

# Hardware vs. Amazon EC2 Cloud

## Performance In the Cloud

### White paper

September, 2009

Ron Warshawsky  
Eugene Skobel

---

**Copyright**

© 2009 Enteros, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express prior written consent of the publisher. Enteros, Inc. retains the right to update or change the contents of this document without prior notice. Enteros, Inc. assumes no responsibility for the contents of this document. Information in this document is subject to change without notice. Enteros, Inc. makes no warranty of any kind with regard to this material, including but not limited to the implied warranties of marketability and fitness for a particular purpose. Enteros, Inc. shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

---

**Content**

- Executive Summary ..... 3
- Introduction ..... 3
- Application Test Description ..... 4
- TPCC Test Description..... 4
- Load Test Description ..... 5
  - Load Test Architectural diagram ..... 5
  - Hardware Description..... 5
- Infrastructure Performance Comparison..... 6
- Load Test Environment Performance Comparison ..... 9
- Node Performance Comparison..... 10
- Transactions per Second Comparison ..... 12
- Maximum Response Time Comparison ..... 13
- Average Response Time Comparison ..... 14
- TPCC Summary Comparison ..... 15
- Conclusions ..... 16

## Executive Summary

Cloud computing, or IaaS – infrastructure as a service has great potential for bringing the IT industry to the next technology level, making infrastructure and software services more attractive and impacting hardware architecture design and procurement.

Cloud environments offer the possibility of significantly reducing issues of resource over-provisioning by introducing a “pay as you go” model. At the same time, a decision has to be made by enterprises whether to use Public clouds or to procure Private ones. Outside of security considerations, this evaluation should be primarily based on performance characteristics that Public clouds would support and their ability to maintain consistent SLA levels during length production operational cycles.

In this whitepaper, we report on the capability of the Public cloud to support operational performance of a typical Web based SME application. Our results show that the performance of the enterprise I/O intensive application still is not comparable to the results derived on hardware equipment. Our tests reveal the following:

- Cloud infrastructures can support similar (and in certain cases even higher) rates of throughput for business transactions and information retrieval requests compared to the infrastructure based on physical hardware for a typical web application of average intensity
- TPC-C results demonstrate that performance of I/O intensive applications on a hardware configuration is much better than on a similar cloud configuration
- Cloud infrastructure database components consumed more CPU (but still remained fairly low), though CPU load on application servers was comparable to the load in the lab environment
- Memory consumption seems more optimal in a cloud environment although network bandwidth consumption is negligible in both environments

Our observations show that caching layers can provide sufficient reduction on I/O for applications of average intensity. However, the cloud environment is unable to sustain high levels of I/O operation and its overall performance demonstrates unpredictability. This was independently confirmed by an earlier white paper (see reference [1])

## Introduction

Two sets of tests were executed to determine performance and scalability of a IBM Lab hardware configuration (San Mateo Innovation Center) vs. Amazon EC2 Cloud environment

Products and Applications used for Benchmarking:

- Enteros Load2Test for Cloud: The test measured CPU/MEM usage, Throughput, IO, Database and Application Server performance behavior of a typical web store application
- TPC-C: The test measured tpmC and transaction count

## Application Test Description

Tests of Business Flows included typical transactions and data retrieval requests, such as:

- a) Adding and removing shopping cart items
- b) Executing purchasing transactions
- c) Multiple types of catalog data searches
- d) "Search & Buy" transactions

## TPCC Test Description

Number of warehouses: 10

Number of Terminals: 1

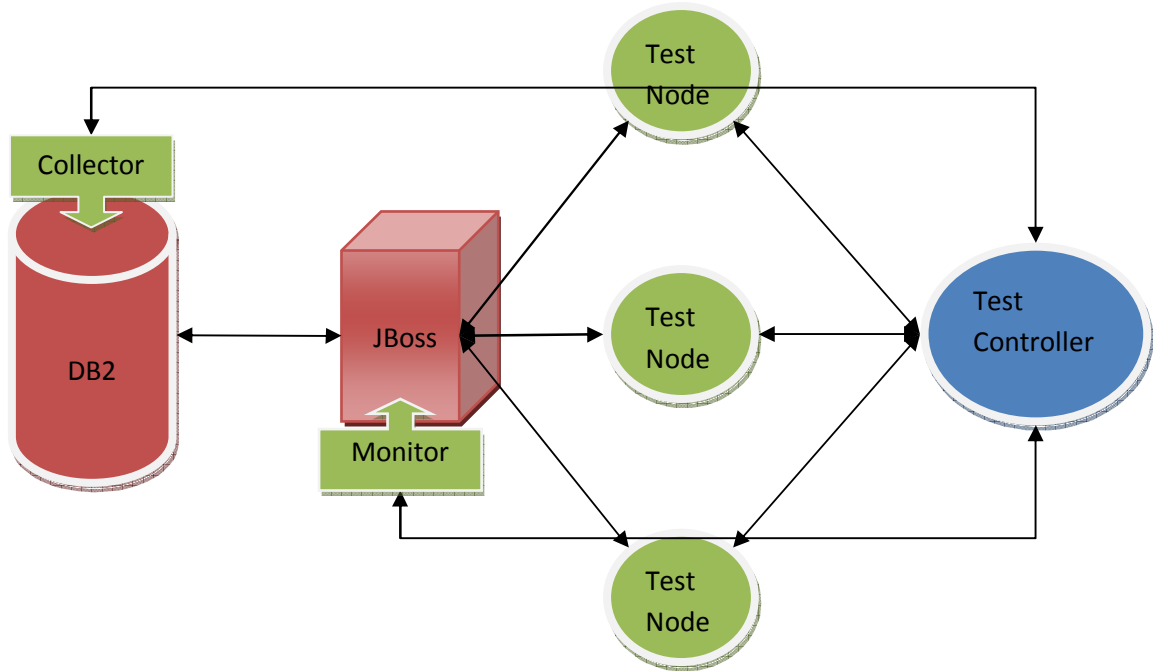
Test executed on EC2 Cloud and in IBM Lab with the following transaction ratios settings.

| <b>Transaction</b> | <b>Distribution</b> |
|--------------------|---------------------|
| New-Order          | 45%                 |
| % Payment          | 43%                 |
| % Order-Status     | 4%                  |
| % Delivery         | 4%                  |
| % Stock-Level      | 4%                  |

## Load Test Description

Test executed on IBM Lab hardware and Amazon EC2 Cloud.

### Load Test Architectural diagram



Number of concurrent users was increased in every phase of the test. CR levels used:

| CR level | Number of concurrent users |
|----------|----------------------------|
| 1        | 3                          |
| 2        | 6                          |

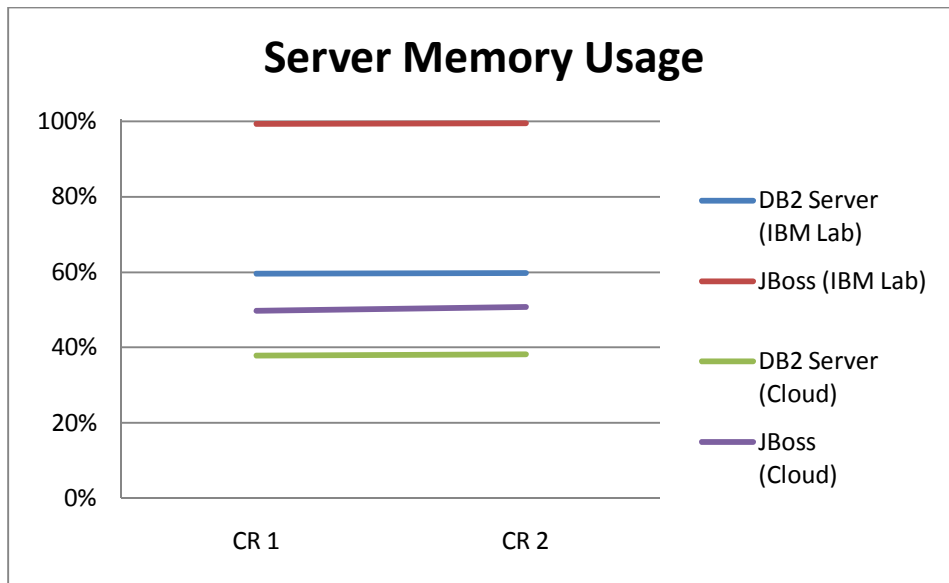
### Hardware Description

|        | IBM Lab        |                |                | Cloud  |   |                      |
|--------|----------------|----------------|----------------|--|---|----------------------|
|        | DB2            | Jboss          | Nodes          | DB2(m1.large)                                      | Jboss(m1.large)                           | Nodes(m1.small)      |
| CPU    | 4x3.60GHz Xeon | 4x3.60GHz Xeon | 4x1.2 GHz Xeon | 2x 2.66GHz Dual-Core AMD Opteron Processor 2218 HE | 2xDual-Core AMD Opteron Processor 2218 HE | 1x2.66GHz Xeon E5430 |
| Memory | 7989MB         | 7989MB         | 8114           | 7680MB   | 7687MB                                    | 1700MB               |


## Infrastructure Performance Comparison

| Infrastructure Monitors |                      |   |   |                    |   |  |  |  |  |
|-------------------------|----------------------|---|---|--------------------|---|--|--|--|--|
| IBM Lab                 |                      |   |   | Cloud              |   |  |  | Delta                                      |  |
| Memory                  | DB2 Server (IBM Lab) | 59.60%                                      | 59.80%                                      | DB2 Server (Cloud) | 37.90%                                    | 38.20%                                     | 21.70%                                     | 21.60%                                     |  |
|                         | JBoss (IBM Lab)      | 99.30%                                      | 99.40%                                      | JBoss (Cloud)      | 49.80%                                    | 50.80%                                     | 49.50%                                     | 48.60%                                     |  |
| CPU                     | DB2 Server (IBM Lab) | Avg./Min/Max<br>2.3/0.0/12.9<br>0 spikes    | Avg./Min/Max<br>2.5/0.0/9.0<br>0 spikes     | DB2 Server (Cloud) | Avg./Min/Max<br>7.6/2.0/20.8<br>0 spikes  | Avg./Min/Max<br>6.8/0.0/13.9<br>0 spikes   | Avg./Min/Max<br>-5.3/-2.0/-7.9<br>0 spikes | Avg./Min/Max<br>-4.3/0.0v/-4.9<br>0 spikes |  |
|                         | JBoss (IBM Lab)      | Avg./Min/Max<br>65.8/42.0/100.0<br>7 spikes | Avg./Min/Max<br>68.8/43.0/98.0<br>13 spikes | JBoss (Cloud)      | Avg./Min/Max<br>68.5/7.0/96.0<br>6 spikes | Avg./Min/Max<br>72.0/55.4/92.1<br>9 spikes | Avg./Min/Max<br>-2.7/35.0/4.0<br>1 spikes  | Avg./Min/Max<br>-3.2/-12.4/5.9<br>4 spikes |  |
| #Users                  | 3                    |   | 6   |                    | 3   |  | 6  |  |  |
| % of baseline           | 100%                 |   | 200%  |                    | 100%                                      |  | 200%                                       |  |  |

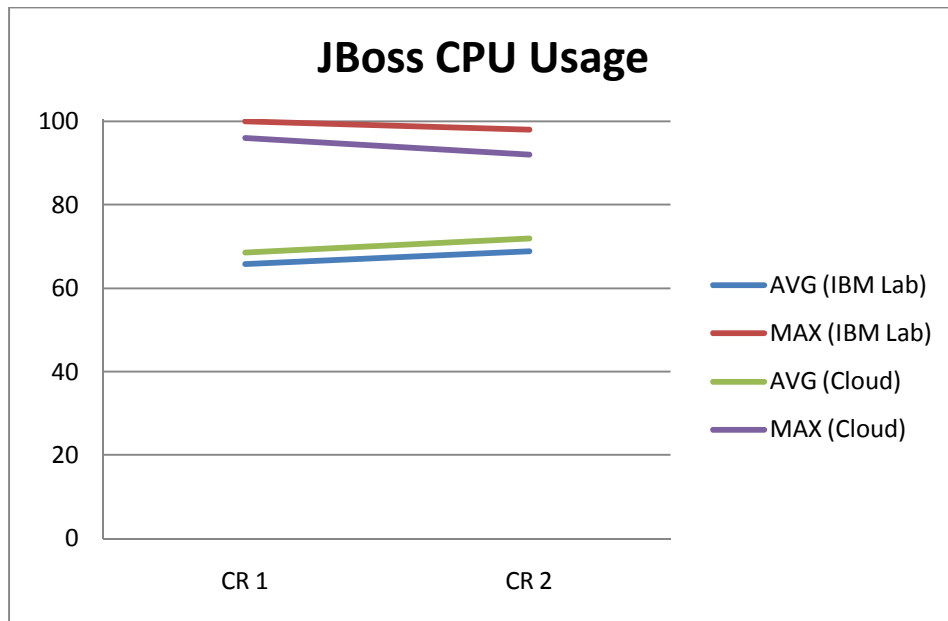
- DB2 Server and JBoss memory utilization was twice as efficient when running on the cloud, than in IBM Lab



- DB2 Server CPU usage was fairly low during cloud and IBM Lab testing, however performance was slightly better in the IBM Lab



- Jboss utilized more CPU on average when running on cloud versus the IBM Lab, however in the IBM Lab CPU spiked more often and with higher peaks



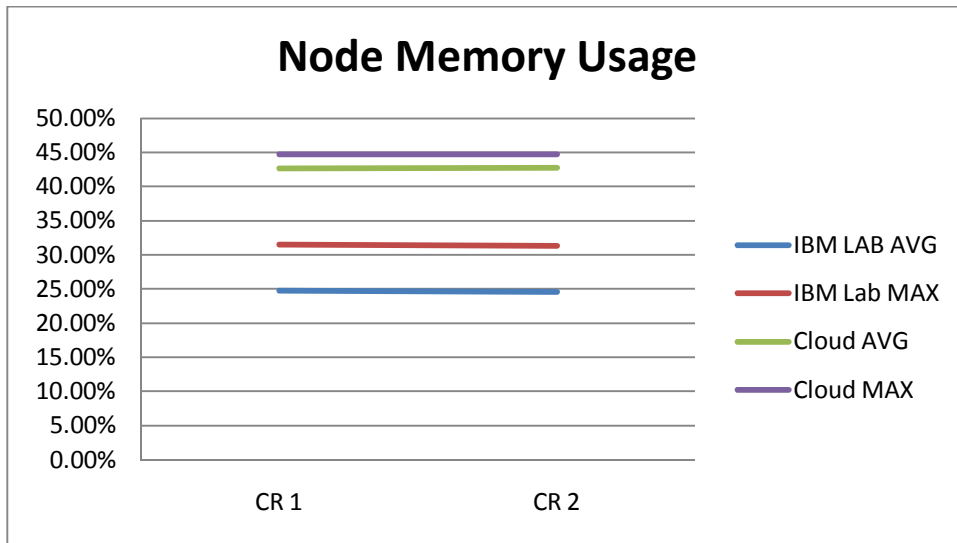


## Load Test Environment Performance Comparison


| System Performance          |                |   |  |                          |   |   |   |  |  |
|-----------------------------|----------------|---|--|--------------------------|---|---|---|--|--|
| IBM Lab                     |                |   |  | Cloud                    |   |   |   | Delta                                      |  |
| <b>Memory</b>               | <b>Sx33618</b> | 31.50%                                  | 31.30%                                 | <b>l2t_ec2_dynamic_1</b> | 44.70%                                  | 44.70%                                  | -13.20%                                   | -13.40%                                    |  |
|                             | <b>Sx33619</b> | 22.60%                                  | 22.40%                                 | <b>l2t_ec2_dynamic_2</b> | 44.50%                                  | 44.60%                                  | -21.90%                                   | -22.20%                                    |  |
|                             | <b>Sx33620</b> | 20.20%                                  | 19.90%                                 | <b>l2t_ec2_dynamic_3</b> | 38.90%                                  | 39.00%                                  | -18.70%                                   | -19.10%                                    |  |
| <b>CPU</b>                  | <b>Sx33618</b> | Avg/Min/Max<br>0.8/0.0/10.0<br>0 spikes | Avg/Min/Max<br>0.8/0.0/6.0<br>0 spikes | <b>l2t_ec2_dynamic_1</b> | Avg/Min/Max<br>3.7/0.0/39.0<br>0 spikes | Avg/Min/Max<br>4.5/0.0/24.0<br>0 spikes | Avg/Min/Max<br>-2.9/0.0/-29.0<br>0 spikes | Avg./Min/Max<br>-3.7/0.0/-18.0<br>0 spikes |  |
|                             | <b>Sx33619</b> | Avg/Min/Max<br>0.6/0.0/3.0<br>0 spikes  | Avg/Min/Max<br>0.8/0.0/4.0<br>0 spikes | <b>l2t_ec2_dynamic_2</b> | Avg/Min/Max<br>3.3/0.0/15.0<br>0 spikes | Avg/Min/Max<br>5.7/0.0/44.0<br>0 spikes | Avg/Min/Max<br>-2.7/0.0/-12.0<br>0 spikes | Avg/Min/Max<br>-4.9/0.0/-40.0<br>0 spikes  |  |
|                             | <b>Sx33620</b> | Avg/Min/Max<br>0.4/0.0/3.0<br>0 spikes  | Avg/Min/Max<br>0.7/0.0/3.0<br>0 spikes | <b>l2t_ec2_dynamic_3</b> | Avg/Min/Max<br>2.5/0.0/11.0<br>0 spikes | Avg/Min/Max<br>6.7/0.0/50.0<br>0 spikes | Avg/Min/Max<br>-2.1/0.0/-8.0<br>0 spikes  | Avg/Min/Max<br>-6.0/0.0/-47.0<br>0 spikes  |  |
| <b>#Users</b>               | 3              |   | 6                                      |                          | 3                                       |   | 6   |  |  |
| <b>% of baseline</b>        | 100%           |   | 200%                                   |                          | 100%                                    |   | 200%                                      |  |  |
| <b>TPS</b>                  | 0.051          |   | 0.103                                  |                          | 0.059                                   |   | 0.118                                     |  |  |
| <b>TPS/Concurrent Users</b> | 0.051          |   | 0.052                                  |                          | 0.059                                   |   | 0.059                                     |  |  |
|                             |                |   |  |                          |   |   |   |  |  |

## Node Performance Comparison


- The memory usage distribution on hardware and cloud nodes varied from machine to machine by about 11% and 6%, respectively
- The average memory usage by nodes deployed in the cloud was around 18% higher than for the IBM Lab



- Overall cloud maximum and average CPU utilization was higher than for the IBM Lab CPU utilization, however both remained fairly low

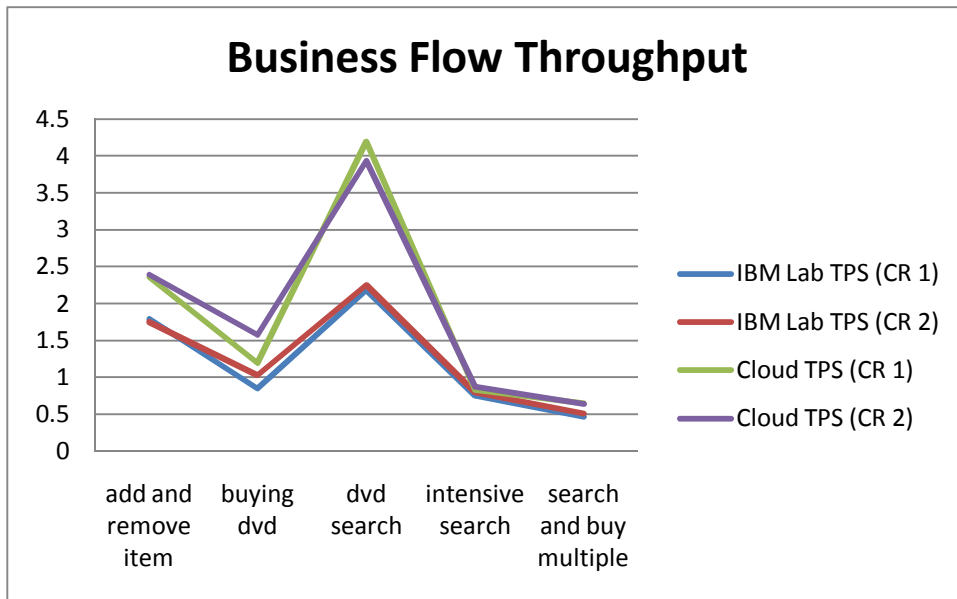


- The number of transactions per second executed using IBM Lab nodes was slightly less than transactions executed using cloud nodes




## Transactions per Second Comparison

|       |                         | Business Flow Throughput  |                           |                           |                           |
|-------|-------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
|       |                         | IBM Lab                   |                           | Cloud                     |                           |
|       |                         | 100%                      | 200%                      | 100%                      | 200%                      |
|       |                         | 3 concurrent active users | 6 concurrent active users | 3 concurrent active users | 6 concurrent active users |
|       |                         | avg                       | avg                       | avg                       | avg                       |
| rbf-1 | add and remove item     | 1.7944                    | 1.7536                    | 2.3559                    | 2.3948                    |
|       | buying dvd              | 0.8472                    | 1.0285                    | 1.193                     | 1.5778                    |
|       | dvd search              | 2.1808                    | 2.2493                    | 4.1957                    | 3.9323                    |
|       | intensive search        | 0.7548                    | 0.7916                    | 0.8314                    | 0.8724                    |
|       | search and buy multiple | 0.4669                    | 0.5099                    | 0.6474                    | 0.6394                    |



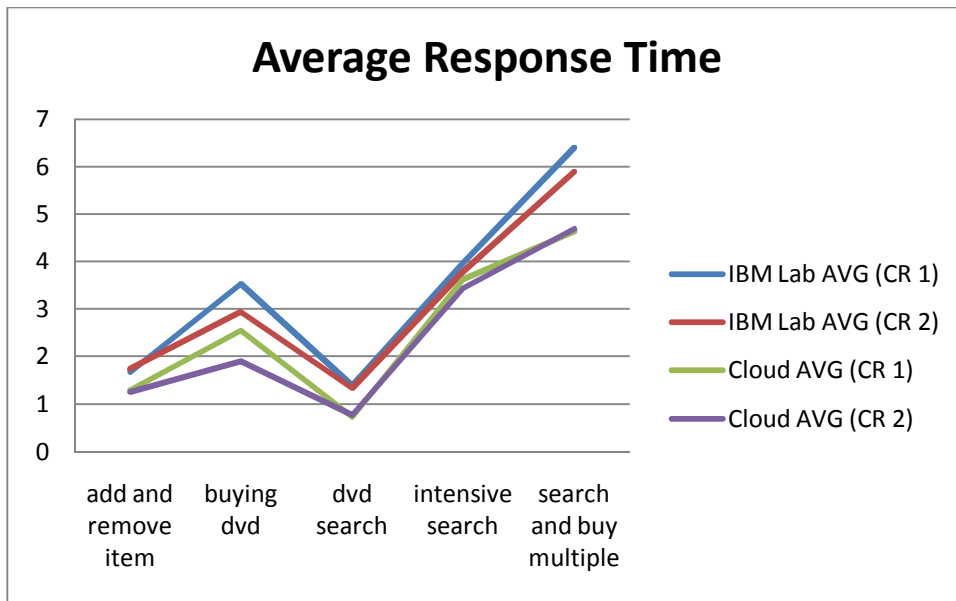
## Maximum Response Time Comparison

|       |                         | Business Flow Response Time |                           |                           |                           |
|-------|-------------------------|-----------------------------|---------------------------|---------------------------|---------------------------|
|       |                         | IBM Lab                     |                           | Cloud                     |                           |
|       |                         | 100%                        | 200%                      | 100%                      | 200%                      |
|       |                         | 3 concurrent active users   | 6 concurrent active users | 3 concurrent active users | 6 concurrent active users |
|       |                         | max                         | max                       | max                       | max                       |
| rbf-1 | add and remove item     | 3.33                        | 2.74                      | 2.44                      | 2.29                      |
|       | buying dvd              | 9.66                        | 6.13                      | 8.46                      | 3.73                      |
|       | dvd search              | 6.23                        | 4.21                      | 0.99                      | 1.67                      |
|       | intensive search        | 7.98                        | 5.05                      | 9.55                      | 4.36                      |
|       | search and buy multiple | 12.74                       | 9.06                      | 7.9                       | 5.83                      |






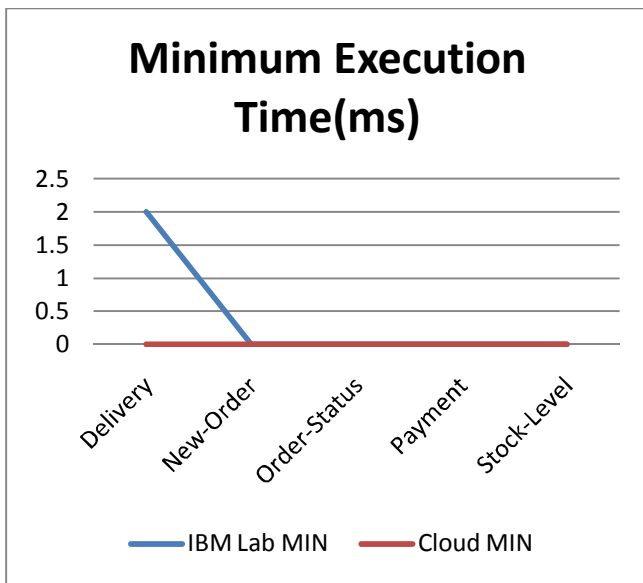
## Average Response Time Comparison

|       |                         | Business Flow Response Time |                           |                           |                           |
|-------|-------------------------|-----------------------------|---------------------------|---------------------------|---------------------------|
|       |                         | IBM Lab                     |                           | Cloud                     |                           |
|       |                         | 100%                        | 200%                      | 100%                      | 200%                      |
|       |                         | 3 concurrent active users   | 6 concurrent active users | 3 concurrent active users | 6 concurrent active users |
|       |                         | avg                         | avg                       | avg                       | avg                       |
| rbf-1 | add and remove item     | 1.67                        | 1.74                      | 1.29                      | 1.25                      |
|       | buying dvd              | 3.53                        | 2.94                      | 2.54                      | 1.9                       |
|       | dvd search              | 1.39                        | 1.33                      | 0.72                      | 0.77                      |
|       | intensive search        | 3.96                        | 3.79                      | 3.62                      | 3.44                      |
|       | search and buy multiple | 6.4                         | 5.89                      | 4.64                      | 4.69                      |



# TPCC Summary Comparison

|               | Hardware |      |       |         | Cloud   |     |       |        |
|---------------|----------|------|-------|---------|---------|-----|-------|--------|
|               | MIN      | MAX  | AVG   | TPS     | MIN     | MAX | AVG   | TPS    |
| Delivery      | 2        | 54   | 2.662 | 34.868  | 0       | 234 | 15.63 | 10.79  |
| New-Order     | 0        | 2156 | 1.051 | 399.490 | 0       | 672 | 3.28  | 120.54 |
| Order-Status  | 0        | 191  | 0.976 | 35.055  | 0       | 328 | 2.72  | 10.51  |
| Payment       | 0        | 210  | 0.946 | 379.895 | 0       | 282 | 3.26  | 115.03 |
| Stock-Level   | 0        | 49   | 0.718 | 35.407  | 0       | 219 | 1.74  | 10.64  |
| Measured tpmC | 7231.92  |      |       |         | 7231.92 |     |       |        |



## **Conclusions**

Benchmarking tests were undertaken between a physical hardware lab configuration and a comparable virtual cloud configuration. Results reveal for a typical web application of average intensity, the cloud infrastructure was able to support similar (and in certain cases even higher) rates of throughput for business transactions and information retrieval requests than does the physical hardware.

At the same time, TPC-C results demonstrate that performance of I/O intensive applications on a hardware configuration is much better than on a similar cloud configuration.

Our tests demonstrate that cloud infrastructure database components consume more CPU (that still remained fairly low) while CPU load on application servers were comparable to the load in lab environment.

Memory consumption appears more optimal in cloud environment, while network bandwidth consumption is negligible in both environments.

References:

[1] "Above the Clouds: A Berkeley View of Cloud Computing"

### **About Enteros**

Enteros solutions lead the industry in Production Performance Management both in the Cloud and on physical hardware. Enteros software and services proactively find and fix performance problems in business-critical applications, Web sites and data centers with unprecedented speed, accuracy, precision, and scope, reducing downtime and increasing customer retention and satisfaction. Established in 2004, Enteros is headquartered in Santa Clara, California with offices in Boston, Israel and the Ukraine. For more information, visit [www.enteros.com](http://www.enteros.com).