

Summary of the 3rd ACM SIGOPS Workshop on Large-Scale Distributed Systems and Middleware (LADIS 2009)

Hakim Weatherspoon*
Cornell University
hweather@cs.cornell.edu

Doug Terry
Microsoft Research
doug.terry@microsoft.com

Gregory Chockler
IBM Research
chockler@il.ibm.com

About the Workshop

The 3rd ACM/SIGOPS workshop on Large-Scale Distributed Systems and Middleware (LADIS'09) was held on October 10-11, 2009, in conjunction with the 22nd Symposium on Operating System Principles (SOSP) in Big Sky, Montana. The workshop was dedicated to hot topics in cloud computing. It lasted for one and a half days and featured keynote talks by industry leaders, peer-reviewed position papers from both academia and industry, and late-night break-out sessions fostering further discussion of important topics. LADIS'09 provided a forum for presenting and discussing cutting-edge ideas, and touched upon topics as diverse as distributed computing infrastructure, programming models, data storage, and security and privacy. This report summarizes key themes and conclusions from the workshop.

Keynote Talks

The keynote talks proved to be the real highlight of the workshop. The presenters were Raghu Ramakrishnan (Chief Scientist for Audience and Cloud Computing, and a Fellow at Yahoo!, where he heads the Community Systems group), Marvin Theimer (Senior Principal Engineer at Amazon), Jeff Dean (a Google Fellow and creator of much of Google's infrastructure), and David Nichols (Microsoft Windows Live). Each speaker discussed their experience with the architectures of the world's most demanding cloud platforms.

Raghu Ramakrishnan's presentation focused on the architecture of Sherpa, the data storage service used by the Yahoo!'s cloud platform. His talk emphasized the importance of a data consistency model that would be useful for supporting a wide range of querying functionality (such as random access and range queries), while being scalable and efficient to address the performance and availability needs imposed by large scales and wide distribution. The second part of Raghu's talk provided an overview of various existing cloud storage platforms (such as BigTable, Dynamo, and Cassandra), and compared their service models and design choices.

The talk by Amazon's Marvin Theimer highlighted the importance of design for failures, and dynamic reconfiguration. He gave the analogy that evolving a cloud infrastructure to a large offering is similar to evolving a Cessna prop-plane into a 747 jumbo jet in-flight. Start with a Cessna prop-plane for cost and timeliness reasons. 4-9's of availability means that you get to land for 52 minutes every year (including scheduled maintenance, refueling, and crash land-

ings). Success implies growth and evolution and rebuilding the plane mid-flight: Passenger capacity goes from 4-person cabin to 747 jumbo wide-body cabin. Support for "scale out" means jet engines are added and propellers removed, testing and safety have to happen while flying!

Jeff Dean gave an up-to-date overview of recent developments in Google's cloud infrastructure. He highlighted the basic principles and lessons learned over the years of designing the Google's cloud, such as the importance of understanding basic factors impacting the system performance at large scale. Towards the end of his talk, Dean offered a sneak peak into Google's new infrastructure for large-scale resource management, called Spanner, currently in alpha testing. The talk generated considerable interest in the community, and is being actively discussed on various Internet forums.

David Nichols of Microsoft discussed Windows Live use of SQL. He argued that SQL can be tamed. SQL has some real issues but they are mostly manageable with some infrastructure tweaks. He stated that it is hard to do better than SQL: it keeps improving and he continually finds that SQL already includes many designs he needs. However, SQL is not always the best solution. Nichols gave an SQL wish list: partitioned data support, easy migration/ placement control, reporting, jobs; supporting aggregated data pattern, improved manageability. More difficult requests included dampening database schema evolution, adding soft delete/versioning support of some kind, and A-D transactions (Atomic and Durable).

Papers and Break-out topics

The 15 presentations of the peer-reviewed papers tackled the challenges being faced by the current and future cloud infrastructure. Topics covered included storage, communication, monitoring and repair of massive cloud systems, as well as cloud programming models and various applications and services. By the end of the workshop, attendees were looking to answer questions impacting tomorrow's datacenter designs and Internet services:

Privacy: Cloud storage systems are increasingly used to store valuable business data and intensely private data, and they even mix data from different individuals on the same servers. When all of a persons (or business) data is stored in the cloud, what steps can be taken to ensure the privacy of that data and to reassure users that their data will not be inadvertently released to others?

Data consistency and replication: Most current cloud- resident storage systems replicate data but have chosen to relax

*The authors are in the reverse alphabetical order.

consistency in favor of increased performance (and availability). What consistency guarantees lie somewhere between strong serializability and weak eventual consistency might appeal to cloud applications? How can they be provided for cloud-based services that serve a globally distributed user population?

Trust: Cloud computing is viewed as risky for various different reasons, thanks to fears of data leakage, outages, lack of long-term business viability of service providers, inability to get data out of the cloud once placed there, and so on. What explicit steps can we take to persuade people that these risks are, in fact, minimal (or make them so), and help them assess them rationally by comparison with their alternatives, which are invariably of the form "keep the cash under the mattress"?

Durability: How can we design cloud-based storage systems to ensure that a users data survives for a hundred years, even as companies come and go?

Service level agreements: The service level guarantees from cloud services are imprecisely specified, often only in the minds of the users. There are certainly many instances of wide spread use of best effort services: IP, email, DNS, etc. Is that good enough? As cloud-based services mature, how should they provide more specific guarantees and what sorts of guarantees will be desired by their clients?

Virtualization: Cloud computing currently relies heavily on virtualized CPU and storage resources to meet elastic demands. Whats the role of virtualization in cloud-based services, is it required, and are current virtualization technologies sufficient?

Programming models: How should computational resources in the cloud be presented to application developers/services, as virtualized hardware or application-specific platforms or something in between?

Mobility: Increasingly, the clients of cloud-based services are not desktop PCs but rather mobile devices, such as cell phones and portable media players. How should mobile devices at the edge of the network interact with cloud-based services to effectively manage data and computation on behalf of users? How does a users location factor into the design of cloud-based services?

Power management: What can be done to reduce the energy demands of cloud-based services?

Principles: Over the past three decades, many fundamental principles of distributed computing have been developed, including techniques for 2-phase commit, consensus, Paxos, replicated state machines, virtual synchrony, Byzantine fault tolerance, etc. Which of these are useful building blocks for cloud services, which are not, and what new principles need to be devised for future cloud computing applications?

Moving forward

The success of the LADIS series of workshop was one of the key factors inspiring the SIGOPS/SIGMOD community to establish the first joint Symposium on Cloud Computing (SOCC) that will be held next year in conjunction with SIGMOD 10. In 2010, the next LADIS workshop will be co-located with the Principles of Distributed Computing Conference (PODC) in Zurich, Switzerland, and maintain its traditional focus on core principles of large-scale distributed systems and middleware. LADIS 2010 will be chaired by Marcos Aguilera (Microsoft Research) and Marc Shapiro (INRIA and LIP6).

LADIS 2009 was made possible by the efforts of the following people:

Workshop Organizers

General Chair, Gregory Chockler, *IBM Research*
Program Co-Chair, Doug Terry, *Microsoft Research*
Program Co-Chair, Hakim Weatherspoon, *Cornell*

Program Committee

Patrick Eaton, *EMC Corp, Cloud Infrastructure Group*
Michael Freedman, *Princeton*
Steve Gribble, *University of Washington*
Rachid Guerraoui, *EPFL*
Benjamin Reed, *Yahoo! Research*
Rodrigo Rodrigues, *Max Planck Institute*
Mendel Rosenblum, *Stanford*
Mike Spreitzer, *IBM*
Franco Travostino, *EBay*
John Wilkes, *Google*