

**EE384A: Network Protocols and Standards
Final Examination Solutions**

Closed Book / Closed Notes

Friday, March 19, 1999

12:15 – 3:15 p.m.

NAME: _____

Honor code observed: _____
(Signature)

Assigned Grade:

1)	_____ / 12
2)	_____ / 12
3)	_____ / 10
4)	_____ / 8
5)	_____ / 9
6)	_____ / 9
7)	_____ / 10

Total _____ / 70

Problem 1: Internetworking between Bridges, Routers and IP (12 points)

Consider the network depicted in Figure 1 below. The network is composed of three Ethernet segments. There is a standard transparent bridge between segments 1 and 2, and a standard IP router between segments 2 and 3. Two subnetworks have been defined: subnet 1 comprises segments 1 and 2, and subnet 2 comprises segment 3. The router is properly configured to route between these two subnetworks. There are four nodes in the network, denoted by A, B, C, and D. The MAC and IP addresses for each node are shown in the figure. For example, for node A, its Ethernet MAC address is Ph.A and its IP address is Net1.A. Note that no IP addresses have been assigned to the bridge since it operates at layer 2. Consider that the network is in the following state:

- The ARP caches of all IP devices are empty.
- The two ports of the bridge are in the forwarding state.
- The forwarding database in the bridge is empty.
- All IP nodes are properly configured, i.e., they know their default router.




Figure 1

Question 1a (3 points):

Assume that Node A decides to send an IP packet to Node B. Assume that Node A knows the IP address for Node B (Net1.B). You are asked to show all the packet transmissions in all segments of this network that will happen because Node A wants to send an IP packet to node B. Use the table provided in the next page and fill all the fields. When filling the table, use the following:

- Use the MAC addresses from Figure 1 in the Source and Destination MAC Address columns. Use the word “Broadcast” for broadcast MAC frames.
- In the “Protocol Type” column, provide the name of the protocol (e.g., IP, ARP, etc.).
- For IP packets, use the IP addresses shown in Figure 1 in the Source and Destination IP address columns.
- For ARP packets, put the Sender IP address in the Source IP column, and the Target IP address in the Destination IP column.

Only show the packets related to this communication; you can ignore routing protocol packets and BPDUs. You do not necessarily need to use all the rows provided in the table.

Segment	MAC Addresses		Protocol Type	IP Addresses	
	Source	Destination		Source	Destination
1	Ph.A	Broadcast	ARP	Net1.A	Net1.B
2	Ph.A	Broadcast	ARP	Net1.A	Net1.B
1	Ph.B	Ph.A	ARP	Net1.B	Net1.A
1	Ph.A	Ph.B	IP	Net1.A	Net1.B

Question 1b (3 points):

Repeat Question 1a, assuming now that Node A wants to send an IP packet to Node C.

Segment	MAC Addresses		Protocol Type	IP Addresses	
	Source	Destination		Source	Destination
1	Ph.A	Broadcast	ARP	Net1.A	Net1.C
2	Ph.A	Broadcast	ARP	Net1.A	Net1.C
2	Ph.C	Ph.A	ARP	Net1.C	Net1.A
1	Ph.C	Ph.A	ARP	Net1.C	Net1.A
1	Ph.A	Ph.C	IP	Net1.A	Net1.C
2	Ph.A	Ph.C	IP	Net1.A	Net1.C

Question 1c (4 points):

Repeat Question 1a, assuming now that Node A wants to send an IP packet to Node D.

Segment	MAC Addresses		Protocol Type	IP Addresses	
	Source	Destination		Source	Destination
1	Ph.A	Broadcast	ARP	Net1.A	Net1.R
2	Ph.A	Broadcast	ARP	Net1.A	Net1.R
2	Ph.R1	Ph.A	ARP	Net1.R	Net1.A
1	Ph.R1	Ph.A	ARP	Net1.R	Net1.A
1	Ph.A	Ph.R1	IP	Net1.A	Net2.D
2	Ph.A	Ph.R1	IP	Net1.A	Net2.D
3	Ph.R2	Broadcast	ARP	Net2.R	Net2.D
3	Ph.D	Ph.R2	ARP	Net2.D	Net2.R
3	Ph.R2	Ph.D	IP	Net1.A	Net2.D

Question 1d (2 points):

Show the bridge's forwarding database after all the activity in questions 1a to 1c (i.e., after Node A has sent an IP packet to Nodes B, C and D). Use the table provided below. You do not necessarily need to use all the rows.

MAC Address	Port
Ph.A	1
Ph.B	1
Ph.C	2
Ph.R1	2

Problem 2: OSPF (12 points)

Consider the Autonomous System depicted in Figure 2 (next page). It consists of a backbone and two areas. Network N1 is a 100 Mb/s Ethernet, while networks N2, N3, N4 and N5 are 10 Mb/s Ethernets. The backbone is comprised of networks N1 and N2. Connected to N1 is an AS Border Router (ASBR) with a route to an external network N6, metric 15 (type 1)¹. Area 1 comprises networks N3 and N4 and is connected to the backbone by router R1. Area 2 comprises network N5 and is connected to the backbone by router R2. Router R1 is the designated router for networks N1, N2, N3 and N4. The default metric is used for routing in this AS (i.e., use a metric of 1 for the 100 Mb/s network and 10 for the 10 Mb/s networks – they are 10 times slower than the 100 Mb/s network).

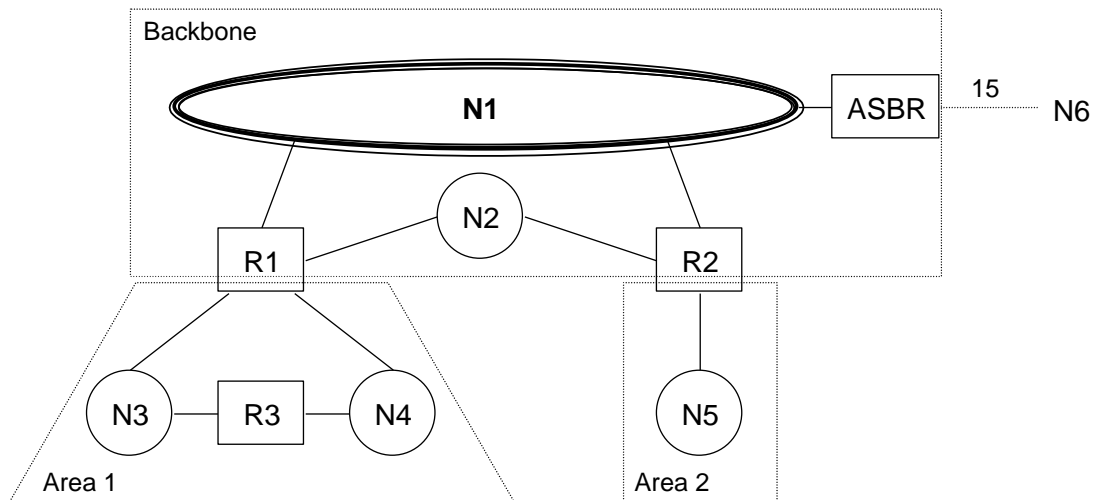


Figure 2

¹ Hint: Type 1 external metrics are equivalent to the link state metrics.

Question 2a (4 points):

What are the link state advertisements issued by router R1 into Area 1 and into the two backbone networks N1 and N2? Specify the type of LSA and most relevant information communicated in each LSA.

R1's LSAs into backbone network N1:

<i>Router Link</i>	<i>connecting R1 to transit network N1</i>	<i>metric = 1</i>
<i>Router Link</i>	<i>connecting R1 to transit network N2</i>	<i>metric = 10</i>
<i>Network Link</i>	<i>for IP network N1 listing routers R1, R2 and ASBR</i>	
<i>Summary Link</i>	<i>for IP network N3 in Area 1</i>	<i>metric = 10</i>
<i>Summary Link</i>	<i>for IP network N4 in Area 1</i>	<i>metric = 10</i>

R1's LSAs into backbone network N2:

<i>Router Link</i>	<i>connecting R1 to transit network N1</i>	<i>metric = 1</i>
<i>Router Link</i>	<i>connecting R1 to transit network N2</i>	<i>metric = 10</i>
<i>Network Link</i>	<i>for IP network N2 listing routers R1 and R2</i>	
<i>Summary Link</i>	<i>for IP network N3 in Area 1</i>	<i>metric = 10</i>
<i>Summary Link</i>	<i>for IP network N4 in Area 1</i>	<i>metric = 10</i>

R1's LSAs into Area 1:

<i>Router Link</i>	<i>connecting R1 to transit network N3</i>	<i>metric = 10</i>
<i>Router Link</i>	<i>connecting R1 to transit network N4</i>	<i>metric = 10</i>
<i>Network Link</i>	<i>for IP network N3 listing routers R1 and R3</i>	
<i>Network Link</i>	<i>for IP network N4 listing routers R1 and R3</i>	
<i>Summary Link</i>	<i>for IP network N1 in Backbone Area</i>	<i>metric = 1</i>
<i>Summary Link</i>	<i>for IP network N2 in Backbone Area</i>	<i>metric = 10</i>
<i>Summary Link</i>	<i>for IP network N5 in Area 2</i>	<i>metric = 11</i>
<i>External Link</i>	<i>for external network N6 advertised by ASBR</i>	<i>metric = 16</i> <i>(type 1)</i>

Question 2c (4 points):

Now, assume that router R1's connection to the 100 Mb/s Ethernet network N1 has failed, as depicted in Figure 3 below. Show the new routing table for router R1, after the routers have realized that the link failed, and the network has reached steady state again.




Figure 3

Type	Destination	Area	Path Type	Cost	Next Hop	Advertising router
N	N3	1	Intra	10	*	*
N	N4	1	Intra	10	*	*
N	N1	0	Intra	11	R2	R2
N	N2	0	Intra	10	*	*
BR	R2	0	Intra	10	*	*
ASBR	ASBR	0	Intra	11	R2	*
N	N5	0	Inter	20	R2	R2
N	N6	*	Type 1 ext.	26	R2	ASBR

Problem 3: Multiple-choice questions (10 points)

Each of the multiple-choice questions will be graded as follows:

- 1 point if you answer it correctly.
- 0 points if you leave it blank.
- -0.25 points if you answer it wrong (i.e., you lose 0.25 points).

Question 3a: The Open Shortest Path First (OSPF) routing protocol is better than the Routing Information Protocol (RIP) because:

- | | |
|-------------------------------------|---|
| <input type="checkbox"/> | a. Since it looks first at the shortest paths, it finds shorter paths than RIP. |
| <input type="checkbox"/> | b. The route computation in the routers is simpler. |
| <input type="checkbox"/> | c. The protocol itself is simpler. |
| <input checked="" type="checkbox"/> | d. Convergence in case of network changes is faster. |
| <input type="checkbox"/> | e. OSPF runs directly over IP, while RIP runs over UDP. |

Question 3b: Regarding to the ICMP Source Quench Message, indicate which of the following statements is *false*:

- | | |
|-------------------------------------|---|
| <input type="checkbox"/> | a. The Source Quench Message is used for congestion control. |
| <input checked="" type="checkbox"/> | b. It can only be sent by the routers in the path of the datagram, never by hosts. |
| <input type="checkbox"/> | c. Hosts are allowed to ignore the Source Quench Message. |
| <input type="checkbox"/> | d. A router is allowed to send Source Quench messages before it experiences congestion. |
| <input type="checkbox"/> | e. A router is allowed to send a Source Quench Message for each datagram it drops. |

Question 3c: Indicate which of the following statements about IP Multicast is *true*:

- | | |
|-------------------------------------|--|
| <input type="checkbox"/> | a. Only members of a multicast group can send to that group. |
| <input type="checkbox"/> | b. A multicast group needs to be created in the router before hosts can join it. |
| <input checked="" type="checkbox"/> | c. From a protocol standpoint, there is no limit to the number of groups a host can join. |
| <input type="checkbox"/> | d. To send a multicast packet on Ethernet, a host first needs to send an ARP packet to find the group's MAC address. |
| <input type="checkbox"/> | e. All members of a group must always respond to a Group Membership Query message from the router. |

Question 3d: On the subject of multicast routing, one advantage of Reverse Path Forwarding (RPF) over Core-Based Trees (CBT) is:

- | | |
|-------------------------------------|---|
| <input checked="" type="checkbox"/> | a. RPF uses multiple trees, which avoids traffic concentration. |
| <input type="checkbox"/> | b. RPF requires the routers to keep less state information than CBT. |
| <input type="checkbox"/> | c. RPF is independent of the unicast routing tables, while CBT uses them. |
| <input type="checkbox"/> | d. Since RPF knows the reverse path, it does not make use of flooding. |
| <input type="checkbox"/> | e. None of the above, there are no advantages of RPF over CBT since they are derived from the same basic algorithm. |

Question 3e: Regarding the Protocol-Independent Multicast (PIM) routing algorithms, indicate which of the statements below is **false**:

- | | |
|-------------------------------------|---|
| <input type="checkbox"/> | a. PIM-Dense Mode is used for situations where the likelihood that an area includes a member is high. |
| <input type="checkbox"/> | b. PIM-Sparse Mode is used for situations where the likelihood that an area includes a member is low. |
| <input type="checkbox"/> | c. PIM-Dense Mode uses Reverse Path Forwarding. |
| <input type="checkbox"/> | d. PIM-Sparse Mode uses an algorithm similar to CBT. |
| <input checked="" type="checkbox"/> | e. Since the PIM is protocol-independent, it can be used for protocols other than IP. |

Question 3f: Of the statements below, select the one that **best** represents the ATM Traffic Contract:

- | | |
|-------------------------------------|---|
| <input type="checkbox"/> | a. When requesting a connection, a node tells the switch its required traffic parameters and quality of service; if the call is accepted, the switch will make sure that the traffic parameters and the quality of service are satisfied. |
| <input type="checkbox"/> | b. When requesting a connection, a node tells the switch its traffic parameters and required quality of service; if the call is accepted, the switch must always guarantee the quality of service. |
| <input type="checkbox"/> | c. When requesting a connection, a node tells the switch its traffic parameters and required quality of service; if the call is accepted, the switch will vary the traffic parameters in order to guarantee the quality of service. |
| <input checked="" type="checkbox"/> | d. When requesting a connection, a node tells the switch its traffic parameters and required quality of service; if the call is accepted, the switch will guarantee the quality of service as long as the host stays within the negotiated traffic parameters. |
| <input type="checkbox"/> | e. When requesting a connection, a node tells the switch its traffic parameters and required quality of service; if the call is accepted, the switch will guarantee that the traffic parameters are met as long as the host stays within the negotiated quality of service. |

Question 3g: Of the ATM functions listed below, select the one that *is not* performed in the ATM Physical Layer:

- a. HEC generation and verification.
- b. Cell Delineation.
- c. Scrambling/descrambling.
- d. Line Coding.
- e. Traffic Shaping.

Question 3h: Of the ATM functions listed below, select the one that *is not* performed in the ATM Layer:

- a. Multiplexing among different ATM connections.
- b. Breaking SDUs into ATM cells.
- c. Congestion control.
- d. Selective cell discarding.
- e. Loss priority indication.

Question 3i: Of the ATM functions listed below, select the one that *is not* a function of Adaptation Layer 1:

- a. Transfer of constant bit rate data.
- b. Transfer of end-to-end timing information.
- c. Forward Error Correction.
- d. Clock recovery.
- e. Transfer of variable-size SDUs at constant bit rate.

Question 3j: Regarding LAN Emulation over ATM version 1.0, indicate which of the statements below is *false*:

- a. The LE Client is required to respond to LE_ARP requests that match its MAC address.
- b. If two LE Clients are connected by a Data Direct VCC, one of them can send a LE_ARP request to the other through this VCC to learn its address.
- c. A LE Client is allowed to drop its Configuration Direct VCC after it has retrieved the LES address from the LECS.
- d. Use of the Control Distribute VCC by the LES is optional; the LES is allowed not to open it and use only the Control Direct VCC.
- e. Multicast and broadcast frames are treated the same; the BUS forwards them to all nodes in the ELAN regardless of group membership.

Problem 4: Comparison of Multicast Routing Algorithms (8 points)

Compare the multicast routing protocols listed below on the issue of how the routes are computed. Make sure to indicate which protocols use algorithms that are optimal in some sense. For example, if you were comparing the unicast routing protocols RIP and OSPF, you would indicate that one of them uses the Bellman-Ford algorithm, and the other the Dijkstra algorithm. Both algorithms are optimal in the sense that they compute the shortest path between any two nodes.

Also indicate which protocols cause extra packets to be transmitted (i.e., packets that are sent into links or networks that do not lead to group members).

The protocols to be compared are:

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast OSPF (MOSPF)
- Core-Based Trees (CBT)
- Protocol-Independent Multicast, Dense Mode (PIM-DM)
- Protocol-Independent Multicast, Sparse Mode (PIM-SM)

***DVMRP** is derived from RIP, and uses Bellman-Ford to compute shortest paths. It uses a different tree for each source and group. The multicast tree built by DVMRP is composed of the shortest paths from the source to each of the destinations. Since DVMRP uses Reverse Path Forwarding with pruning, at times packets get transmitted into regions of the network where there are no group members. **MOSPF**, like OSPF, uses Dijkstra's shortest-path algorithm, and builds a tree composed of the shortest path from the source to each destination. If the AS is a single area, there are no extra packet transmissions, because each router has a complete picture of the network, and pruning happens in memory. If the AS is divided into areas, then extra packet transmissions may happen since Inter-Area routers will receive all multicasts. **CBT** builds a "core" tree per group, based on a "core" router. This tree is the shortest-path tree from the core router, but it is not optimal in any sense for the sources and destinations in the group. In CBT there are no extra transmissions into networks that do not have group members. **PIM-DM** is similar to DVMRP, but it is not tied to the unicast routing protocol; if that protocol computes shortest-paths, then it (eventually) builds the shortest-path tree by pruning. Like DVMRP, it sends multicast packets into regions of the network where they are not needed; however, since it has even less information than that protocol (since it does not touch the routing tables), it will flood more packets. **PIM-SM** is similar to CBT; instead of a "core", it defines a "rendezvous point", and uses the shortest paths to that point. Therefore, the routes are not optimal in any sense. The main difference between PIM-SM and CBT is that traffic can be re-routed around the rendezvous point if the volume justifies it. No packets are sent into networks that do not contain members.*

Problem 5: AAL-1 Clock Recovery (9 points)


Question 5a (2 points): Explain what is AAL-1 Source Clock Recovery and why it is needed.

One of the functions provided by AAL-1 is transfer of timing information between source and destination. In other words, it receives a bit pipe with its clock, and must transfer both the bits and the clock itself. Source Clock Recovery is the method by which the clock signal is regenerated at the destination. This needs to happen because the rate of bits into the network must match exactly the rate of bits out of the network; otherwise, the receiver will either underflow (run out of bits) or overflow.

Question 5b (3 points): Briefly explain the Synchronous Residual Time Stamp (SRTS) Method for AAL-1 Source Clock Recovery.

The SRTS Method is based on the existence of a synchronized network-wide clock both at the transmitter and at the receiver. The basic idea is to use this clock at the transmitter to measure the service clock, and transmit the measurement to the receiver using a bit in the AAL-1 header. This way, the receiver can synchronize its clock. The exact method is as follows: count the number of network clock cycles that happen during a certain number of cycles of the service clock. Subtract the nominal value from the actual count, and send this difference. The receiver will adjust its (locally generated) service clock to match.

Question 5c (4 points): Another way of doing Source Clock Recovery in AAL-1 is the Adaptive Clock Method, depicted below:



Let us assume that, on a certain implementation, the receiver has only two possible values for the bit clock: a slow value and a fast value. If the buffer falls to 25% full, the receiver switches to the slow clock; if it goes past 75% full, the receiver switches to the fast clock. Let T_{slow} be the period of time the receiver uses the slow clock, and T_{fast} be the period of time it uses the fast clock. Let F_0 be the nominal bit clock frequency in bits/second for this channel, and let δ be the range on either side of F_0 for the receiver clock (i.e., the frequency of the slow receiver clock is $F_0 - \delta$ and the frequency of the fast receiver clock is $F_0 + \delta$). Let F_s be the actual service clock for this channel at the transmitter; obviously, we need to have $F_0 - \delta \leq F_s \leq F_0 + \delta$. Let us define $\Delta = F_s - F_0$. If the total size of the buffer is B bits, derive the values of T_{slow} and T_{fast} as a function of Δ , δ and B . You can assume that, at startup, the buffer is filled to 25% and the receiver is using the slow clock².

During T_{slow} the buffer is filling, because F_s is higher than $F_0 - \delta$. It will fill from 25% to 75%; the amount of data is $(0.75 - 0.25)B = 0.5B$. Since the buffer is being emptied at F_0 , the rate at which the buffer is being filled is $F_s - (F_0 - \delta)$. The time this will take is:

$$[F_s - (F_0 - \delta)]T_{slow} = 0.5B$$

or

$$T_{slow} = 0.5B/(\delta + \Delta)$$

Conversely, during T_{fast} the buffer is emptying at a rate $(F_0 + \delta) - F_s$, from 75% to 25%. The time this will take is:

$$[(F_0 + \delta) - F_s]T_{fast} = 0.5B$$

or

$$T_{fast} = 0.5B/(\delta - \Delta)$$

² Note: this problem is much simpler than it looks.

Problem 6: RFC1483 and Classical IP over ATM (9 points)

Question 6a (2 points): Explain the difference between VC-based and LLC-based protocol multiplexing over ATM.

In LLC-based multiplexing, there is a single VC for all protocols, and individual frames have an LLC protocol type to indicate what protocol they are carrying. In VC-based multiplexing, there is a separate VC for each protocol, and the protocol type is inferred from the VC carrying the frame.

Question 6b (3 points): Explain the purpose of Inverse ARP (InARP). How is InARP different from Reverse ARP (RARP)?

*Inverse ARP is a protocol used to request the IP address of the node on the other side of a VC. In other words, the “hardware address” (the ATM address, or the VCI/VPI pair) is known, and the node is requesting the IP Address. This is the “inverse” of ARP, where the IP address is known and the node is requesting the corresponding hardware address. Reverse ARP is a mechanism for the host to find out **its own IP** address. In summary: in Inverse ARP, you have some other node’s hardware address, and want to find out its IP address; in Reverse ARP, you have your own hardware address and want to find out your own IP address.*

Question 6c (4 points): In ATMARP a new ARP message, the ARP NAK, was defined. This message does not exist in the traditional ARP. What is its purpose? Why wasn’t a similar message defined in the traditional ARP? Do you think that ATMARP would work without it? Justify your answer.

The ARP NAK is the message used by the ATMARP server as a response to a request to indicate that it does not know the address. In the traditional ARP, there is no centralized server; the fact that a host does not know the mapping does not mean that it does not exist, someone else in the network may know it. In traditional ARP, a failure is detected by a timeout – nobody responds. ATMARP would work just fine without the NAK, by using the same timeout mechanism; the NAK just makes it faster.

Problem 7: ATM LAN Emulation version 1.0 (10 points)

In this problem, you are asked to show the protocol interchange when two ATM-attached nodes communicate via LAN emulation. For this problem, node A knows the IP address of node B, and wants to send an IP packet to it. Assume the following:

- The IP and ATM ARP caches of all LAN Emulation Clients are empty.
- All LECs have joined the ELAN and are in the Operational State.
- There are no Data Direct VCCs open.
- The LECs implement the Flush Protocol.
- If the upper-layer software asks the LEC to send a packet with an unknown MAC address (i.e., with a destination MAC address that is not in the LE ARP cache), the LEC will forward that packet to the BUS. It will also start the LE_ARP process.
- Each transmission takes one unit of time, regardless of the source and destination.
- The time required to create a Data Direct VCC is negligible.
- To make it easier to draw, assume that a LEC will only send one frame per time unit. In other words, if at a certain stage, you get to a point where the LEC should send two frames, space them by one unit of time. Frames that go to the BUS should be sent first.
- To make it easier to draw, assume that all LECs need one unit of time to “think” about any frame they get. In other words, if they receive a frame that needs a response, the response goes out in the next time slot. This ***does not*** apply to the BUS or the LES; assume that they are infinitely fast and respond right away (but the response still needs one unit of time to get to the destination).

Use the diagram provided in the next page to show the protocol messages being exchanged. Each horizontal dashed line on that diagram corresponds to one time unit. You are required to show all frames exchanged at the IP, ARP, LE_ARP and FLUSH levels. There is no need to specify the frame in detail, just indicate the type (e.g., IP, LE_ARP Request, etc.). You are also required to show which virtual circuit is used for each message; use the following numeric key for this:

1. Control Direct VCC
2. Control Distribute VCC
3. Multicast Send VCC
4. Multicast Forward VCC
5. Data Direct VCC

Make sure to flag in your diagram the point(s) in time when ***new*** VCCs are created.

