


Note:

- Overwritten answers will not be accepted for rechecks
- Write all parts of a question together.

Q1. Consider a traditional network topology of a datacenter given below. Now that datacenter is moving to FatTree (6,3) topology, explain how it affects (i) path diversity (ii) bandwidth oversubscription. Assume that all links are of 10 Gbps bandwidth. Show calculations comparing the two topologies.

[10M]



- Q2.** State whether the following statements are true/false with justification. [2*5=10M]
- (a) 2PC and Paxos are same except for Paxos uses majority criteria for decision making.
 - (b) Zookeeper gets blocked when there are network failures.
 - (c) In hardware-assisted virtualization, when critical instruction is encountered, CPU performs vmexit operation giving control to VMM.
 - (d) In leaderless replication, if read and write quorum is followed by all clients, then all replicas will be consistent.
 - (e) Paxos is more efficient compared to ZAB protocol.
- Q3.** A university has decided to move all its course labs in computer science to a Cloud setup. Computer science department has 50 courses with lab component. A student has 3 courses with lab components. Each course has 200 students. The lab sessions are conducted once every week. At one time, a minimum of 10 courses and a maximum of 30 courses can have lab. Students access labs using SSH client from their own systems. They are required to login at the prescribed time and carry out the lab tasks. You are tasked to design a Cloud system to takes care of
- a) a proper scheduling of resources to minimize energy consumption
 - b) dynamic migration of lab sessions

- c) provide real-time backup services for the student work during lab sessions. Student can restore to any earlier point.
- d) restore student session in case of network or system failures
- e) provide workspace for group-based lab sessions where multiple students can login into a single system to work collaboratively
- f) system should be 99.99% available
- g) system should be scalable

In tabular format, identify challenges, solutions available, and your choice of the solution with justification.

[12M]

Q4. The following data is from black box monitoring of a datacenter. Assume CPU, Swap rate and Network utilizations are equal. VM1, VM2 and VM3 are on PM1.

Time	VM1	VM2	VM3	PM1	PM2	PM3
0	0.17	0.25	0.05	0.93	0.34	0.67
5	0.19	0.97	0.16	0.59	0.51	0.07
10	0.07	0.89	0.64	0.03	0.96	0.93
15	0.68	0.78	0.02	0.7	0.92	0.91
20	0.08	0.48	0.89	0.76	0.33	0.53
25	0.16	0.44	0.35	0.14	0.6	0.3
30	0.65	0.2	0.68	0.21	0.31	0.88
35	0.15	0.87	0.93	0.87	0.76	0.11
40	0.68	0.56	0.25	0.76	0.81	0.29
45	0.45	0.21	0.3	0.51	0	0.69
50	0.17	0.73	0.66	0.01	0.41	0.97
55	0.14	0.85	0.69	0.05	0.92	0.36
60	0.73	0.51	0.54	0.51	0.51	0.88
65	0.47	0.7	0.55	0.8	0.28	0.67
70	0.87	0.44	0.25	0.29	0.55	0.59
75	0.95	0.84	0.18	0.9	0.69	0.62
80	0.24	0.1	0.31	0.39	0.04	0.61
85	0.36	0.45	0.88	0.77	0.41	0.28
90	0.35	0.52	0.27	0.61	0.18	0.15
95	0.29	0.35	0.33	0.01	0.47	0.24
100	0.97	0.83	0.31	0.26	0.8	0.7

Compute the following over window size W (W=last 20 readings): [8M]

- (a) For the three VMs, compute peak CPU, and network resource requirements using distribution profile
- (b) Identify if there are any hotspots.
- (c) Determine which VMs to migrate and to where.

--&--