

1 Multiple Choice

- 1.1 Which of the following is true of datacenters?
- (a) Many applications that run in datacenters have extreme performance requirements.
 - (b) The hardware cost of the network is more expensive than that of the compute.
 - (c) CPU cycles are especially important since a cloud provider may rent them to customers for additional profit.
 - (d) Datacenter operators have full administrative control over their datacenters.
- (a), (c), (d)
- 1.2 Why would one use kernel bypass?
- (a) To customize networking behavior
 - (b) To avoid making changes to kernel code
 - (c) For performance optimization
 - (d) To save money on hardware
- (a), (b), (c)
- 1.3 T/F: Stateless offloads are easier to implement since all information needed is contained in each individual packet.
- (a) True (b) False
- True
- 1.4 T/F: NIC offloads can be easily changed.
- (a) True (b) False
- False
- 1.5 T/F: Segmentation offload is a stateless offload.
- (a) True (b) False
- True
- 1.6 Which of the following are true of congestion control signals?
- (a) ECN contains more information than loss.

- (b) ECN contains more information than delay.
 - (c) Delay can be difficult to measure at high precision.
 - (d) Receive rate (bandwidth) contains more information than ECN.
- (a), (c)

1.7 Which of the following are true of QoS?

- (a) It only matters when a link is at 100% utilization.
- (b) Only one flow can be in each class.
- (c) It is used to prioritize some types of traffic.
- (d) The queue can enforce the policy with weighted fair queueing.

(a), (c), (d)

2 Software-Defined Networking

You are a chief network engineer at CS168DB, a company that specializes in databases. Currently, your networking infrastructure still relies on hardware switches that converge on routing with various L2/L3 protocols. You have tens of thousands of servers that need to communicate with each other. As the chief network engineer, you are thinking about making the switch to software-defined networking.

2.1 What are your current control plane abstractions?

None!

2.2 What are the required control plane's abstractions if you make the switch to SDN?

There are three. The first is an abstraction of the forwarding model of switches, OpenFlow is an example. The second is an abstraction of network state, in the form of a global network view via a Network Operating System. The third abstraction is an abstraction of the specification for network behavior, which is used to specify the **goals** of an operator on an abstract view of the network.

2.3 What happens when a link goes down in your new SDN setup?

The software switches connected to the link will communicate the topology change to the controller. The controller will then recompute routing state and disseminate that information to the routers.

2.4 Is OpenFlow equivalent to SDN?

No.

3 Network Virtualization Offload

Suppose you are a cloud provider implementing network virtualization. You want to allow a single server to host multiple VMs with potentially the same private IP addresses. In order for the servers hosting VMs to know which VM to send a packet to if there are, the packets must have some sort of identifier other than

the destination IP. Accordingly, you want to add a virtual network header to each packet to distinguish it from other customer's virtual networks.


3.1 Why would you want to allow this addressing overlap?

For better scalability (not limited by the number of private IPs in how many VMs you can host) and so that tenants do not have to worry about others' address spaces.

3.2 Where could you implement this functionality?

This can be done in the hypervisor/software switch or the NIC.

3.3 Assume the setup below. What rules should your SDN controller program into S1 in order for the below network to be routable? Fill in the table with the necessary match-action rules.



Assume the packets in the underlay network look like the one below. The virtual network header has one field, the ID, which has the name of the customer in it.



Header Field	Value	Action
Virtual Network ID	A	Forward to VM 1
Virtual Network ID	B	Forward to VM 2

3.4 What other functionality might you put in this MA table?

- (a) Bandwidth limiting
 - (b) HTTP caching
 - (c) Access control
 - (d) Accounting
- (a), (c), (d)

4 RDMA

Assume your application is about to use RDMA to transfer some data. In this question, we will walk through the steps to complete the process. Assume that we will use the RDMA Send operation discussed in lecture.

4.1 First, what benefits come from using RDMA?

- (a) Reduce complexity
- (b) Better performance
- (c) Save CPU cycles
- (d) Faster design iteration

(b), (c)

4.2 Which sides must register buffers before the transfer?

Both

4.3 How does the sender application start the transfer?

It places a pointer to the memory to send in the send queue.

4.4 Does the receiver have to do anything for the transfer to work? If so, what?

It must place a pointer to the buffer in the receive queue.

4.5 Which component of each host creates a completion queue entry?

The NIC

4.6 How does the application know that the transfer is complete?

It checks the completion queue.