

Review Questions

Instructions: You are not required to turn in answers to these review questions, but you should do them on your own after you read this section's papers to test your knowledge. Use Piazza or the TAs if you have questions or get stuck on any of these.

Giant Scale services

- 1) Discuss the pros and cons of load management at different levels in giant scale services.
- 2) Discuss the tradeoffs between replication and partitioning in architecting the data repositories of giant scale services.
- 3) Define the availability metrics: "uptime", "yield", and "harvest".
- 4) Explain the DQ principle and how it helps in handling graceful degradation of giant scale services.
- 5) Discuss the pros and cons of the three approaches to online evolution of giant-scale services: "fast reboot", "rolling upgrade", and "big flip".

Map-Reduce

- 1) Write a Map-Reduce solution for ranking the pages of web pages. Be explicit in identifying the