CloudTran

Data Management for Large-scale Applications

A White Paper

William Payne
Date: 19th January 2010

Overview

The lack of an effective mechanism for very large scale transaction processing over distributed, grid and cloud architectures limits many organisations' ability to build complex, enterprise-wide, scalable applications.

Primary impacts on organisations include a loss of business agility, lack of value chain integration, difficulty in handling large trading volumes, and limited customer responsiveness, leading to potential loss of market share.

Secondary impacts on organisations include higher IT hardware and software costs, poorer ROI on IT assets, costly migrations, software project time and cost over-runs, and major IT programme failures.

CloudTran is a middleware technology that enables highly responsive transaction processing over the largest distributed and cloud computing environments. It provides for significant improvements in the performance, availability, reliability and cost-efficiency of very large scale transaction processing.

The Drivers : Business Impact

The limitations of the current technologies for managing very large scale transaction processing over distributed environments can affect, and curtail, a wide range of business operations.

Reduced cost and risk in investment banking

In the financial services industry, investment banks typically operate across multiple time zones and regions, with trading operations in major financial centres across the globe.

Current technologies mean high cost and complexity of IT infrastructure coupled with increased vulnerability both to outages and risk management failures. An effective solution for high performance, high volume distributed transaction processing – such as offered by CloudTran – would not only reduce costs and simplify operational data architectures, but increase the quality and timeliness of data available to financial sector central oversight and management functions.

Improved web and branch response times in retail banking

In the retail banking arena, businesses can juggle hundreds of gigabytes of customer records in cache. Storage I/O bottlenecks often impose long response times, leading to slow performance on customer websites and branches during

busy times. Due to lack of a viable alternative, a typical response is to curtail the range and sophistication of services available to customers to reduce peak loads on the system.

More sophisticated and innovative services

Major improvements in high volume distributed transaction processing would allow retail banks to provide their customers not only with improved web and branch service response times, improving customer experience and customer retention rates. It would also enable them to manage and peaks of demand, and to introduce more sophisticated and innovative services. In particular, it would help banks reduce the transactional cost of operating general and basic retail accounts, while expanding the scope and range of their retail wealth management services aimed at middle segment customers.

Public Sector

An improved high volume distributed transaction processing architecture would also support the development of modern applications and services integration in the public sector and within health service organisations.

Integrated records

Service delivery integration and consolidation is a major objective across the public sector. Delivering integrated electronic services, based on shared electronic records systems and resources, has frequently been hampered by the prevalence of legacy systems and the highly distributed structure of much public service and health service delivery. In addition, the complexity of many public sector and health service functions means that applications can be highly complex, with very high processing and data volume loads.

An efficient architecture supporting high volume transaction processing over highly distributed environments would remove a critical bottleneck to the development of successful programmes for integrated public sector and health service IT applications, as well as reducing the cost of existing IT installations.

Faster and easier compliant production

Poor support for high volume distributed transaction processing can also adversely affect firms with multinational manufacturing operations, especially those with significant regulatory compliance, or batch record requirements, such as manufacturers of semiconductors, electronics, pharmaceuticals and medical devices.

In particular, this can affect such firms' production agility and operational flexibility, making it more difficult to load-balance or switch production of complex or regulated products between different plants across their multinational

operations.

An effective platform for high volume distributed transaction processing would allow firms operating in such industries to increase the speed and ease of fully compliant production setup from one plant to another across any timezone, and create a single centralised environment for fully agile manufacturing control, Quality, compliance control and reporting for all global production operations.

The Response : CloudTran

CloudTran is a unique middleware technology that delivers robust, durable, isolated and highly responsive transaction processing over very large scale distributed, grid and cloud enterprise architectures. It ensures major performance improvements in large scale transaction processing over conventional technologies, and incorporates a mechanism that is inherently robust as well as highly performant over the most distributed, heterogeneous or integrated application architectures.

Based on Java enterprise technologies, CloudTran brings major performance and reliability improvements to high volume transaction environments ranging from legacy host-based systems through to the latest grid and cloud architectures.

Applications that can benefit from CloudTran include web customer facing systems, e-commerce systems, SOA and web services, large scale trading environments, ERP, CRM, and logistics systems.

CloudTran Pedigree

CloudTran has been developed by NT/e, a leading J2EE development specialist. NT/e has a track record of successful large scale systems and transaction processing technology development going back over twenty years, with a major involvement in both European and North American markets. The company was one of the first integration specialists to adopt the J2EE platform, and is one of the largest European providers of technical consultancy and integration services for BEA Systems and Oracle enterprise application infrastructure and SOA solutions. The company has a particular expertise in building high performance, high volume enterprise solutions in the financial and banking industry sectors.

Background : The Shift to Distributed Computing

Large scale enterprises and public sector organisations are increasingly moving from hierarchical computing models to flatter, more distributed architectures. This is being driven by a combination of business and technical factors :

- As enterprises have sought to consolidate their business operations onto single unified platforms, data and process volumes have increased significantly. Distributed hardware and software architectures, rather than hierarchical architectures, provide a better solution for the large volumes that result from consolidating multiple operations;

Corporate systems handling increased volumes

- Organisational data is also becoming more complex. Organisations are seeking greater functionality from their applications and the volumes of data that must be handled. Coupled with certain technical issues (such as those associated with relational databases), these factors are greatly increasing the processing loads and data volumes that must be handled by corporate systems. Distributed hardware and software architectures offer the promise of a better solution for these volumes than hierarchical architectures;

Integration and interoperability

- The need for greater integration and interoperability between core enterprise applications and data is an additional driver. Whether for increased financial and operational control, better value chain integration, or the development of new cross function products and services, greater integration between existing business systems, applications and data is a key factor driving the migration towards more distributed architectures;

Customer applications inherently distributed

- Web facing customer applications are inherently distributed, as they typically involve a number of different business applications – frequently legacy systems – as well as web mediation and content management systems. Web facing customer applications are typically very sensitive to system response times, and many organisations often have to make trade-offs between service functionality and poor response times. Moreover, this sensitivity is increasing. In 2006, a survey found that 28% of the public would abandon a web transaction after a three second delay in the transaction. The same survey repeated in 2009 found that figure had grown to 40%.

Scalable, flexible, architecture

There are also significant technical reasons for organisations moving to more distributed architectures :

- Scaling up through distributed hardware and software architectures can avoid the need to migrate applications and data to larger, more expensive and

complex systems;
- Distributed architectures also allow greater resilience and fault tolerance, enabling load balancing across a grid;
- Distributing applications across a grid or cloud is a highly cost effective way to maximise performance in integrated, multiple application or heterogeneous data environments, such as trading, ERP, CRM, logistics or engineering applications.

The Requirement : Distributed Transaction Processing

Distributed, grid or cloud infrastructures offer organisations clear benefits in terms of improved processing and data volume handling, scalability and ROI. Yet the lack of an effective mechanism for distributed transaction processing is severely limiting their effectiveness

Any business application that utilises data to transact a business operation or process requires a mechanism for transaction processing. An obvious example is a web based e-commerce application. Another is a customer booking or reservations system. But most enterprise business applications incorporate transaction processing at their core, including financials, accounts, ERP, CRM, manufacturing and logistics.

Today's prevalent transaction processing technologies were developed for mainframe and host server architectures. They have proved not to scale well in distributed, grid and cloud environments.

The result is that business applications migrated to large scale distributed environments have suffered from highly fragile transaction processing. One result of this fragility has been long and unpredictable response times in transaction execution. Another is the loss of transactions altogether, as conventional phased commit mechanisms fail to execute over complex distributed application layers.

Current Approach : Extreme Transaction Processing

The industry response to the challenges posed by transaction processing over distributed computing environments such as grids and clouds is the creation of extreme transaction processing technologies (XTP).

XTP as a class provides a range of technologies from different vendors that allow for transaction processing across many different applications spread across distributed or grid environments.

Typically, transaction processing within an XTP technology is

created within a local grid memory "space", or distributed across different grid memory spaces. Connections to database systems are provided through a mirroring function that provides an approximate transactionality. This approach at present suffers from a number of drawbacks :

- distributing transaction processing across complex environments with this approach can incur a steep degradation in performance and transaction response time;
- programming applications to execute XTP transactions can be complex and intuitively different from conventional programming approaches, creating problems with migrating applications into an XTP environment. Fine-tuning an enterprise application for an XTP environment also poses significant challenges;
- there is an issue with mirrored transactionality as a replacement for conventional two phased committed transactions.

For durable persistence, the transaction that exists in the grid or cloud memory space must be reliably committed to the database in a way that is fully ACID (atomic, consistent, isolated, durable).

XTP processes vulnerable

However, none of the XTP solution providers currently provide a robust transaction solution that addresses the gap between the in-memory space transaction and the database. This leaves XTP based transactional processing applications vulnerable to power outages, memory cache failures or other disasters. These problems can affect even the very largest cloud based infrastructures. Without credible ACID commitment processes for XTP transactions, business systems utilising such technologies are significantly more vulnerable to business interruption and reputational loss than their competitors still employing conventional hierarchical host-based transactional processing.

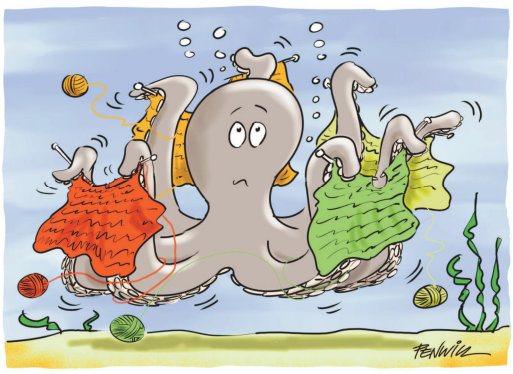

The Solution : The CloudTran Layer

CloudTran provides a number of unique features that allow the creation of extremely robust, high performance transaction processing over very large grid and cloud infrastructures with full ACID durability and persistence.

Based on Java enterprise application technology, CloudTran is a middleware product that enables enterprise application developers to build scalable applications for distributed, grid and cloud environments very quickly and easily.

CloudTran design

CloudTran :

- includes an Object-Relation Mapper layer that provides a simple Java view of distributed data. It finds data in the grid, collects it for the application and then redistributes any changes;
- caters fully for In-Memory Storage, automatically transfering data between storage services and in-memory data formats. Relational databases are supported by default and plug-in points are provided for non-relational database and other file storage systems;
- is the first Transaction Monitor specifically designed for clouds, providing ACID guarantees coordinated across distributed nodes, messaging and storage services.

The CloudTran solution is highly reliable as it provides a buffer between the customer-facing application and back-end resources, such as databases or transaction systems. This means that the application runs even when there are spikes in traffic or back-end resources are unavailable.

The Result : Robust, high performance transactions

CloudTran provides a platform for high performance transaction processing across grid or cloud computing infrastructures with full transactional commit to enterprise databases or other data stores.

It is able to handle any number of nodes and any amount of data in a single transaction, providing for almost unlimited, scale-free distributed transactional processing.

Five times faster

CloudTran can accelerate database performance: although with CloudTran usual 'secure transactions to disk' mode can be used, developers can also utilise a 'secure transactions to backed-up space' technology. This mode is about 5 times faster than the 'disk-first' mode and allows a single-CPU machine to support thousands of transactions per second.

Handles messages and non-RDBMS

CloudTran supports scenarios to coordinate message passing with persistence operations, whilst it plug-in architecture can support any type of persistent store. JDBC connections to RDBMSs are provided as standard. The technology is able to handles any number of persistent stores, with consistency supported across all databases.

Continuous performance

CloudTran buffers transactions that are to be persisted, allowing the main application to handle demand spikes without being delayed by database commits, and continue to run at full speed even when the database is down. In-flight transactions are restored from a backup copy at the transaction controller if a primary node fails over to a secondary.

Summary

Large scale organisations have an increasing need for robust, high performance distributed transaction processing capable of scaling over the largest grid or cloud software architectures.

Failure to implement secure distributed transactional systems entails a number of business and operational risks for a wide range of organisations.

CloudTran provides a robust, high performance middleware system that allows organisations to rapidly and easily integrate applications and systems across multiple nodes and software systems with high performance degradation and full ACID compliance.

"Cloud computing provides a low cost, high performance infrastructure ... CloudTran adds the data management capabilities that enterprise applications need. Integrated ORM, highly scalable transactions, and unique DBMS integration features make CloudTran the leader in cloud-based data management."
Dan Stone, Scalable Application Strategies (http://blog.scapps.co.uk/)

- ends –

For further information :
www.cloudtran.com
Matthew Fowler, NT/e 020 7529 9797