

Aims and Scope

More and more individuals, companies and organizations are relying on the cloud to store and manage their data, which translates into increasing pressure on the cloud infrastructure. Cloud data can be very diverse, including a wide variety of personal data collections, very large multimedia content repositories and very large datasets. Users and application developers can be in very high numbers, with little DBMS expertise. Data-intensive applications can be very diverse too, with requirements ranging from basic database capabilities to complex analytics over big data. In particular, the pay-as-you-go model makes the cloud attractive for supporting novel large-scale elastic applications.

NoSQL solutions for the cloud, for instance, have traded consistency and transactional guarantees for scalability. However, the grand challenge for a data-intensive cloud infrastructure is to provide ease of use, consistency, privacy, scalability and elasticity, simultaneously, over cloud data. Addressing this challenge requires novel solutions across the spectrum of data management techniques, including massive data storage, elastic parallel query processing, transactions over data replicated at geographically distributed sites, security and privacy, and efficient data loading and access. This special issue focuses on recent advances in research and development in data-intensive cloud infrastructures.

Topics of Interest

The VLDB Journal solicits contributions to a Special Issue on Data-Intensive Cloud Infrastructure. Relevant topics include, but are not limited to, the following:

- Novel data-intensive cloud architectures
- Data sharing in heterogeneous clouds
- Energy-efficient data-intensive clouds
- Massive data storage, distributed file systems
- Data models and query languages for cloud data
- Transactional models for cloud data
- Large-scale distributed and parallel data management, including partitioning, caching, replication, and transactional mechanisms to synchronize them
- Elastic query processing, including parallelization and optimization
- Big data analytics
- Multi-tenancy and workload isolation

- Data security, data privacy and access control
- Metadata management, data integration, data sharing, recommendation
- Data provenance, data pricing
- Data availability, fault-tolerance
- Data management on virtual machines and virtual storage
- Hardware-aware data management, including GPU, NUMA and flash memory
- Benchmarking and testing of data-intensive clouds
- Database administration, self-management and automatic tuning
- Resource and workload management for data-intensive clouds
- Real-world applications and system-oriented use cases

We solicit submissions of high-quality, original research and are particularly interested in system-oriented papers and practical use cases involving new technologies. We are also happy to consider well-written survey articles.

Guest Editors

Ashraf Aboulnaga, University of Waterloo, Canada (ashraf@uwaterloo.ca)
Beng Chin Ooi, National University of Singapore (ooibc@comp.nus.edu.sg)
Patrick Valduriez, INRIA and LIRMM, France (Patrick.Valduriez@inria.fr)

Important Dates

- Paper submissions: September 15th, 2013
- First round notification: January 15th, 2014
- Revised versions: April 1st, 2014
- Second round notifications: May 1st, 2014
- Final version: June 1st, 2014
- Publication in August or October, 2014

Special Issue on Data-Intensive Cloud Infrastructure

Written by VLDB Journal Team

Monday, 25 March 2013 00:00 - Last Updated Monday, 25 March 2013 14:30

Submission Website: <http://www.editorialmanager.com/vldb/default.asp> (under "Article Type" choose "Data-Intensive Cloud Infrastructure").