Kommunikationssystem FK, Examination 2G1305
Date: March 13th 2004 at 14:00 – 17:00

SOLUTIONS

1. General (5p)

Place each of the following functions/protocols in the correct TCP/IP layer (Application, Transport, Network, or Link/Physical layer): logical host addressing, Cyclic Redundancy Check (CRC), port addressing, end-to-end reliability, network management, name space lookup, SLIP, TCP, IGMP, Spanning Tree Protocol. (5p)

Application: name space lookup
Transport: port addressing, end-to-end reliability, TCP
Network: logical host addressing, IGMP
Link/Physical: CRC, SLIP, Spanning Tree Protocol

Network management can be made in more or less every layer: any layer is an acceptable answer.

2. *Link Layer - VLAN (5p)*

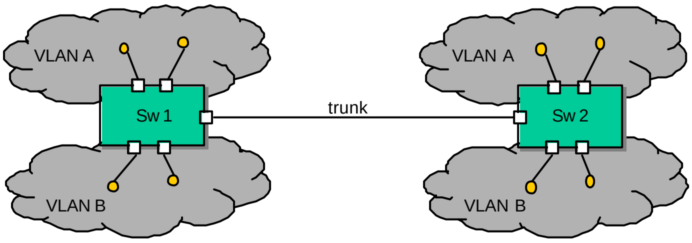

The figure above illustrates two VLAN switches (Sw 1 and Sw 2), two VLANs (A and B), and a VLAN trunk between the switches.

a) How do the two switches distinguish between frames belonging to VLAN A and frames belonging to VLAN B on the VLAN trunk between the switches? (1p)
By examining the VLAN tag in the 802.1Q encapsulation format.

b) How many broadcast domains exist in the network depicted above? (1p)
2

c) What is needed in order to send a packet from VLAN A to VLAN B? (1p)
A router connected to both VLANs.

d) How does the multiple spanning tree protocol differ from the original spanning tree protocol? (2p)
Multiple STP: calculates a spanning tree for each VLAN.
Original STP: calculates one common spanning tree covering all VLANs

3. *IPv4 Addressing (5p)*

For this problem you may need to know the following translations between binary and decimal numbers:

Binary	***Decimal***
1000 0000	128
1100 0000	192
1110 0000	224
1111 0000	240
1111 1000	248
1111 1100	252
1111 1110	254
1111 1111	255

An ISP (Internet Service Provider) is granted a block of addresses starting with 149.70.0.0/16. The ISP creates sub-blocks to customers as follows:

?? 200 medium sites, each needing 200 addresses
?? 80 small sites, each needing 30 addresses
?? 256 households, each needing 5 addresses

a) Design the sub-blocks and give the slash-notation (CIDR notation) for the first and last sub-block in each category (3p)
149.70.0.0/24 – 149.70.199.0/24
149.70.200.0/27 – 149.70.209.224/27
149.70.210.0/29 – 149.70.217.248/29

b) What is the interval of unallocated addresses that the ISP has after the allocation above? (2p)
The interval is 149.70.218.0 – 149.70.255.255 (corresponding to 9728 free addresses)

4. Basic IP routing (5p)

Construct an IPv4 network satisfying the routing table below. The routing table shows the state of one router A. Your assignment is to draw the complete network, including networks, routers and hosts. You may not place a host on a subnet that does not match the host's IP address. Note that there may be many solutions, and you only need to give one.

Destination	***Nexthop***	***Flags***	***Interface***
80.4.5.5/24	-	U	eth0
201.4.0.0/17	80.4.5.1	UG	eth0
193.14.5.160/27	-	U	eth1
201.4.2.4/32	193.14.5.161	UGH	eth1
33.2.3.0/24	-	U	eth2
0.0.0.0/0	33.2.3.1	UG	eth2

Answer:

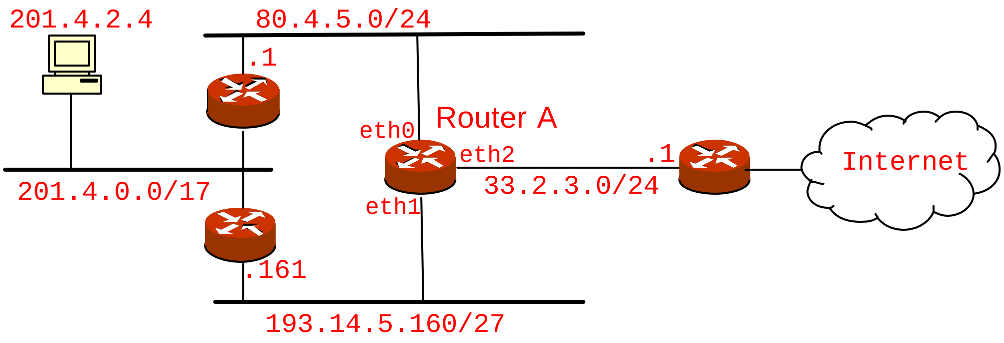

5. ARP (5p)

a) Briefly describe proxy ARP and its purpose? (3p)
Proxy ARP is when a machine responds to an ARP request on another machine's request. ARP requests are broadcast, and will not be forwarded by routers. However, a router replies with its own MAC address to an ARP request on behalf of another node. The router will then forward IP packets destined to that other node. Proxy ARP allows subnets to be hidden.

b) What is the role of the ARP cache (2p)
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 weren't stored, an ARP request/reply would be needed for every IP packet transmitted.
An entry in the ARP cache times out when it has not been used within a certain amount of time. This works fine because of the correlation in use of IP addresses (a sender typically sends many IP datagrams in a row to the same destination).

6. UDP and fragmentation (5p)

For the following problems you may need to know the following:

?? Ethernet MTU is 1500 bytes
?? UDP header length is 8 bytes
?? IP header length is 20 bytes (you can assume there are no IP options)

a) What action does the receiving side of UDP take if the checksum calculation of a UDP datagram results in 0xDEAD (hex)? (2p)
The datagram is silently discarded due to checksum error. The checksum calculation of a correctly received datagram results in 0.

b) A sending UDP protocol peer sends a datagram with 2048 bytes of user data. Sender and receiver are directly connected via Ethernet. How many datagrams will arrive to the receiving UDP protocol peer as a result of the transmission? (2p)
One datagram will arrive to the receiving UDP protocol peer. IP will fragment the datagram at the sending side, but also take care of reassembly at the receiving side before delivering the datagram to UDP.

c) Assume a client uses UDP to send 28 bytes of data to a server. What will be the efficiency (ratio of useful bytes to total bytes) of this transmission at the IP level? (1p)
Efficiency = data/(data + UDP header length + IP header length) = 28/(28 + 8 + 20) = 0.5

7. TCP (5p)

a) The TCP header contains a header length field, but no information about the total length of the TCP segment. How does a receiving TCP protocol peer determine the total length of a TCP segment? (2p)
The IP header holds both a total length field and an IP header length field. Thus, the IP protocol module can compute the length of the TCP segment (total length – header length) and pass this information along with the TCP segment itself to the TCP protocol module.

b) A TCP connection is using a window size of 8,000 bytes. The sending side receives a segment with acknowledgement number 32,001. Give the window at the sending side as a range of byte numbers after the reception of this ACK. (1p)
The range is 32,001 – 40,000.

c) TCP sends a segment at 4:30:20. It does not receive an acknowledgement. At 4:30:28, it retransmits the previous segment. It receives an acknowledgement at 4:30:30. Give the values of both the RTT (Round Trip Time) and the RTO (Retransmission Time-Out) after reception of the ACK according to Karn's algorithm. When the original TCP segment was sent, the RTT was 4 seconds. (2p)
*According to Karn, RTT estimation of a retransmitted segment should not be considered. RTT should not be updated until an ACK is received without the need for retransmission. In addition, exponential back off should be used meaning that the RTO will be doubled for each retransmission. The value of RTO was 8 seconds when the original TCP segment was sent, given by 4:30:28 – 4:30:20 (also given by the formula RTO = 2*RTT = 2*4 = 8 seconds). Thus, the new values of RTT and RTO are as follows:*
RTT_{new} is 4 seconds
*RTO_{new} is $2*RTO_{prev}$ = 2*8 = 16 seconds*

8. Dynamic routing (5p)

OSPF is an intra-domain routing protocol for dynamic routing within an autonomous system

a) How does OSPF handle network topology: how does OSPF partition the network, and what are the limitations of this partitioning? (1p)
OSPF partitions network topologies by *areas*. The backbone area (area 0) may have sub-areas (but sub-areas may not be further partitioned). All traffic must pass through the backbone area. Routers connecting areas are called Area Border Routers.

b) OSPF is said to converge faster than RIP. What does this mean, and why is this the case? (2p)
OSPF is based on link-state routing, whereas RIP uses distance-vector. Link-state routing distributes original link information by flooding to every other node. Thus, every node has complete link information fast. The system can reach a correct routing state in a short time. In contrast, distance-vector uses periodic updates between neighbours to distribute information. In addition, distance-vector recomputes routes – nodes do not have access to original data. Therefore, RIP takes longer time to reach a correct state. Thus, the system may be inconsistent causing routing loops and black holes during a relatively long time.

c) A network consists of three routers running OSPF connected by an Ethernet. Which OSPF link type represents the Ethernet? Which OSPF LSAs do the three routers advertise? (2p)
The Ethernet is a transient link.
The three routers advertise four LSAs. Each router advertises a router link LSA. In addition, the designated router of the transient link advertises a network link LSA.

9. Autoconfiguration – DHCP and DNS (5p)

a) What is the major enhancement with DHCP compared to BOOTP? (2p)
DHCP can provide dynamic configuration of IP addresses, while BOOTP uses a

predetermined mapping between physical addresses and IP addresses. The dynamics means that DHCP can give a client a temporary IP address from a pool of available IP addresses, by issuing a lease for a specified period of time. When the lease expires, the client must either stop using the IP address or renew the lease. DHCP can also provide static IP addresses, like in BOOTP. DHCP is backward compatible with BOOTP.

b) Which are the two DNS message types defined? (1p)
DNS query and DNS response.

c) In which DNS message type are resource records used? (1p)
In the DNS response message.

d) What is a DNS pointer query? (1p)
It is a reverse lookup to find out the hostname of a given IP address.

10. IP QoS (5p)

Describe the two IP QoS (Quality of Service) models developed by IETF (Internet Engineering Task Force). You should:

?? Cover basic characteristics of the two models

?? Point out main differences between the two models when it comes to e.g., resource reservation, signalling, flow granularity, and complexity

Integrated Services and Differentiated Services

int-serv: Model for end-to-end resource reservations for application flows.
Includes signalling protocol (RSVP) and defines 2 (3) service classes;
guaranteed service, controlled load service (and best-effort).
Rather complex approach due to reservation set-up on an end-to-end basis

diff-serv: simpler model, addressing drawbacks of int-serv.
Traffic is divided into a small number of classes and resources reservations are made on a per-class basis. SLAs are used, and there is no signalling involved.
Bits from the ToS field in the IP header are used to mark packets with their traffic class.
Flows are aggregated and all flows within an aggregate are given the same treatment when it comes to priority, dropping, etc.

11. IP Security (5p)

One can classify IP security into four aspects: integrity, authentication, privacy and non-repudiation.

a) Briefly describe each security aspect. (1p)
Integrity – Proof that a message is received exactly as it was sent.
Authentication – The receiver is sure of the sender's identity.
Privacy – The transmitted message is only readable by the receiver.
Non-repudiation – Proof that the sender actually sent the message.

b) Secret key versus Public key encryption/decryption: Give an advantage of each. (1p)
Secret-key – Efficient encryption and decryption algorithms
Public-key – Easier to exchange keys: the public key can be distributed freely; a

smaller number of keys since the same keys can be re-used by everyone communication with the node.

c) Public key encryption is typically made using the global key for encryption and the private key for decryption. But in some scenarios the opposite method is employed: encryption with private key and decryption with public key. Give an example of such a scenario. (1p)
Digital signatures. The sender encrypts the message (or its digest) using the private key. The receiver decrypts the signature using the public key and compares with the original message (or digest).

d) There are two IPSec modes of operation. Briefly describe each mode and in which scenarios the two different modes are useful. Also describe how the original IP header, the ESP and AH headers are handled in the two modes. (2p)
Transport mode and tunnel mode. In transport mode, the security association is established between end-nodes and gives end-to-end security. This useful in single-node scenarios where only the end-nodes are trusted. The original IP header is modified to accommodate the ESP and/or AH headers, and the IP header is not encrypted (in the case of ESP).
In tunnel mode, the security association is established between routers. This is useful in VPN scenarios when trusted sub-networks are connected over an un-trusted internet. A new IP header is added in front of AH and/or ESP and the original IP header. In this way, the original IP header can be encrypted (in the case of ESP)

12. IPv6 (5p)

Three transition strategies for deploying IPv6 on the Internet have been devised by the IETF. Describe each strategy, and which scenario each strategy is intended for. Describe how the address mapping between IPv4 and IPv6 addresses is handled in each case. (5p)

1. Dual stack is used in an initial phase. Every node has both an IPv4 and IPv6 implementation. An application can use one of them depending on who it communicates with. The node then uses either an IPv4 address, or an IPv6 address.
2. Tunnelling is used in a “later” phase when two IPv6 nodes communicate over an IPv4-only network. The IPv6 packet is encapsulated within an IPv4 packet when it reaches an IPv4/IPv6 boundary router. In automatic tunnelling, IPv4-compatible IPv6 (::<IPv4addr>) are used. The boundary router automatically creates an IPv4 header from the IPv6 addresses. In configured tunnelling, the encapsulated IPv4 header uses the IPv4-boundary routers as source and destination addresses.
3. Header translation. Used when a majority of Internet nodes have moved to IPv6, only a few nodes do not understand IPv6. The sender sends an IPv6 packet and a new IPv4 header needs to be created from the IPv6 header. IPv4-mapped IPv6 addresses are used in this case (::FFFF:<IPv4addr>)