

We'll post an outline of each recitation after the fact so you know what was covered in case you had to miss class that day. The recitation outlines are not a full replacement for recitation!

This week's recitation was largely a problem-solving session (good for exam prep). These are the problems we went over; some sections may not have gotten to every problem.

Performance — Latency, Throughput, Utilization

Question 1: Consider Katrina's Amazing Web Suite, with multiple clients sending requests to a single server. The clients are spread across the Internet, but their requests ultimately reach a single switch that is directly connected to the server.

Katrina measures a variety of quantities to assess the performance of the system. *In theory*, which of the following quantities (if any) could increase without bound? For the quantities that are bounded (if any), what bounds them?

- The average time it takes a client to receive a reply after sending a request, measured in milliseconds
- The average throughput of requests at the server, measured in requests per second
- The utilization of the network link into the server

Performance — HDDs vs. SSDs

Question 2: For each of the following questions, select whether the statement applies to hard disk drives (HDDs), solid state drives (SSDs), both, or neither. Explain your answers.

1. Batching writes is an effective technique for improving performance
2. Cost effective for storing large amounts of data across hundreds or thousands of machines
3. Provides volatile storage

Question 3: Which of the following is the best option to improve performance on a workload that consists of many random reads from a few large files?

- Lay the files out contiguously on the disk
- Add a cache that stores the value of previous reads. Assume that the cache is large enough to hold 10% of the total data, and that it evicts data using a least-recently-used policy
- Replace the HDD with a solid state drive

Encapsulation

As part of Katrina's Amazing Web Suite, Katrina has designed an application-layer protocol called KAWS. As packets generated by KAWS travel across the Internet, they're encapsulated by lower-level protocols. You observe one of these packets as it traverses an Ethernet link, and are able to see the original packet data as well as four headers, one for each layer. Packets, then, look like this, where each element (A-E) is either a packet header or the original packet data:

[A | B | C | D | E]

Question 4: Which element corresponds to the Ethernet header? Which is the KAWS header?

In hopes of improving performance for this amazing web suite, Katrina decides to utilize a different addressing scheme. Instead of IP addresses, Katrina develops KP addresses, which have an entirely different format. Instead of using IP as the network-layer protocol, the machines in Katrina's amazing web suite use KP.

Question 5: What else would Katrina need to do in order to successfully send these packets across the Internet?

- A. Make sure the destination endpoint was also using KP, but make no changes to the switches on the path.
- B. Make sure that the destination endpoint and some of the switches along the path were using KP
- C. Make sure that the destination endpoint and all of the switches along the path were using KP
- D. Nothing; KP will work fine without any changes.

Ethernet

Question 6: Pick True or False for each of the following statements

- Ethernet receives discard duplicate packets
- Ethernet transceivers implement distributed packet switching so as to avoid a single central point of failure
- Ethernet's efficiency does not depend on how large the packets are.