

SUMMARY:

Lei Gao, Mike Dahlin, Amol Nayate, Jiandan Zheng, and Arun Iyengar.

Application specific data replication for edge services.

In *Proc. Int'l Conf. on World Wide Web (WWW'03)*, pages 449-460, 2003.

DATE: February 8, 2010

The work presented in this paper introduces a model of application-specific data replication for edge services. The standard setup for edge services that contains a centralized database used to service the requests of the remote edge servers is not able to provide the availability desired by the applications due to the network latency and partitioning, while providing the high consistency required. The model suggested by Gao et al. trades off some of the consistency to provide high availability by replicating the data of the application at the edge servers in an application-specific manner.

Gao et al. propose that the application developers define distributed objects encapsulating subsets of the application data with similar consistency requirements. For the prototype presented in the paper, a variation of the TPC-W benchmark, a distributed bookstore application, is implemented. The data for the application is divided into several distributed objects, for example, *catalog* object that is updated on the backend server and read at the edge servers, *order* object that is written at an edge server and read at the backend server only, *profile* object that can be (but is not likely to be) concurrently updated at multiple edge servers, and for which the data can be reconciled using field-specific rules. For each such distributed object the system defines rules that control data propagation and conflict resolution, if required. Each edge server in the system stores all the data that is required for the application, providing very high availability - since all the requests are satisfied using local data - while not sacrificing consistency beyond what can be tolerated by the particular application.

Results presented in the paper show that the proposed model provides much higher availability than the model using a single backend data store. At the same time data high level of consistency, defined in terms of the particular application and not some generic system, can be maintained. Additionally, distributed objects allow web-applications to maintain copies of data at the edge servers, thus providing a much better performance and throughput than that of systems with a single backend database.

SUMMARIZED BY: Evguenia Eflou