

CloudTran

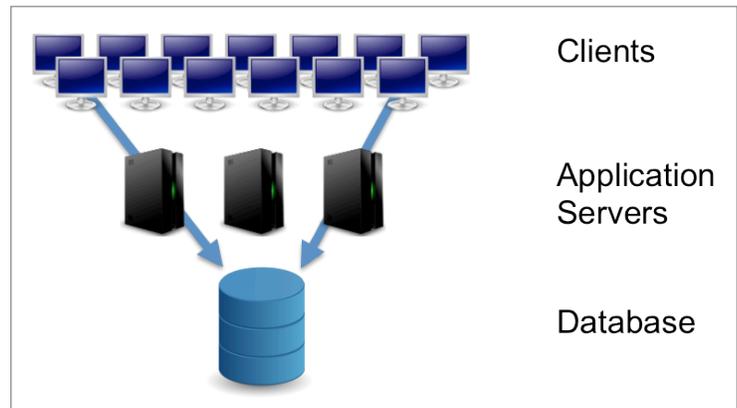
Data Layer Server for Web Applications

CloudTran, Inc.
www.CloudTran.com
4000 Pimlico Drive, Suite 114
Pleasanton, CA 94588
925-271-5512

Is Your Database the Bottleneck?

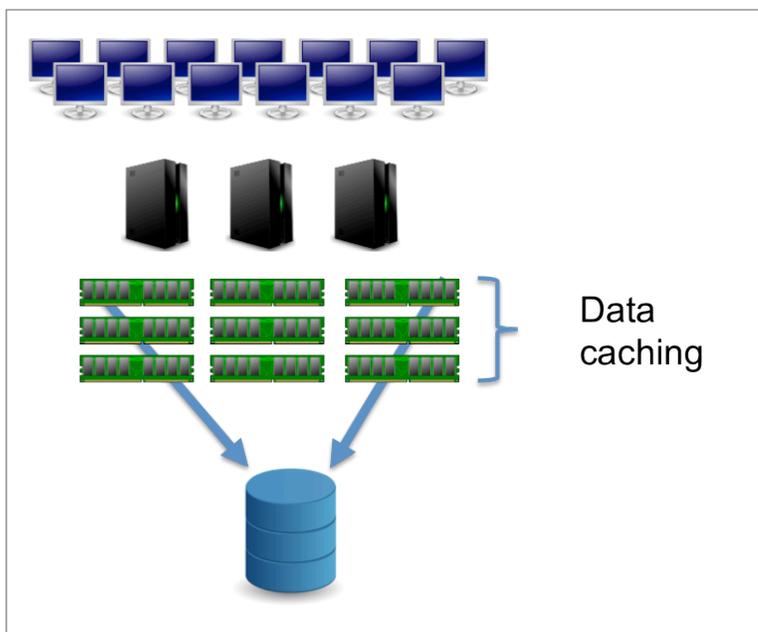
CloudTran helps you write scalable applications for distributed and cloud computing environments. Our software eliminates the performance bottleneck that typically occurs at the database level by making it easy to use several, loosely coupled data stores in your application.

Common application architectures support scale-out of clients and application servers, but when it comes to the data layer, the architecture still relies on a single database. This can result in severe performance problems, especially with write-intensive applications that perform frequent updates, inserts, and deletes to the database. A single database gets serialized with requests for data, and the resulting disk I/O becomes an overall limiting factor.



But Doesn't Caching Solve the Problem?

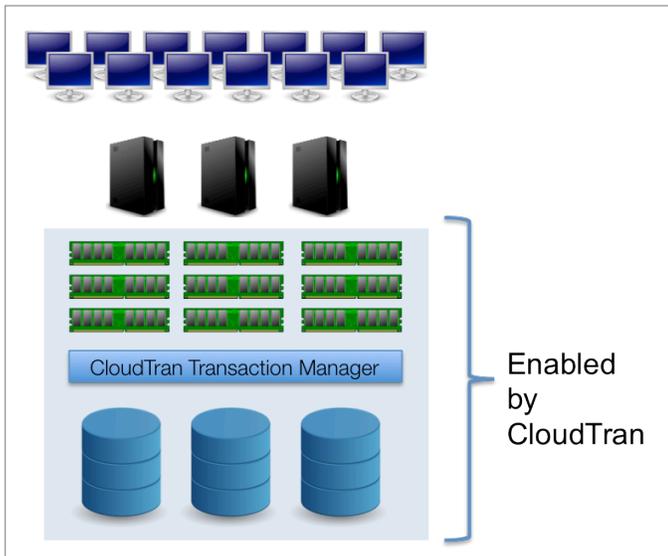
A better approach is to create scalability at the data level too, just like there is at the application level. The first step that many developers take is to use a caching system like Memcached or Ehcache for storing frequently used data in memory. In some cases, the entire database can be stored in memory. More elaborate architectures use an elastic caching platform like GigaSpaces or Coherence that provide extra services such as failover recovery and backup to disk.



Putting caching in place is a great start, but for applications that must persist data to disk, there is still a bottleneck when it comes time to handle inserts, updates, and deletes to the database. What applications need persistence? Just about every, business-critical use case from eCommerce to inventory processing, supply chain management, micropayments, and even online gaming.

CloudTran Breaks the Bottleneck

CloudTran removes the database bottleneck by allowing you to easily use multiple databases located on different machines as a persistent data store for in-memory data. Our data layer server coordinates the movement of data between databases and the in-



memory data grid and shields you from having to code for ACID transactionality or location of data.

It's important to note that CloudTran is NOT an implementation of distributed transactions. CloudTran uses a proprietary method for persisting data to disk that doesn't require the typical two-phase commit protocol. As a result, transaction rates of several thousand per second can be achieved with just a single CloudTran transaction controller. CloudTran gives you a complete, out-of-the-box architecture for creating scalable data applications. We handle all the scatter, gather, failover recovery, logging, and data mapping to make life as a developer easier.

FAQs

Can't I use Oracle RAC or MySQL Cluster to scale-out datastores?

Scale-out applications use multiple application nodes to achieve their performance goals. When these applications need to commit transactions, they will have to use distributed transactions, which are notoriously slow. As a result, performance issues will arise quickly when more than a few transactions per second are required. CloudTran is optimized for data grids and commits transactions at grid speed.

But data grid products these days have automatic data persistence – why not use that?

These systems do not provide the ability to scale-out at the database level and do not provide ACID transactionality between the data grid and the database. While many applications can still benefit from this feature, business-critical applications that are write-intensive will suffer in performance.

What databases and caches does CloudTran support?

The goal of CloudTran is to be agnostic of the data store and of the caching layer. Today, we support GigaSpaces XAP system as an elastic caching platform, and support for Oracle Coherence is coming soon. For data stores, you can use Oracle, SQL Server, MySQL or any JDBC database. We also support MuleSoft's ESB and Amazon's SQS, and support for NoSQL products like MongoDB is coming soon.

How do I write code for CloudTrain – is there an API?

CloudTran provides a distributed ORM, but you can also use JPA. The API is similar to an ORM because CloudTran takes care of distribution across the data grid and the conversion to database format for persistence. CloudTran includes an Eclipse plug-in you use to develop applications that work with the CloudTran transaction manager. The plug-in is used to specify your data model as well as deployment configurations. We use codegen to build the routines that perform data access. You add your business logic to these routines.