Overview of Internet Protocol

Ben Schultz

UNH InterOperability Lab

June, 2001

Overview

- Understanding the Internet Protocol Problem
- Solving the Problem
- Overview of Internet Protocol
- Alternative Technologies
- Routing and how it works

Layer 2 Review

- MAC Addressing is a 6 Byte addressing Scheme
- This is a Local Area Network (LAN) solution – it solves the problem of direct communication
- What if you need to communicate with a node in a network beyond the LAN on which your station resides?

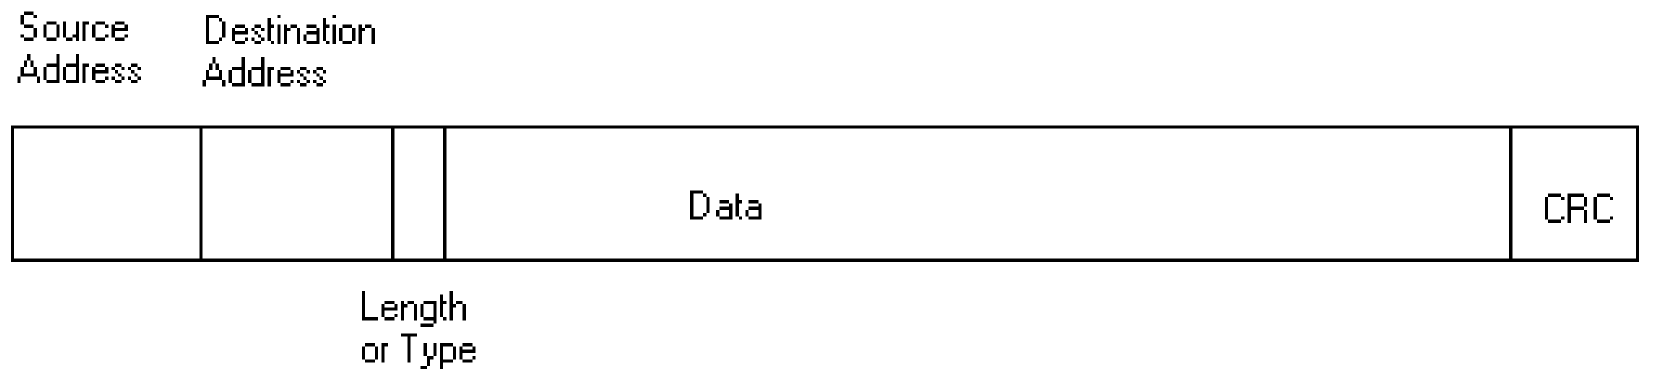

Inside The MAC Frame

- The Internet Protocol addressing information is contained in the “data” field of the MAC frame
- This field is referred to as “data” from the Layer 2 perspective because the MAC entity views this information as data

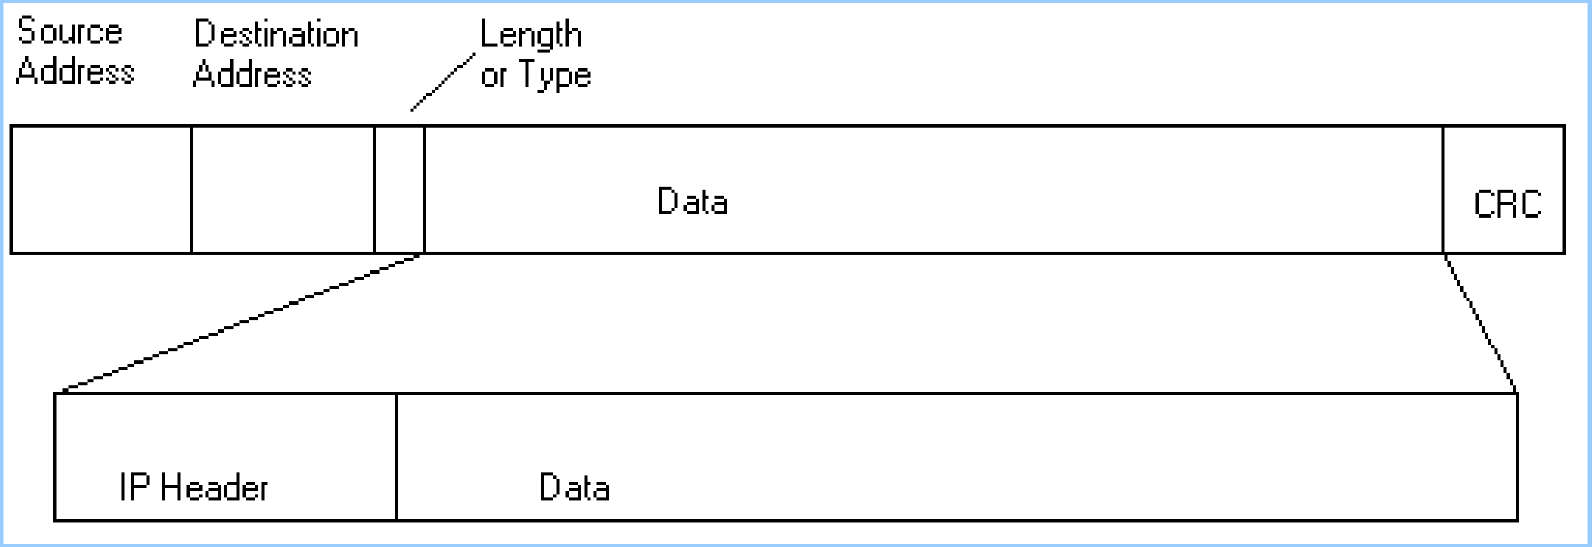

Network Layer Addresses

The Network Layer provides end to end transport. This addressing scheme allows routing of data packets between Local Area Networks (LANs).

The transmission scheme is connectionless. This means that station A sends a packet to station B without checking if station B received the information.

Internet Protocol

The IP addressing scheme consists of 4 Byte addresses. Unlike MAC addresses, these addresses are represented in decimal numbers, with dots separating the bytes.

Some Example Addresses:

132.177.121.222

10.10.10.20

How does IP Interface with Layer 2?

Through the Layer 2 interface, IP sends a broadcast query message.

This asks “Who has the unknown IP address of 10.10.10.20?”

The response message contains the proper MAC address and the station can now send the IP data packet to the correct MAC address.

This is called Address Resolution Protocol or **ARP**

Beyond the LAN

If a device must send to an IP address not on the local area network it must ARP the router, obtaining the router’s MAC address.

The IP Packet is then transmitted to the proper destination address with the MAC address of the router.

The router then forwards this packet to the destination network, using ARP if necessary.

Address Resolution Protocol

Or ARP, is the interface between Layer 2 and Layer 3. Remember MAC addresses?

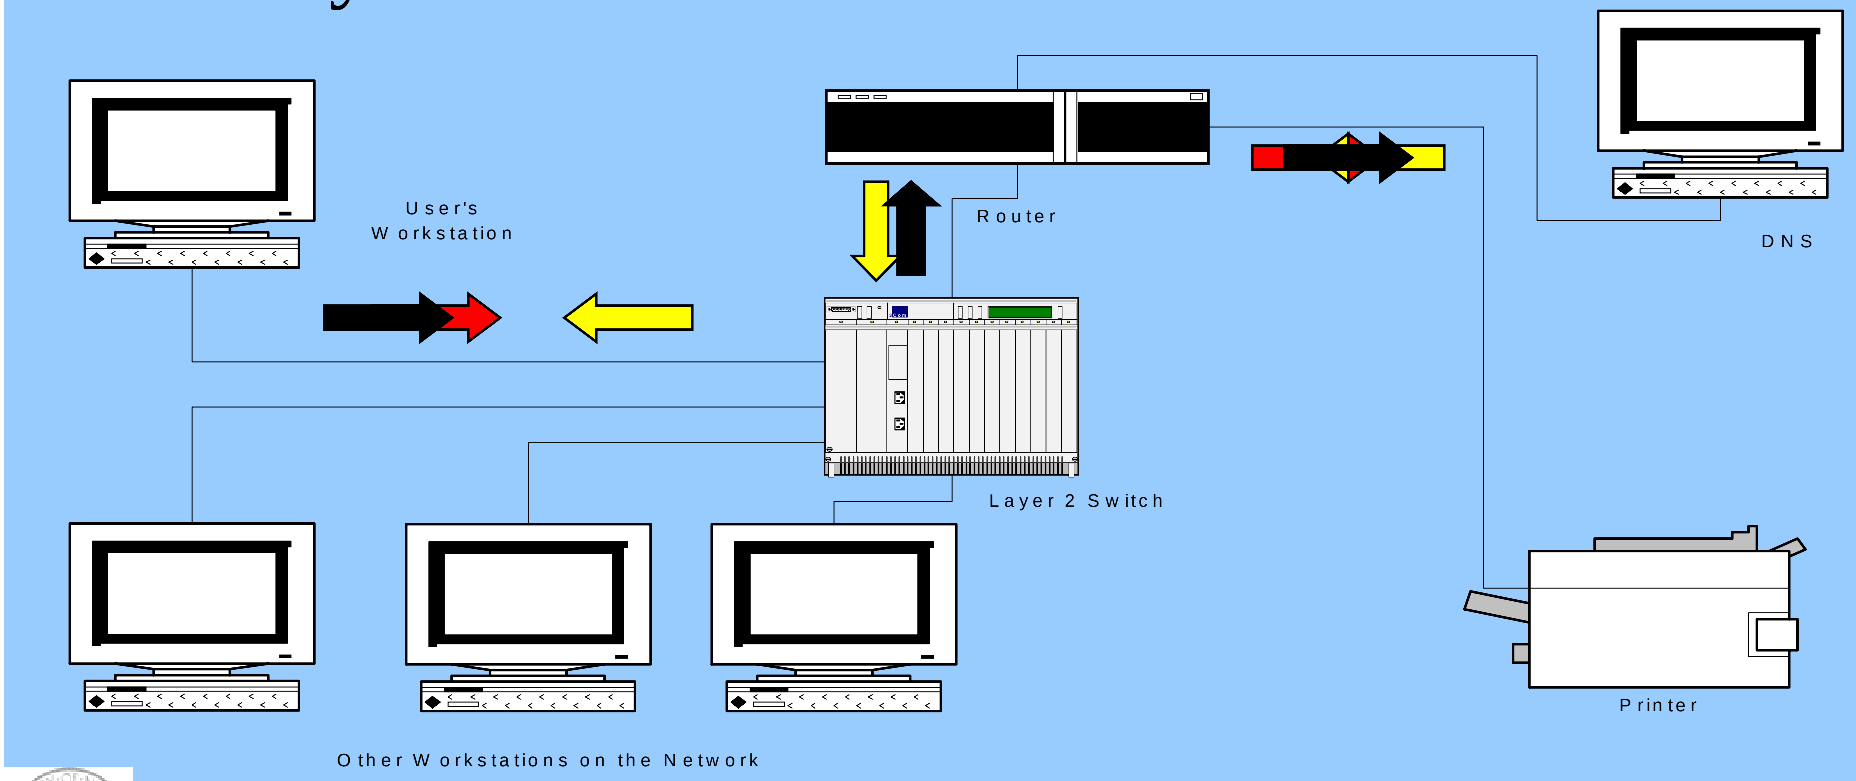

IP Header Information

Like the Layer 2 header, the IP header includes source and destination addresses.

The header includes a checksum that is run on the header and has a similar function to the CRC at layer 2.

The IP Header includes a Protocol Type field, similar to the Type field in a Layer 2 packet.

A total length field of the data and header information

IP Header Information That Differs from Layer 2

A Type of Service Field for routing service information

Fragmentation Information

Extra Options

So… Why do we need IP?

Layer 2 addresses only can communicate within LANs.

Layer 2 wide area networking technologies use different addressing schemes, so how would the UNH network be able to communicate with the CNN network?

Internet Protocol, an Example

Before the specific details are disclosed, you should have a reference point as to what Internet Protocol includes.

The following problem is experienced:

•A user wishes to print a file from his host system to a printer driven by another system on the same network.

•The user types the print command on his system with the file name and receives the response, " 132.177.128.4 unknown", and the file does not print.

What Now?

We need to figure out:

1. What happened to deny the user’s computer access to the printer?
2. What is wrong?

Environment and Parameters

a. Network structure – The network topology and configuration

b. Protocols used

c. Physical connections

d. Host systems

e. Printer system and the commands to make it work

Network Topology

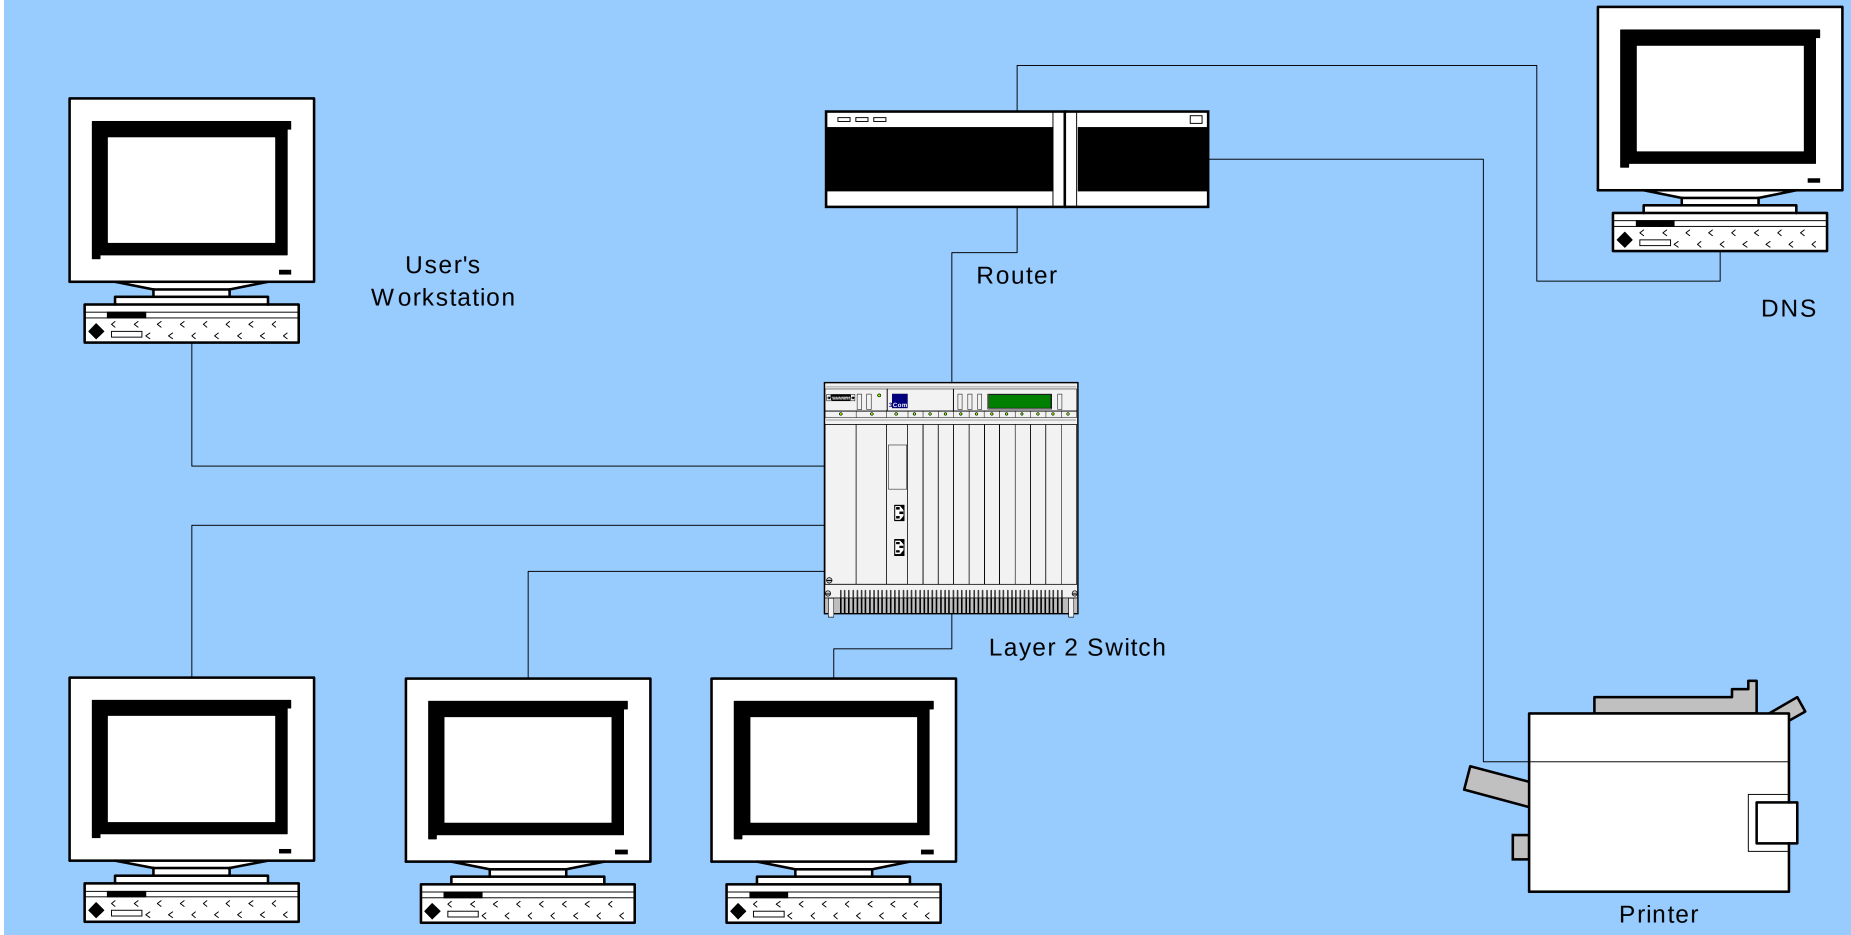

Other Workstations on the Network

How Does Remote Printing Work?

- Start from the user request from the UNIX machine …. Lpr -P hp4ip
- hp4ip is the printer name
- Lpr command checks a queue by the name of “hp4ip”, submits a request to the lpq daemon.

A New Process

- Lpq daemon spawns a new process to deal with the print job.
- If the printer entry is in the /etc/printcap file, the 2nd daemon reads the printer definition. If there is an "rm" line in the definition, lpd knows to send the job to a remote machine.

The Printcap File

- A typical definition may look like this:

```
hp4ip|lp0|hp printer:\
        :if=/var/spool/printers/hp4ip-filters/hp4ip-ps:\
        :rp=hp4ip:\
        :rm=fleck.iol.unh.edu:\
        :ct=network:lp=:\
        :sd=/usr/spool/printers/hp4ip:\
        :lf=/usr/adm/printers/hp4ip-errs
```

Name Resolution

The next thing that happens is that the name “fleck.iol.unh.edu” needs to be mapped to a network address. How does this happen?

- The host system has a configuration file called /etc/resolv.conf and associated libraries for domain name service (DNS).
- The resolver queries the domain server for iol.unh.edu and gets back an address (132.177.118.20)

Passing the Job to the Printer

After the lpd daemon determines where to send the file it will pass it through a filter (defined by the "if=" tag) to format the print job properly.

Once that is done, it will send the formatted print job to the remote machine and spool it to the proper queue on that system (as defined by the "rp=" tag).

The print server may then take the job (and apply it to its own filter) and then send it directly to the printer.

This did not happen, instead an error occurred.

The Problem Revealed

- The server system, fleck, is also running name service, and must do a lookup in inaddr-arpa to map an IP address to a hostname so it can accept a connection from the spooling system. The resolv.conf file points to 132.177.123.46, the primary name server.
- The resolv.conf file points to 132.177.123.46 as the primary name server.
- But, the configuration file named.rev had an error in it!

The Solution

- The address field for the host sending the file was 4.128.177.132
- It should have been 41.128.177.132

There is of course no entry for this bad address, so the lookup failed. Therefore, the message was sent “132.177.128.4 unknown”

Reasons for Internet Protocol

1. Provides a name to address mapping with the Domain Name Service
2. Provides a flexible address hierarchy
3. Provides a way to route the data packets between Local Area Networks

Domain Name System

Network addresses are fine for computers, but people have an easier time remembering names. Names for each station must be unique on the same IP network.

IP networks also have names, like unh.edu. There can be the same name on different IP networks.

The distributed database that keeps track of the name to address mapping is called DNS.

DNS Hierarchy

- Each piece of the name gives more information about the station. From right to left the information is more and more specific.

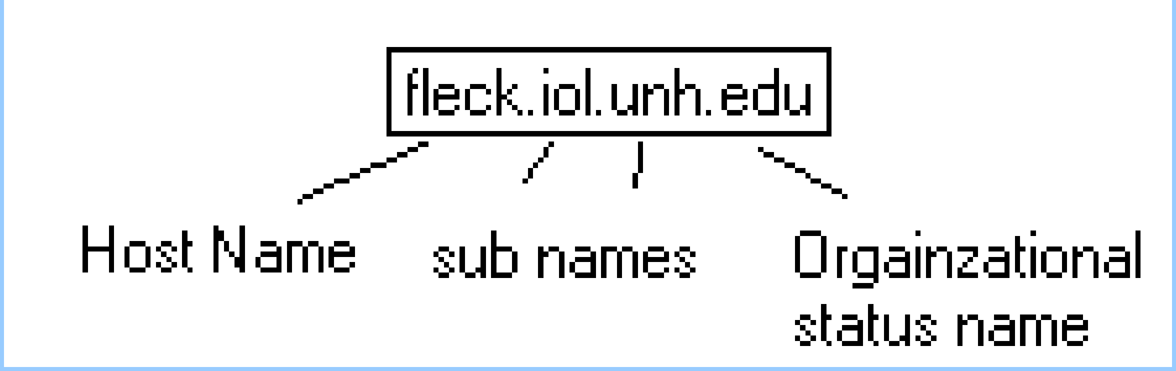

Top Level

The top level domains include the following:

.edu for educational institutions

.com for commercial organizations

.org for non-profit organizations

.gov for government organizations

.net for network provider organizations

There are also country domains at the top level, such as .us for the United States, .fr for France, etc.

Servers in the DNS

DNS is the translator between host/network names and IP addresses.

No one server can keep track of this mess – hosts are added, changed and deleted at a constant rate.

The Solution

- Each organization must have at least 2 DNS servers.
- Each server contains a list of host addresses and names for that organization.
- The secondary server backs up and checks the primary name server from time to time.

How it works

Say your computer is attempting to connect to hera.micro.biology.wisc.edu

- This results in your system sending a query to the local DNS server.
- The local DNS server then queries the root .edu server asking for the IP address of the DNS server for .wisc.edu.
- The local DNS server then queries the wisc.edu name server for biology.wisc.edu.
- Once it obtains the IP address for the biology.wisc.edu, the local DNS server queries the biology.wisc.edu name server for micro.biology.wisc.edu

The Tree

When the local DNS server obtains the IP address for micro.biology.wisc.edu it can then query that server for the proper IP address.

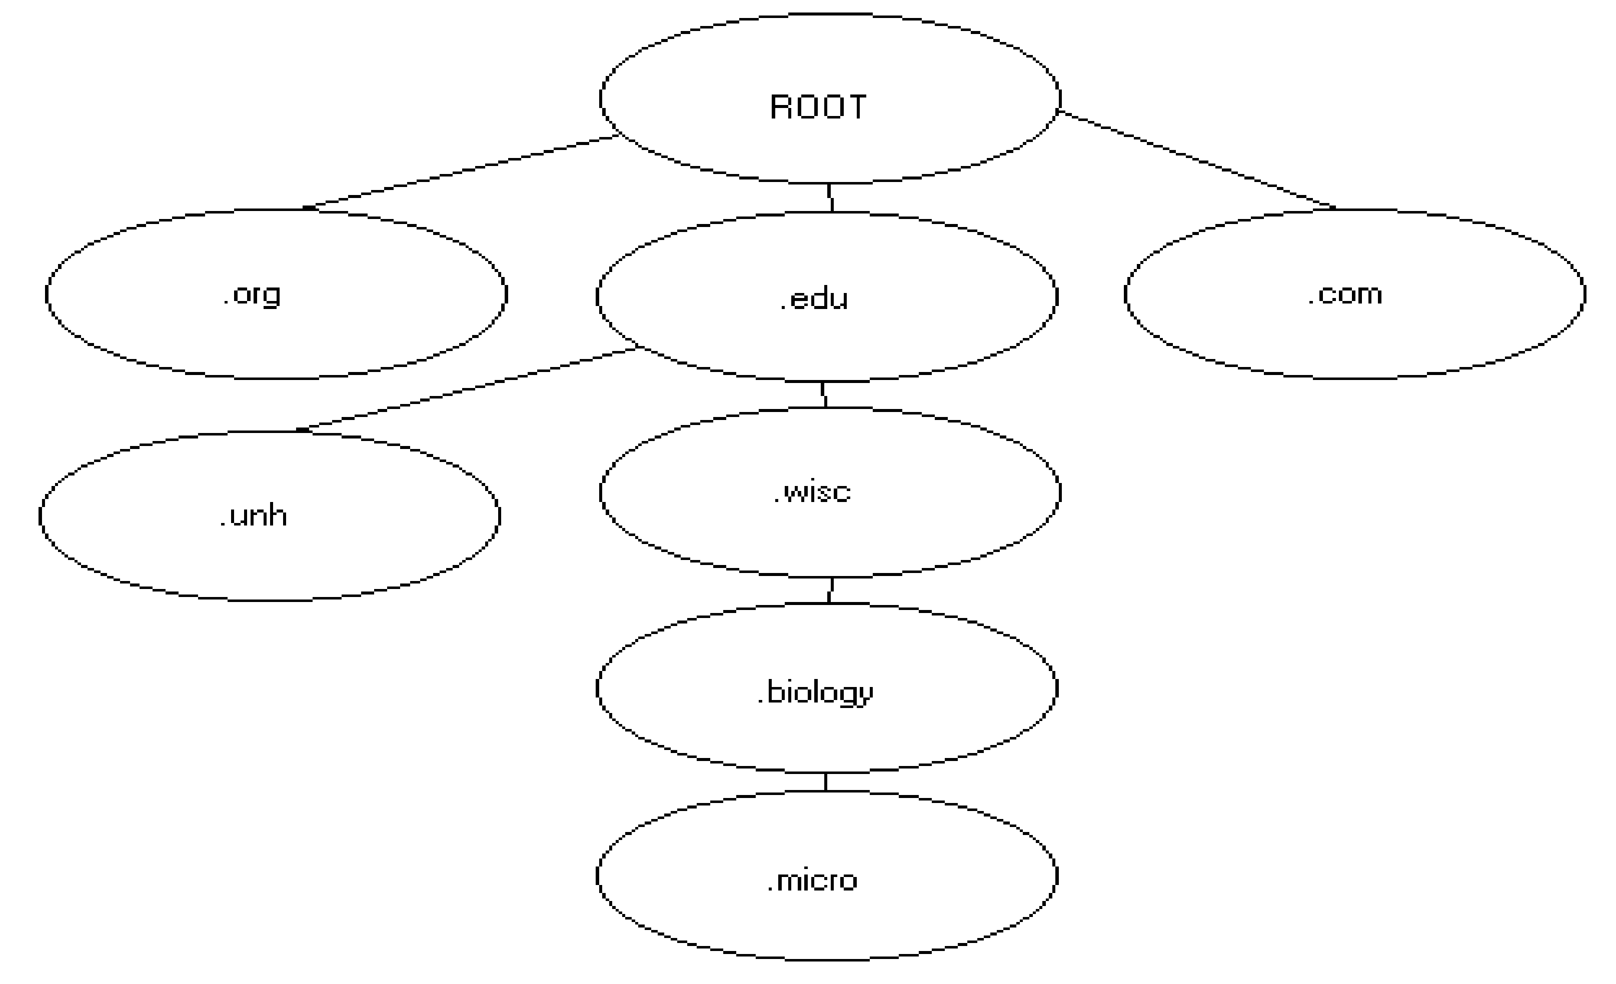

Address Hierarchies

- As stated earlier, IP addresses are 4 byte addresses, normally represented in dotted decimal format. 0.0.0.0 and 255.255.255.255 are the maximum and minimum values that can be assigned.
- Originally there were 5 classes of addresses defined

Address Classes

Network Class	Address Range	Networks in each class	Maximum number of hosts on a network
Class A	0.0.0.0 to 127.255.255.255	126	Over 16 Million
Class B	128.0.0.0 to 191.255.255.255	16,384	65,534
Class C	192.0.0.0 to 223.255.255.255	2,097,152	254
Class D	224.0.0.0 to 239.255.255.255	Reserved for Multicast Protocols	N/A
Class E	240.0.0.0 to 247.255.255.255	Reserved for future use	N/A

Network Classes and Subnetting

- In the above scheme, distinguishing which class a network was in was as easy as checking the first byte.
- But what if you want have a Class A address block and you do NOT want 16 million stations on your LAN??
- You must divide your allocated area into subnets.

Subnets – how they work

The division is called Subnetting, the overall network consisting of smaller sub-networks.

All IP addresses contain a network piece and a host piece. This is designated by something called a subnet mask. Subnet masks are usually notated as dotted decimal, like IP addresses. A common subnet mask is the equivalent of a Class C address: 255.255.255.0.

Subnet Mask Example

Here is an example of a Host address (top) and a subnet mask (bottom). The mask tells the host what stations are on the local area network. In this example, all stations that have address prefixes of 132.177.118.X are on the local area network. Thus the host can send frames with a direct layer 2 address paired with the destination IP address.

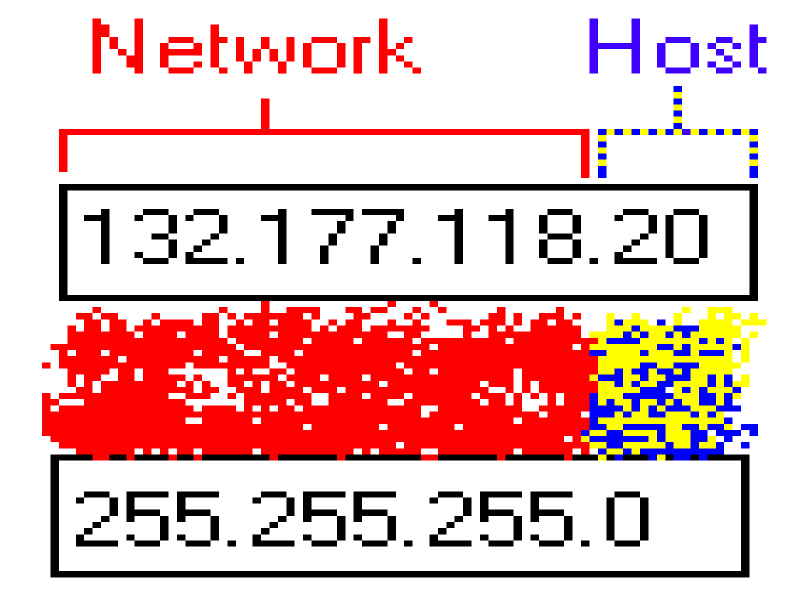

Outside the Subnet

What occurs when the host wants to send to an address outside 132.177.118.X? It will send the data with the desired IP address to the MAC address of the local router or **GATEWAY**.

When configuring a host device, there are 3 direct parameters you must configure – IP address, subnet mask and gateway.

Broadcast Addresses and Router Conventions

- Like layer 2, there is an IP broadcast address for a subnet, but NOT for the entire internet. In the above example it would be 132.177.118.255. The broadcast address is the maximum address number permitted by the subnet.
- The minimum address permitted by the subnet (in the above case 132.177.118.0) is reserved.
- The router is usually one more than the minimum address permitted by the subnet (132.177.118.1). This is not required, but it is usually the case.

BREAK

Introduction to Routing Concepts

How do the packets in the above examples get to their destinations across the internet? How are the paths that they travel in the network optimized?

Routing exists to solve these problems.

Distance Vector Routing

- Each network is assessed a cost that it takes data packets to travel across it. Slow networks usually have a higher cost, fast networks usually have a lower cost.
- Each router must remember the distance from its networks to each possible destination.

Distance Vector Example

Ignoring routes with more than 4 routers, there are multiple ways to travel from source to destination:

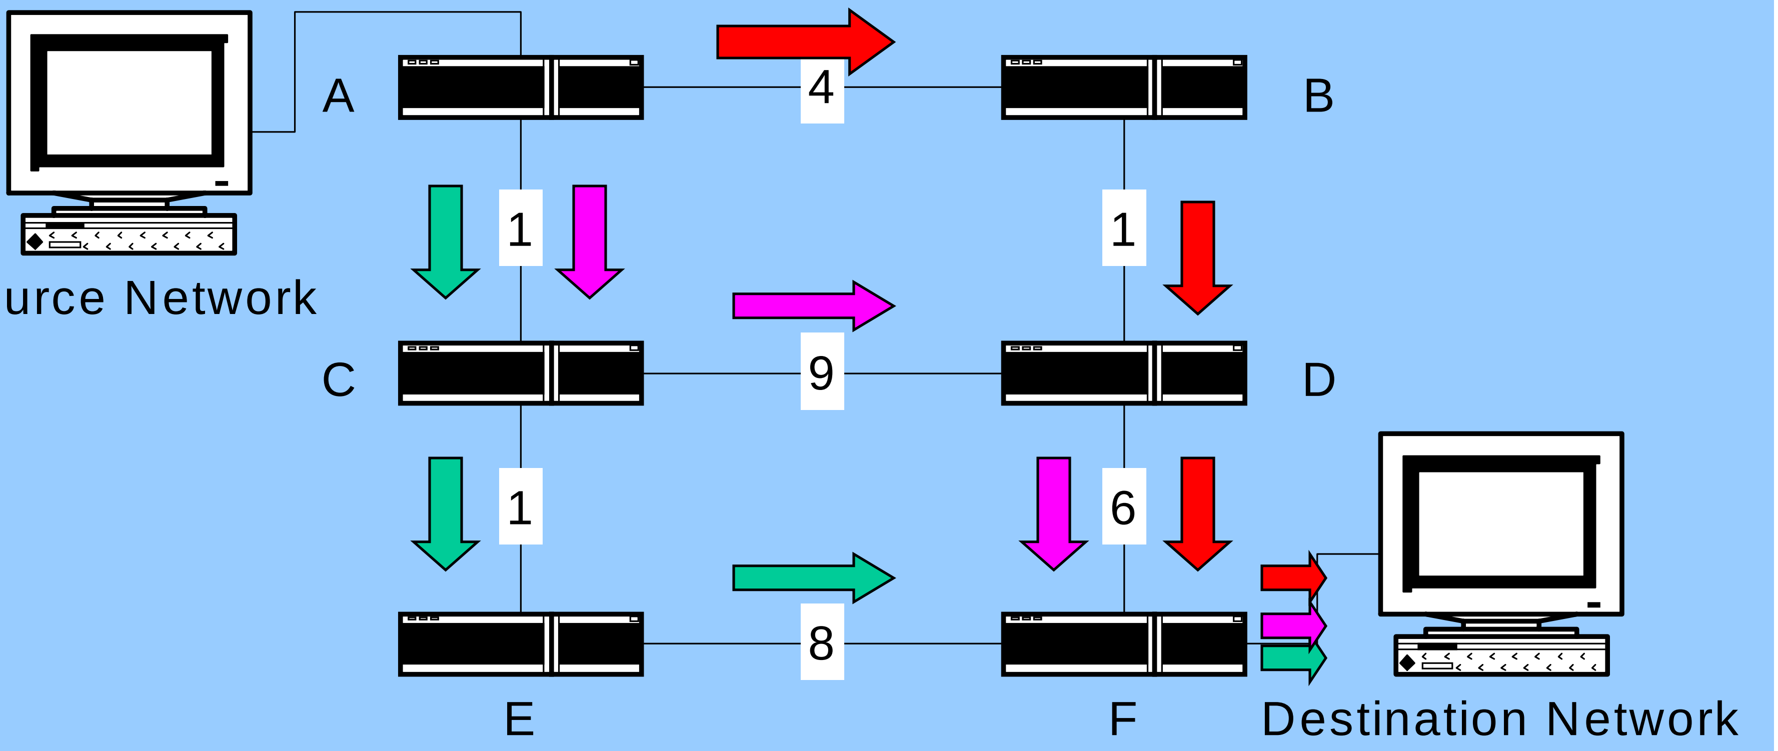

Routing Path Cost

The Path cost is passed around the network and incremented at each router to properly display the distance to each network.

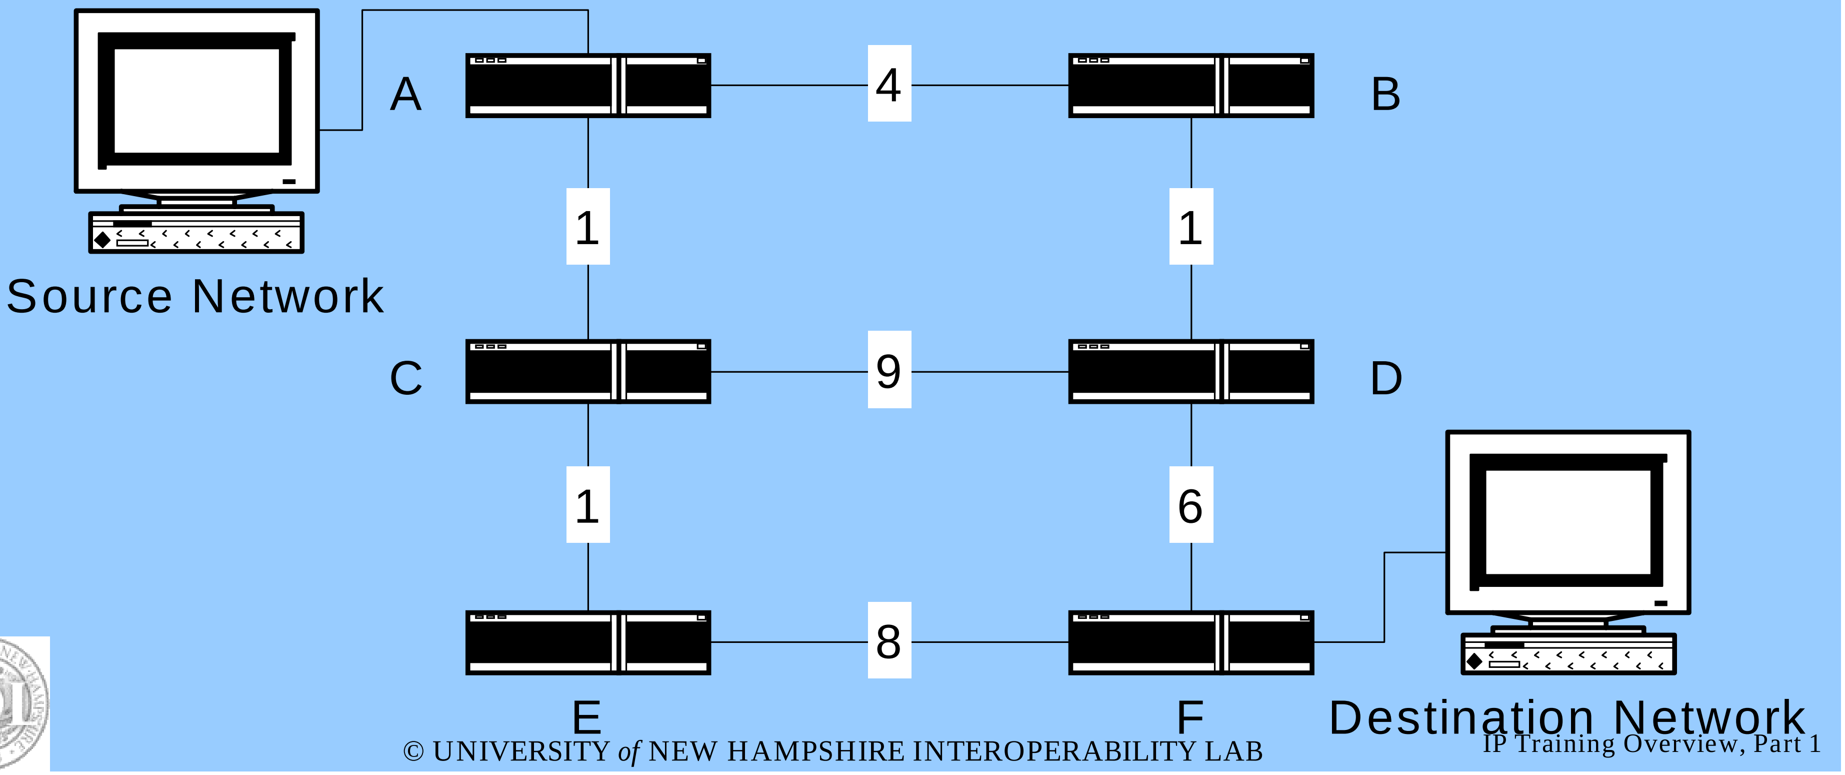

Packet Forwarding

Router A sees the following routes to the destination network: Path 1 to Router B with a cost of 11. Path 2 to Router C with a cost of 16. Path 3 to router C with a cost of 10. The packets are forwarded out the bottom interface to Router C.

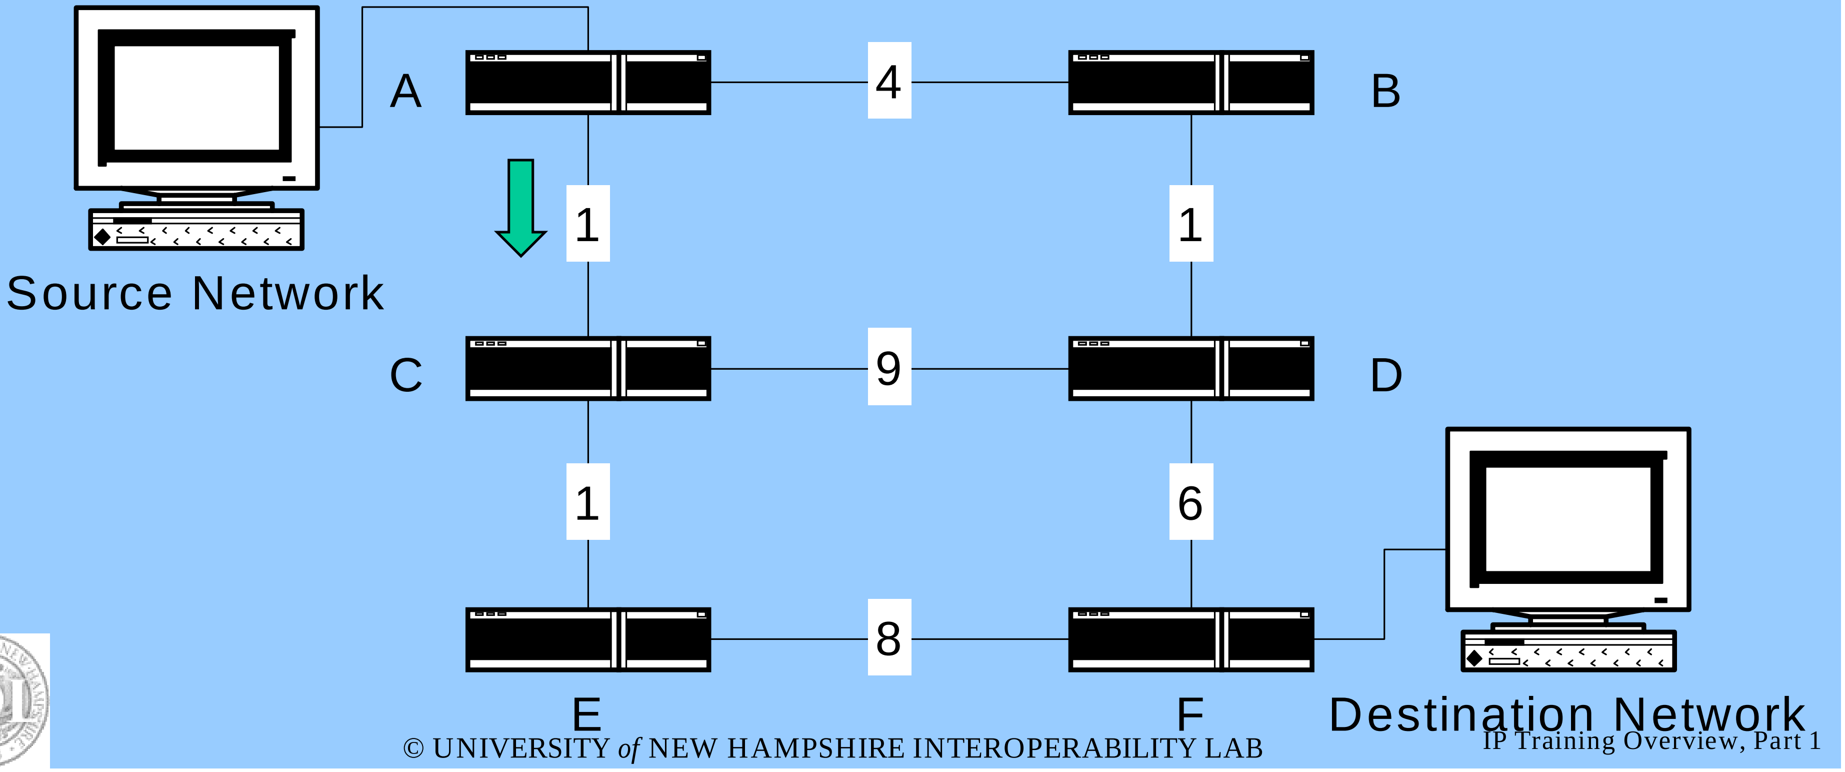

How Do Routing Tables Work?

On a more simple example, Router A is interested in making a routing table. As each network is advertised, A increments the advertisement with its connected networks. For example, the routes it receives from Router B are incremented by 4 and the routes it receives from Router C are incremented by 1.

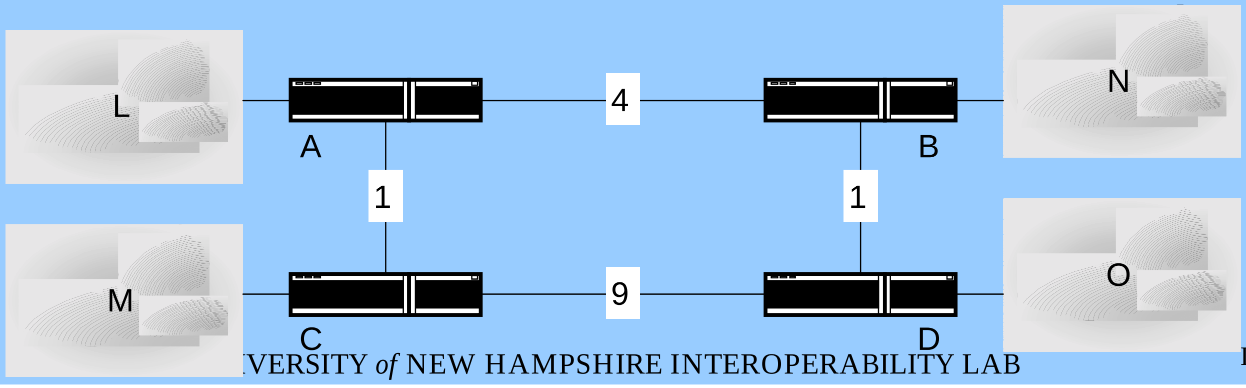

Routing Table for Router A

Destination	Cost	Next Hop
Network L	0	None
Network M	1	Router C
Network N	4	Router B
Network O	5	Router B

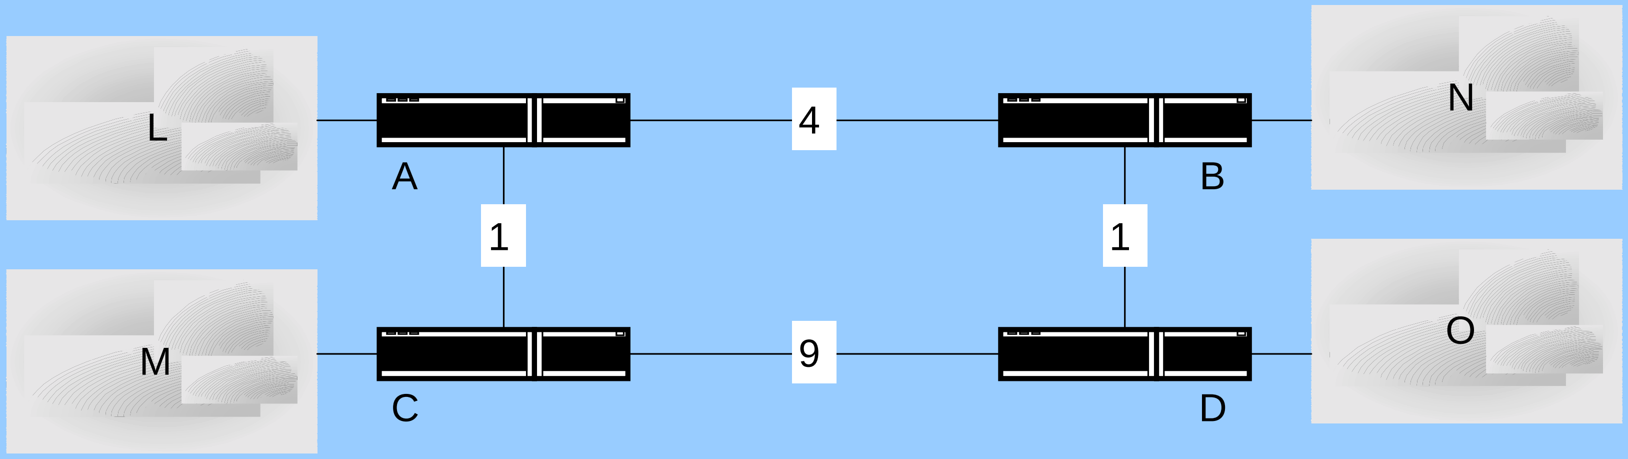

Link State Routing

- Like Distance Vector Routing, each network is assessed a cost.
- Each router must meet the neighboring routers.
- Each router must transmit the status of each network they are connected to and the cost.
- The above status reports are used to create a complete map of the network topology.

Link State Example

The routing tables and calculations as done by Router A.

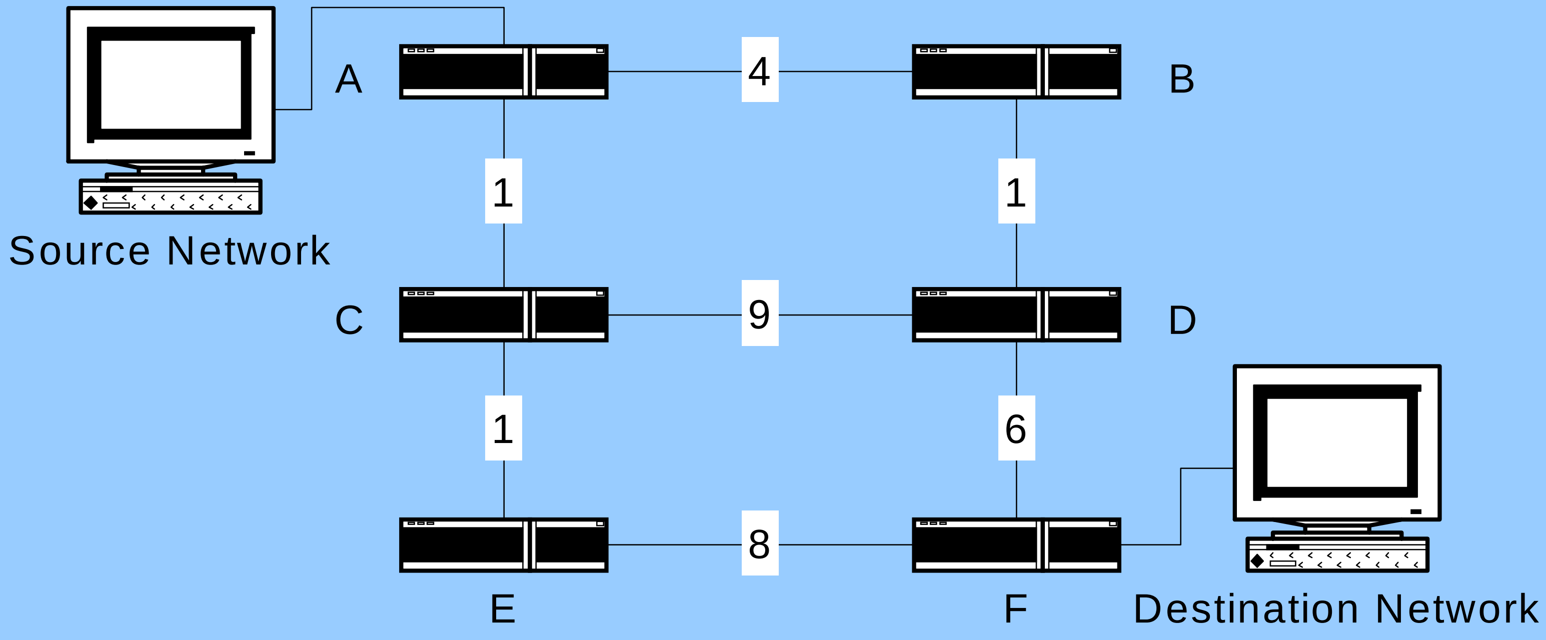

Router A’s Link State Database

Router A	Router B	Router C	Router D	Router E	Router F
B 4	A 4	A 1	B 1	C 1	D 6
C 1	D 1	D 9	C 9	F 8	E 8
		E 1	F 6		

Once the link state database is constructed, the router chooses the best path, with the PATH status, or investigates a possible best path with the TENT status. Once this is done the forwarding database is constructed from the results.

Dijkstra’s Algorithm

- The Router starts by placing itself in PATH
- It examine’s A’s LSP, adding B and C to TENT.

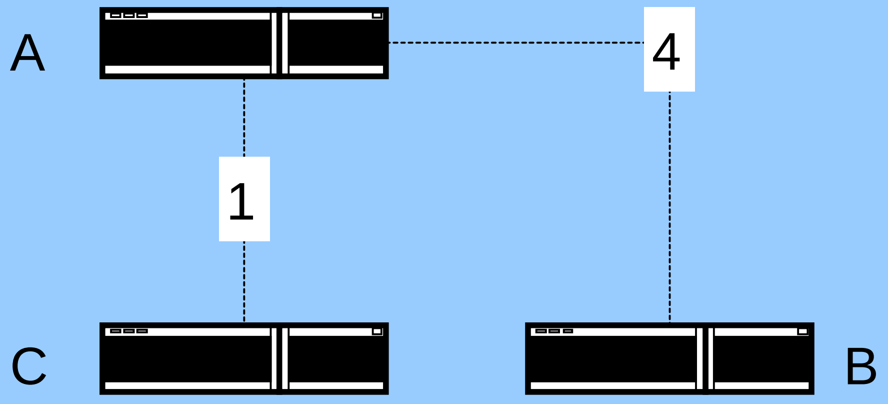

Routes Calculated

The router then places C in PATH and examines its entry in the link state database.

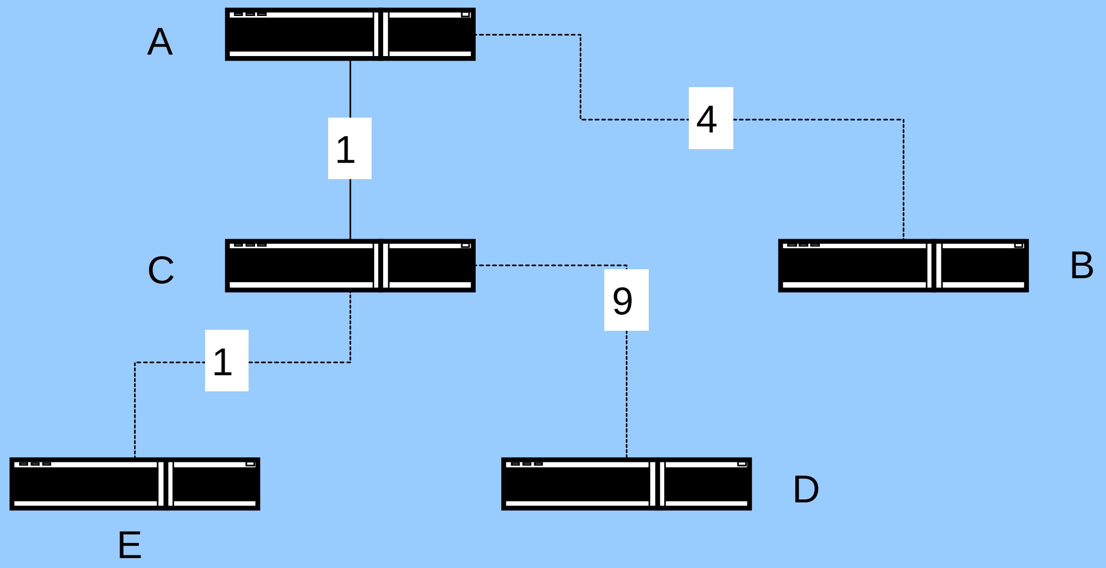

Routes Calculated

The Router then places B in PATH and examines its entry in the link state database.

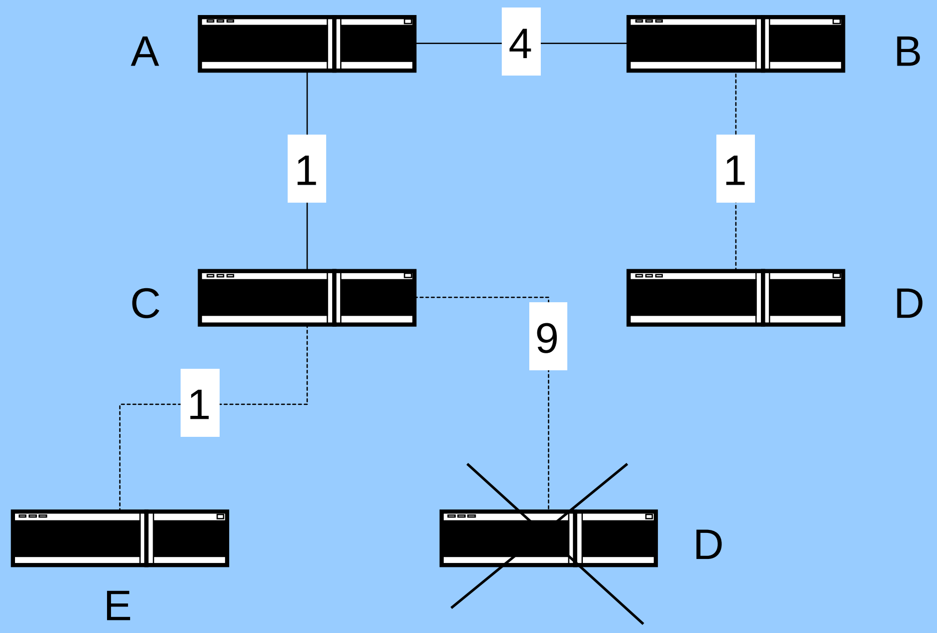

Routes Calculated

The Router then places E in PATH and examines its entry in the link state database.

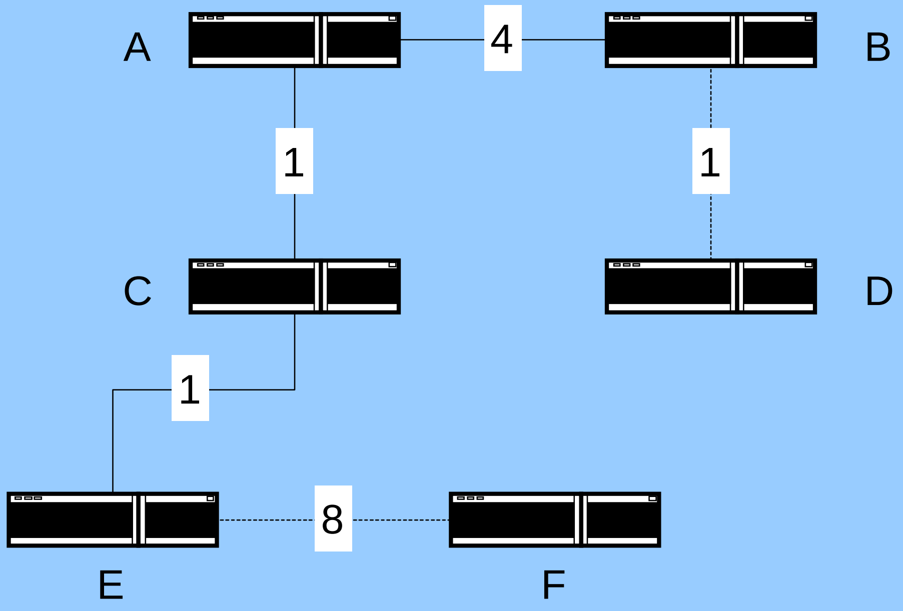

Routes Calculated

The Router then places D in PATH and examines its entry in the link state database.

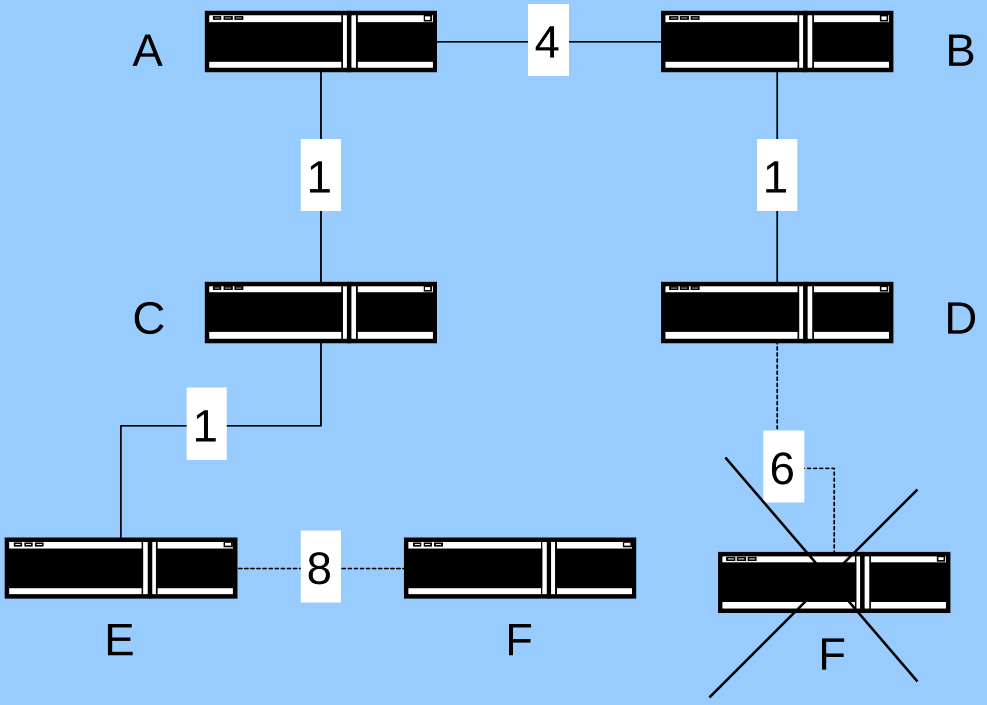

Routes Calculated

The Router then places F in PATH. Because there are no nodes left to calculate routes for, the process terminates.

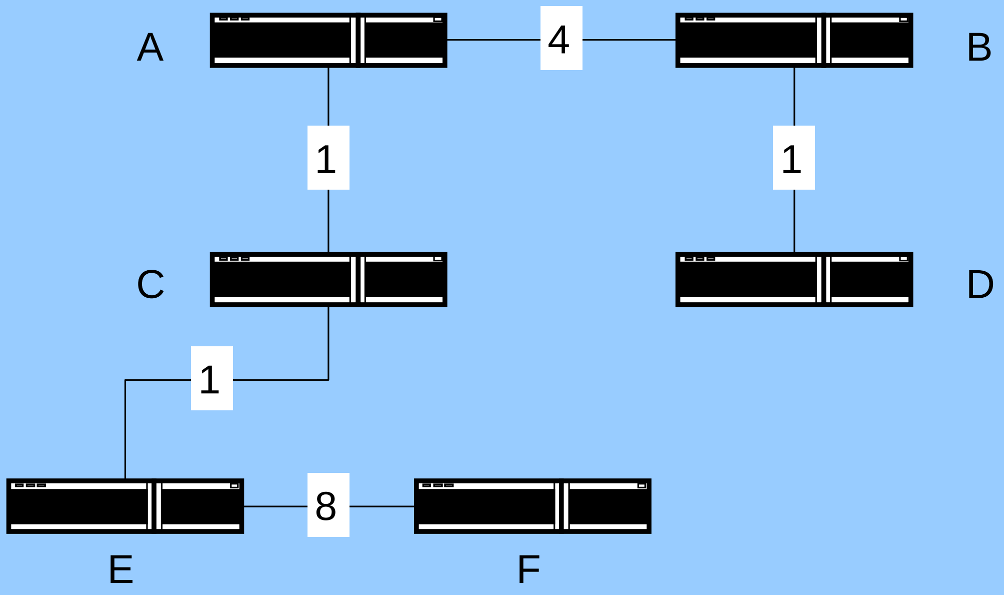

Which is better? Neither…

Distance Vector Routing

- Is easier to implement
- Takes up less memory

Link State Routing

- Makes it easier to discover the network topology and troubleshoot failures because each router has the network topology in memory
- Convergence time is faster.

Any Questions?