

VLDB Workshop on Design, Implementation, and Deployment of Database Replication

in conjunction with the 31st International Conference on Very Large Data Bases (VLDB 2005)
August 28th, Trondheim, Norway

<http://gorda.di.uminho.pt/wdiddr>

Important Dates

Papers due: June 1, 2005.
Notification: July 1, 2005.
Camera ready: July 15, 2005.
Workshop: August 28, 2005.

Program Chair

R. Oliveira
U. Minho, Portugal

Organizing Chair

J. Pereira
U. Minho, Portugal

Program Committee

E. Cecchet
EMIC/ObjectWeb, Finland

J. Goldstein
Microsoft Research, USA

B. Kemme
McGill U., Canada

T. Lane
RedHat Inc., USA

H. Madeira
U. Coimbra, Portugal

M. Patiño-Martínez
UPM, Spain

F. Pedone
USI, Switzerland

L. Rodrigues
U. Lisboa, Portugal

J. Stanton
George Washington U., USA

J-B. Stefani
INRIA, France

L. Thalmann
MySQL AB, Sweden

J. Wieck
Affiliis Inc., USA

Context

Database replication is widely used to improve both the performance and resilience of database management systems. Although most commercially available solutions and the large majority of deployments use asynchronous updates in a shared nothing architecture, there is an increasing demand for additional guarantees, configuration flexibility, and manageability, which are currently provided only, if at all, by costly high-end solutions.

An example is the desire to upgrade current fail-over clusters to an active-active configuration, thus leveraging the additional computational resources for additional performance. It should also be possible to configure eager update replication over a MAN or WAN for disaster recovery, without resorting to volume replication, thus improving manageability and performance. The combination of both local clusters and wide-area systems in grid-style large scale systems poses new challenges in performance and manageability. The drive for self-manageable systems also discourages the need for human intervention in recovery and conflict solving, thus favoring eager update approaches. Finally, typical loads of current middleware and applications might allow novel trade-offs between resilience and performance that are not viable with OLTP loads.

This demand has been addressed by a flock of proposals for eager update replication that do not use the traditional 2PC-based approach. These range from middleware based approaches, which often also provide load balancing of read-only queries, to integrated solutions as demonstrated in the Postgres-R prototype. Finally, there is a flourishing academic research trend in group communication based replication, which offloads many complexity and performance issues to group communication toolkits and aims at scalable eager update everywhere replication.

Goal

The goal of the workshop is to bring together researchers and practitioners from the database and fault-tolerant distributed systems communities to discuss the current state of the art, pending challenges and trends, and novel solutions in the design, implementation and deployment of database replication.

Topics of interest include, but are not limited to, the following:

- cluster, MAN and WAN replication protocols;
- shared nothing and shared disk eager update replication;
- replication middleware;
- novel applications and loads for replicated databases;
- interfacing the DBMS with replication protocols;
- replication transparency and client interfaces;
- performance and dependability benchmarks for replicated databases;
- fault models and required guarantees;
- self-manageable and autonomic replicated databases;
- management interfaces and tools for replicated databases;
- recovery protocols and tools for replicated databases;
- group communication based replication;
- group communication protocols for database replication.

Paper submission

Authors are invited to submit short position papers (no longer than 6 pages in "VLDB format") describing original, previously unpublished work. Work-in-progress and experience reports, that foster discussion and bridge the gap between target communities, are specially welcome.

All submitted papers will be reviewed by the program committee. At least one author of an accepted paper is expected to attend the workshop in order to present the work. The authors are requested to send their contributions to datarep@lsd.di.uminho.pt no later than June 1st, 2005, as a postscript or PDF file.

Additional Information

The workshop is organized in association with the VLDB conference (<http://www.vldb2005.org>) by the GORDA project (FP6-IST2-004758) (<http://gorda.di.uminho.pt>).

For more informations contact the organization at datarep@lsd.di.uminho.pt.

