

1 Datacenter True/False

- (1) Achieving low latency for small jobs (mice) is critical for datacenter networks.

Solution: True: It is necessary to prevent large slow flows from starving small latency sensitive tasks.

- (2) Clos topology enables the use of commodity hardware in datacenter networks.

Solution: True


- (3) ECMP leads to reordering of packets in the same flow.

Solution: False. Packets from the same flow take the same path.

- (4) Over-subscription is the ratio of worst-case achievable aggregate bandwidth between end-hosts to total bisection bandwidth.

Solution: True

2 Clos-based Topology



Given the datacenter network above:

- (1) How many paths from an arbitrary server upward to an arbitrary root switch? Consider only minimal-length paths.

Solution: 1.

- (2) How many paths can be set up between M1 to M3 in this topology? M1 to M5? Consider only minimal-length paths.

Solution: 2 ; 4.

- (3) How many of the M1 to M5 paths go through switch S2?

Solution: 1.

3 Bisection Bandwidth and ECMP

Consider the same network topology in Q2. Assume all links have capacity 10 Gbps and ECMP is used for multipath routing.

- (1) What is the bisection bandwidth of this topology?

Solution: 80 Gbps.


- (2) Assume ECMP makes the best possible hashing/load balancing. If switch S1 and S3 fail, what is the available aggregate bandwidth from each four-server Pod to any other such Pod (a Pod contains four servers closest to each other)?

Solution: 20 Gbps.

- (3) Assume M1, M2 are sending flows to M5; and M3, M4 are sending flows to M6. TCP makes optimal use of the available bandwidth and ECMP makes the worst possible hashing. What is the approximate average data rate server M1 can send to server M5?

Solution: 2.5 Gbps.

4 Encapsulation



- (1) If routing a packet from VM 1 to VM 6, what is the expected path?

Solution: S1-R1-S4

- (2) What encapsulation and decapsulation should the servers/router on the path from VM 1 to VM 6 perform (if any)? Fill in the following table with the destination address (both overlay and underlay) of the packets at each step.

Server/Router	Adds/Deletes Header	Overlay Address	Underlay Address

Solution:

Server	Adds/Deletes Header	Overlay Address	Underlay Address
Server 1	Adds	192.0.2.3	4.4.4.4/32
Router 1	Neither	192.0.2.3	4.4.4.4/32
Server 4	Deletes	192.0.2.3	4.4.4.4/32

- (3) How many destinations does Router 1 have to keep track of without encapsulation? How about with encapsulation? Consider how this isolation impacts VM scalability in a data center!

Solution: Without = all of them = 10, With = just the underlay / physical addresses = 4