

1 True/False


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

- 1.2 Clos topology enables the use of commodity hardware in datacenter networks.

- 1.3 ECMP leads to reordering of packets in the same flow.

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


2 STP



The diagram above shows a L2 network where all links have equal cost. We want to apply the Spanning Tree Protocol. Choose the switch with the lowest alphabetical ID as the root, and when there are multiple shortest paths to the root, choose the path that uses the neighbor switch with the lowest alphabetical ID.

- 2.1 Let XY denote the link that connects switch("X") and switch("Y"). List all links that are a part of the spanning tree after the STP protocol converges.
- 2.2 Assume that at some point during STP, switch("F") believes that switch("A") is the root and switch("D") believes that switch("B") is the root. Is it possible that switch("C") also believes that switch("B") is the root?
- 2.3 After the STP protocol has converged, switch("A") and all links directly connected to switch("A") go down. List all links that are a part of the new spanning tree after the STP protocol converges again.
- 2.4 switch("A") and all links connected to switch("A") finally come back up, and the STP protocol begins to run again. However, switch("F") forgets to send any update messages to its neighbors. After the protocol converges, does the resulting set of links form a valid spanning tree?

3 Clos-based Topology & ECMP




Given the datacenter network above:

- 3.1 How many paths from an arbitrary server upward to an arbitrary root switch? Consider only minimal-length paths.
- 3.2 How many paths can be set up between M1 to M3 in this topology? M1 to M5? Consider only minimal-length paths.
- 3.3 How many of the M1 to M5 paths go through switch S2?

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

- 3.4 What is the bisection bandwidth of this topology?
- 3.5 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)?
- 3.6 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?

4 Encapsulation



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

4.2 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
Server 1			
Router 1			
Server 4			

4.3 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!