

Seminar “Cloud Computing”

Organization

Klaus Berberich¹, Katja Hose¹, Jörg Schad²

¹Max-Planck-Institut für Informatik

²Information Systems Group

October 19, 2010

Meetings and Agenda

- Regular meetings
Tuesdays 14-16 s.t. in Room 001, Building E1.7 (MMCI)
- Kickoff meeting and registration
October 19 (today)
- Agenda for today
 - Requirements for the certificate
 - Registration and paper assignment

Requirements for the Certificate

- Attend all talks - not just your own
- Read your paper and other related literature
- Contact your tutor at least two weeks before your talk and present an intended outline of your talk
- Prepare a 45 minute talk about your topic that introduces the matter to your fellow students, use your own slides
- Present and discuss your slides with your tutor by the Friday before your talk (by 16:00)
- Presentation and slides must be given in English

Requirements for the Certificate

- Each presentation is followed by a discussion, in which all students participate
- For each presentation, a second student will be selected as an opponent
 - Asks tough questions to challenge the paper presented in the talk
 - Receives a preliminary version of the slides on the Friday before the talk
- Two weeks after the talk, the presenter submits a written summary of the assigned topic
 - Pointing out strengths and weaknesses of the approach presented in the paper
 - Template on our website: <http://bit.ly/ccseminar>
 - The summary has been reviewed by the opponent before it is handed in!

Requirements for the Certificate

- Final meeting with your tutor to give you feedback on your talk and your summary
- Final grade influenced by the following factors:
 - Oral presentation
 - Knowledge about your topic (discussion after the presentation)
 - The questions you asked as an opponent
 - General participation in the seminar
 - Written summary

Papers and Topics

Infrastructure I

09.11.2010 The Eucalyptus Open-Source Cloud-Computing System K. Berberich

Infrastructure II

16.11.2010 Energy Management for MapReduce Clusters J. Schad

Infrastructure III

23.11.2010 Multi-tenant databases for software as a service: schema-mapping techniques K. Berberich

Storage & Indexing I

30.11.2010 Towards Elastic Transactional Cloud Storage with Range Query Support K. Hose

Papers and Topics (cont'd)

Storage & Indexing II

07.12.2010 Efficient B-tree Based Indexing for Cloud Data Processing K. Hose

Replication I

14.12.2010 Volley: Automated Data Placement for Geo-Distributed Cloud Services K. Berberich

Replication II

04.01.2011 A self-organized, fault-tolerant and scalable replication scheme for cloud storage K. Hose

Transactions I

11.01.2011 Consistency rationing in the cloud: pay only when it matters J. Schad

Papers and Topics (cont'd)

Transactions II

18.01.2011 G-Store: a scalable data store for transactional multi key access in the cloud K. Berberich

Applications I

25.01.2011 Behavioral Simulations in MapReduce K. Hose

Applications II

01.02.2011 Efficient parallel set-similarity joins using MapReduce J. Schad

Applications III

08.02.2011 Max-cover in MapRedcue K. Berberich

Recommended Reading

- *Above the Clouds: A Berkeley View of Cloud Computing*
M. Armbrust et al.
Technical Report, EECS Department, UC Berkeley, 2009.
- *MapReduce: Simplified data processing on large clusters*
J. Dean and S. Ghemawat
Communications of the ACM 51(1), 2008.

Registration & Paper Assignment

- Send e-mail to `kberberich@mpi-inf.mpg.de` by *tomorrow noon* with
 - your name
 - your student number (“Matrikelnummer”)
 - your favorite three topics in ranked order
(e.g., Applications II, Replication I, Infrastructure II)
- Paper assignment will be posted online tomorrow evening