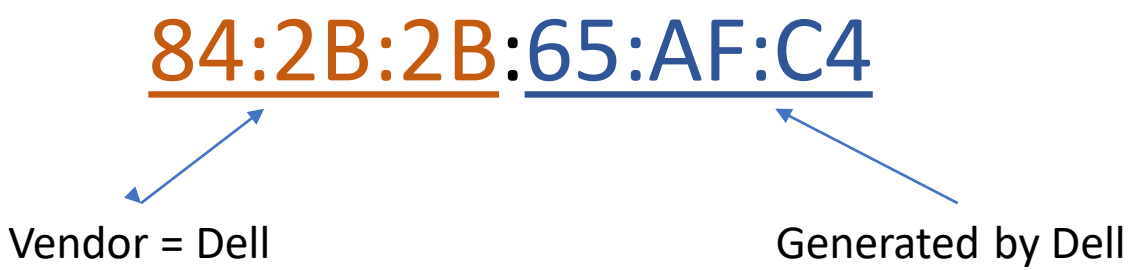

- Thus Ethernet addresses are globally unique!

MAC Address Assignment

- Although not commonly done, it is possible to manually assign a MAC address
  - A past VTC network administrator once configured the interface on his office machine to have a MAC address of DEAD:BEEF:CAFE.
- Fun game: How many words can you spell using only the letters A-F?

Sending IP Packets 101

- You have an IP packet. You want to send it…
- Is the packet going to a machine on the same link?
  - If yes… send it directly to that machine.
  - If no… send it to your configured “default gateway” (aka router) that is attached to the link.
  - No default gateway? Oh well, can’t communicate.

On The Same Link?

- Example: Lemuria at 155.42.107.97 wants to send to 155.42.107.85.
- Is that on the same link? Use the netmask!

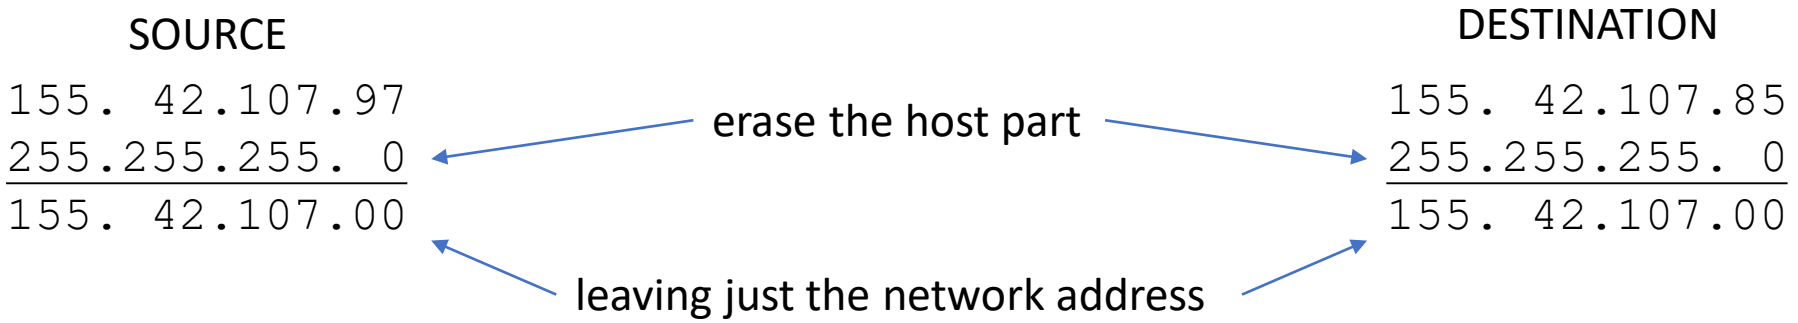

- In this case the networks are the same, so destination on same link!

On Different Link?

- Example: Lemuria at 155.42.107.97 wants to send to 155.42.113.85.
- Is that on the same link? Use the netmask!

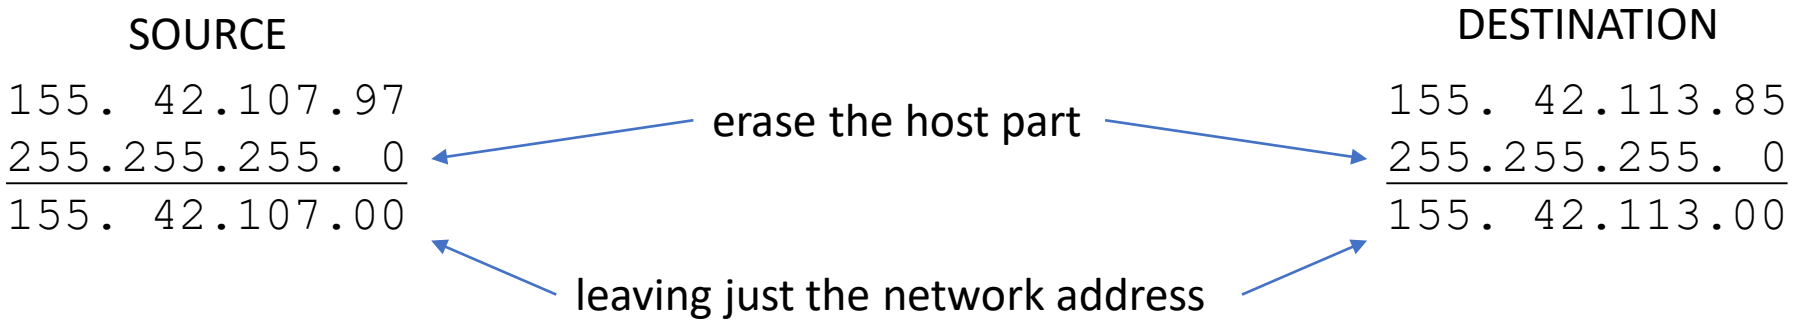

- In this case the networks are different, so send to the default gateway

Default Gateway?

- In Lemuria's case the default gateway is 155.42.107.1
  - A router connected to the 155.42.107.0/24 network managed by the IT department.
  - The x.x.x.1 address is commonly used for the default gateway, but not at all necessary. Any IP address on the link is fine, as long as all the hosts on that link know about it.
  - The default gateway is involved for any traffic that goes off the link.

Where Does Ethernet Come In?

- Any traffic on the link, either directly to the destination or to the default gateway, must be sent in an Ethernet (for Lemuria) Frame.
  - But what MAC address should be used??
- We need to “resolve” an IP address to a MAC address.
- We need “Address Resolution Protocol”… **ARP**

Send an IP Packet… The Steps

1. Send to the destination IP directly if on the same link; otherwise send to the default gateway for all other destinations.
2. If the MAC address associated with the IP is known, use it. Done.
3. Otherwise… broadcast an ARP request that says, “Hey! Who has such-and-such an IP address? Send your response to me.”
4. The machine with that IP address sends a frame to the requester containing the its MAC address.
5. The requester caches the IP -> MAC mapping in the “ARP cache.”
6. Sends a frame to the discovered MAC address containing the IP