CS 50011: Introduction to Systems II

Prof. Jeff Turkstra

Computer Science Department
Purdue University

Copyright 2017

Internet Review and Socket Programming

History of the Internet

- In the late 1960s the Department of Defense Advance Research Project Agency (DARPA) created a nationwide network to allow computer access to the different research centers.
- The alternative was to give a computer to each center/university but this was –very- expensive.
- Research in the Internet continued in the 70s and 80s
- The Internet became a commercial success in the 90s.

History of the Internet

- The Internet has been doubling in size every nine to twelve months.
- Some people attribute the increase in productivity in the last 10 years to the existence of the Internet. People produce more in less time.

Internet Architecture

- The Internet is a collection of
  - Networks
  - Routers interconnecting networks
  - Hosts connected to networks.

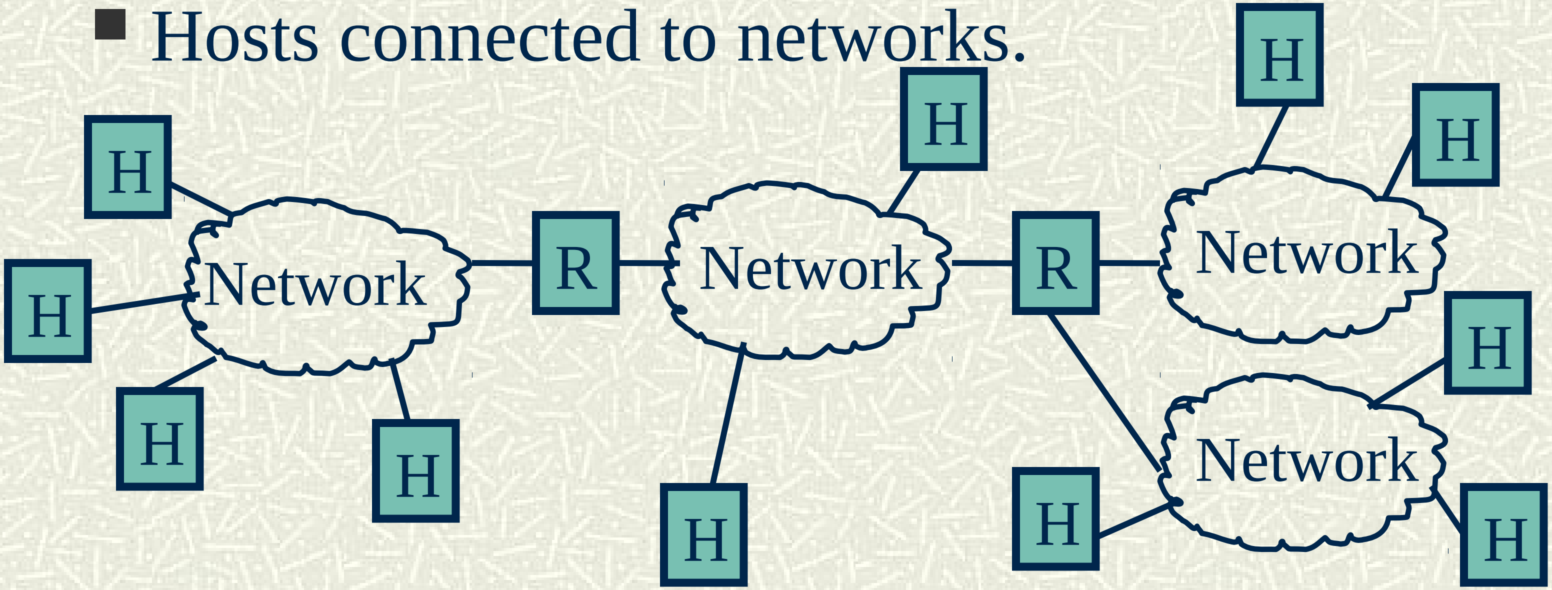

Internet Architecture

- The networks may be implemented using different kinds of hardware: Ethernet, Token Ring, Serial Line, Apple Talk, etc.
- The goal of the Internet is to hide all this heterogeneity to the user and user programs.

Internet Architecture

- The internet is a virtual network with its own addressing and name scheme.

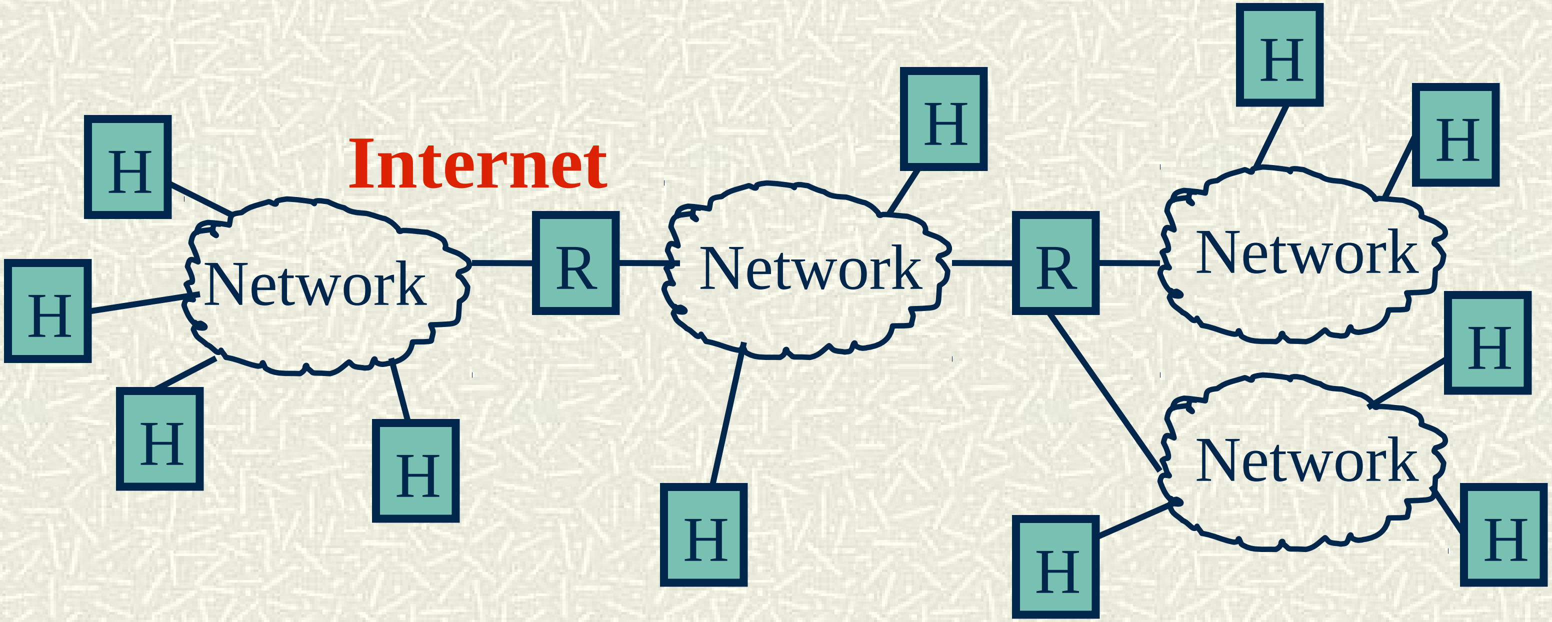

Internet Layering

- It reflects the layering used by the TCP/IP protocols
- Closer to reality than ISO-OSI Layering

Application	- Individual Application Programs (HTTP)
Transport	- Program to Program (TCP and UDP)
Internet	- Packet Forwarding. Machine to Machine. (IP)
Network Interface	- Local Area Network (Ethernet, RS232, etc)
Physical	- Basic Network Hardware

IP Addressing

- It is an abstraction to hide the network internals.
- It is independent from hardware addressing.
- IP addresses are used for all communications.
- IP addresses use 32 bits.
- There is a unique value for each host

IP Addressing

- Important:

  An IP Address does not specify a specific computer. Instead, each IP address identifies a connection between a computer and a network.

- An IP address identifies a network interface.
- A computer with multiple network connections (like a router) must be assigned one IP address for each connection.

IP Addressing

- It has two parts:
  - The prefix identifies a network
  - The suffix identifies the host in that network.

Network Number	Host Number

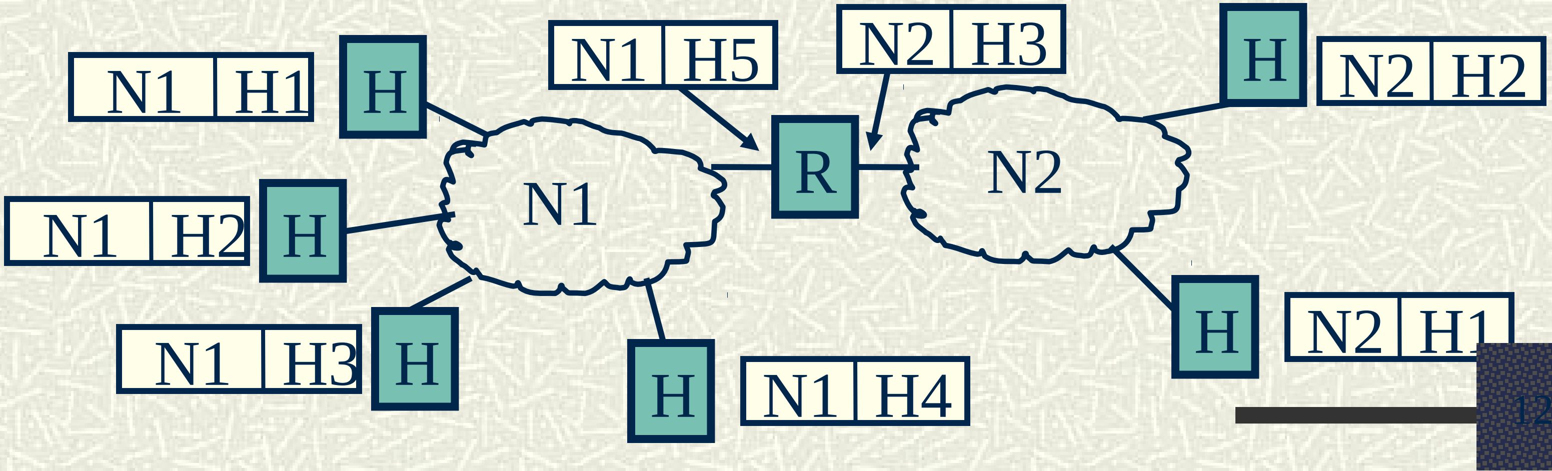

IP Addressing

- A global authority assigns a unique prefix for the network.
- A local administrator assigns a unique prefix to the hosts.
- The number of bits assigned to the prefix and suffix is variable depending on the size of the number of hosts in each network.
- The subnet mask is a parameter in the interface that tells the number of bits used for the network number.

Routing

- The routing table gives the next router necessary to reach the destination network.
- The source of the table information can be:
  - Manual:
    - By hand
    - Small networks
    - OK if routes never change
  - Automatic
    - Software creates/updates the routing table using information from neighboring routers.
    - It is needed for lager nets
    - It changes routes if failure.

IP packet from A to M

A: Routing Table

Target Net	Net/Subnet Mask	Next Hop
40.0.0.0	255.0.0	Directly
128.10.3.0	255.255.255.0	40.0.0.11(R1)
128.10.5.0	255.255.255.0	40.0.0.11(R1)
128.10.4.0	255.255.255.0	40.0.0.11(R1)
216.109.112.0	255.255.255.0	40.0.0.11(R1)
Default:	255.255.255.255	40.0.0.1(R3)

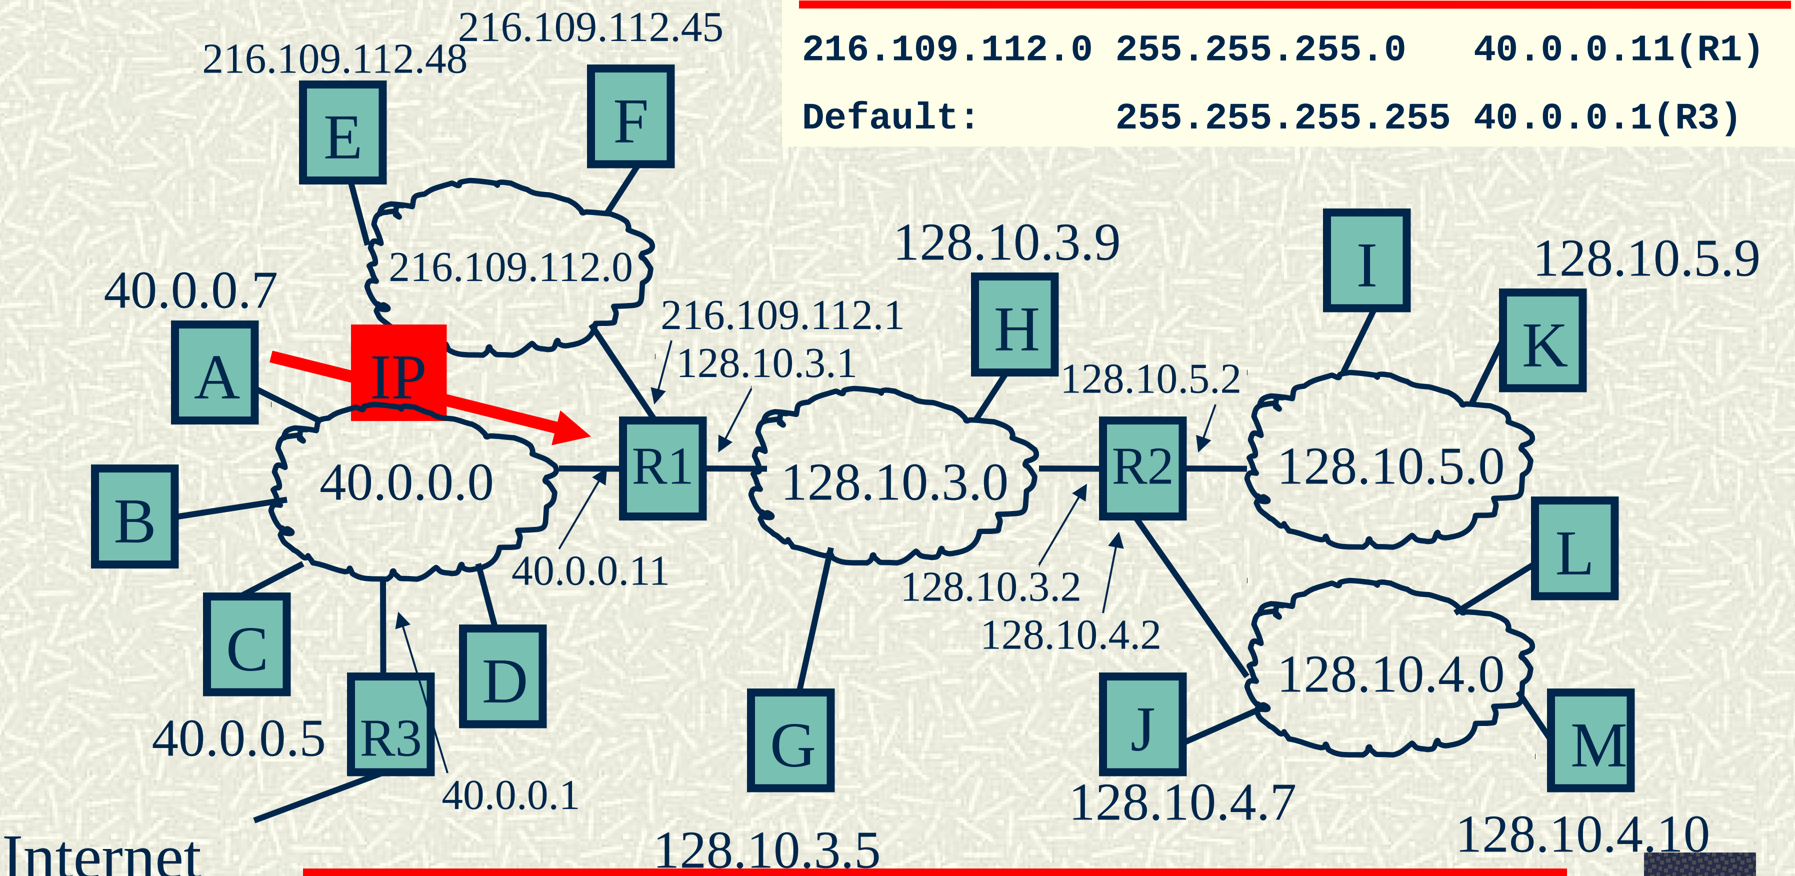

IP: $E_{src}=E_a$, $E_{dst}=E_{R1}$, $IP_{src}=A$, $IP_{dst}=M$

IP packet from A to M

R1: Routing Table

Target Net	Net/Subnet Mask	Next Hop
40.0.0.0	255.0.0	Directly
128.10.3.0	255.255.255.0	Directly
128.10.5.0	255.255.255.0	128.10.3.2(R2)
128.10.4.0	255.255.255.0	128.10.3.2(R2)
216.109.112.0	255.255.255.0	Directly
Default:	255.255.255.255	40.0.0.1(R3)

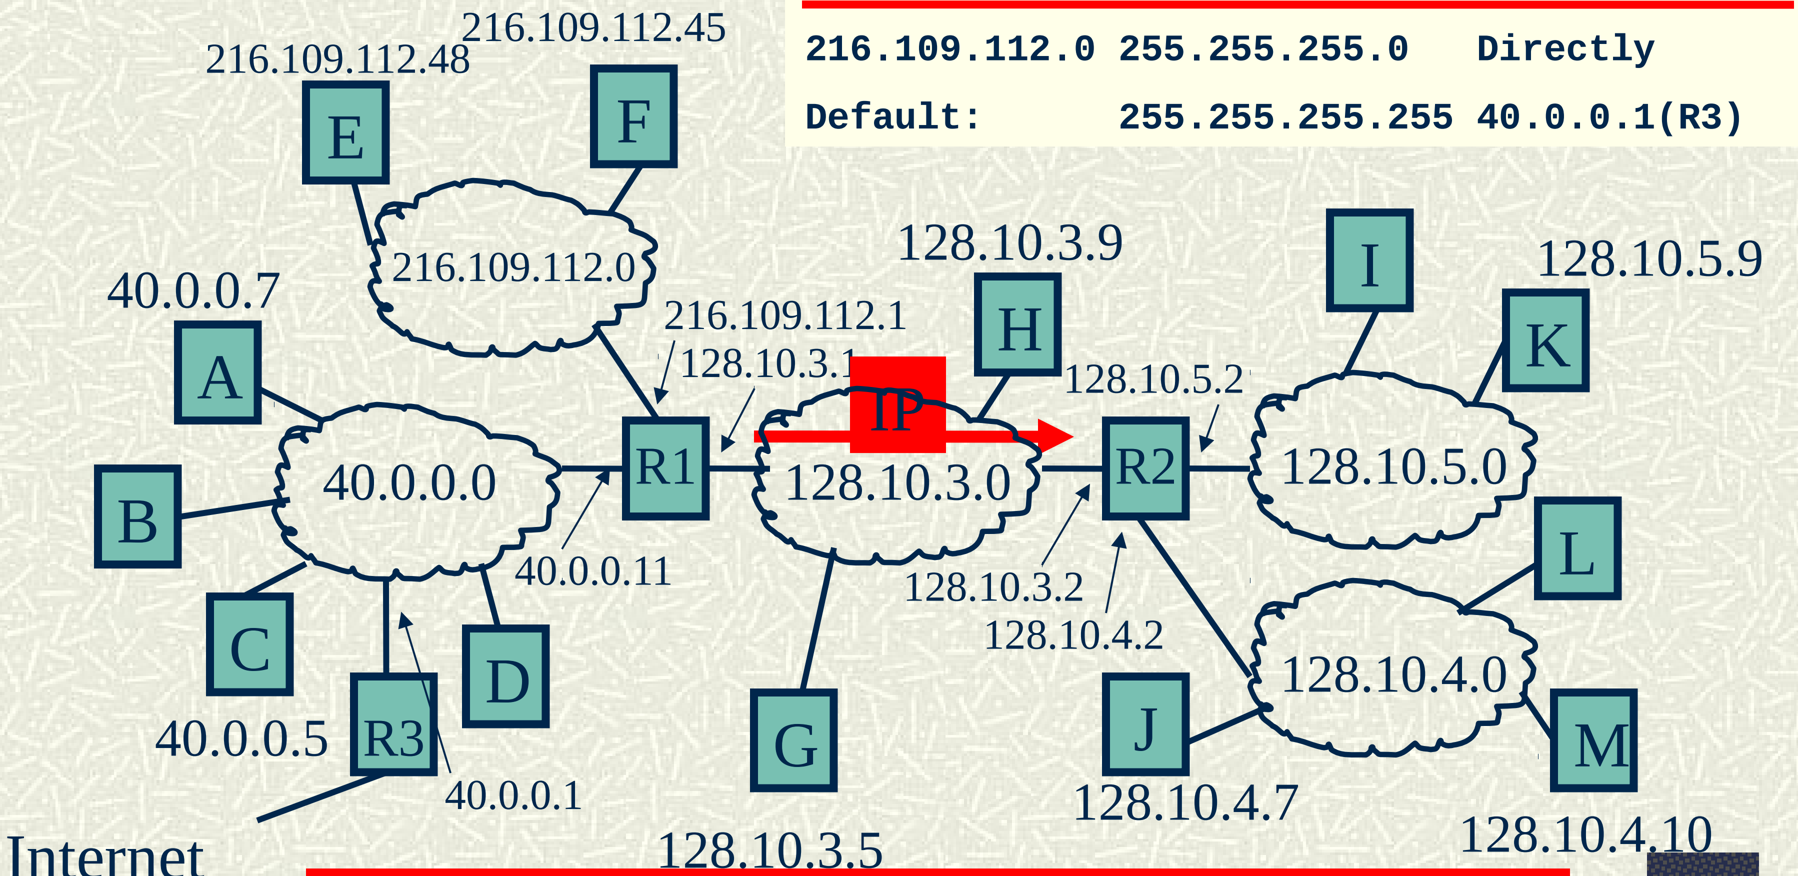

IP: E_{src}=E_{R1}, E_{dst}=E_{R2}, IP_{src}=A, IP_{dst}=M

IP packet from A to M

R2: Routing Table

Target Net	Net/Subnet Mask	Next Hop
40.0.0.0	255.0.0	128.10.3.1(R1)
128.10.3.0	255.255.255.0	Directly
128.10.5.0	255.255.255.0	Directly
128.10.4.0	255.255.255.0	Directly
216.109.112.0	255.255.255.0	128.10.3.1(R1)
Default:	255.255.255.255	128.10.3.1(R1)

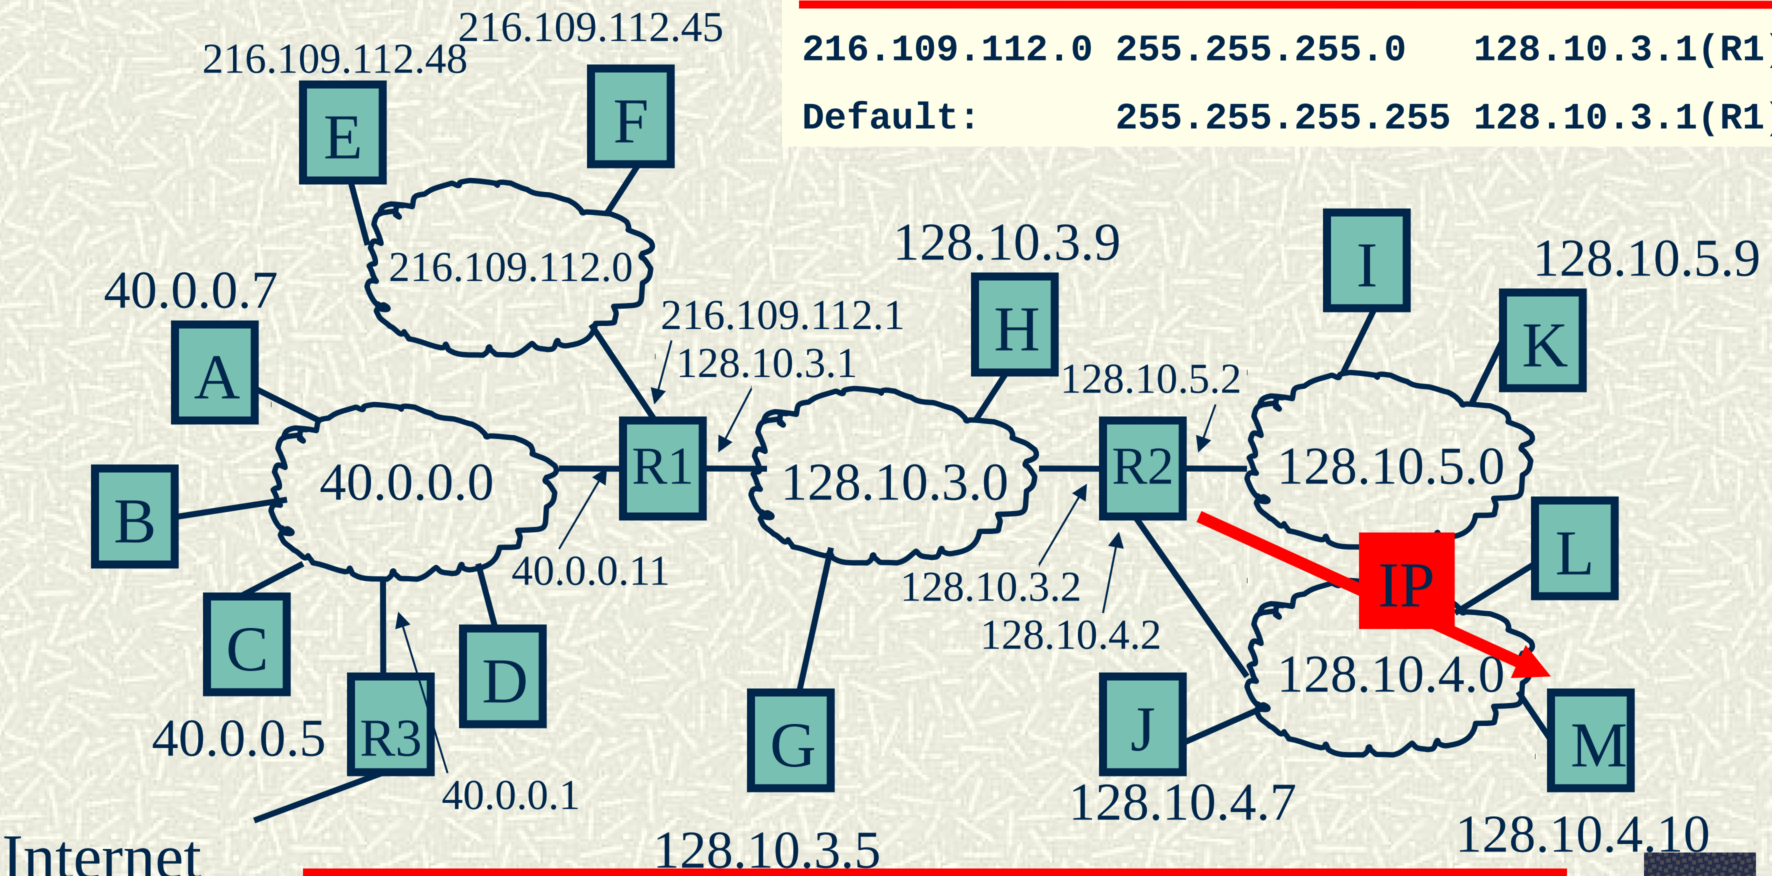

IP: $E_{src}=E_{R2}$, $E_{dst}=E_{M}$, $IP_{src}=A$, $IP_{dst}=M$

Addresses in IP packet

- The IP source and destination address of the packet is the same during the transit of the packet.
- The hardware source and destination address will be different every time the packet is forwarded.
- The source host or some of the routers also may require ARP requests if the hardware destination address is not in the ARP cache.

ARP Address Resolution Protocol

- When it is time for the router or host to deliver a packet directly, it is necessary to convert the IP address to a hardware address.
- For example in an Ethernet LAN , the target IP address in the IP packet has to be translated to the Ethernet address of the destination machine.
- ARP does this translation.

ARP Address Resolution Protocol

- ARP Input and Output:
  - Input: IP address C of computer in locally connected network N
  - Output: Ethernet address for C.
- ARP keeps bindings (IPAddr, EtherAddr) in a table called ARP table or ARP cache.
- ARP builds the table as needed.

ARP Command

```
C:\Users\owner>arp -a
Interface: 10.184.105.105 --- 0xd
  Internet Address      Physical Address      Type
  10.184.96.1           00-24-c4-c0-fe-c0     dynamic
  10.184.111.255        ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  224.0.0.252           01-00-5e-00-00-fc     static
  239.255.255.250       01-00-5e-7f-ff-fa     static
  255.255.255.255       ff-ff-ff-ff-ff-ff     static
```

DNS: Domain Name Server

- Humans prefer to use computer names instead of IP addresses.
- Example: www.yahoo.com instead of 204.71.200.68
- Before DNS the mappings name to IP address where stored in a file /etc/hosts
- The Net administrators used to exchange updates in the /etc/hosts file.
- This solution was not scalable..

DNS: Domain Name Server

- One host name may map to multiple IP addresses. Why? One host will have multiple IP addresses if it has multiple network interfaces.

DNS: Domain Name Server

- DNS is a service that translates host names to IP addresses.
- DNS uses a distributed lookup algorithm and contacts as many servers as necessary.
- Host names are divided in domains: host.dom2.dom1.dom0.
- Example: ector.cs.purdue.edu with the most general domain at the right.

DHCP – Dynamic Host Configuration Protocol

- Allows connecting computers to the Internet without the need of an administration.
- Before DHCP, an administrator had to manually configure the following parameters to add a computer to the Internet:
  - The local IP address – Current address
  - The subnet mask – Used to send packets to hosts in same LAN
  - The default router – Used to send packets to hosts outside the LAN
  - The default DNS server – Used to convert names to IP addresses.
- In UNIX the command used to set these parameters is ifconfig. In Windows is ipconfig or the Control Panel.

Transport Protocols

- Two transport protocols available in the TCP/IP family
  - UDP – User Datagram Protocol
  - TCP – Transmission Control Protocol

UDP- User Datagram Protocol

- Unreliable Transfer. Applications will need to implement their own reliability if necessary.
- Minimal overhead in both computation and communication.
- It is best for LAN applications
- Connectionless – No initial connection necessary. No state in both ends

UDP- User Datagram Protocol

- Message Oriented
- Each message is encapsulated in an IP datagram.
- Size of message is restricted by the size of the MTU of the directly connected network. (Maximum Transfer Unit =1500bytes for Ethernet networks)
- The UDP header has ports that identify
  - Source application (Source Port)
  - Destination application (Destination Port)

TCP – Transmission Control Protocol

- It is the major transport protocol used in the Internet
- It is:
  - Reliable – It uses acknowledgement and retransmission to accomplish reliability
  - Connection-Oriented - An initial connection is required. Both end points keep state about the connection.
  - Full-Duplex – Communication can happen in both ways simultaneously.
  - Stream Interface – Transfer of bytes look like writing/reading to a file.

TCP Reliability

- How does TCP achieves reliability?
- It uses Acknowledgments and Retransmissions
- Acknowledgement-
  - The receiver sends an acknowledgement when the data arrives.
- Retransmission
  - The sender starts a timer whenever the message is transmitted
  - If the timer expires before the acknowledgement arrives, the sender retransmits the message.

TCP Reliability

Normal Exchange

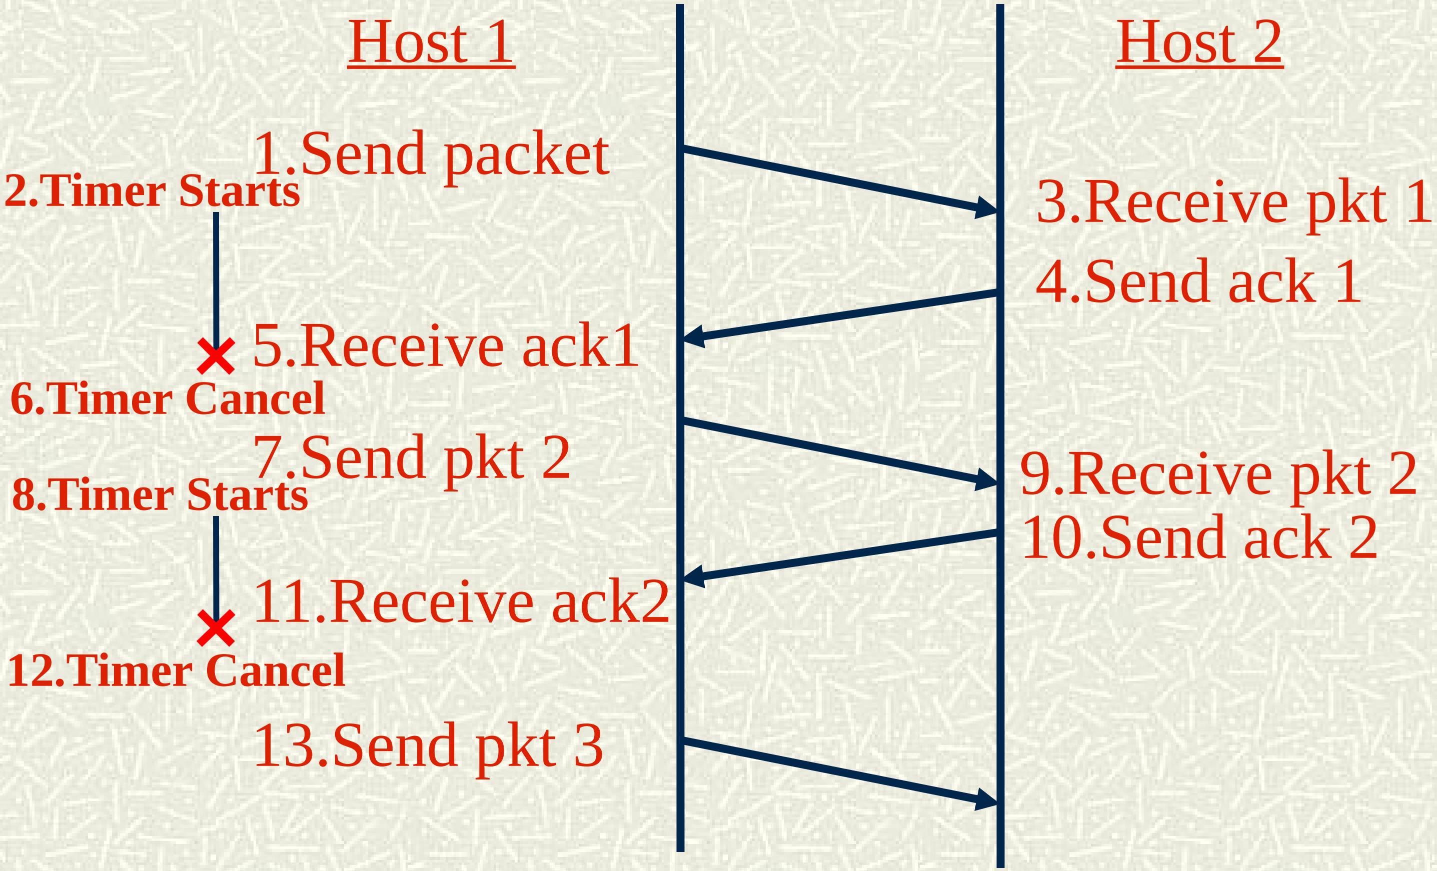

Host 1

Host 2

1.Send packet

2.Timer Starts

3.Receive pkt 1

4.Send ack 1

5.Receive ack1

6.Timer Cancel

7.Send pkt 2

8.Timer Starts

9.Receive pkt 2

10.Send ack 2

11.Receive ack2

12.Timer Cancel

13.Send pkt 3

TCP Reliability

Packet Lost

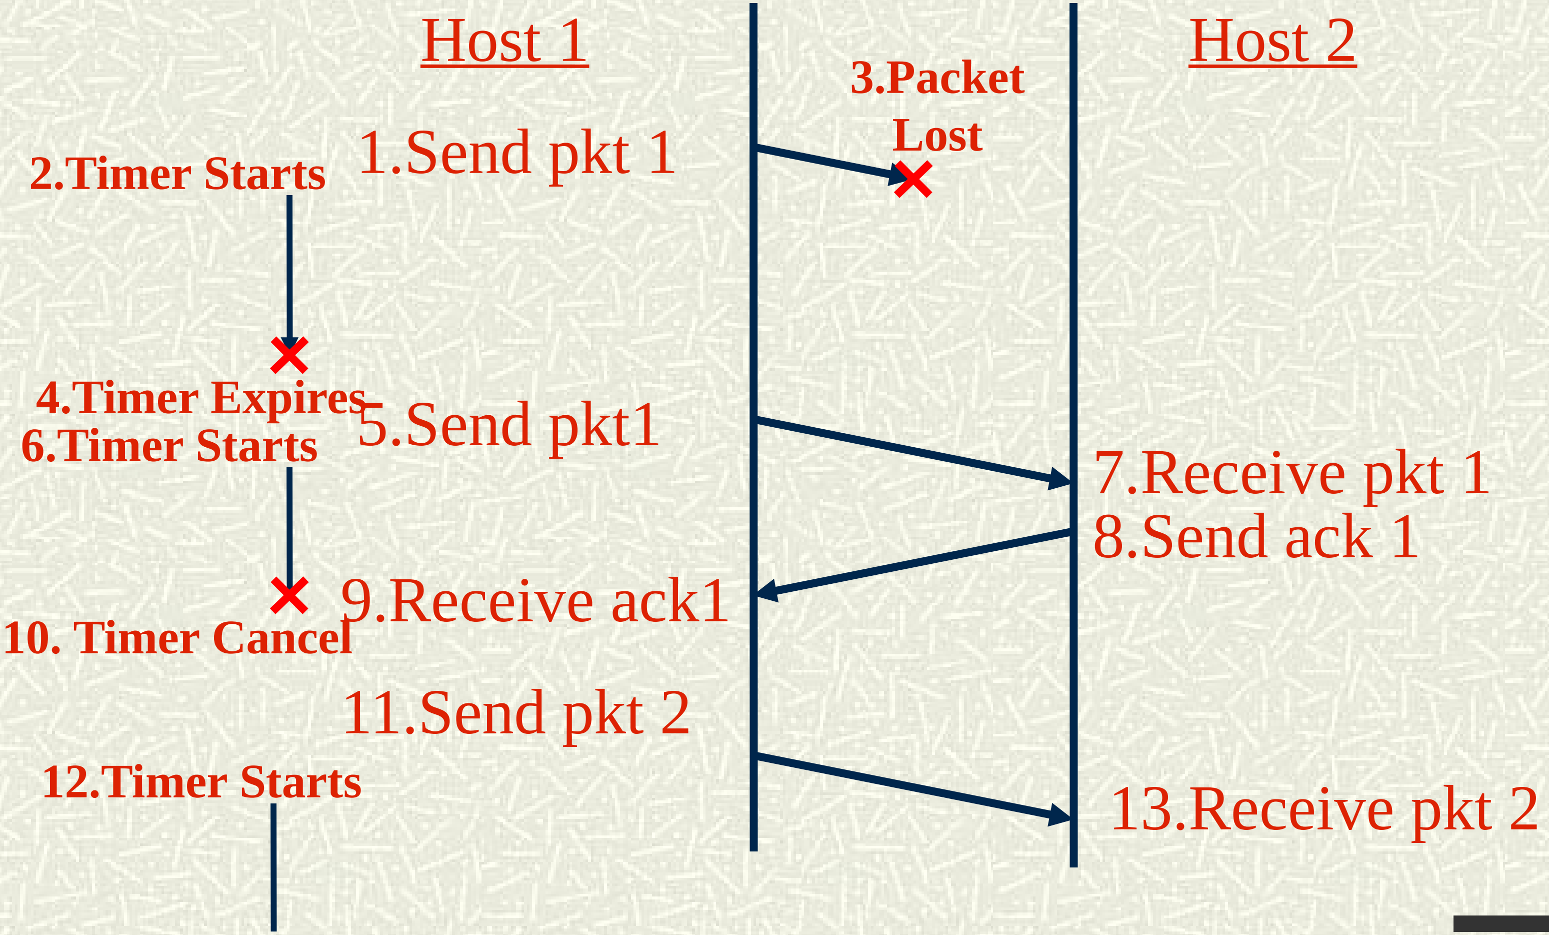

Host 1

Host 2

1.Send pkt 1

2.Timer Starts

3.Packet Lost

4.Timer Expires

5.Send pkt1

6.Timer Starts

7.Receive pkt 1

8.Send ack 1

9.Receive ack1

10. Timer Cancel

11.Send pkt 2

12.Timer Starts

13.Receive pkt 2

TCP Reliability

Ack Lost

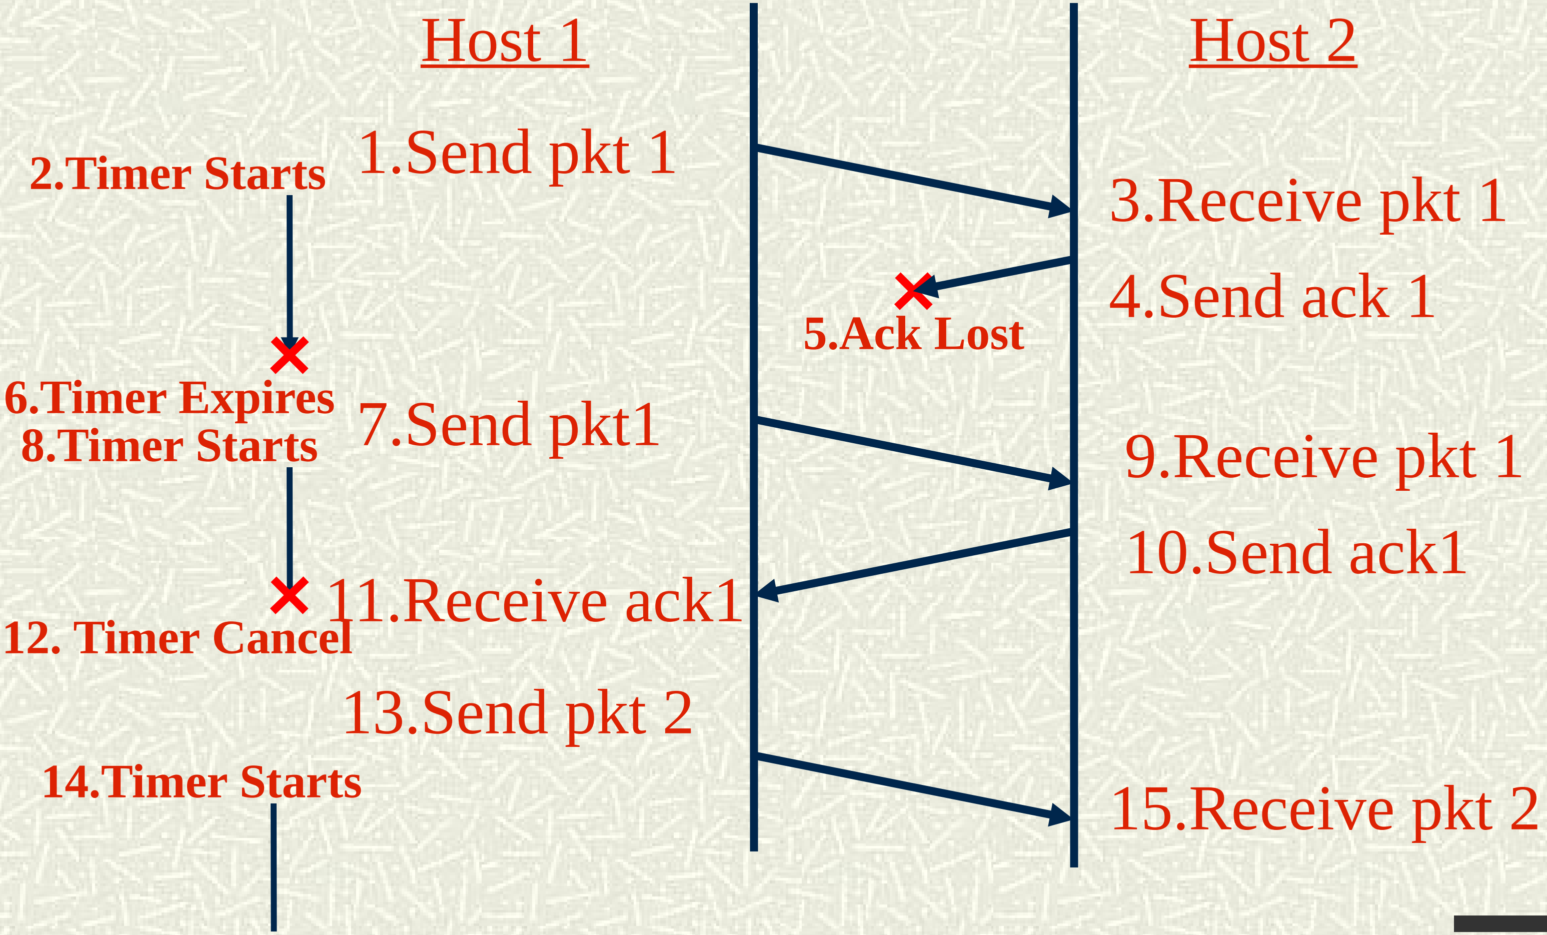

Host 1

Host 2

1.Send pkt 1

2.Timer Starts

3.Receive pkt 1

4.Send ack 1

5.Ack Lost

6.Timer Expires

7.Send pkt1

8.Timer Starts

9.Receive pkt 1

10.Send ack1

11.Receive ack1

12. Timer Cancel

13.Send pkt 2

14.Timer Starts

15.Receive pkt 2

TCP Summary of Features

- ***1. Adaptive Retransmission***
  - The retransmission timer is set to RTT+4*RTTSTDDEV where RTT is estimated. This allows TCP work in slow and fast networks.
- ***2. Cumulative Acknowledgments***
  - An acknowledgment is for all the bytes received so far without holes and not for every packet received.
- ***3. Fast Retransmission***
  - It is a heuristic where a duplicated acknowledgment for the same sequence is signal of a packet lost. The data is retransmitted before the timer expires.

TCP Summary of Features

- ***4. Flow Control***
  - It slows down the sender if the receiver is running out of buffer space. The window (receiver's buffer size) is sent in every acknowledgment.
- ***5. Congestion Control***
  - For TCP a lost packet is signal of congestion.
  - Instead of aggressively retransmit, it will slow down the retransmission. It will use first ***"Slow Start"*** and then a ***"Congestion Avoidance"*** where the window of retransmitted data is reduced in size.

TCP Summary of Features

- ***6. Reliable Connection and Shutdown***
  - TCP uses a Three way Handshake to close connections.
  - Three packets are enough to make sure that lost packets or host crashes will not interfere future connections.

When to Use UDP or TCP

- If you need reliable communication in your application use TCP.
- Only use UDP in the following cases:
  - Broadcasting: that is the computer needs to reach all or part of the computers in the local network.
    - Example: Find out of the existence of a server (Example DHCP, or finding a printer).
    - Multicasting data to several machines simultaneously.
  - Real Time Data: Applications where packets arriving on time with minimum delay is more important than reliability where retransmission can add to the delay.
    - Example: Voice over IP, teleconferencing.

NAT: Network Address Translation

- *Network Address Translation (NAT)* is used when you want to connect multiple computers to the Internet using a single IP address.
- The NAT software can run on a computer or specialized device (NAT box) that has two network interfaces: one connected to the private network and the other one to the Internet.

NAT: Network Address Translation

- As a side effect, NAT provides protection.
- Packets will be allowed into the private network only if they belong to a connection that started by a machine in the private network.
- A NAT box is also called a ***Firewall.***
- NAT also mitigates the problem of running out of Assigned Network Numbers.
- Potentially you could have another Internet behind a NAT box.

NAT: Network Address Translation

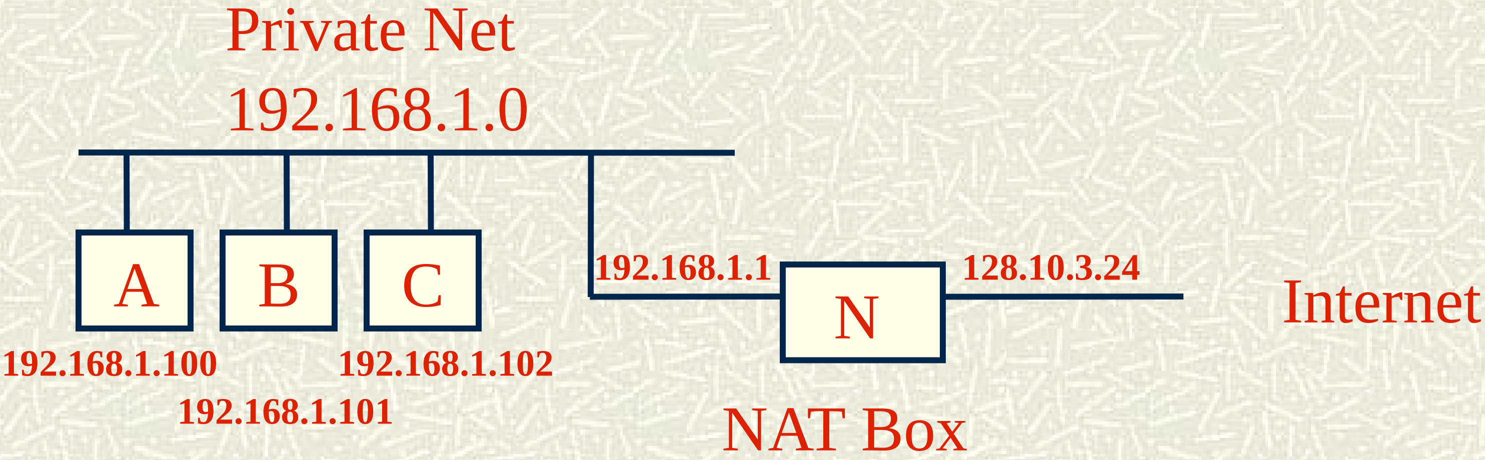

From the point of view of the Internet, all computers in the private network have the address 128.10.3.24

The NAT box is also a DHCP server that assigns IP addresses and it is the default router.

NAT: Network Address Translation

- A TCP connection is defined uniquely in the entire Internet by four values:

  <src-ip-addr, src-port, dest-ip-addr, dest-port>

- The NAT box will work as follows:

1. The machines in the private network use the NAT box as the default router.
2. When a TCP packet with header

   <IPsrc, PORTsrc, IPdest, PORTdest>

   goes from the private network to the Internet, the NAT box will change the header to

   <IPnat, PORTrand, IPdest, PORTdest>

   Where ***IPnat*** is the shared IP address and ***PORTrand*** is a random unused port in the NAT box.

NAT: Network Address Translation

3. The NAT box will also add the NAT mapping to the NAT table

   (PORTrand, IPsrc, PORTsrc)

4. When a packet

   < ***IPdest, PORTdest , IPnat, PORTrand>***

   comes from the Internet to the NAT box, the NAT box will lookup ***PORTrand*** in the NAT table and it will change the header to

   < ***IPdest, PORTdest , IPsrc, PORTsrc>***

   And it will forward the packet to ***IPsrc.***

5. A similar translation is done for UDP packets.

NAT Example

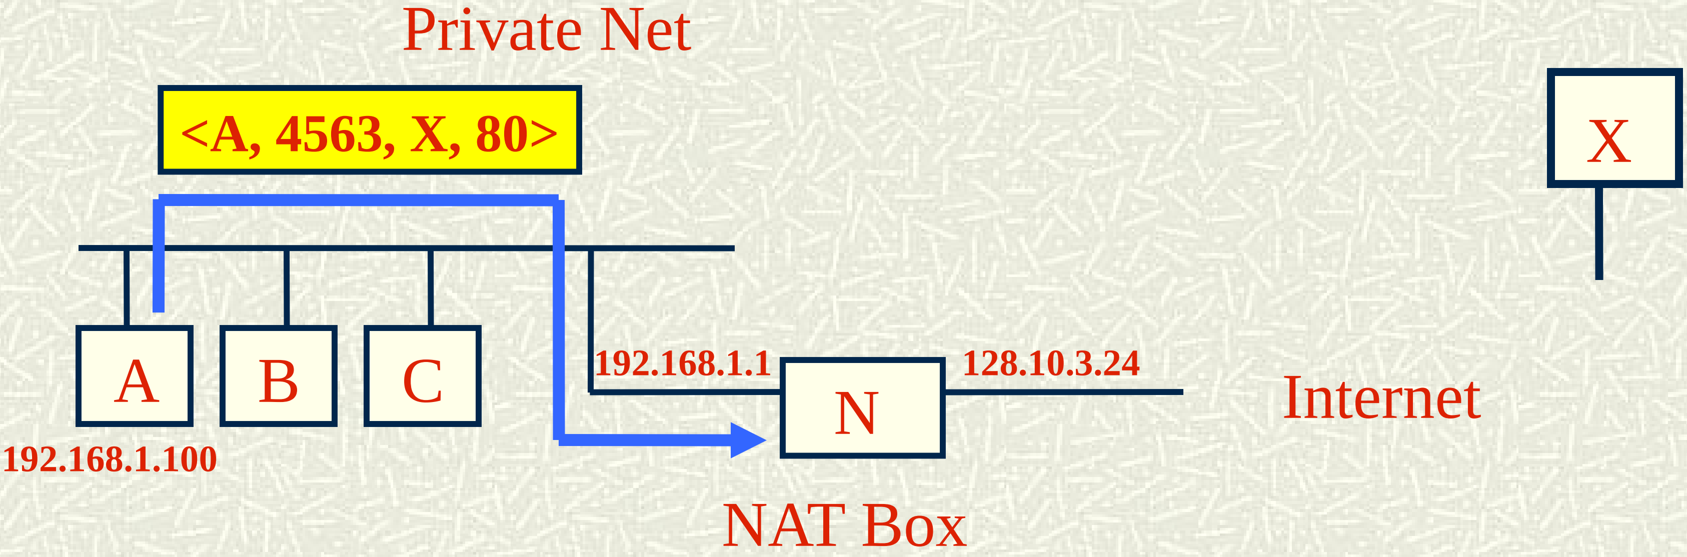

1. Computer A wants to establish a TCP connection with an HTTP server X in the Internet.
2. The NAT box is the default router for A so A sends the TCP packet to the NAT box.

NAT Example

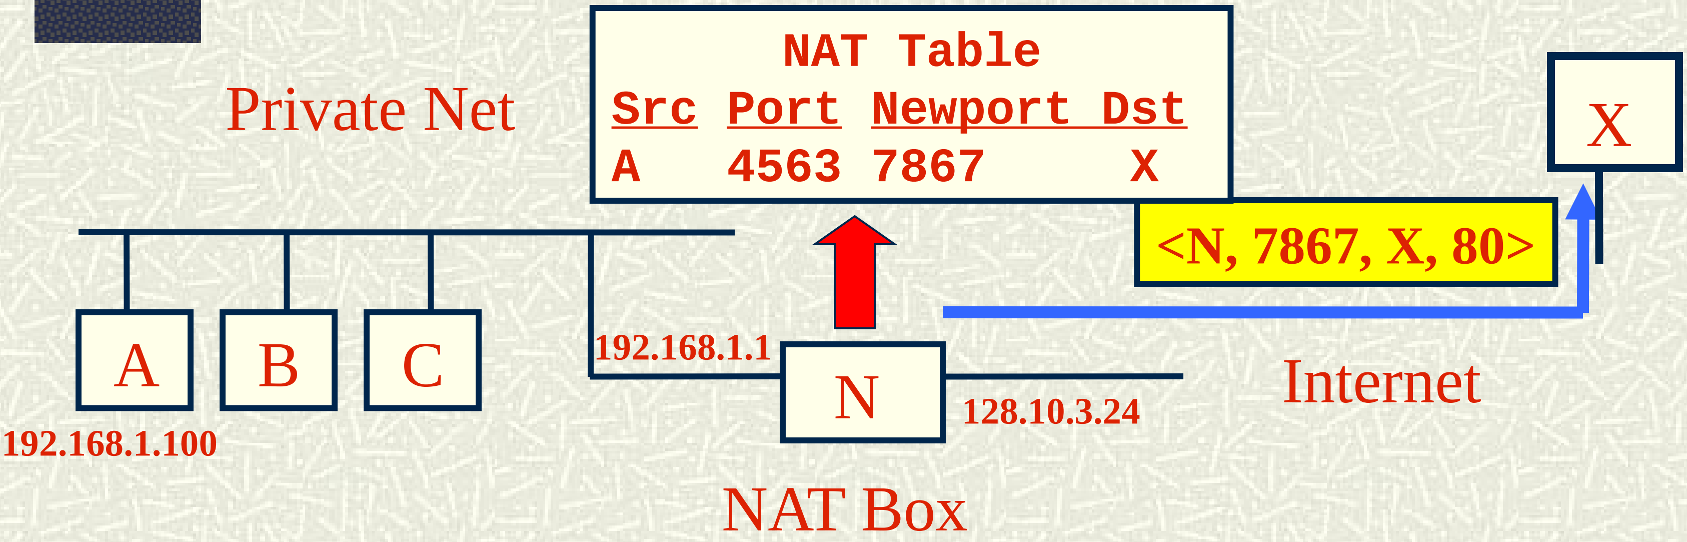

NAT Table

Src	Port	Newport	Dst
A	4563	7867	X

3. The NAT box chooses an unused random port (7867) and substitutes the source port in the packet as well as the source address with its own IP address. The new packet is sent to X.

4. The old port (4563), old source address (A), and new port are added to the NAT table for use when packets of the same connection come back.

NAT Example

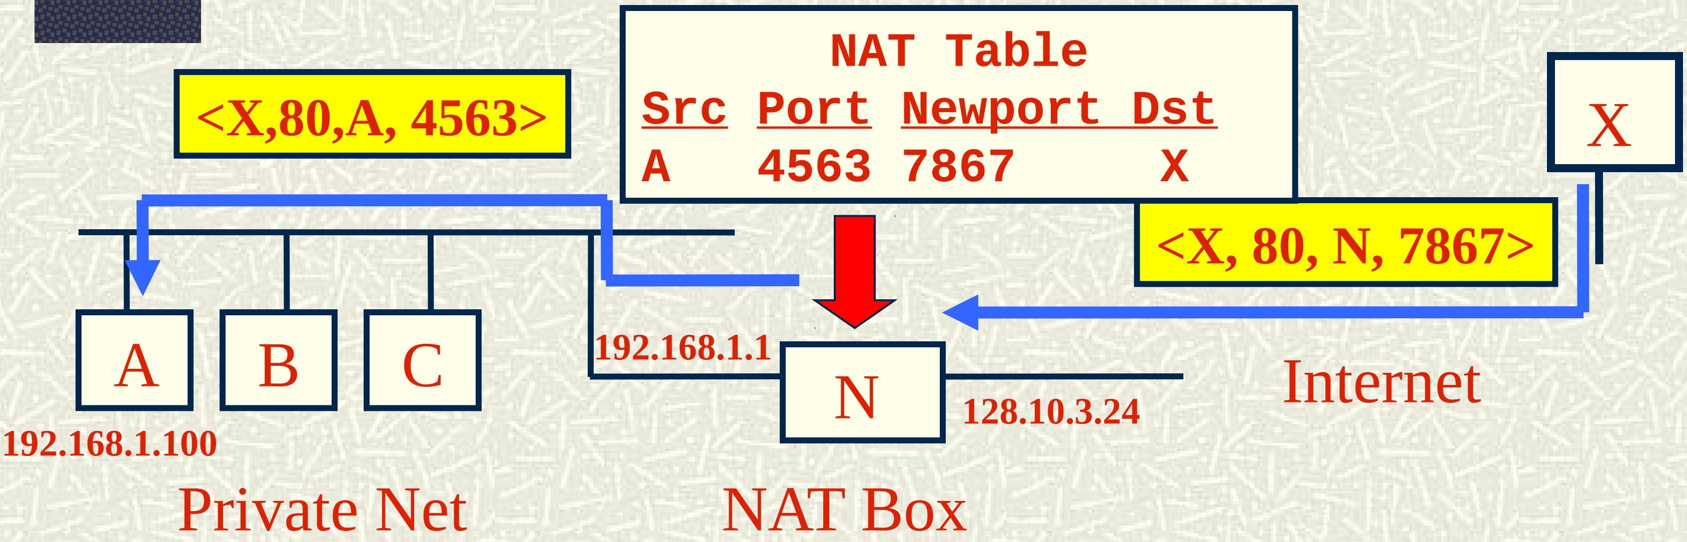

NAT Table

Src	Port	Newport	Dst
A	4563	7867	X

5. When a packet comes back from X, the NAT box will lookup the destination port (7867) in the NAT table and it will substitute the destination address and destination port with the original values.

6. The packet is forwarded to A. Everything will be transparent for both X and A. They will never know that a translation took place. The same mapping will be used for all packets of the same connection until the connections closed or the entry times out..

NAT and Firewalls

- The NAT box will pass into the private network only the packets that belong to a connection that started from the inside the private network.
- This is similar to the recommendation “Never give your credit card through the phone if you did not start the call”.
- If a packet is received that does not have a mapping in the NAT table, the packet is discarded.

NAT and Firewalls

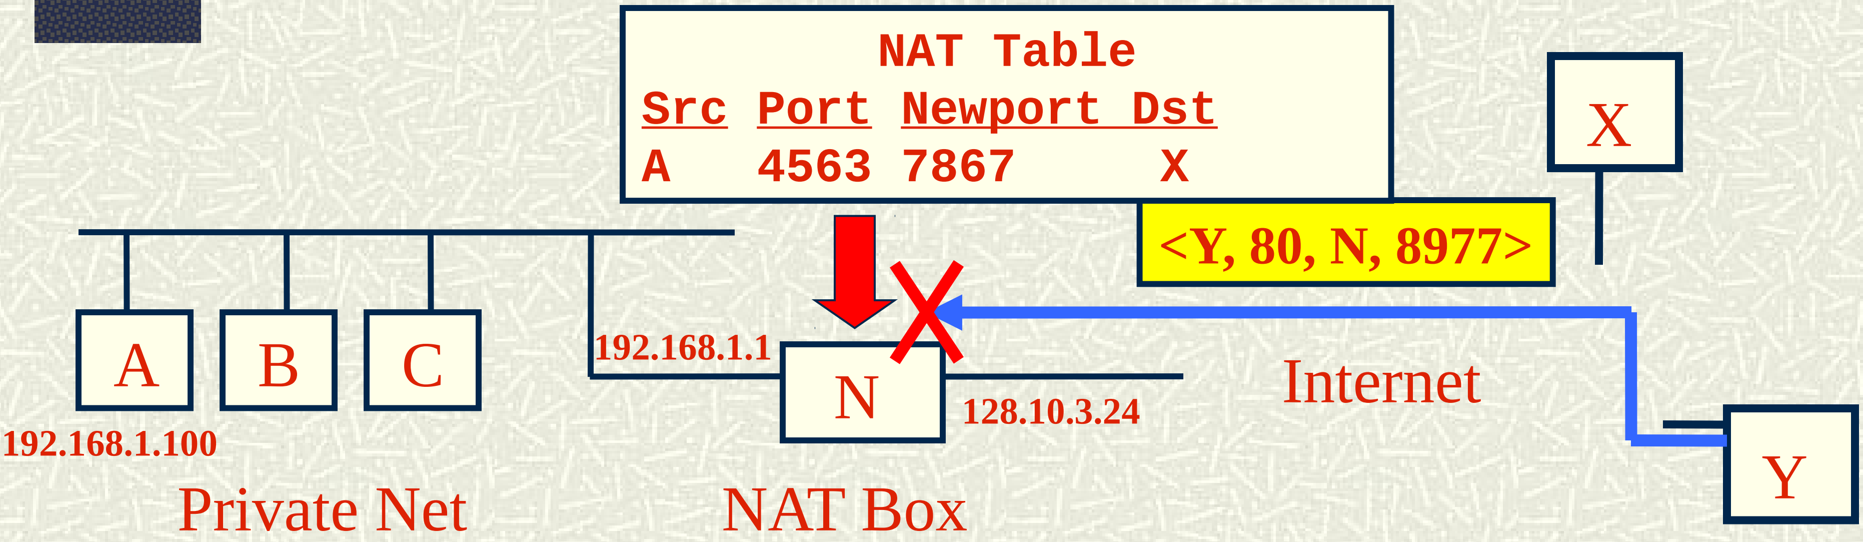

NAT Table

Src	Port	Newport	Dst
A	4563	7867	X

1. If a packet arrives to the NAT box that does not have an existing entry in the NAT table, the packet will be discarded by the NAT box.
2. This protects the private network from hacker attacks. It is recommended to have a NAT box/Firewall even when you have one computer connected to the Internet.

NAT: Problems with Firewalls

- What will happen if you want to run a server at home?
- The connections will start from the outside and there is no mapping in the NAT table.
- The packets will be dropped before reaching the server.
- Solution: Add static entries to the NAT table manually.

NAT: Static Entries

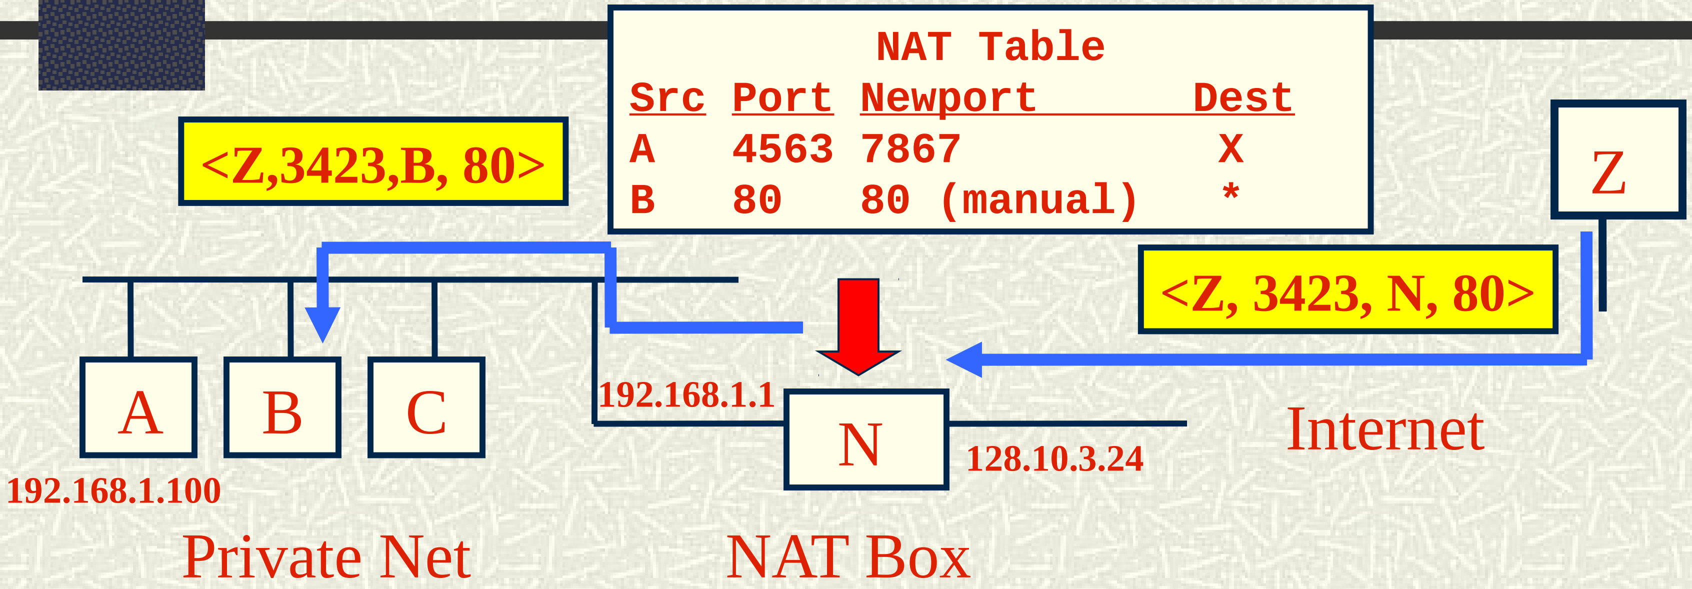

NAT Table

Src	Port	Newport	Dest
A	4563	7867	X
B	80	80 (manual)	*

1. B is running an HTTP server in the private network.
2. To allow connections from the Internet, the administrator adds a static entry in the NAT table to redirect any request to port 80 to B.

NAT: Problems with Firewalls.

- Some application eliminate this problem by having a proxy server forward the data.
- The computer in the private network makes a connection to the proxy server. Computers in the internet that want to connect to the computer in the private network contact the proxy server.

NAT Configuration

- NAT boxes have a built in HTTP server that you can use to configure it or add static entries.
- For protection, any request to add a static entry in the NAT table or any configuration in the NAT has to come from a computer in the private Network.

Wireless Router Web Interface

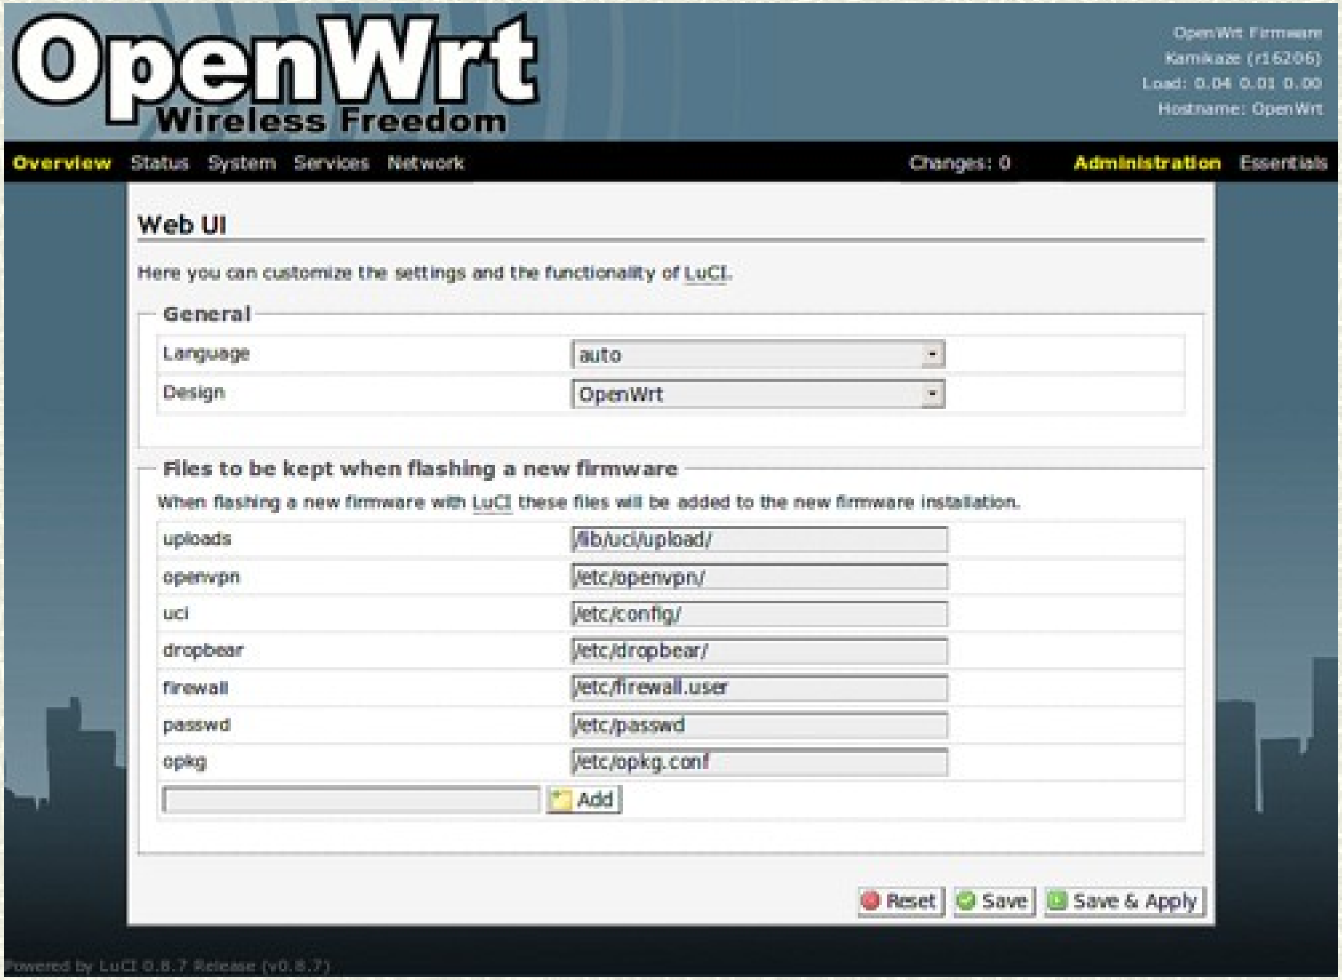

OpenWrt Firmware
Kamikaze (r16206)
Load: 0.04 0.01 0.00
Hostname: OpenWrt

Changes: 0

Web UI

Here you can customize the settings and the functionality of LuCI.

General

Language	auto
Design	OpenWrt

Files to be kept when flashing a new firmware

When flashing a new firmware with LuCI these files will be added to the new firmware installation.

uploads	/lib/uci/upload/
openvpn	/etc/openvpn/
uci	/etc/config/
dropbear	/etc/dropbear/
firewall	/etc/firewall.user
passwd	/etc/passwd
opkg	/etc/opkg.conf

Powered by LuCI 0.8.7 Release (v0.8.7)