Internet Security [1]

VU 184.216

Engin Kirda engin@infosys.tuwien.ac.at
Christopher Kruegel chris@auto.tuwien.ac.at

The Internet

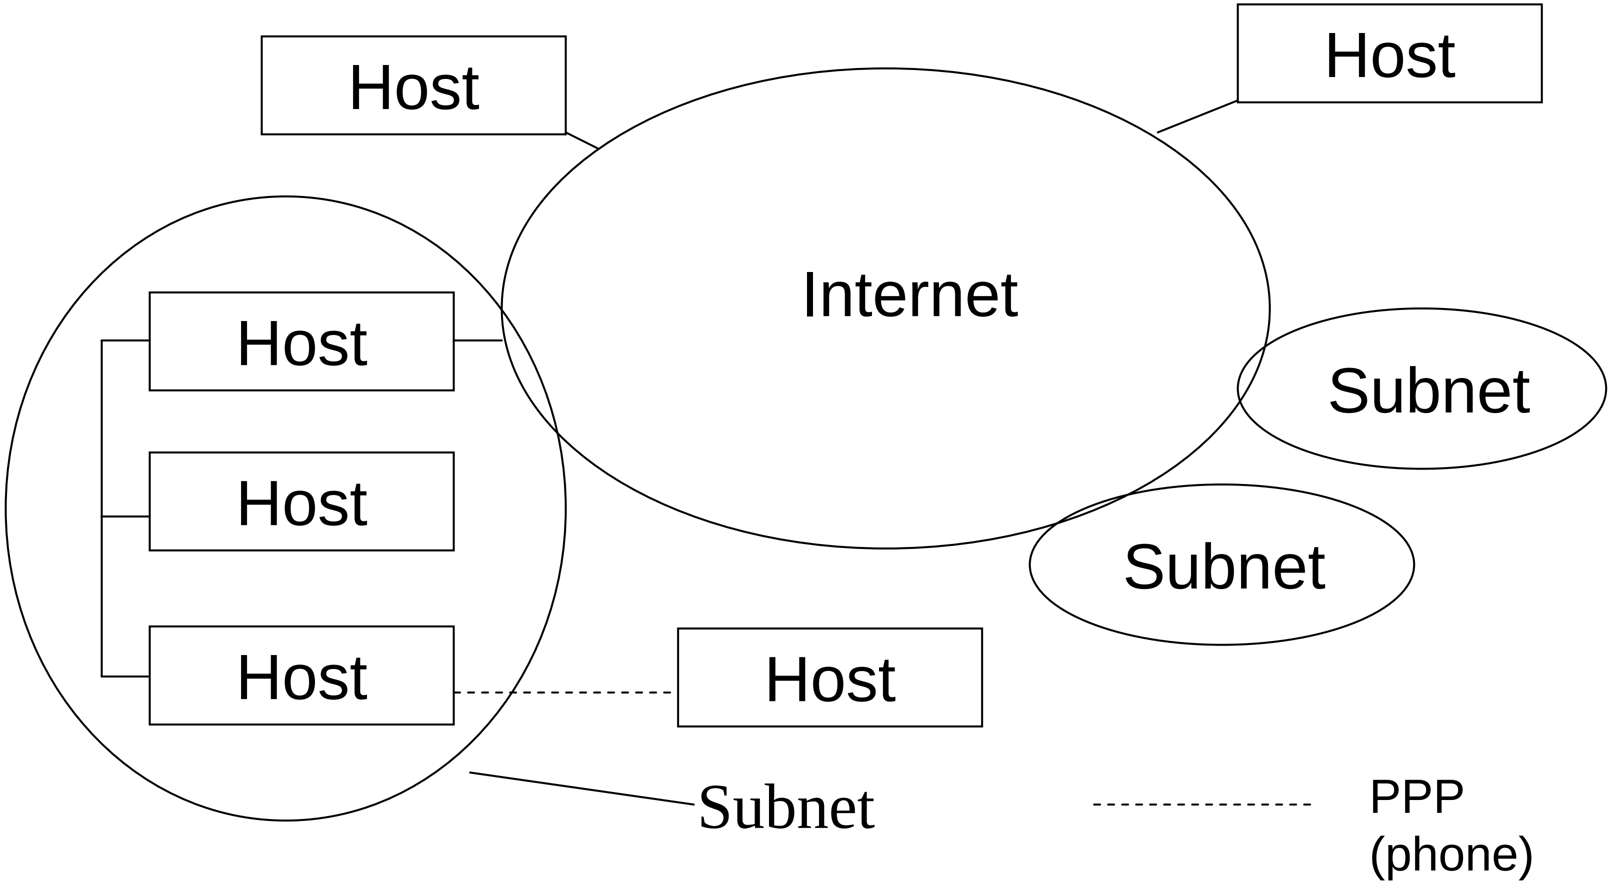

Direct IP delivery

- If two hosts are in the same physical network the IP datagram is encapsulated and delivered directly

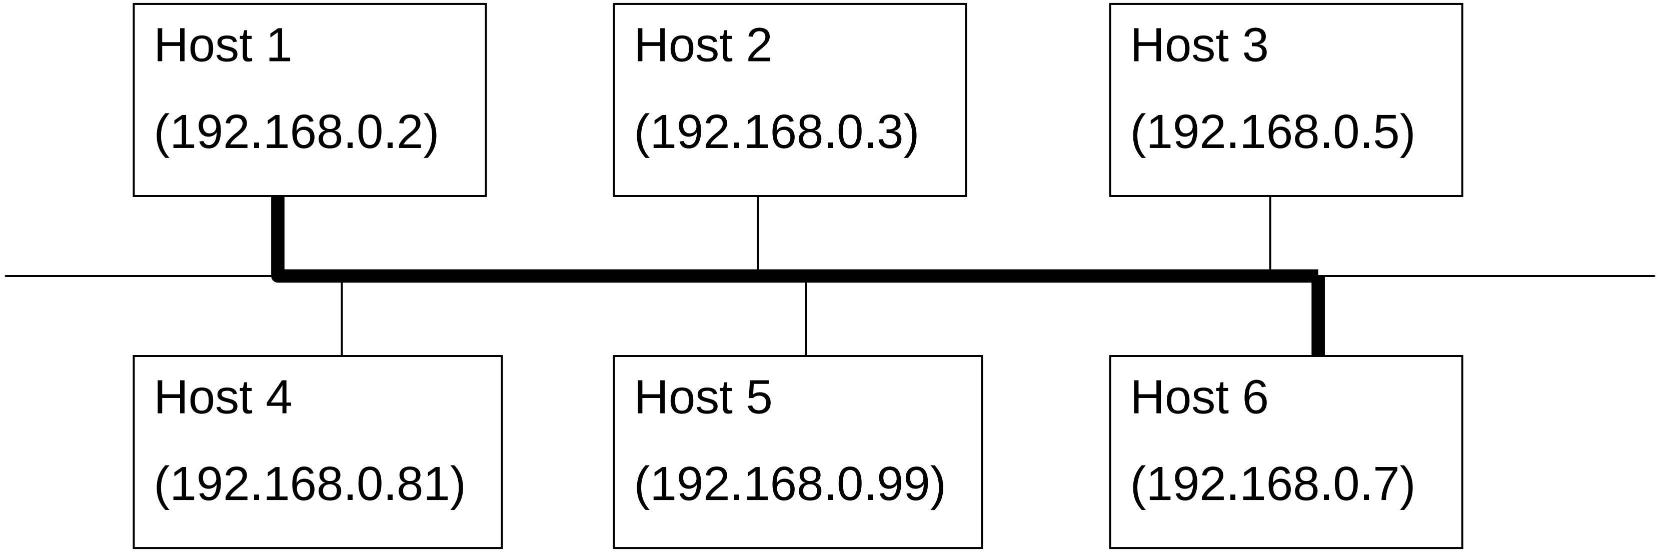

Ethernet

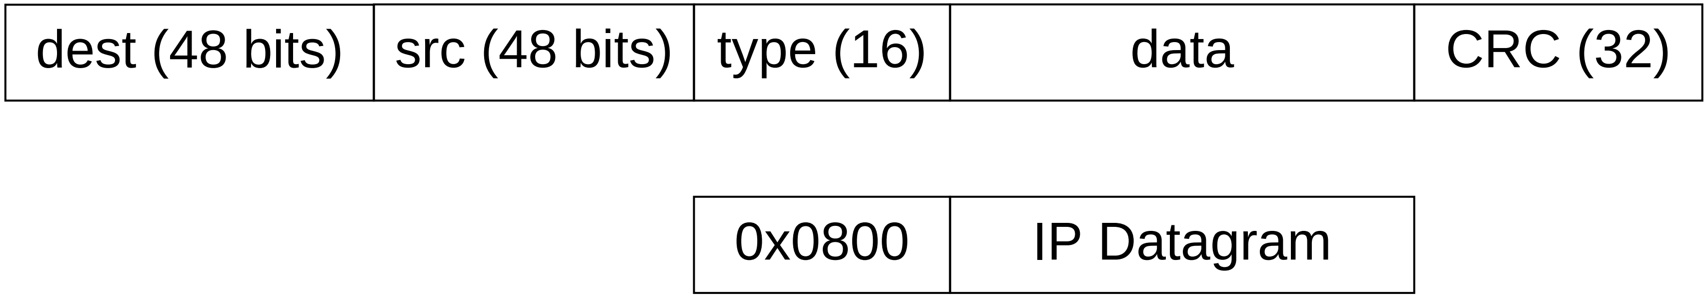

dest (48 bits)	src (48 bits)	type (16)	data	CRC (32)
		0x0800	IP Datagram	

Ethernet

- Widely used link layer protocol
- Carrier Sense, Multiple Access, Collision Detection
- Addresses: 48 bits (e.g. 00:38:af:23:34:0f), mostly
  - hardwired by the manufacturer
- Type (2 bytes): specifies encapsulated protocol
  - IP, ARP, RARP
- Data:
  - min. 46 bytes payload (padding may be needed), max 1500 bytes
- CRC (4 bytes)

Direct IP delivery

Problem:

- Ethernet uses 48 bit addresses
- IP uses 32 bit addresses
- we want to send an IP datagram
- but we only can use the Link Layer to do this

ARP

ARP (Address Resolution Protocol)

- Service at the link-level, RFC 826
- maps network-addresses to link-level addresses
- Host A wants to know the hardware address associated with IP address of host B
- A broadcasts ARP message on physical link
  - including its own mapping
- B answers A with ARP answer message
- Mappings are cached: arp -a shows mapping

RARP

RARP (Reverse Address Resolution Protocol)

- maps link-level addresses to network-addresses
- for diskless stations to obtain their own IP address
- Service at the link-level, RFC 903

Host A wants to know its IP address (which is IP_A)

- A broadcasts RARP message on physical link
- RARP server answers with RARP answer
  - containing IP_A

(R)ARP Message

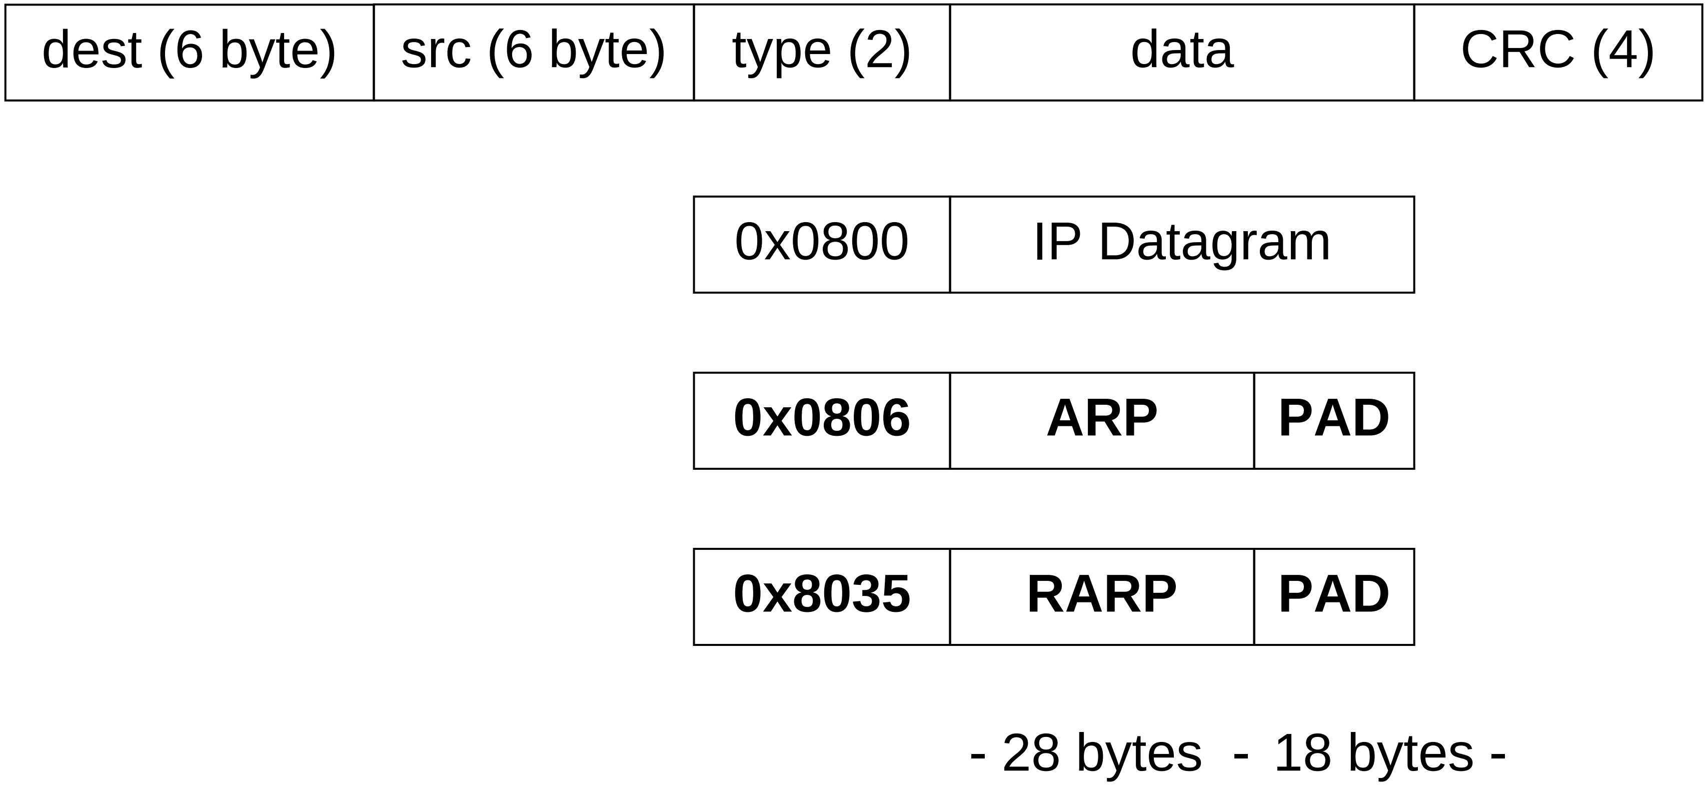

dest (6 byte)	src (6 byte)	type (2)	data	CRC (4)

0x0800	IP Datagram

0x0806	**ARP**	**PAD**

0x8035	**RARP**	**PAD**

- 28 bytes - 18 bytes -

(R)ARP Message

<table>
<tr><td colspan="2">hardware type (2 byte)</td><td>protocol type (2 byte)</td></tr>
<tr><td>hw.adr.size
(1 byte)</td><td>prot. adr. size
(1 byte)</td><td>opcode (2 byte)</td></tr>
<tr><td colspan="3">sender Ethernet address (6 byte)</td></tr>
<tr><td colspan="3">sender IP address (4 byte)</td></tr>
<tr><td colspan="3">target Ethernet address (6 byte)</td></tr>
<tr><td colspan="3">target IP address (4 byte)</td></tr>
</table>

(R)ARP Message

- use same message format

- contain:
  - types and address sizes of hardware and protocol
  - type of message (=opcode, (R)ARP request/reply)
  - link-level and network level addresses of sender and target.

- depending on type, different fields are empty
  - ARP: target link level address
  - RARP: everything except source link-level address

Sending an IP Packet

- Assume host A wants to send an IP packet to host B and that all ARP caches are empty

Procedure

- A sends ARP request for IP-B.
- B sends ARP answer to A
- ARP caches on A+B are filled
- A sends encapsulated IP datagram on link level to B
- datagram is delivered

Direct IP Delivery

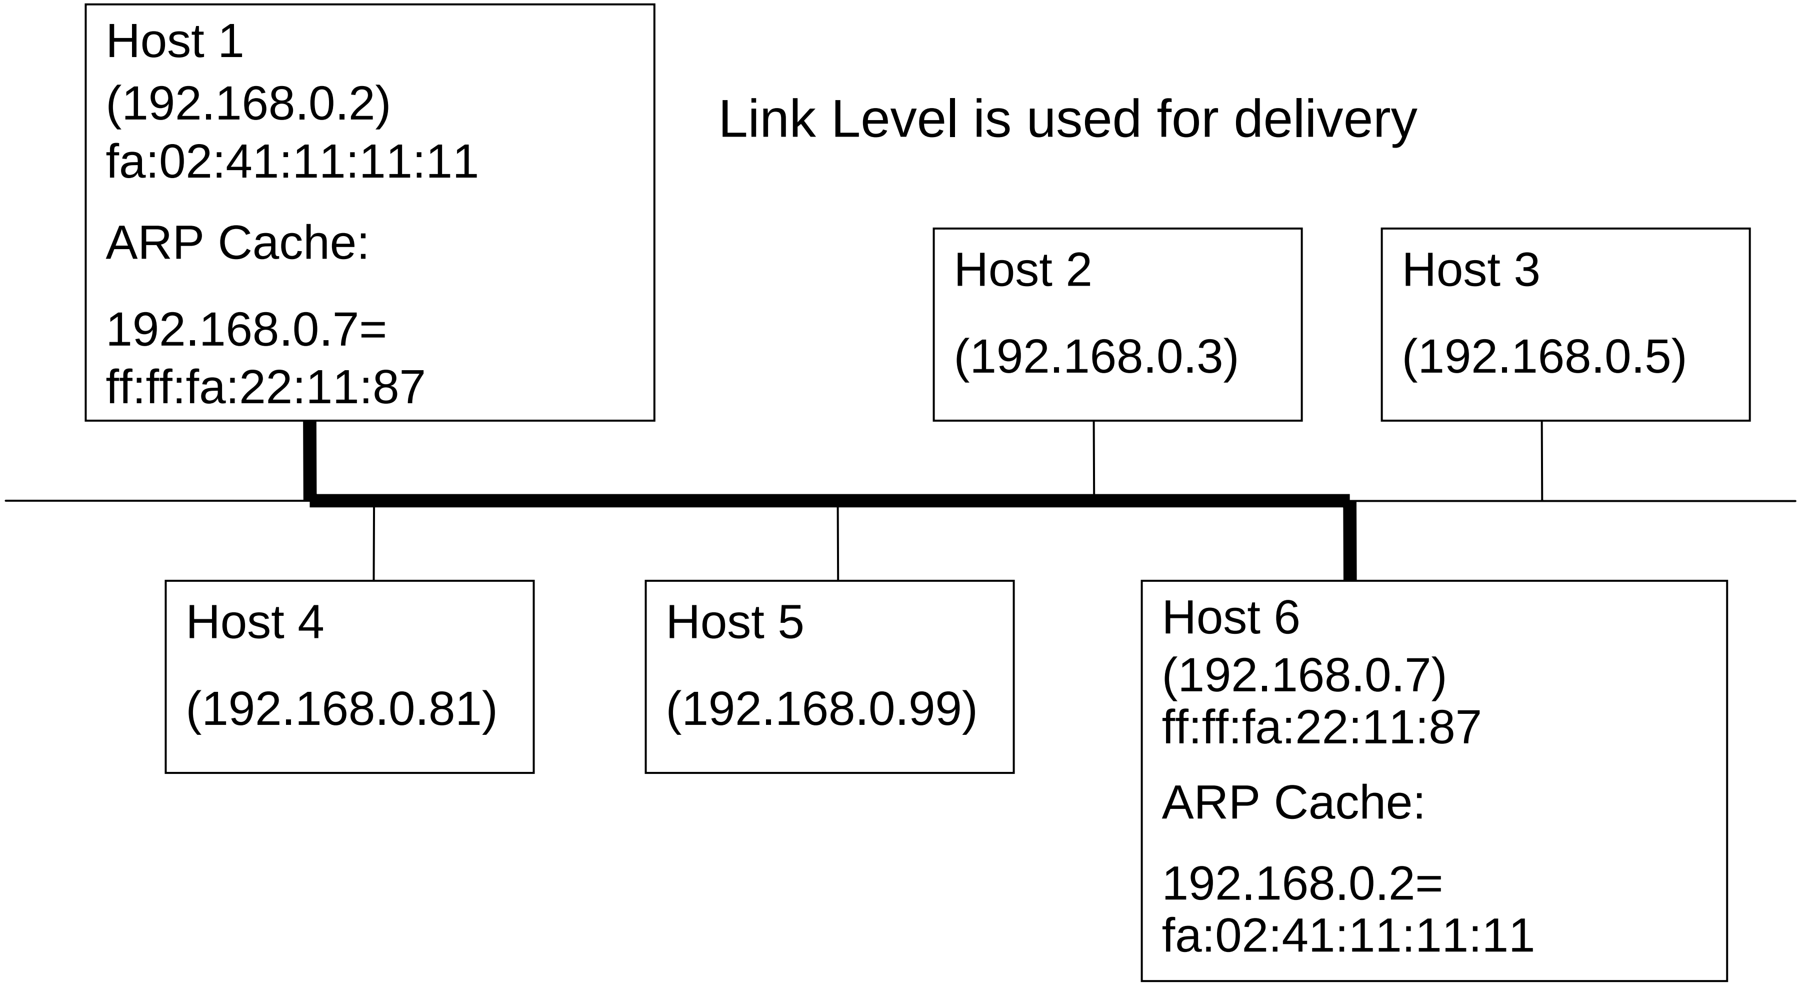

Link Level is used for delivery

Host 1
(192.168.0.2)
fa:02:41:11:11:11

ARP Cache:

192.168.0.7=
ff:ff:fa:22:11:87

Host 2

(192.168.0.3)

Host 3

(192.168.0.5)

Host 4

(192.168.0.81)

Host 5

(192.168.0.99)

Host 6
(192.168.0.7)
ff:ff:fa:22:11:87

ARP Cache:

192.168.0.2=
fa:02:41:11:11:11

LAN Attacks

- Goals:
  - Information Recovery
  - Impersonate Host
  - Tamper with delivery mechanisms

- Methods:
  - Sniffing
  - IP Spoofing
  - ARP attacks

IP Spoofing

= impersonating another host by sending a datagram with a faked IP-address

- used to impersonate sources of security critical info
- explicit address-based authentication
  - RPC, DNS
  - „r-“ commands (rsh, rcp, etc).

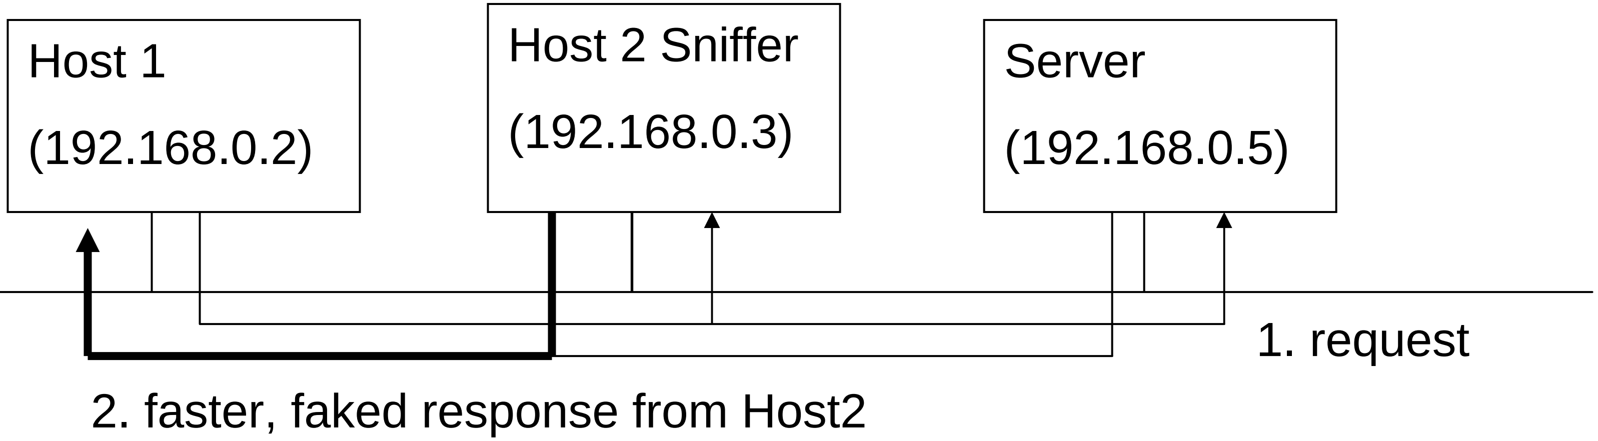

IP Spoofing

How can you do it on your own?

- open a RAW socket
  - socket(AF_INET, SOCK_RAW, IPPROTO_RAW)
- craft the packet
  - with faked IP address
  - including all headers with all attributes set correctly
  - including data
  - including checksums (TCP: required, UDP: recommended)
- send the packet using the RAW socket

ARP Attack 1/2

ARP does not provide any means of authentication

Attacks

- Racing against the queried host is possible
  - provide false IP address/link-level address mapping
- Fake ARP queries
  - used to store wrong ARP mappings in a host cache

=> result in a redirection of traffic to the attacker

ARP messages are sent continuously to have caches keep the faked entries

ARP Attack 2/2

- can be used to impersonate the gateway and filter ALL the traffic
- OR: use ARP to map gateway IP to non-existent MAC address (denial-of-service)

- Tools:
- e.g. WinARP: denial of service against Windows
  - requires the victim to click on many modal dialogs
  - or reboot the machine (1 dialog / ARP packet)

Some Tools

everyone should know (do `man` <toolname> on UNIX/Linux)

- arp
  - service program for the ARP service
- ping
  - check whether a host is alive
- tcpdump
  - check what is going on on the net down to the packet level
- IPtraf
  - check what is going on on the net connection tracing, GUI
- nslookup (dig / host)
  - DNS resolving

The Internet

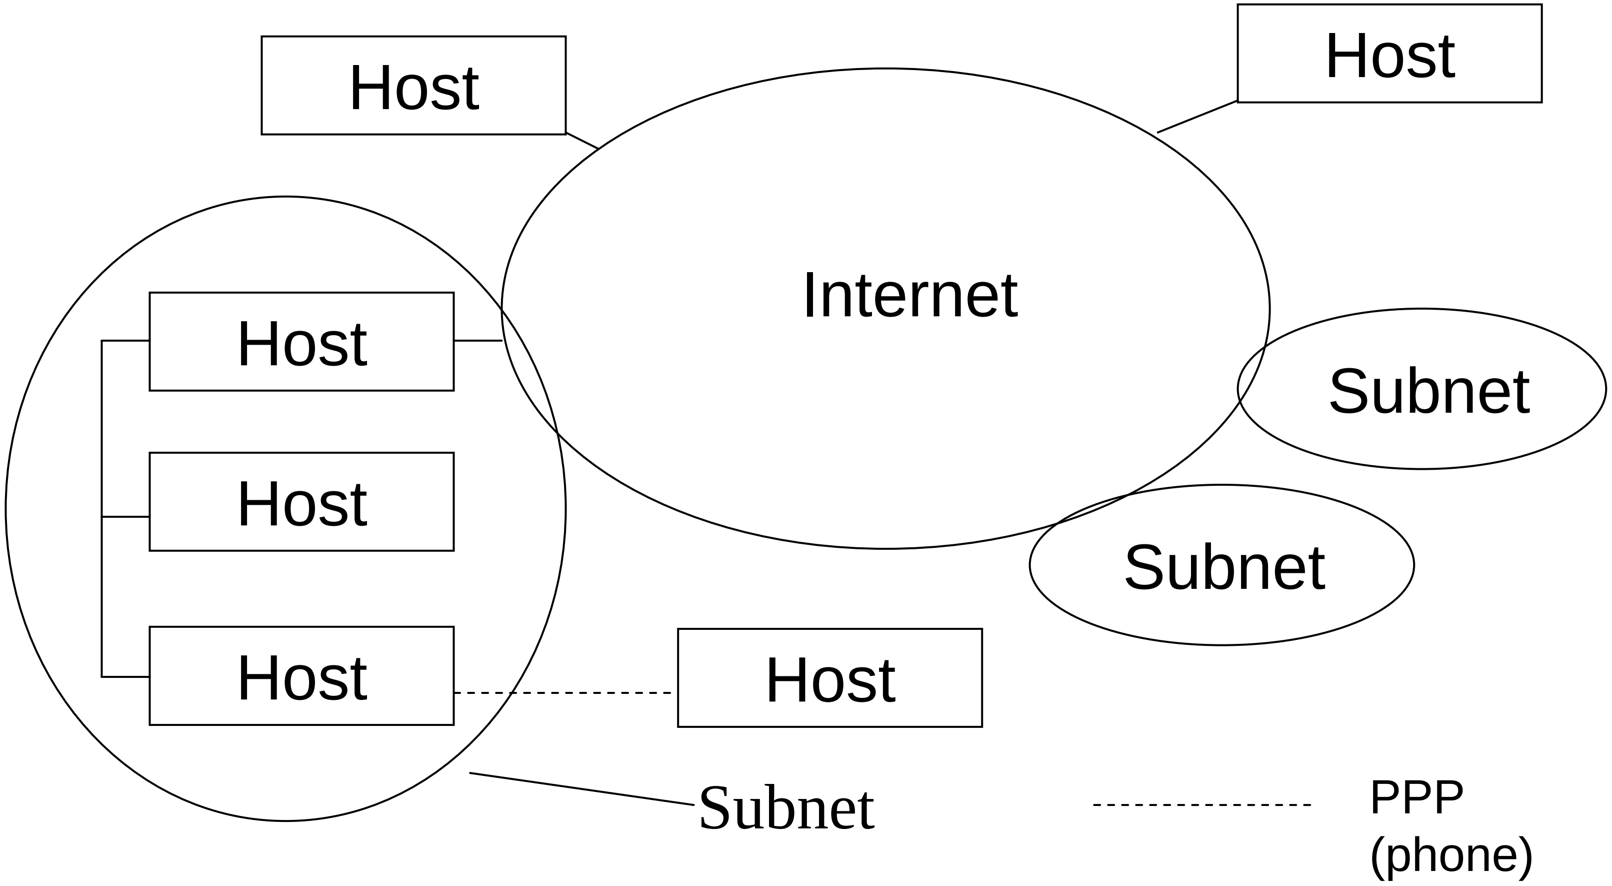

Indirect Delivery: Routing

If hosts are in different physical networks packet can‘t be delivered directly

- Packet is forwarded to a **gateway**
  - has access to other network(s)
  - decides upon destination where to send the packet next
  - this is repeated until packet arrives at network with target host
  - then direct delivery is performed
  - link level addresses change at every step, also TTL field

Indirect Delivery: Routing

Store and forward communication

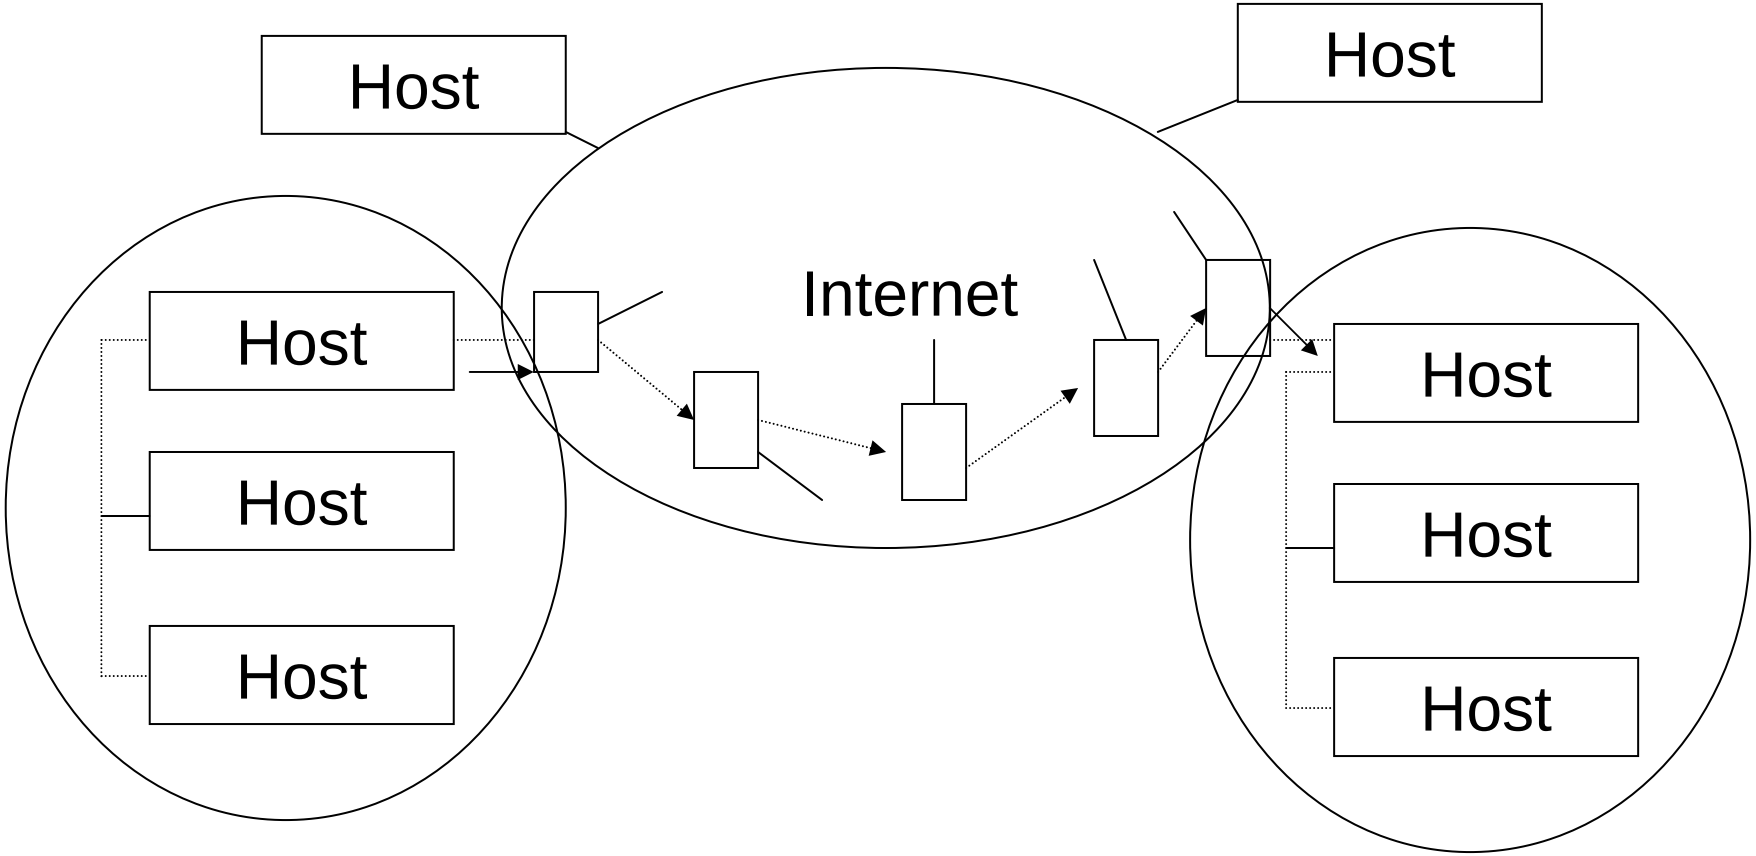

The Routing Table

contains information how to do hop-by hop routing

```
% route -n
Kernel IP routing table
Destination         Gateway             Genmask             Flags     Iface
192.168.1.0         0.0.0.0             255.255.255.0       UH        eth0
127.0.0.1           0.0.0.0             255.0.0.0           U         lo
0.0.0.0             192.168.1.1         0.0.0.0             UG        eth0
```

- Flags:
  - U: the route is up
  - G/H: destination is a gateway/host

Types of Routing

- Hop by hop routing
  - delivery route is determined by the gateways that participate in the delivery process

- Source routing
  - originator of a datagram determines route
  - uses IP source routing option
  - strict source routing: only the specified hosts may be used
  - loose source routing: prefer routing over stated hosts

Routing Mechanism

- Route-daemon searches for
  - matching host address
  - matching network address
  - default entry

- If no route can be found: ICMP message
  - „Host unreachable“ is sent back to originator
- Routing tables can be set
  - statically
  - dynamically (using routing protocols)

Routing Protocols

- automatically distribute information about delivery routes
- hierarchically organized with different scope

- divided in
  - exterior gateway protocols (EGPs)
    - distribute information between different autonomous systems
    - e.g. Border Gateway Protocol (BGP) for Internet backbone
  - interior gateway protocols (IGP)
    - distribute information inside autonomous systems
    - e.g. in LANs
    - e.g. Routing Information Protocol (RIP)

- autonomous means: under a single administrative control

Routing Protocols

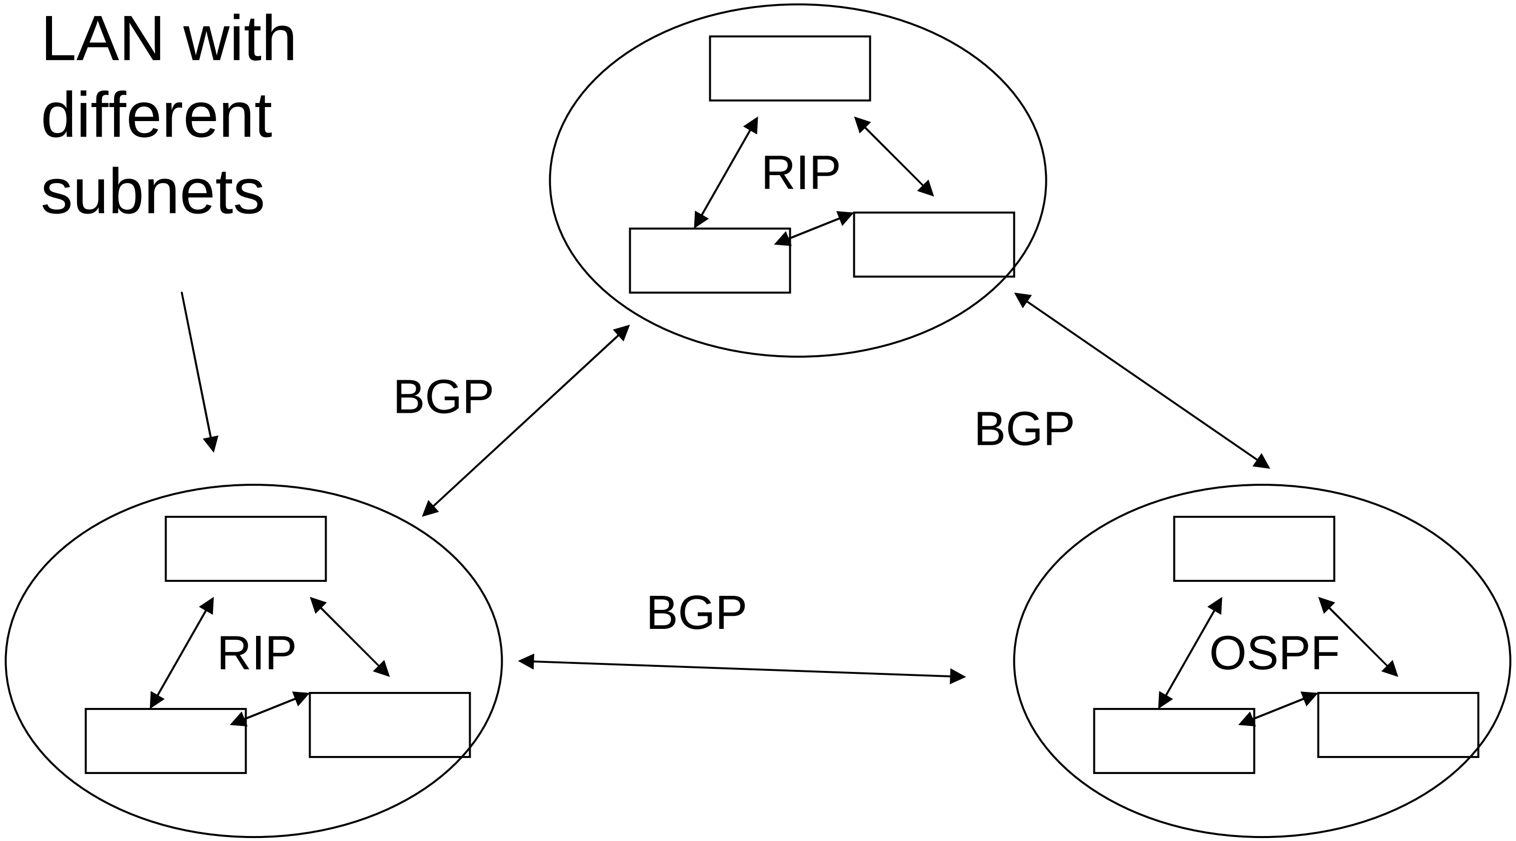

RIP

RIP (Routing Information Protocol)

- uses UDP to transport messages (520)
- no authentication (RIPv1), password in the clear (RIPv2)

- router knows which nets it is connected to
- routers broadcast RIP messages every 30 seconds
  - contains 1-25 advertisements (all its knowledge)
  - each advertisements contains metric: hop count
- only route with smallest hop count is stored in the router
- timeout for routes (3 minutes) if not advertised again

RIP

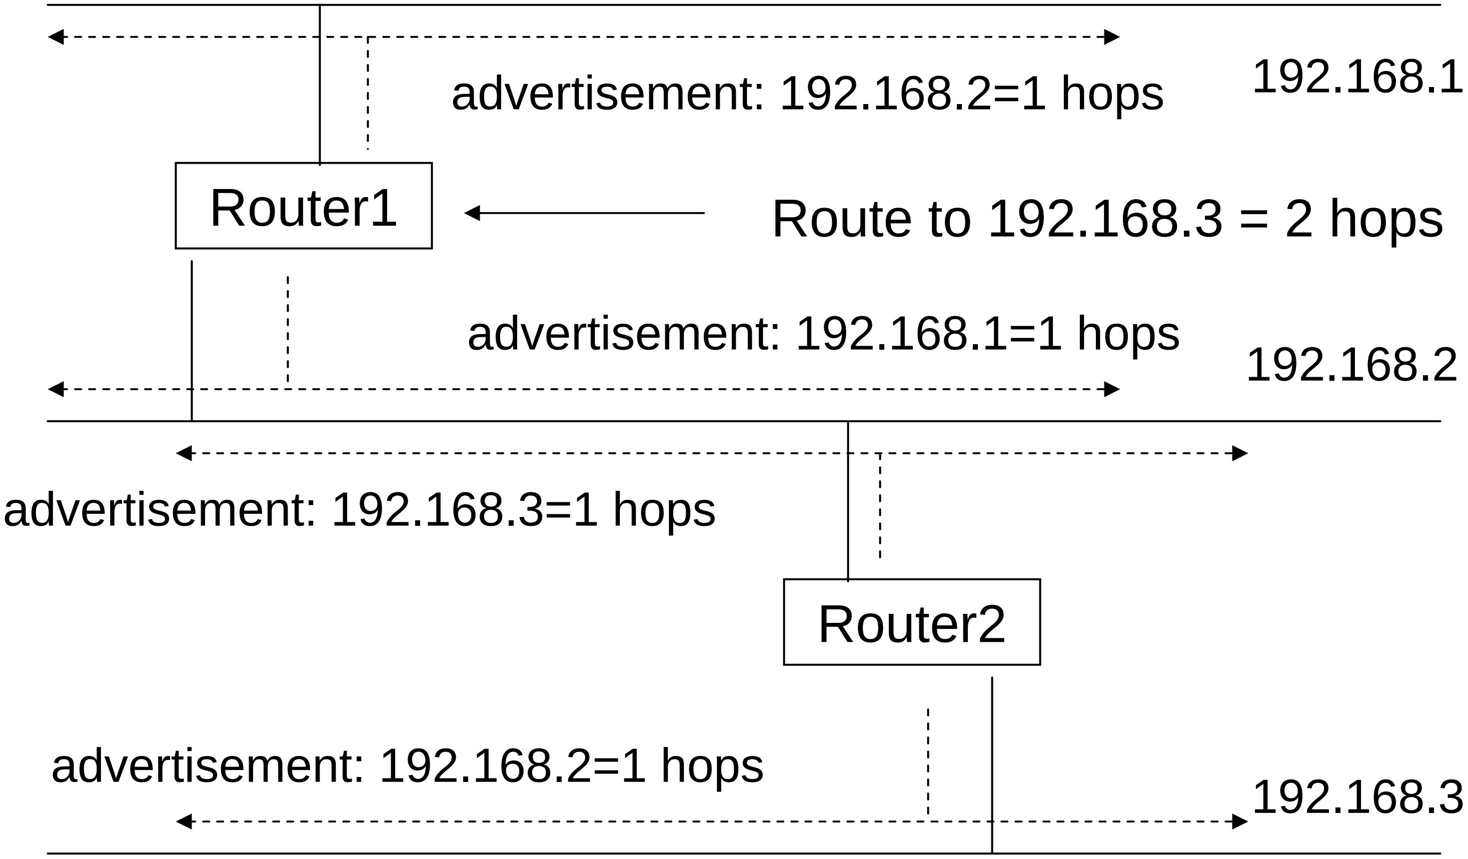

advertisement: 192.168.2=1 hops
192.168.1
Router1
Route to 192.168.3 = 2 hops
advertisement: 192.168.1=1 hops
192.168.2
advertisement: 192.168.3=1 hops
Router2
advertisement: 192.168.2=1 hops
192.168.3

OSPF

OSPF (Open Shortest Path First)

- uses IP datagrams directly
- instead of hop counts, it uses a link-state information
  - each router tests the status of its link to each neighbor
  - then it sends a summary to each of its neighbors
- uses multicast (not broadcast!) for traffic delivery
- It provides a cleartext password authentication

Attacks involving Multiple Networks

- Blind IP spoofing
- Man-in-the-middle-attacks
- Attacks concerning the routing mechanism
  - e.g.RIP attacks

Blind IP Spoofing

- usually the attacker does not have access to the reply, abuse trust relationship between hosts
  - e.g.
    - Host C sends an IP datagram with the address of some other host (Host A) as the source address to Host B
    - attacked host (B) replies to the legitimate host (A)

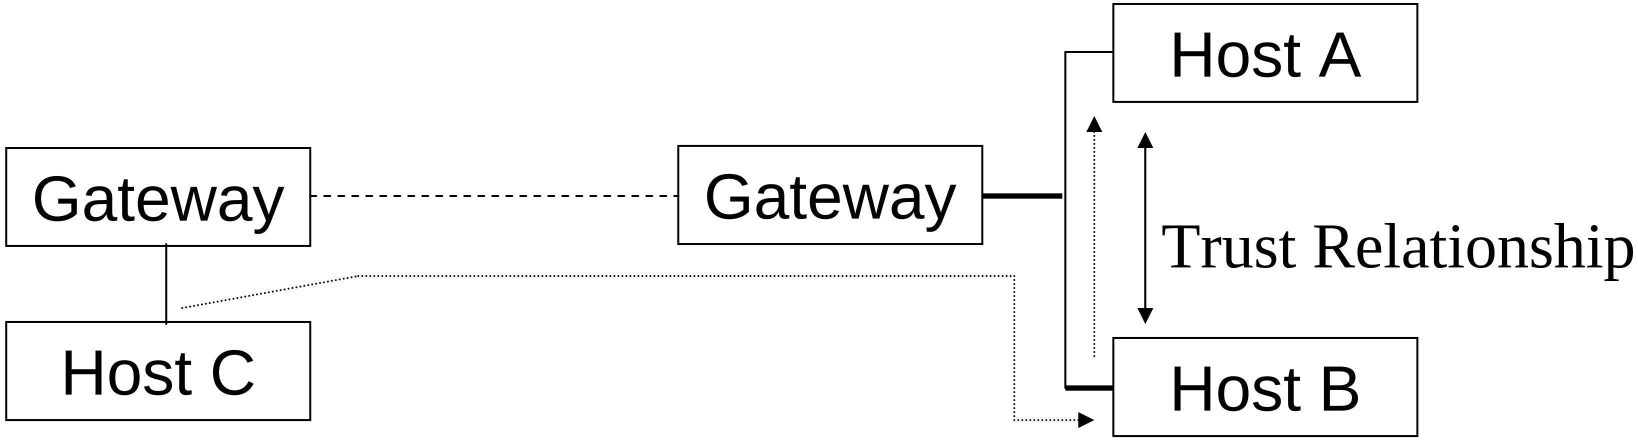

Man-in-the-Middle Attack

Attacker controls a gateway that is used in the delivery process can

- sniff the traffic
- intercept/block/delay traffic
- modify traffic

works only properly if attacker is on „best“ route

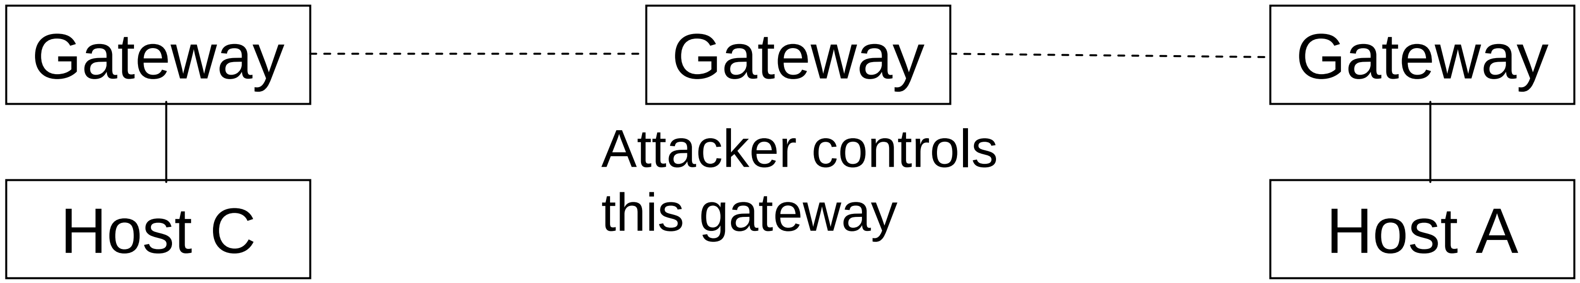

Man-in-the-Middle Attack

- not easy in the Internet because of hop-by-hop routing
- unless you control one of the backbone hosts
- or source routing is used
  - The IP source routing option can be used to specify the route to be used in the delivery process independent of the „normal“ delivery mechanisms
  - the traffic can be forced through specific routes (=specific hosts)
  - if the reverse route is used to reply to traffic, a host on the route can easily impersonate another host
  - can be used to abuse a trust relationship

RIP Attacks

A host can send spoofed RIP packets in order to

- „inject“ routes into a host (requires only IP/UDP spoofing)
- a route with a smaller hop count would be used

- This attack can be used for
  - Hijacking
  - DOS

- On a LAN with RIPv2 passwords have to be used for updating routes, but plaintext passwords are used
  - can be sniffed

Layer 4 Protocols

Many protocols use IP as the underlying network layer

- Important ones are

  ICMP (Internet Control Message Protocol)

  UDP (User Datagram Protocol)

  TCP (Transmission Control Protocol)

ICMP

ICMP (Internet Control Message Protocol)

- is used to exchange control/error messages about the delivery of IP datagrams
- ICMP messages are encapsulated inside IP datagrams

- ICMP messages can be:
  - Requests
  - Responses
  - Error messages
    - includes header and first 8 bytes of offending IP datagram

ICMP Message Format

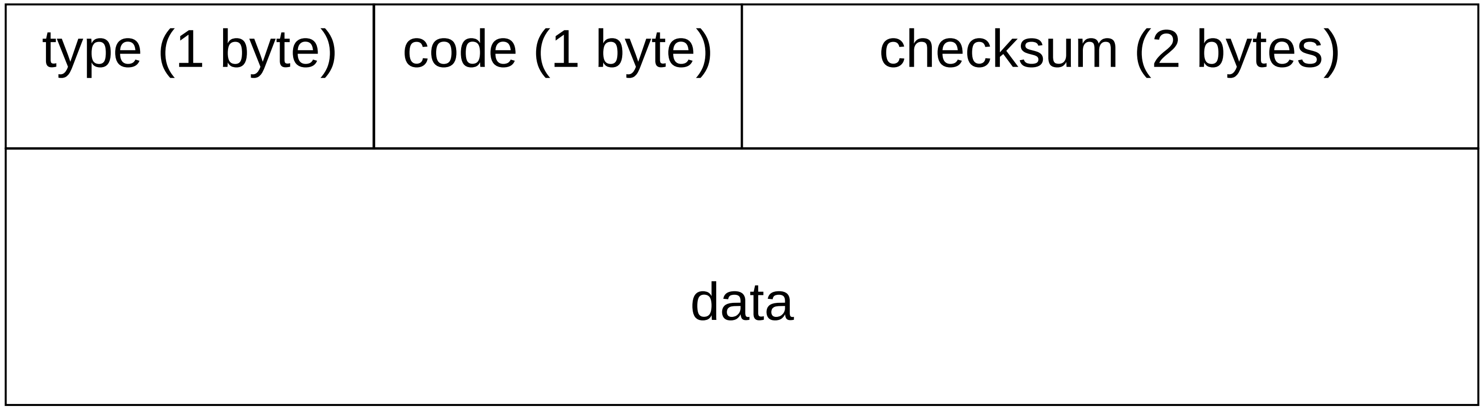

type (1 byte)	code (1 byte)	checksum (2 bytes)
data		

type field: specifies the class of the ICMP message

code field: specifies the exact type of the message

ICMP Messages

- Address mask request/reply
  - used by diskless systems to obtain the network mask at boot time
- Timestamp request/reply
  - used to synchronize clocks
- Source quench
  - used to inform about traffic overloads
- Parameter problem
  - used for inform about errors in the IP datagram fields

ICMP Messages

- Echo request/reply
  - used to test connectivity (ping)
- Time exceeded
  - used to report expired datagrams (TTL=0)
- Redirect
  - used to inform hosts about better routes (gateways)
- Destination unreachable
  - used to inform a host that it is impossible to deliver traffic to a specific destination

ICMP Echo

- Used by the ping program

<table>
<tr><td>type (1 byte)</td><td>code (1 byte)</td><td colspan="2">checksum (2 bytes)</td></tr>
<tr><td colspan="2">identifier (2 bytes) =
Process ID</td><td colspan="2">sequence number (2 bytes)</td></tr>
<tr><td colspan="4">data</td></tr>
</table>

identifier is used by „ping“ to deliver back the packet to the right process (allowing more than one ping to run concurrently)

remember: in ICMP (based on IP) there are no ports

ICMP Echo Attacks

- map the hosts of a network
  - ICMP echo datagrams are sent to all the hosts in a subnet
  - attacker collects the replies and determines which hosts are alive

- denial of service attack (SMURF attack)
  - send spoofed (with victim‘s IP address) ICMP Echo Requests to subnets
  - victim will get ICMP Echo Replies from every machine

Smurf Attack

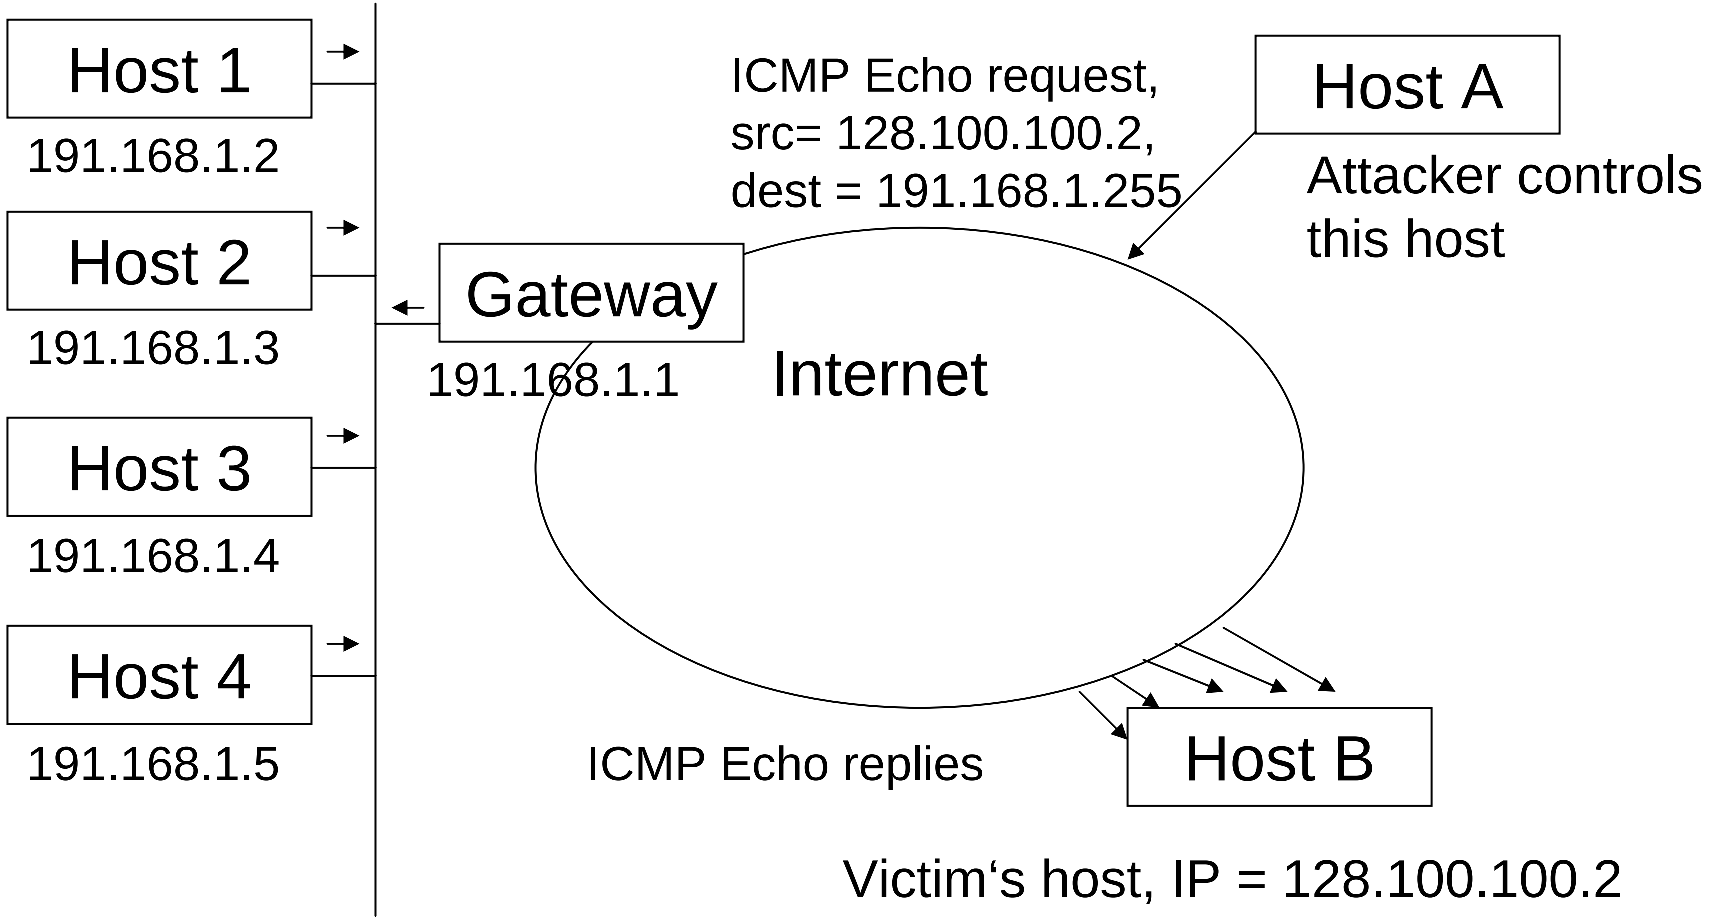

Victim's host, IP = 128.100.100.2

ICMP Redirect

- is used for stating that there is a better route to a host/net
- is sent by a router that routes a packet over the same interface that was used for receiving this packet

type (=5)	code	checksum (2 bytes)
IP address of the router that should be used		
IP header + first 8 bytes of the original datagram		

ICMP Redirect - Example

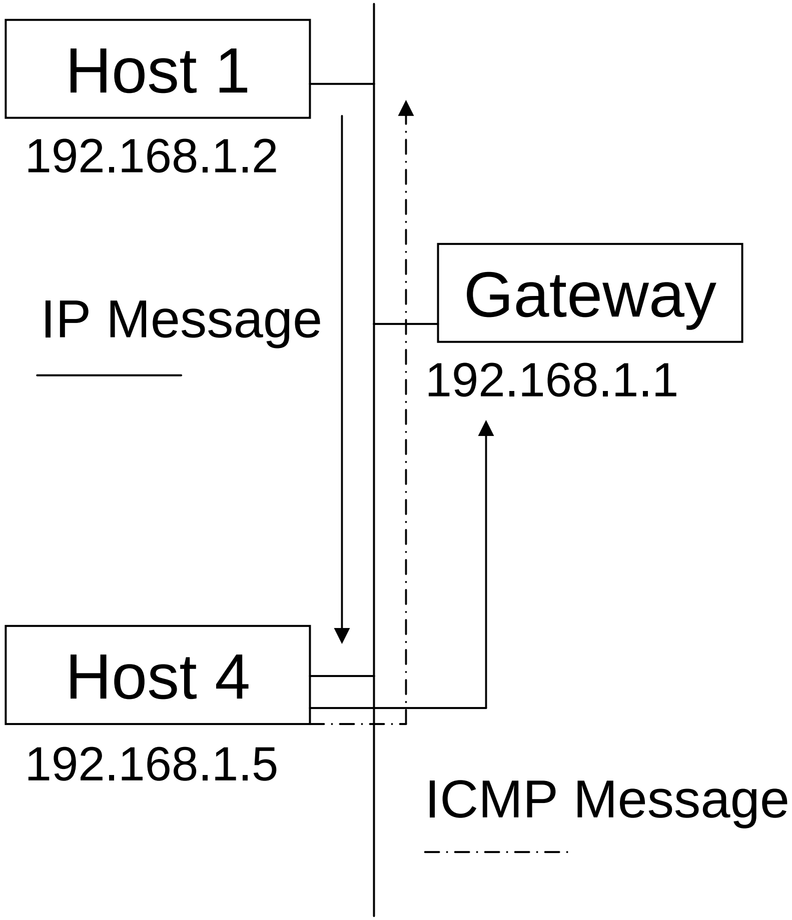

1) In Host1's configuration, it is stated to use Host4 as a gateway. So when Host1 sends a packet outside the subnet, this is forwarded to Host4.

2) Host4 gets the packet, but has to forward the packet to Gateway.

3) Additionally Host4 sends Host1 an ICMP redirect message. „The net xxx can be reached better via Gateway yyy.

ICMP Redirect

- A host that receives an ICMP redirect message checks:
  - whether the new router is directly connected to the network
  - the redirect must be from the current router for this destination
  - the redirect can‘t tell the host to use itself as the router
  - the route that is being modified has to be an indirect route

- What is not checked
  - is message really sent by the current router?
  - is the target host (the new router) a router?

ICMP Redirect Attacks

- ICMP redirect messages can be used to re-route traffic on specific routes or to a specific host that is not a router at all
- The attack is very simple: just send a spoofed ICMP redirect message that appears to come from the host‘s default gateway
- Can be used to
  - Hijack traffic
  - Perform a denial of service attack

ICMP Dest. Unreachable

- ICMP message used by gateways to state that the datagram cannot be delivered
- Many subtypes
  - Network unreachable
  - Host unreachable
  - Protocol unreachable
  - Port unreachable
  - Fragmentation needed but don‘t fragment bit set
  - Destination host unknown
  - Destination network unknown etc.

Dest. Unreachable Attack

- Can be used to „cut“ out nodes from the network
- is a denial of service attack (DOS)

Example

An attacker injects many forged destination unreachable messages stating that 100.100.100.100 is unreachable into a subnet (e.g. 128.100.100.). If 128.100.100.2 net tries to connect to 100.100.100.100, he will immediately get an ICMP Time Exceeded from the attacker‘s host. For 128.100.100.2, this means that there is no way to contact 100.100.100.100, and therefore communication fails.

Dest. Unreachable Attack

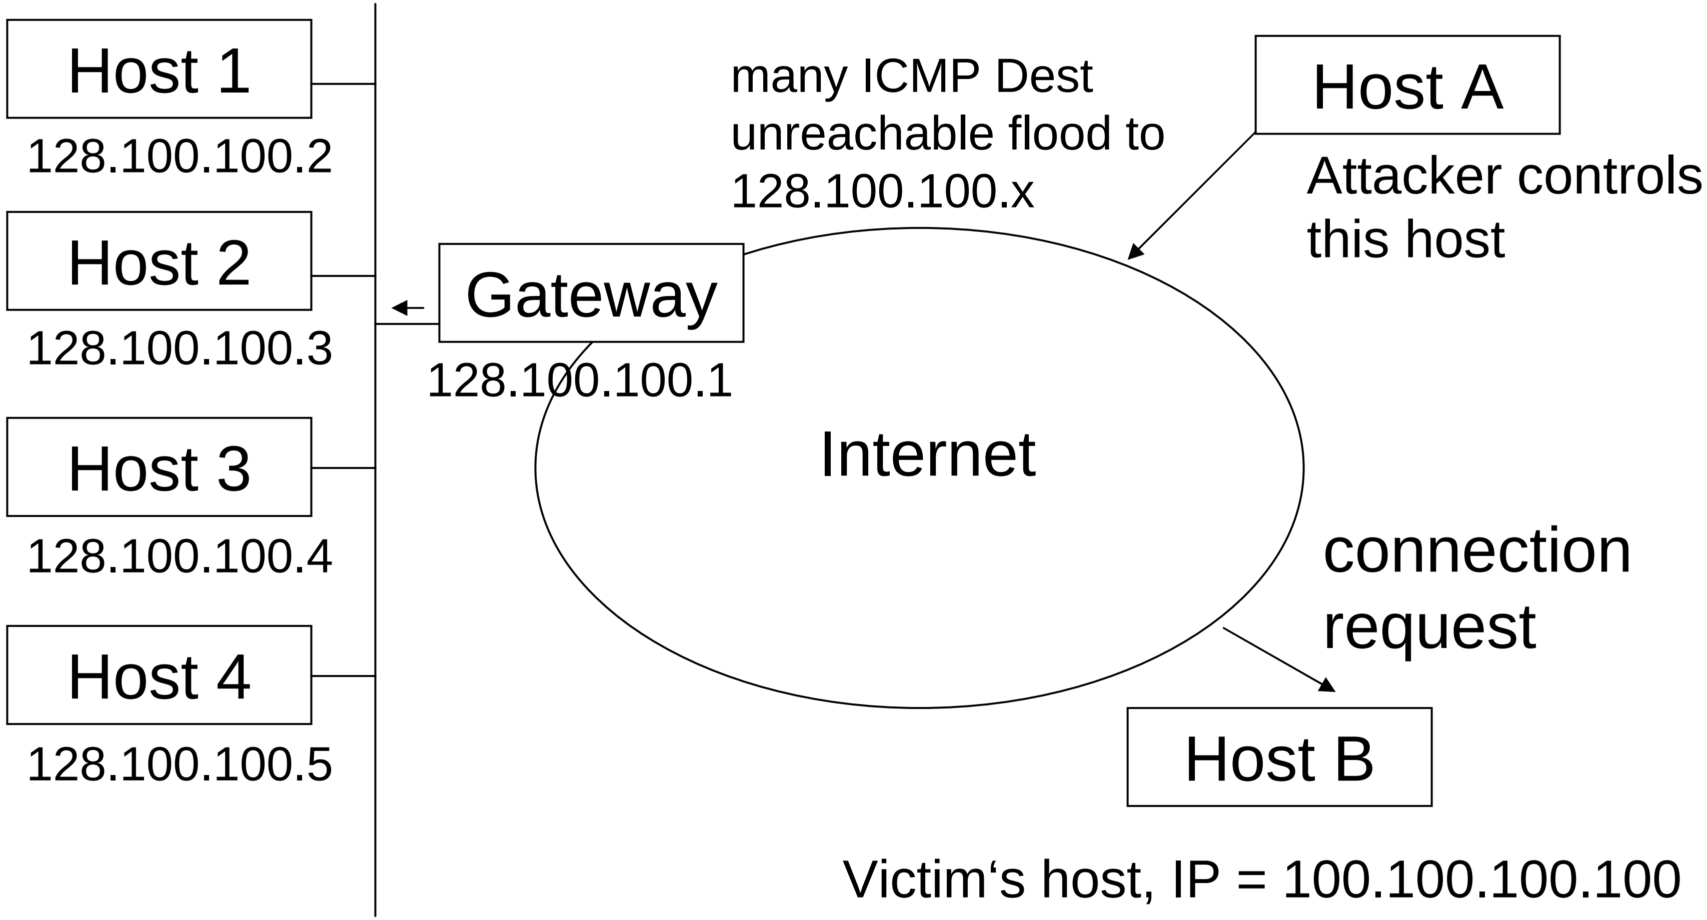

Victim‘s host, IP = 100.100.100.100

ICMP Time Exceeded

Used when

- TTL becomes zero (code =0)
- The reassembling of a fragmented datagram times out (code=1)

type (=11)	code (0 or 1)	checksum (2 bytes)
unused (4 bytes)		
IP header + first 8 bytes of the original datagram		

Traceroute

- Program to determine the path to a specific host/net by evaluating ICMP Time Exceeded messages
- Does this by
  - sending a series of IP datagrams to the destination node
  - each datagram has an increasing TTL field (start=1)
  - gets back ICMP Time Exceeded messages by the intermediate gateways
  - so the full path can be reconstructed by Traceroute
- Traceroute also allows to use loose source routing
- Useful tool for topology mapping