A HURWITZ WHITE PAPER

The Challenge of Managing On-line Transaction Processing Applications in the Cloud Computing World

Marcia Kaufman,
COO and Principal Analyst

HURWITZ & ASSOCIATES
Insight to Action

Sponsored by CloudTran

The Challenge of Managing On-line Transaction Processing Applications in the Cloud Computing World

In highly competitive markets, companies are increasingly turning to cloud computing as a way to improve their ability to react to market changes. Whether we are looking at financial services, healthcare, or retail, cloud computing is becoming a strategic approach. Many industries are grappling with the complexities of ensuring that online transaction applications can be effectively and safely managed in highly distributed cloud environments. Processes and technologies designed for lower transaction volumes operating in siloed environments cannot easily support unbounded transaction volumes operating in cloud and service oriented environments. Organizations with a requirement to run mission-critical online transaction processing (OLTP) applications need to manage these extreme volumes effectively in order to provider optimal service to customers.

The increasingly distributed nature of IT environments makes it challenging to optimize many of the performance requirements including transaction speed and near real-time responsiveness.

In this paper, we will present an overview of the requirements for running OLTP applications in multi-processing and cloud environments using a retail ecommerce example. In addition, we will provide insight into evolving technology requirements for overcoming scalability challenges in non-mainframe distributed environments. In addition, we will provide a perspective on managing near real-time access to data required to support OLTP transactions in non-mainframe distributed environments.

The increasingly distributed nature of IT environments makes it challenging to optimize many of the performance requirements including transaction speed and near real-time responsiveness. While all industries have specific challenges and requirements, retailers exemplify the complexities of managing OLTP in a distributed environment. Retailers with very large online channels, for example, need to collect and analyze extremely large amounts of data in real time to ensure consistent and high quality results for their customers. Data covering all aspects of the transaction, from product availability to product pricing and customer credit and payment process, must be accessible at the right time to be actionable. A system designed to handle these data requirements for a typical sales cycle may not have enough compute power, memory, or bandwidth to accommodate a surge in demand due to a special promotion or the introduction of a high-demand product. These systems lack scalability or the ability to expand IT resources as needed to meet the demand for those resources.

In order to have the scalability needed to maintain service levels with customers at all levels of the sales cycle, retailers running OLTP applications would like to leverage cloud technology. The cloud can help organizations to scale up (add more compute resources in a single node) and to scale out (add more compute resources using many nodes). However, many retailers have found it difficult to transition to the cloud. One of the key challenges they face in moving to the cloud is due to the bottlenecks that occur within their data architecture. These bottlenecks are created when access to memory and databases are insufficient to meet the real time performance requirements for complex transactions.

HURWITZ & ASSOCIATES
Insight to Action

Requirements for delivering optimal OLTP performance

Companies encounter many performance challenges when dealing with multiple nodes and multiple databases required for mission-critical transactions. The situation becomes even more complicated when organizations themselves are highly distributed and they are interacting with business partners that are located across the globe. In this type of situation, traditional "distributed transactions" can be slow and unreliable. This is primarily due to the fact that distributed transaction capabilities were largely created for smaller demand loads on fewer computer nodes. For example, a retailer that developed a point-of-sale system to manage in-store transactions only, now has to ensure that online and mobile transactions have consistent and reliable access to customer and product data as well. In other situations, a large-scale ecommerce site may need to scale up and out to accommodate higher transaction volumes at holiday time. In addition, enterprises need to have seamless transaction management with their partners.

... traditional "distributed transactions" can be slow and unreliable. This is primarily due to the fact that distributed transaction capabilities were largely created for smaller demand loads on fewer computer nodes.

With this vast range of transactions coming from a multitude of systems, Service Oriented Architecture (SOA) services, devices, and users, the negative impact on data sources is significant. Accommodating this demand, and the irregularity of the demand, requires new algorithms and techniques to be successful.

An E-commerce scenario

Consider the example of the CIO of Big Retail Corporation. Sam was recently hired to help transform the company's IT infrastructure so that the company's changing business model can be implemented quickly and changed at the same speed. Only 10 years ago the retailer focused on selling high end products with a personalized customer service approach in its chain of globally distributed stores. New competitors entered the market, offering comparable products at a fraction of the price. In order to compete, this retailer had to streamline its store operations and increasingly offer its products online.

Traditionally, retailers like Big Retail have not had the cash to strategically invest in IT. As a result, many retail corporations are faced with a siloed and inefficient approach to managing their channels. Simply put, IT did not plan for the possibility that the online channel would become paramount to long-term business success.

However, today's retail customer expects a seamless transition between in-store, online, and mobile channels. Customers don't distinguish between the channels; they assume the physical asset they want to purchase will be available whenever and wherever they makes a decision – whether in-store, online, or with a mobile device. They demand a high level of service in the store and also require a fast and high-quality shopping experience online.

HURWITZ & ASSOCIATES
Insight to Action

Sam recognizes that Big Retail's customers are the driving force behind the need for IT transformation. Customer use of the Internet for product evaluation, selection, and purchase has dramatically changed the competitive landscape for Big Retail. Many of the company's customers research products online to compare features and prices. And while they may walk into the store to touch a physical product, they will often make their eventual purchase though an online or mobile channel. The challenge for Big Retail is to provide a seamless experience for its customers across all channels. Over the years as the company developed the infrastructure to support its online channel, it never anticipated the degree to which multiple product and customer databases would need to be integrated and shared across the enterprise.

There is a huge database bottleneck that causes major problems for retailers and other organizations with highly distributed online environments.

Managing data in distributed OLTP environments

There is a huge database bottleneck that causes major problems for retailers and other organizations with highly distributed online environments. Databases containing customer and product data must be read from and written to constantly and in near real-time to support the quality and timeliness of each transaction. For example, after a customer places his selections in a shopping cart, data on product price and availability in addition to customer address and credit must be instantly accessed, analyzed and updated before the transaction can be completed. In a distributed IT environment, various data elements required for this transaction are likely to be located in multiple nodes across the entire system. Making sure that the right data is available at the right time to provide a seamless and high quality experience for the user is a goal that gets increasingly harder to achieve for some companies as they grow their business.

There is actually a lot at risk for the seller from the time a buyer places his selections in a shopping cart until his credit card is charged. Products may be unavailable when the customer is ready to buy. If there is a sudden surge of interest in a product, it may be difficult to manage the transaction volume and the company could potentially sell the same product to two customers who hit the purchase key at the same time. Sometimes a customer is ready to buy only to determine that shipping costs are much higher than they had anticipated putting the purchase at risk. If product inventory data across multiple warehouses is not well integrated, a buyer might be told the product is on backorder when it is actually available in an alternative warehouse. This type of problem happens because there is too much data that needs to be analyzed, accessed and coordinated in real time.

How can companies deal with these complex real time transaction management issues? In essence, companies follow one of the three available options described below:

Option 1. Use of in-memory data grids. This method supports the caching of frequently used data in a distributed, in-memory data grid. The quantity and variety of data that must be cached in-memory is often expansive including

HURWITZ & ASSOCIATES Insight to Action

data related to customers, individual products, shipping locations, third party suppliers, warehouse, and inventory. While this approach helps to eliminate data bottlenecks and can lead to significant improvements in transaction performance, it comes with the real risk of data loss when nodes fail. Because data and transactions are stored in memory only, loss of a node can mean failed transactions and loss of revenue. Even with the deployment of elaborate backup and recovery plans, most companies find this approach insufficient to support business-critical OLTP applications.

This approach can take many months to implement in a distributed environment and requires significant development knowledge and ongoing manual software code updates.

Option 2. Use of in-memory data grids combined with manual management of persistence. This method uses advanced software development designed to improve a company's success with in-memory data grids. One important aspect of this method is the implementation of the two-phase commit protocol to eliminate some of the risk of data loss. The two-phase commit protocol is a technical algorithm that provides the required coordination between all the processes that must take place to persist a transaction accurately to permanent disk. This approach can take many months to implement in a distributed environment and requires significant development knowledge and ongoing manual software code updates. The software engineers need to know the server location of all the data elements required for a transaction so that the data can quickly be read, reviewed, and then written in the appropriate location to keep up with the desired transaction speed. This approach requires highly skilled engineers to successfully handle moderate transaction loads with software that is written for a specific use case. Most companies find that, persisting data in this way results in systems insufficient to support more extreme loads.

Option 3. Keep with the status quo and react to exceptions. Surprisingly, this approach is very common in retail and other distributed transaction environments. Some times retailers in highly competitive markets will cut corners when it comes to designing fault tolerant architectures and quality control measures for data because they are under pressure to get to market quickly with new features and functionality. They may also eliminate some system safeguards that might slow down the system because they want to ensure that the online shopping environment performs well during spikes in demand. As a result, the system may maintain performance levels, but there are lots of places where errors may occur. In systems like this, the engineering teams analyze the product and customer databases every night looking for places where the integrity of the data has been disturbed and they make corrections as needed. Companies following this approach will obviously find it difficult to scale or react quickly to unanticipated market changes.

What's missing in the current approach to managing OLTP applications?

All of the options described above require significant manual intervention to monitor exceptions and maintain the performance levels required to compete. What companies need is an integrated approach that enables in-memory

HURWITZ & ASSOCIATES
Insight to Action

and physical database technology to work in tandem with fast and reliable persistence of data. While many companies are successful with the use of in-memory data grids for data that must be read, these companies need a way to automate the manual aspects surrounding the management of data that must be inserted, updated, or deleted. For example, if 1000 customers purchase a hot new designer bag on a retail site, the in-memory data grid software is designed with an escape route allowing the retailer's IT managers to manually make the connection between data related to the sales of this new product and the persistent database storing the customer and product information. While lengthy and technically sophisticated programming may get the job done, a more automated approach would improve scalability, repeatability, business flexibility, and bottom line results. Companies with business critical online applications need an automated way to persist the data – store data quickly on data and provide a backup – rather than relying on extensive custom coding.

Companies with requirements for business critical OLTP environments need to a way for their IT engineers to persist data to a database while using an in-memory data grid.

For large retailers and other organizations with highly distributed OLTP environments, two key questions loom large for their IT executives. How can we make sure that all database actions commit at the same time – customer information, billing, ordering of like items – and maintain quality of performance? And how can we keep up with customer expectations if we need to take a manual approach to exceptions and update? Companies with requirements for business critical OLTP environments need to a way for their IT engineers to persist data to a database while using an in-memory data grid. In order to compete effectively, the business needs to ensure its IT developers can focus more time on business specific logic and less on custom code related to exceptions, logs, and transaction specific data requirements.

Moving to a cloud model will require unprecedented escalation in the automation and repeatability in the underlying architectural processes supporting OLTP applications. All the transactions that are dependent on each other need to be linked in a way that enables changes to be handled efficiently and for performance to be tuned to meet customer expectations.

HURWITZ & ASSOCIATES Insight to Action

About Hurwitz & Associates

Hurwitz & Associates is a consulting, market research and analyst firm that focuses on how technology solutions solve real world business problems. The firm's research concentrates on disruptive technologies, such as Cloud Computing, Service Oriented Architecture and Web 2.0, Service Management, Information Management, and Social and Collaborative Computing. We help our customers understand how these technologies are reshaping the market and how they can apply them to meet business objectives. The team provides direct customer research, competitive analysis, actionable strategic advice, and thought leadership. Additional information on Hurwitz & Associates can be found at www. hurwitz.com.

HURWITZ & ASSOCIATES
Insight to Action

175 Highland Avenue, 3rd Floor • Needham, MA 02494 • Tel: 617-597-1724
www.hurwitz.com