Examination
IK2218 Protocols and Principles of the Internet
EP2120 Internetworking

Date: 07 January 2014 at 08:00–13:00

a) ***No help material is allowed - You are not allowed to use dictionaries, books, or calculators!***
b) *You may answer questions in English or in Swedish.*
c) *Please answer each question on a separate page (not sheet).*
d) *Please write concise answers!*
e) *Put a mark in the table on the cover page for each question you have addressed.*
f) *The grading of the exam will be completed no later than 28 January 2014.*
g) *After grading, EP2120 exams will be available for inspection at STEX (Q-building) and IK2218 exams will be available for inspection online.*
h) *Deadline for written requests for grading review is 11 February 2014.*
i) *Course responsible IK2218 is Peter Sjödin, phone 08-790 4255.*
j) *Course responsible EP2120 is György Dán, phone 08-790 4253.*

Important note!

Your grade is F in any of these two cases:
- if you do not reach at least 10 (ten) points out of 20 for problems 1-4 or
- if you reach less than 30 points in total.

We advise you to start with problems 1-4.

Part I (Problems 1-4)

1. *IP and addressing (5p)*

a) You connect your computer to a LAN in which there is no DHCP server, and there is only one router (the LAN is not a transit network). You use Wireshark to capture the traffic on the network, and observe a number of packets with the following source and destination addresses:

Datagram 1	*S*: 213.130.24.15	*D*: 66.249.91.103
Datagram 2	*S*: 87.248.113.14	*D*: 213.130.29.120

Based on this information you try to manually configure an IP address, default gateway, and netmask for the wireless network interface. What is the longest possible netmask for this subnet? (1p)

b) Give the network address of your subnet in CIDR notation! Choose an IP address for your computer. Try to avoid IP address collision. What could be a reasonable guess for the default gateway? (1 p)

c) What is the directed broadcast address of the subnet assuming the longest possible netmask you identified? (1 p)

d) Explain the difference between IPv6 link local addresses and IPv6 unique local addresses. Can a network interface have both kinds of addresses simultaneously? (1 p)

e) Consider that you have an *IPv6* network running, with a link MTU of 2MB. You would like to make use of the large MTU offered by the link layer technology. What support is there in *IPv6* that you can rely on? (1 p)

a)The addresses 213.130.24.15 and 213.130.29.120 are on your subnet, so the longest netmask is /21.
b) The network address is 213.130.24.0/21. The IP address of your host could be any in the given range, excluding the router address (possibly 213.130.24.1), the directed broadcast and the network address.
c) The directed broadcast address is 213.130.31.255.
d) Link local address is unique on a link, and is not routable. Unique local address is unique within a private network and is routable in that network, but is not routable in the global Internet. Every host must have a link local address, hence a host can have both of these addresses.
e) The IPv6 Hop-by-hop extension header can contain the Jumbo payload option, which allows one to send datagrams of length up to 4GB.

2. *Delivery and address resolution (5p)*

ARP – the Address Resolution Protocol – is primarily used to resolve IPv4 addresses to link layer addresses. ARP typically works with a cache and a number of timeouts.

a) What is the purpose of the ARP cache? What would happen if no cache were used in the ARP protocol? (1 p)

b) The entries in the ARP cache are controlled by timers. There are two variants of ARP cache timeouts that are used in two different situations. Which are the two variants and what is the purpose of each? What happens when the timers expire? (1 p)

Consider the following *IPv4* network consisting of 1 router, 2 bridges and 6 hosts. Hosts H_1 to

H_6 have one interface each. B_1 and B_2 are learning bridges. The router R_1 has an appropriate routing table. All ARP caches and the bridges' learning tables are empty. Assume that ARP snooping is used.

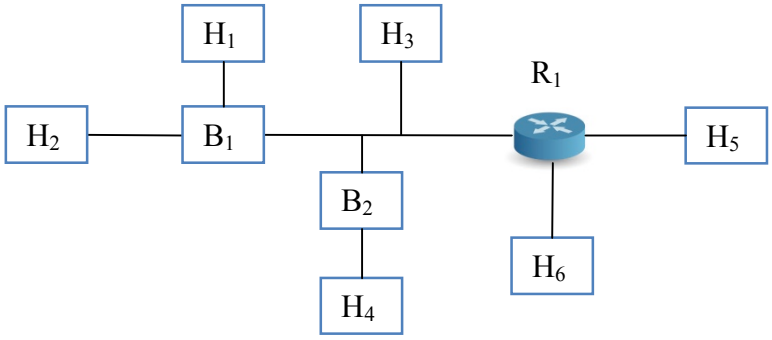

c) Consider the table below. Which of the physical (MAC) and logical (IP) addresses are *not* needed? Identify the subnets of the network! (1 p)

	IP	MAC
H_1	A	a
H_2	B	b
H_3	C	c
H_4	D	d
H_5	E	e
H_6	F	f
R_1	G (*West*) H (*East*) I (*South*)	g (*West*) h (*East*) i (*South*)
B_1	J (*North*) K (*West*) L (*East*)	j (*North*) k (*West*) l (*East*)
B_2	M (*North*) N (*South*)	m (*North*) n (*South*)

d) A process on Host H_5 sends 100 bytes via UDP to a process on host H_2. Show the contents of the learning tables and the ARP caches after the packet has been delivered. Assume that the process on Host H_5 knows the IP address of Host H_2. (1 p)

e) Which protocol is used to perform address resolution in IPv6? (1 p)

a) The role of the ARP cache is to store resolved addresses in order to limit the amount of ARP traffic on a subnet. If resolved addresses were not stored, an ARP request/reply would be needed for every IP datagram transmitted.

b) First, a timer is set for an incomplete entry, when an ARP request has been sent. If no ARP reply is received, the timer expires, a new ARP request is sent, and if no response is received after multiple tries then the entry is cleared eventually. The purpose of this timer is to give up if there is no host on the link with the given IP address.
Second, a timer is set for a complete entry. When the timer expires, the entry is cleared. The purpose of this timer is to reinitiate address resolution in case the remote host crashes, or changes address.

c) Bridges B_1 and B_2 do not need an IP and a MAC address.
Subnet 1: A B C D G
Subnet 2: H E
Subnet 3: I F

d)
H_1 : G-g
H_2: G-g
H_3: G-g
H_4: G-g
H_5: H-h
H_6: -
R_1: E-e, B-b
B_1: g-East, b-West
B_2: g-North, b-North

e) In IPv6 address resolution is done using the neighbor discovery protocol, which is part of ICMPv6 (neighbor solicitation and advertisement messages).

3. *IP forwarding (5p)*

a) Consider an IPv4 network as shown in the Figure below. Hosts A and B are connected by an Ethernet network and they belong to the same subnet. Host C belongs to another subnet, and it is reachable directly from Hosts A and B via the router. The MTU is the same in the entire network and it equals to 1500 bytes.
Host A sends a datagram of size 840 bytes with TTL=64 to Host B and a datagram of size 1200 bytes with TTL=64 to Host C. Hosts B and C receive these datagrams. What is the value of the TTL field that Host B and Host C observe? Has any other field of the base header changed in any of the packets after they left Host A, and if yes, which one(s)? (1p)

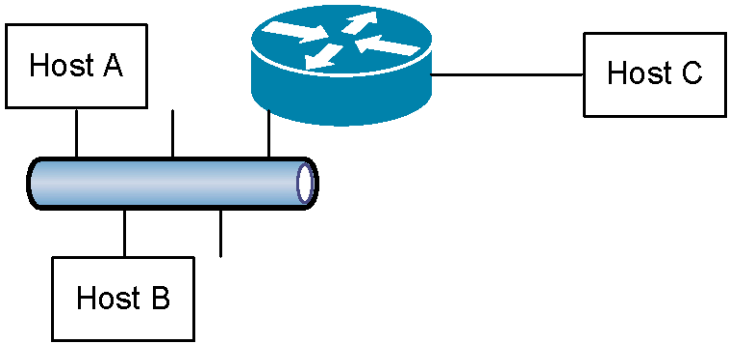

A router has the IPv4 forwarding table shown below. Determine the next-hop address and the outgoing interface for the packets arriving to the router with destination addresses as given in points (b)-(e).

Destination	Next hop	Flag	Interface
133.15.16.0/24	-	U	m0
142.13.0.0/16	-	U	m1
82.93.192.0/18	-	U	m2
171.171.80.0/20	133.15.16.2	UG	m0
160.43.12.0/23	82.93.193.161	UG	m2
82.93.224.0/20	133.15.16.131	UG	m0

160.43.14.0/23	142.13.0.52	UG	m1
0.0.0.0/0	142.13.42.9	UG	m1

b) 171.171.97.134 (1p)
c) 82.93.225.78 (1p)
d) 160.43.16.78 (1p)
e) 82.93.240.189 (1p)

a) Host B observes TTL = 64, Host C observes 63. CRC will be different in the datagram received by Host C due to the change in TTL value; the datagram received by Host B will will be unchanged.
b) 142.13.42.9 on m1 (default route)
c) 133.15.16.131 on m0
d) 142.13.42.9 on m1 (default route)
e) 82.93.240.189 on m2 (direct delivery)

4. *TCP (5p)*

a) During TCP connection establishment the sender and the receiver agree, among others, on the initial sequence number (ISN). Which party chooses the ISN, how does it choose it, and how does the party that chooses the ISN know that the other party is aware of its choice? Assume that a TCP implementation would always choose the same ISN. Give an example of a situation when this could corrupt the data transmission. Support the example with a drawing that shows the segments exchanged between the two processes! (2p)
b) TCP sets a retransmission timer for every sent segment. TCP Tahoe retransmits an unacknowledged segment upon the expiration of the timer. TCP New Reno uses fast retransmit and fast recovery to improve the throughput in the presence of losses. Describe how fast retransmit improves over timer-based error recovery and how it interacts with flow control. How does fast recovery change the way congestion control reacts to losses compared to the case when the retransmission timer expires? (2p)

Consider two hosts, A and B, connected by a network running IPv4. The capacity of all links is 10Mbps and the round trip time is 200ms. A process P_A on host A would like to transmit 40000 bytes to a process P_B on host B using TCP. The path MTU is known to be 1540 bytes. The sender uses a value of 65535 for sshthresh for congestion control. Delayed acknowledgements are not used. The receiver can process the data as fast as they arrive.
c) You observe an average throughput of 6Mbps between the two processes using TCP. Assuming there are no losses in the network, what appears to be the maximum receiver window size used by TCP on host B? (1p)

a) The ISN is chosen by the sender at random, and is included in the SYN segment. The sender knows that the receiver is aware of the ISN through the SYN-ACK segment, which contains ISN+1 as the next byte expected (the ACK field).
If a sender always uses the same ISN then a segment sent in a connection that gets delayed could easily be accepted by the receiver in a subsequent connection. This would lead to the received data being compromised. An example of such a scenario is shown in Fig. 2 in C.A. Sunshine, Y.K. Dalal, "Connection Management in Transport Protocols", Computer Networks, Vol 2, No. 6, 1978, pp. 454-473
b) Fast retransmit: Upon 3 duplicate acknowledgements (ACK with same sequence number) the sender retransmits the segment following the ACKed sequence number. If the 3 duplicate ACKs arrive before the timer would expire this means that data is still arriving at the receiver and only some segments were lost. Following the retransmission, the sender sends one

segment with new data for every duplicate ACK (as the duplicate ACKs signal that segments were received by the receiver, they are not in the network any longer).
Fast recovery: Upon 3 duplicate ACKs the sender sets sshthresh=CWND/2 and once the retransmitted segment is ACKed, the sender sets CWND=sshthresh. The CWND is thus not reduced to 1 MSS (or 3), and therefore the loss event does not trigger slow start.
*c) The maximum receiver window size appears to be $6*10^6$Mbps*0.2s=1.2Mb=150kbyte.*

Part II (Problems 5-12)

5. *UDP and fragmentation (5p)*

a) What are the two major differences between the way fragmentation is implemented in IPv4 and in IPv6? What is the reason for the difference? (hint: where can fragmentation be done and where is the necessary information transmitted) (1p)

b) Consider the IPv4 network shown in the figure below. An application on Host A transmits 2377 bytes of data via UDP to Host B. The UDP header is 8 bytes long, there are no IP options used. Consider the two following scenarios:
   i. Host A knows the path MTU (1000 bytes). (2p)
   ii. Host A assumes that the path MTU is 1500 bytes. (2p)
   In each of the scenarios, how many IP fragments arrive at the router and how many at Host B? Give the segment sizes, the fragmentation offset and the more fragments (MF) bit of all fragments.

a) In IPv6 fragmentation can only be done in the end hosts (not in the routers). In IPv6 the fragmentation related information is carried in an extension header (not in the base header). Fragmentation should be avoided as much as possible, hence the related information is not part of the base header in IPv6.

b) Total data to be sent is 2377+8 = 2385 bytes.
Scenario i)
Both the router and Host B observe the following fragments.
Fragment 1: 976 bytes IP payload, offset=0, MF=1
Fragment 2: 976 bytes IP payload, offset=122, MF=1
Fragment 3: 433 bytes IP payload, offset=244, MF=0

Scenario ii)
The router observes
Fragment 1:1480 bytes IP payload, offset=0, MF=1
Fragment 2: 905 bytes IP payload, offset=185, MF=0
Host B observes
Fragment 1:976 bytes IP payload, offset=0, MF=1
Fragment 2: 504 bytes IP payload, offset=122, MF=1
Fragment 3: 905 bytes IP payload, offset=185, MF=0

Offset given in bytes is acceptable, but it should be clear if the unit is bytes or multiple of 8 bytes (as in the offset field).

6. *Routing (5p)*

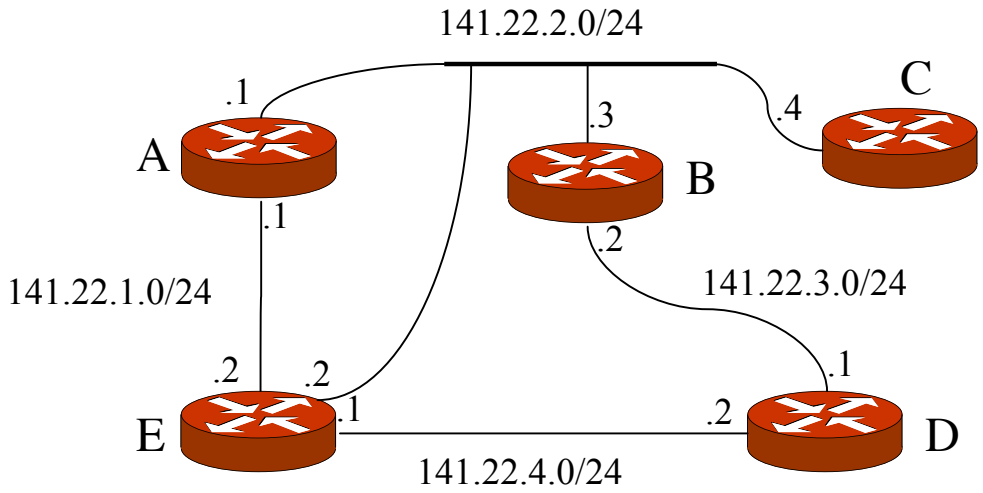

In the IPv4 network shown in the figure all routers A-E run RIPv2 and all link metrics are 1. The addresses of the IPv4 networks and the associated interface addresses are given in the figure. Note that the letters A-E do *not* denote addresses. Assume an initial state for all routers, where only the addresses of the directly connected networks are present in the routing tables. The destinations in the network are the /24 prefixes. Assume also that all RIP implementations support Equal-cost-multi-path (ECMP). All routers implement split-horizon and poison reverse.
Express routes as 'destination, metric, next-hop'. If the destination is a directly connected network, the route is given as 'destination, metric, -'.

a. What is the initial routing state of router A? (1p)

b. Assume that router B starts by sending a RIP response to its neighbours. What is the routing state of E after it has received the initial distance-vector from B? (1p)

c. Assume that the second event that happens in the network is that router E sends RIP responses to its neighbours. Which RIP response messages does E send, and which distance-vectors do they contain? You should indicate source and destination address of each RIP message, on which interface they are sent out (and to where) and which distance-vector (destination, metric tuples) are contained in each message. (1p)

d. What are the routing states of C and D after they have received the distance-vector from E in the previous step (c)? (1p)

e. Assume that the link between E and D fails and that the next event that happens in the network is that router E sends RIP responses to its neighbours. Which RIP response messages does E send, and which distance-vectors do they contain? What is the routing state of D after it has detected the link failure? (1p)

a) 141.22.2.0/24, 1, - #north interface
141.22.1.0/24, 1, - #south interface

b) 141.22.1.0/24, 1, - #north interface
141.22.4.0/24, 1, - #east interface
141.22.2.0/24, 1, - #north-east interface
141.22.3.0/24, 2, 141.22.2.3
c) On the north interface, E sends a RIP response message with src address 141.22.1.2 and destination address 224.0.0.9. Alternatively, if the link between A and E is point-to-point, the destination address may be 141.22.1.1. The distance-vector of this message using split-horizon with poison reverse is:
141.22.4.0/24, 1
141.22.2.0/24, 1
141.22.3.0/24, 2
(141.22.1.0/24, 16 # RIP implementations may announce this network but it is not necessary since all connected routers have this as a directly connected network: it is accepted both to have this route and to omit it)

On the east interface, E sends a RIP response message with src address 141.22.4.1 and destination address 224.0.0.9. Alternatively, if the link between E and D is point-to-point, the destination address may be 141.22.4.2. The distance-vector of this message using split-horizon with poison reverse is:
141.22.1.0/24, 1
141.22.2.0/24, 1
141.22.3.0/24, 2
(141.22.4.0/24, 16 # Same comment as above)

On the north-east interface, E sends a RIP response message with src address 141.22.2.2 and destination address 224.0.0.9. The distance-vector of this message using split-horizon with poison reverse is:
141.22.1.0/24, 1
141.22.4.0/24, 1
141.22.3.0/24, 16
(141.22.2.0/24, 16 # Same comment as above)
d) Routing state of C:
141.22.2.0/24, 1, - # west interface
141.22.3.0/24, 2, 141.22.2.3
141.22.1.0/24, 2, 141.22.2.2
141.22.4.0/24, 2, 141.22.2.2

Routing state of D:
141.22.3.0/24, 1, - #north interface
141.22.4.0/24, 1, - #west interface
141.22.2.0/24, 2, 141.22.3.2
141.22.2.0/24, 2, 141.22.4.1
141.22.1.0/24, 2, 141.22.4.1
e)
On the north interface, E sends a RIP response message containing the following distance-vector:
141.22.4.0/24, 16
141.22.2.0/24, 1
141.22.3.0/24, 2
(141.22.1.0/24, 16)

On the north-east interface, E sends a RIP response message containing the following distance-vector:
141.22.1.0/24, 1
141.22.4.0/24, 16
141.22.3.0/24, 16
(141.22.2.0/24, 16 # Same comment as above)

Routing state of D:
141.22.3.0/24, 1, - #north interface
141.22.4.0/24, 16
141.22.2.0/24, 2, 141.22.3.2
141.22.2.0/24, 16
141.22.1.0/24, 16

7. *Electronic Mail (4 p)*

Suppose that you want to send email messages to three different users: user1@domain1.com, user2@domain2.com, and user3@domain3.com. You could either do this 1) by sending three separate emails, or 2) by sending a single email to three recipients.

a) Explain what the differences would be between the two cases: describe the TCP connections your email client would establish in each case, and explain the communication that takes place over them. (2 p)
b) In case 2), you want to conceal the recipients, so that one recipient cannot see who the other recipients are. Discuss how you could do this, and explain the mechanisms in email protocols that make it possible not to reveal information about who is receiving an email. (2 p)

Note that this question is *not* about how you would use some existing email client (Outlook, Thunderbird, gmail, etc) to achieve the results. Instead, assume that you are acting as the email client yourself and perform the email protocols manually, for instance by using a text-based interface (such as Telnet or PuTTY) in order to communicate with mail servers.

Solution

a) In case 1), three emails are sent from your client to the outgoing SMTP server. So there will be three TCP connections, with one SMTP session over each. (Alternatively, it could also be a single TCP connection, with three SMTP sessions.)

   In case 2), there will be a single transaction with the SMTP server, with one TCP connection and one SMTP session.

b) The delivery of the email is controlled by the "RCPT TO:" SMTP command, which is not visible in the email message delivered to the recipients. The recipients will only see what is in the email header (i.e., "To:" and "Cc:" fields, etc), and since you are using a "raw" telnet interface you can put whatever you like in those fields. You could even leave them empty.

8. *Remote Access (4 p)*

You are working as a network programmer, and you are designing a scheme for communication with a specific server application. The server application is accessed with Telnet over TCP, and requires the user to login with username and password. This is not

acceptable from a security point of view, since the username and password are transmitted in cleartext over the network. Your job is to design a way for clients to access the server application in a secure way, where all communication with the server over the network is encrypted.
You do not have access to the source code for the server application, and you cannot modify it in any way. Hence, you are required to run it exactly as it is. You do, however, have full access to the system where the server application is running. For instance, you can configure the IP address and the port number at which the server application is accepting incoming TCP connections, and you can install and configure other applications to run on the same system.

a) How you would design your solution? Explain how your solution would achieve the desired goal: encrypted communication over the network without modification of the server application. (2 p)
b) To further improve security, you want to prevent external Telnet access to the server application. In other words, it should not be possible access the server over the network using unencrypted communication over Telnet. Hence, your solution should be the only way of accessing the server application. Explain how you could achieve this. (2 p)

Note that the question is about network applications, not encryption, so you should not discuss encryption algorithms or any other cryptographic aspects of the protocols involved.

Solution

a) Implement an application program, *proxy*, on the server that acts like a relay and works as follows: The proxy has a server side that accepts incoming TCP connections from the network. The incoming connections are encrypted. The proxy also has a client, and for each incoming connection it creates a Telnet connection to the server application. The proxy receives data on the incoming TCP connection, decrypts the data, and relays it to the server application over the Telnet connection. Hence, the unencrypted Telnet connection is between two processes on the same machine, and never travels over the network. (You could for instance use ssh in port forwarding mode to accomplish this.)

b) You could configure the server to only accept incoming connections on the loopback interface (IP address 127.0.0.1). Hence, it will not accept incoming connections on external interfaces. You could also install a firewall filter that blocks incoming connections to the server application on external interfaces (but not on loopback). It is sufficient that your solution includes one approach.

9. *DNS (5p)*

You are using the dig tool to do DNS lookups. You run the command `dig www.hackenbush.info` and get a response with the following information in the ANSWER section:

```
www.hackenbush.info   2600   IN    A       130.237.12.18
www.hackenbush.info   2600   IN    A       130.237.12.15
www.hackenbush.info   2600   IN    A       130.237.12.12
```

You run the same command again. Now you get this response:

```
www.hackenbush.info   2600   IN    A       130.237.12.15
www.hackenbush.info   2600   IN    A       130.237.12.12
```

```
www.hackenbush.info   2600   IN    A       130.237.12.18
```

a) Explain what happens here. Why are there several IP addresses for the same name? Why are they returned in different order? (2 p)
b) DNS also makes it possible to have multiple names for the same IP address. Explain why this could be useful. (2 p)
c) Suppose that you instead send a DNS request for a non-recursive query, for instance through the command `dig +norecurse www.hackenbush.info`. This time you get a response with an empty ANSWER section, and where the AUTHORITY and ADDITIONAL sections contain information about root name servers. Why did you get this response instead? (1 p)

Solution

a) The administrators for the domain want to achieve load sharing, most likely. There are three servers. The assumption is that a client will first try with the first IP address on the list, so by rotating the order, different clients will use different IP addresses and thereby different servers.

b) Having multiple names on the same IP address can be used to maintain a level of indirection that is useful for naming services rather than physical servers. For example, the host names "www", "mail", "ns", could be used to represent a web service, mail, and DNS. The names would then translate to the IP addresses of the server(s) where those services are currently located. If a service is moved to a different server, for some reason, the name is kept, but the mapping from name to IP address is changed.

c) The name server to which you send the request (your local name server) does not know the name. Since you have requested the server not to do the complete lookup for you, the server just refers you to some other server that might know more, a root server in this case.

10. *Autoconfiguration (6p)*

d) What is the purpose with a DHCP relay agent? Your answer should describe a scenario where a DHCP relay agent can be useful and explain how it can simplify the administration of the network. (3p)
e) In IPv6 stateless autoconfiguration, the client can create an IP address based on its MAC address instead of requesting it from a DHCP server. Discuss advantages and problems with using an IPv6 address generated from the MAC address and explain how IPv6 privacy extensions address the problems. (3p)

SOLUTION

a) DHCP uses limited broadcast (255.255.255.255) and IP packets destined to this address can't be forwarded by routers from one subnet to another. A DHCP relay agent can relay (tunnel) DHCP messages to/from a DHCP server so that one DHCP server can serve several subnets. This simplifies the administration of large networks since it won't be necessary to maintain a distributed database of valid address mappings.
b) A MAC-derived IPv6 address is a straight forward way to generate a unique IP address automatically and L3/L2 address translation can be done locally by the sender (no ARP needed). The problem is that the MAC address reveals information about the interface card, such as identity and vendor of the interface card so that e.g. potential

bugs could be exploited. IPv6 privacy extensions solve this problem by using a randomly assigned interface ID instead and this number can change over time (temporal address).

11. *IP Gateways—NATs and firewalls (4p)*

a) Consider the figure below. Assume that host 10.0.0.3 on a private network (10.0.0.0/24) sends an HTTP request through its NAT box to a web server (port 80) on address 130.237.15.10 and that this web server answers with an HTTP response back to the host. Fill in source address, source port, destination address, and destination port in the IP packets 1-4 in the figure. Also, fill in the NAT table as it will look when the four packets have been exchanged. (2p)

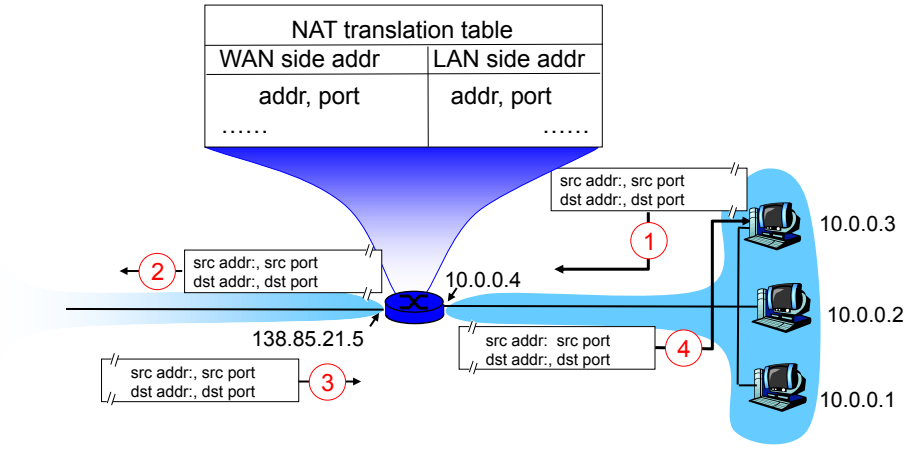

b) Briefly explain how a stateful packet filter works. (2p)

SOLUTION

a)

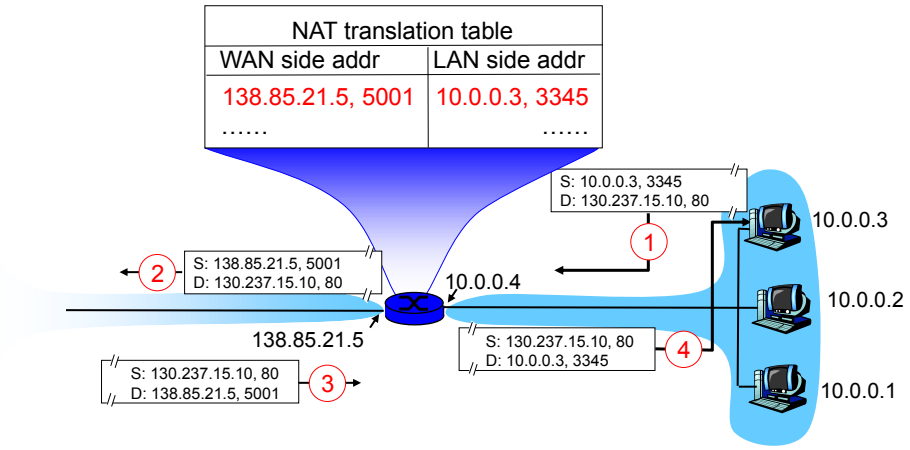

b) A stateful packet filter can identify that a connection was initiated from *s* (internal) to *d* and then allow (for some period of time) connections from *d* to *s*.

12. *IPsec and IKE (7p)*

a) Draw an IP packet where IPsec ESP (Encapsulated Security Payload) is used in *transport* mode for both encryption and authentication. You don't need to show any header fields, just headers/trailers and payload. Mark the parts of the IP packet that are encrypted and the parts that are authenticated. (2p)

b) Explain why ESP encapsulation in *tunnel* mode can cause problems if it used in combination with private addresses and NAT (Network Address Translation). (2p)

c) The following picture illustrates the general idea for IKE phase-1 protocols, main mode. Note that the message exchange is simplified in several ways.

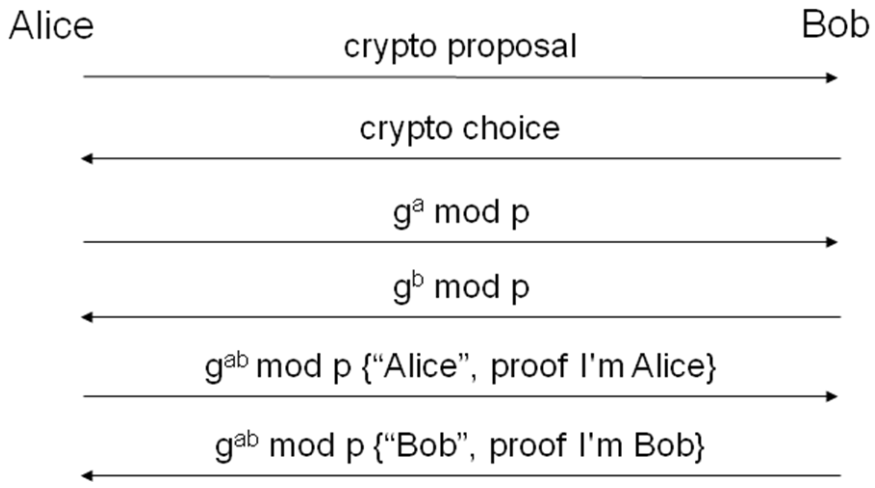

To protect against certain attacks, cookies and nonces are used in IKE. Redraw the figure and show where to add cookies and nonces. Against what type of attack are cookies used? Against what type of attack are nonces used? (3p)

SOLUTION

a) IPsec ESP in transport mode:

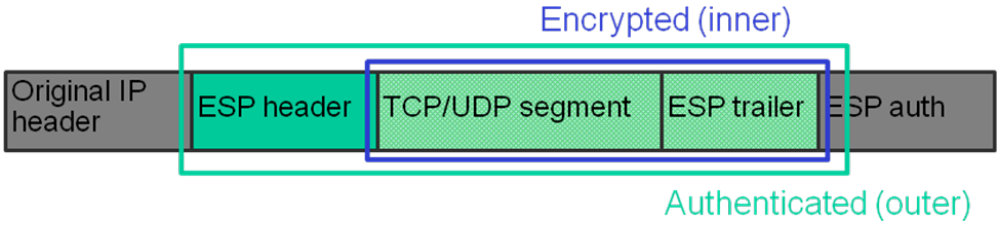

b) If we set up a tunnel between our host and a public gateway, it won't work, since our private addresses will be in the original IP header.

c) Cookies are used to protect against denial-of-service attacks. Nonces are used to protect against replay attacks.

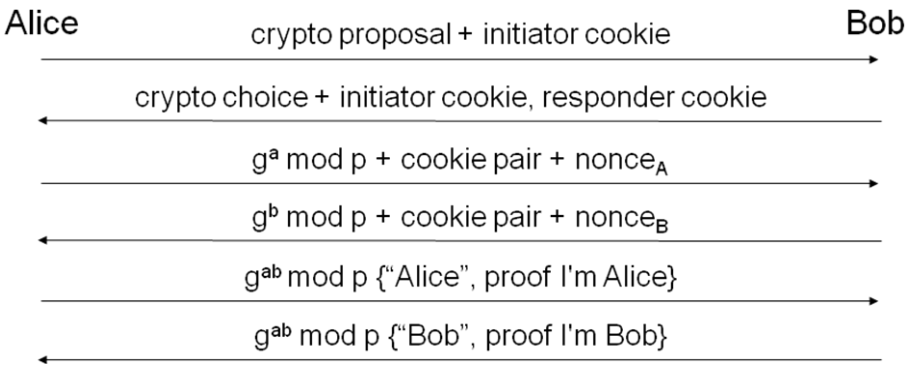
Alice
Bob
crypto proposal + initiator cookie
crypto choice + initiator cookie, responder cookie
g^a mod p + cookie pair + nonce_A
g^b mod p + cookie pair + nonce_B
g^ab mod p {"Alice", proof I'm Alice}
g^ab mod p {"Bob", proof I'm Bob}