

Homework C

Due by Midnight (IST) of Sat 25 Apr, 2015. Total of 100 points.

Submit the homework solution as a PDF file by email with the subject line (and file name)

SE252_JAN2015_HW-C_StudentName.

Be succinct. The points allocated to each question gives you a sense of the depth of explanation required.

State assumptions you are making. Always provide citations for external sources of information and avoid verbatim use of external content!

(I) ILO5: Performance Benchmarking and Consistency

- 1) Big Data Nuggets Ltd is an enterprise that offers analytics as a service using Hadoop in the backend. They are currently using local commodity clusters within their internal data centre and are considering proposals from several IaaS Cloud vendors to host their service. You are asked to evaluate the performance of these IaaS Clouds. Discuss the process you will follow in putting together a system performance benchmark to help Big Data Nuggets make a selection. Follow the guidelines offered by Folkerts, et al.¹ [25 points]
- 2) Say a large university were designing several online services to be made available to students and faculty as given below. In each case, discuss which of the CAP properties that the service should support and why? State your assumptions [20 points]
 - a) Email service
 - b) Student Vacation Approval service
 - c) Salary & Reimbursement payment service
 - d) Classroom Discussion/Mailing List service
 - e) Student Admissions Form Submission service
 - f) Course Lecture Video Hosting Service
 - g) Emergency SMS Notification Service
 - h) Collaborative Document Editing Service
- 3) Give one example of a data storage service, software or platform for each combination of CAP (i.e. CA, AP, CP). Discuss briefly how they support the two properties and not the third. [25 points]
- 4) Compare the BASE properties of Amazon Dynamo Key-Value Store² with the ACID properties of a relational database like MySQL. Discuss in detail the techniques used by Dynamo to ensure a weak form of consistency? [15+ 15 points]

¹ Benchmarking in the Cloud: What it Should, Can, and Cannot Be, Enno Folkerts, Alexander Alexandrov, Kai Sachs, Alexandru Iosup, Volker Markl, and Cafer Tosun, *TPCTC*, 2012

² Dynamo: Amazon's Highly Available Key-value Store, Giuseppe DeCandia, et al, *SOSP*, 2007