

CloudTran Unbundles Transaction Services for Cloud Applications

Background

While the academic world has portrayed the benefits of [unbundling transaction services](#) in a cloud computing architecture, the practical reality of doing so has only recently become available in a commercial product.

[CloudTran](#) is the premiere developer of transaction management software for cloud computing. We've bridged the gap between theoretical studies and real-world requirements for running mission-critical, OLTP applications in a cloud architecture. This paper discusses how an independent transaction management system has in fact been "unbundled" resulting in key benefits to mainstream application development.

Stepping Up

The exciting thing about cloud architectures — apart from instant deployment and pay-as-you-go — is the jump in scale: some applications now serve millions of online users in fractions of a second with information refined from huge amounts of data. This category of application has large data, heavy processing and high transactions-per-second. At a certain level, the cheapest and best solution is an in-memory architecture (see [RAMClouds](#)), which explains why every industry player is betting on some sort of “in-memory database/data grid” product.

The Problem

However, the problem with cloud architectures for mission-critical applications revolves around ACID transactions — specifically, the lack of them. Brilliant engineers have tried to use existing techniques such as distributed transactions (e.g., XA) to provide scalable applications with ACID support. This effort has failed because distributed transactions are too slow and unreliable.

The CloudTran Approach

The CloudTran solution is to unbundle transaction management from the data stores (i.e., databases, document stores). The first change resulting from this is a central transaction coordinator that sits between the in-memory data grid and the data stores. The coordinator can handle changes from any number of nodes in the grid and can send data to any number of stores.

The second change made by CloudTran is to distribute transactional control — particularly constraint and isolation handling — to achieve maximum performance. The client, the “ORM” (object relational mapper), the in-memory nodes, and the transaction coordinator all handle different aspects. This allows even an entry-level configuration to handle **thousands of update transactions per second**.

The path to the data stores is supremely important for durability of course, and, secondarily for links into data warehousing and other BI that feed off the database. With CloudTran, however, the performance requirements of the data stores changes and gives opportunities:

- data affinity and foreign-key constraints are handled in the grid, so tables can, for example, be sent to one physical database each, which may avoid complicated sharding
- latency is no longer important, which means the application front-end can be in the cloud and the durable data in the datacenter, reducing security concerns
- real-time analytics can be fed from the transaction coordinator, rather than via the database, to improve the timeliness of the information

The Way Forward

Successful companies in the future will use in-memory architectures to create faster websites delivering on more refined intelligence from deeper analysis of personal preferences and social trends. The upside of the move is increased competitiveness and the ability to serve a global customer base. The challenge is the uncertainty as new tools are adopted and the risks in strategy and execution. This paper has focused on the CloudTran approach to providing strong, scalable transactions linking to standard databases gives architects and developers a familiar reference point.

Following papers will highlight CloudTran solutions for simplifying distributed development and managing deployments.

About CloudTran

CloudTran is the premiere developer of transaction management software for cloud computing. We make it possible to run mission-critical OLTP applications in public, private, and hybrid clouds by adding ACID transactions to data grids and distributed persistent data stores.

Until now, it has been extremely difficult and costly for IT to develop scalable, write-oriented (OLTP) applications that take full advantage of cloud/grid computing. This is because there is significant application architecture and coding work required for which IT has neither the time nor skill to deliver. CloudTran is addressing this with its first products to market: a transaction manager for the cloud and a data modeling and architecture tool for Java applications.