

# **Digesting an Open-Source Fair-Use TPC-E Implementation: DBT-5**

2007-May, Ottawa

**Mark Wong  
Rilson Nascimento**

PGCon 2007

# Introductions

- Mark Wong
  - Database Performance specialist
  - Previously worked at OSDL (presently, Linux Foundation)
- Rilson Nascimento
  - MSc Candidate in the Federal University of Pernambuco, Brazil
  - Previously worked at Itautech Performance Lab

# Table of Contents

- TPC-E Overview
- Workload Architecture
- Experimental Results
- Research
- Future Work

# Question Policy

- Interrupt us if something is unclear
- Keep long generic questions to the end
- Approach us during the conference
- Write us
  - [markwkm@gmail.com](mailto:markwkm@gmail.com)
  - [rilson.nascimento@gmail.com](mailto:rilson.nascimento@gmail.com)

## Why TPC-E?

- <http://www.tpc.org/tpce/spec/TPCEpresentation.ppt>
- TPC-C is over 14 years old
- Not practical to modify existing workload
- Transaction are too lightweight by today's standards
- CPU performance grew according to Moore's Law
- Disk latency did not
- Reduce cost/complexity of running benchmark
- Encourage DB uses which is more representative of what customer do

## TPC-E vs. DBT-5

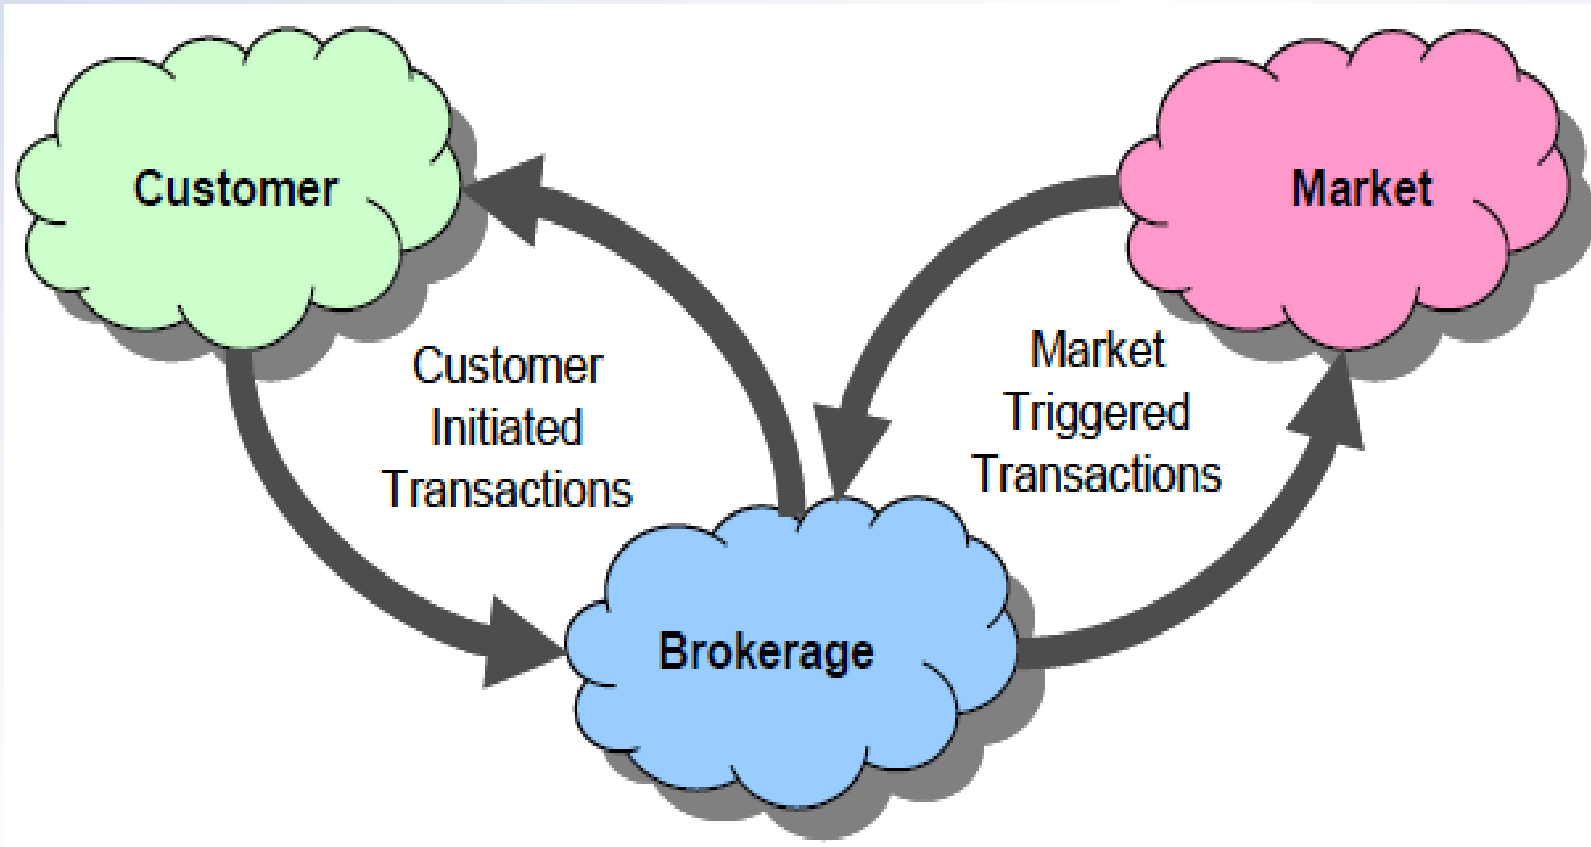
- TPC-E is a benchmarking specification for generating marketing collateral
- DBT-5 a test kit to help improve PostgreSQL

# TPC-E Overview

- TPC-E Goals
  - OLTP Database-centric workload
  - Comparability of results
  - Familiar business model – easy to understand
  - Reduce cost of running benchmark
  - Enhance schema complexity
  - Be more representative to what customers do

## TPC-E Overview (2)

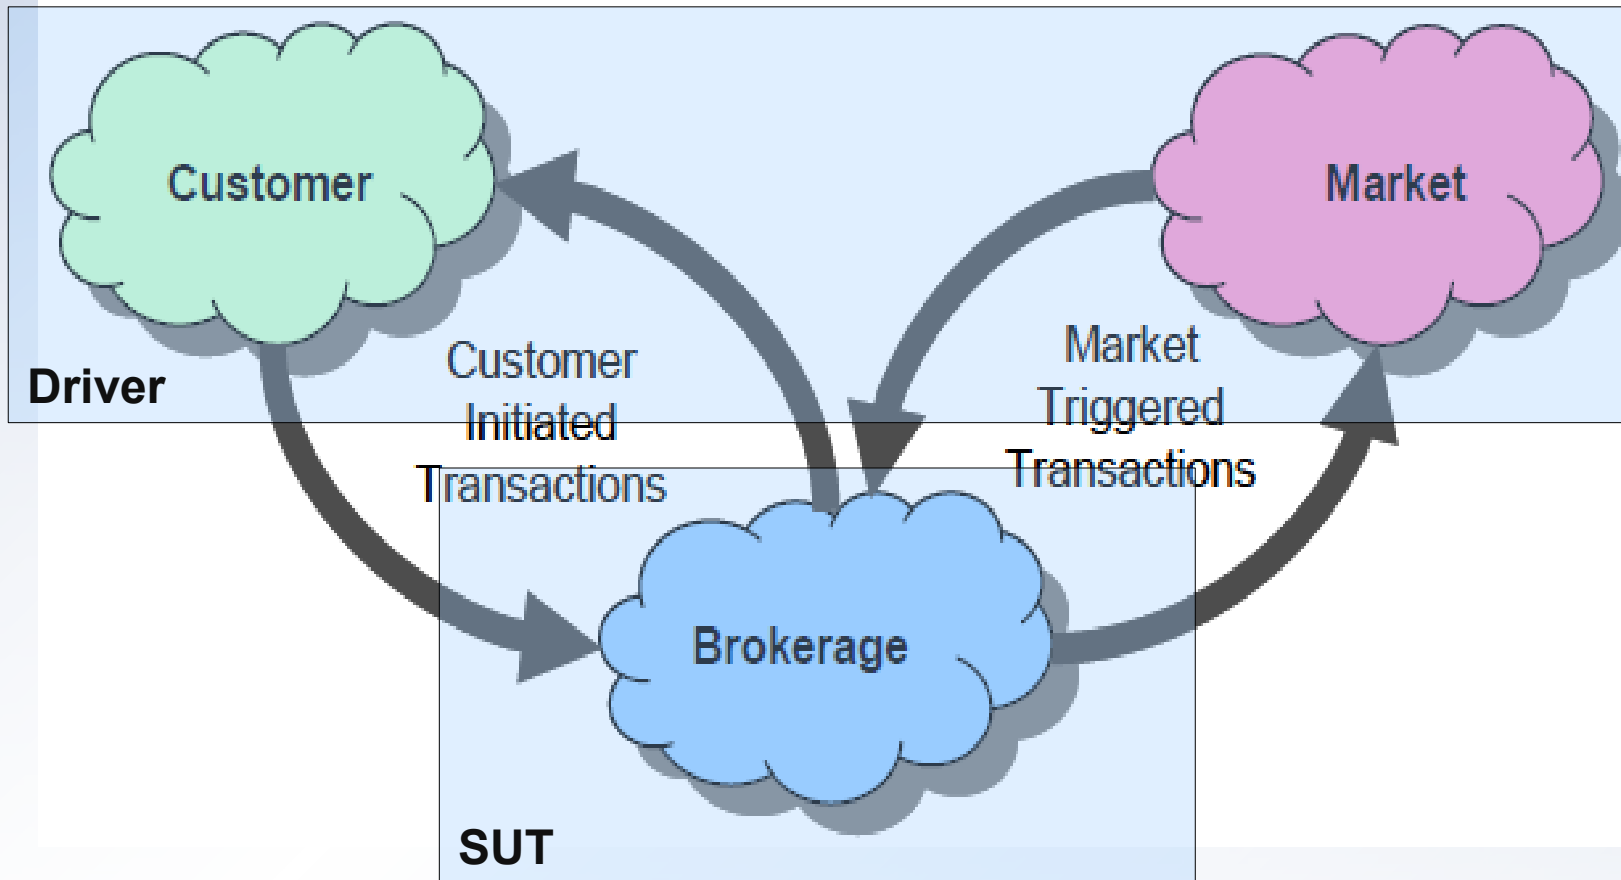
- Business Model – Financial Market





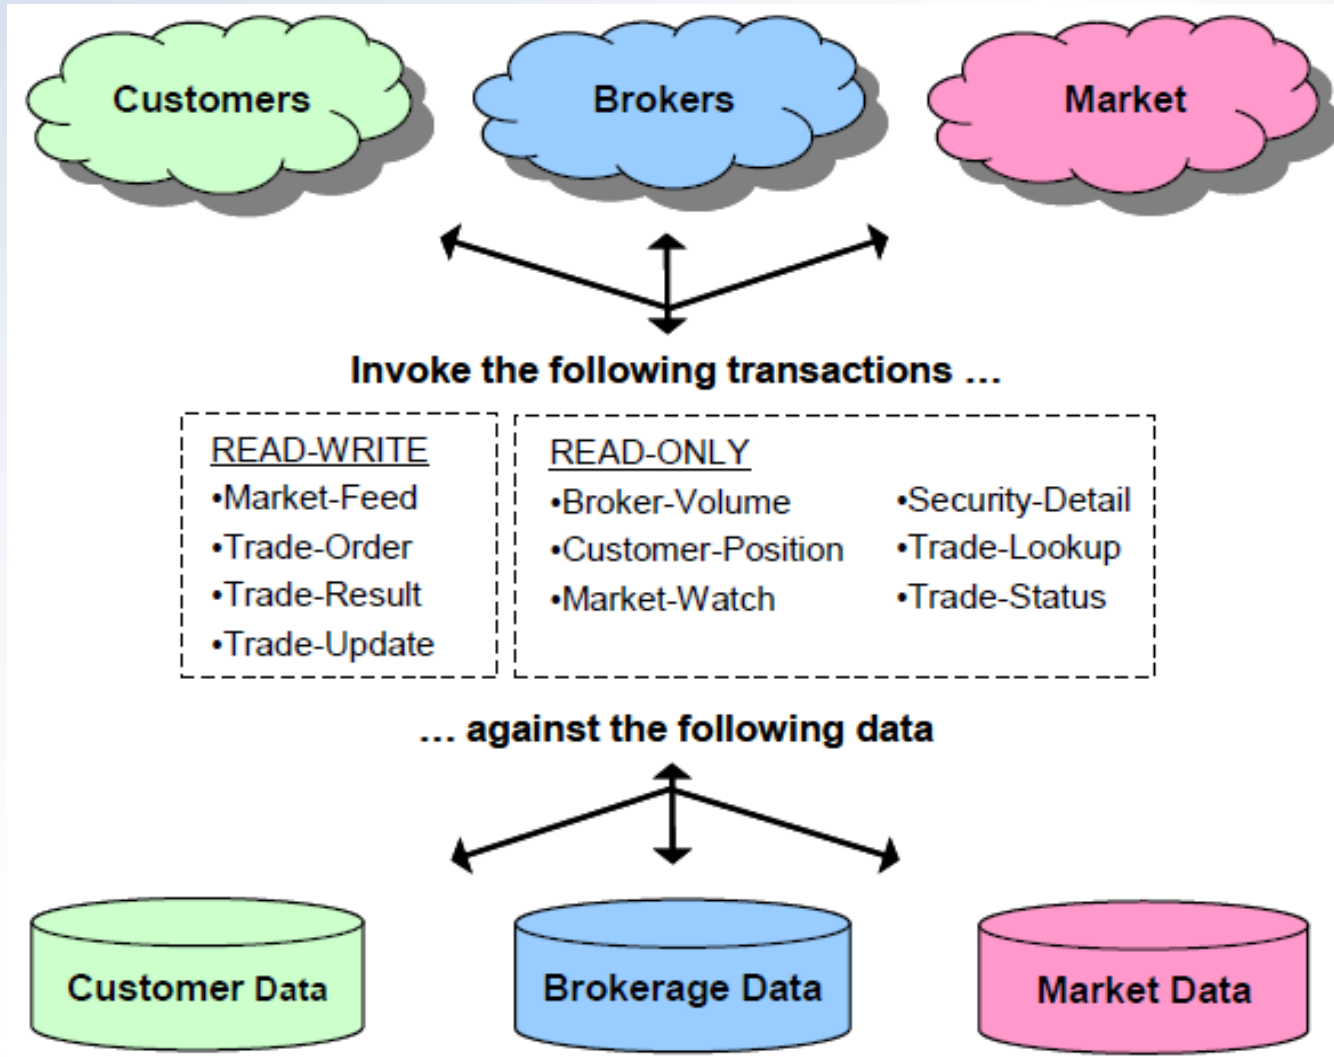
# TPC-E Overview (3)

- Business Model – Financial Market



# TPC-E Overview (4)

- Business Model – Financial Market



## TPC-E Overview (5)

- Business Model – Comparison with TPC-C

### TPC-C

- Wholesale supplier
- Organized by
  - Warehouses
    - Districts
    - Customers

### TPC-E

- Brokerage House
- Organized by
  - Customers
    - Accounts
    - Securities
    - Companies

## TPC-E Overview (6)

- Database Schema – Comparison with TPC-C

Characteristic	TPC-E	TPC-C
Tables	33	9
Columns	188	92
Min Cols / Table	2	3
Max Cols / Table	24	21
Data Type Count	Many	4
Data Types	UID, CHAR, NUM, DATE, BOOL, LOB	UID, CHAR, NUM, DATE
Primary Keys	33	8
Foreign Keys	50	9
Tables w / Foreign Keys	27	7
Check Constraints	22	0
Referential Integrity	Yes	No

# TPC-E Overview (7)

- Database Tables and Scaling

Customer	Broker	Market	
ACCOUNT_PERMISSION	BROKER	COMPANY	SECTOR
CUSTOMER	CASH_TRANSACTION	COMPANY_COMPETITOR	SECURITY
CUSTOMER_ACCOUNT	CHARGE	DAILY_MARKET	
CUSTOMER_TAXRATE	COMMISSION_RATE	EXCHANGE	
HOLDING	SETTLEMENT	FINANCIAL	Dimension
HOLDING_HISTORY	TRADE	INDUSTRY	ADDRESS
HOLDING_SUMMARY	TRADE_HISTORY	LAST_TRADE	STATUS_TYPE
WATCH_ITEM	TRADE_REQUEST	NEWS_ITEM	TAXRATE
WATCH_LIST	TRADE_TYPE	NEWS_XREF	ZIP_CODE
Legend:			
	Fixed Tables	Growing Tables	Scaling Tables

# TPC-E Overview (8)

- Transactions

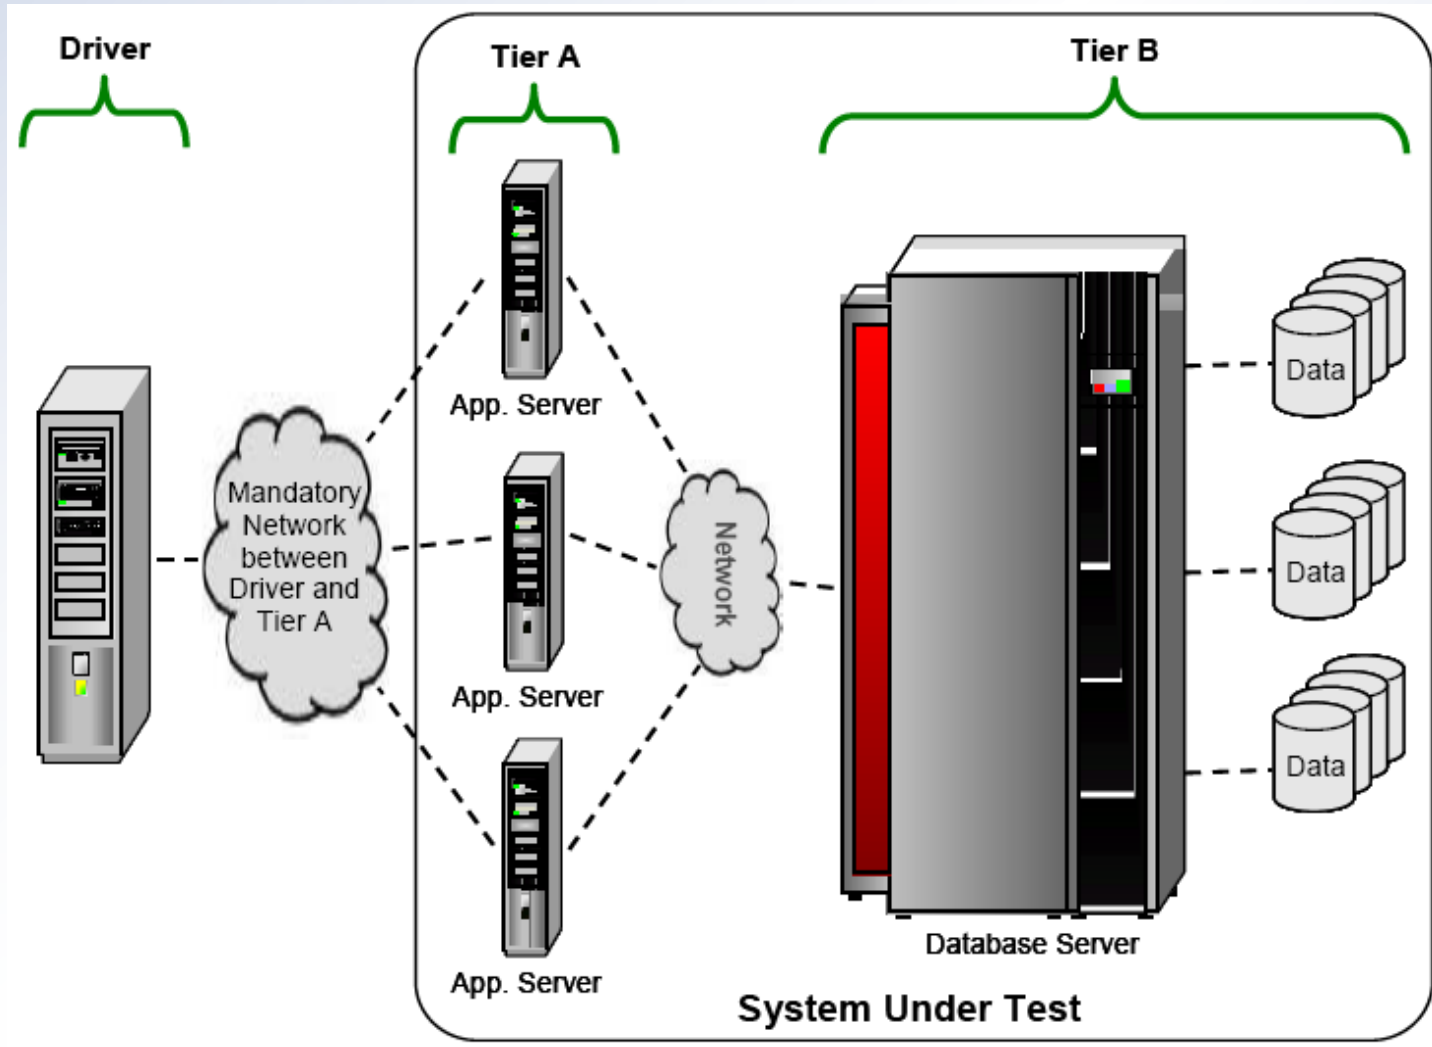
Transaction	Weight	Access	Mix%	90% Response Time Constraint
Trade-Order	Heavy	R/W	10.1	2 sec.
Trade-Result	Heavy	R/W	10	2 sec.
Trade-Lookup	Medium	R/O	8	3 sec.
Trade-Update	Medium	R/W	2	3 sec.
Trade-Status	Light	R/O	19	1 sec.
Customer Position	Mid-Heavy	R/O	13	3 sec.
Broker Volume	Mid-Heavy	R/O	4.9	3 sec.
Security Detail	Medium	R/O	14	3 sec.
Market Feed	Medium	R/W	1	2 sec.
Market Watch	Medium	R/O	18	3 sec.
Data Maintenance	Light	R/W	-	-
Trade-Cleanup	Medium	R/W	-	-

# TPC-E Overview (9)

- Metrics
  - Performance (throughput), expressed in tpsE
  - Price/Performance, expressed in price/tpsE
  - Availability Date, when all products necessary to achieve the stated performance will be available

# TPC-E Overview (10)

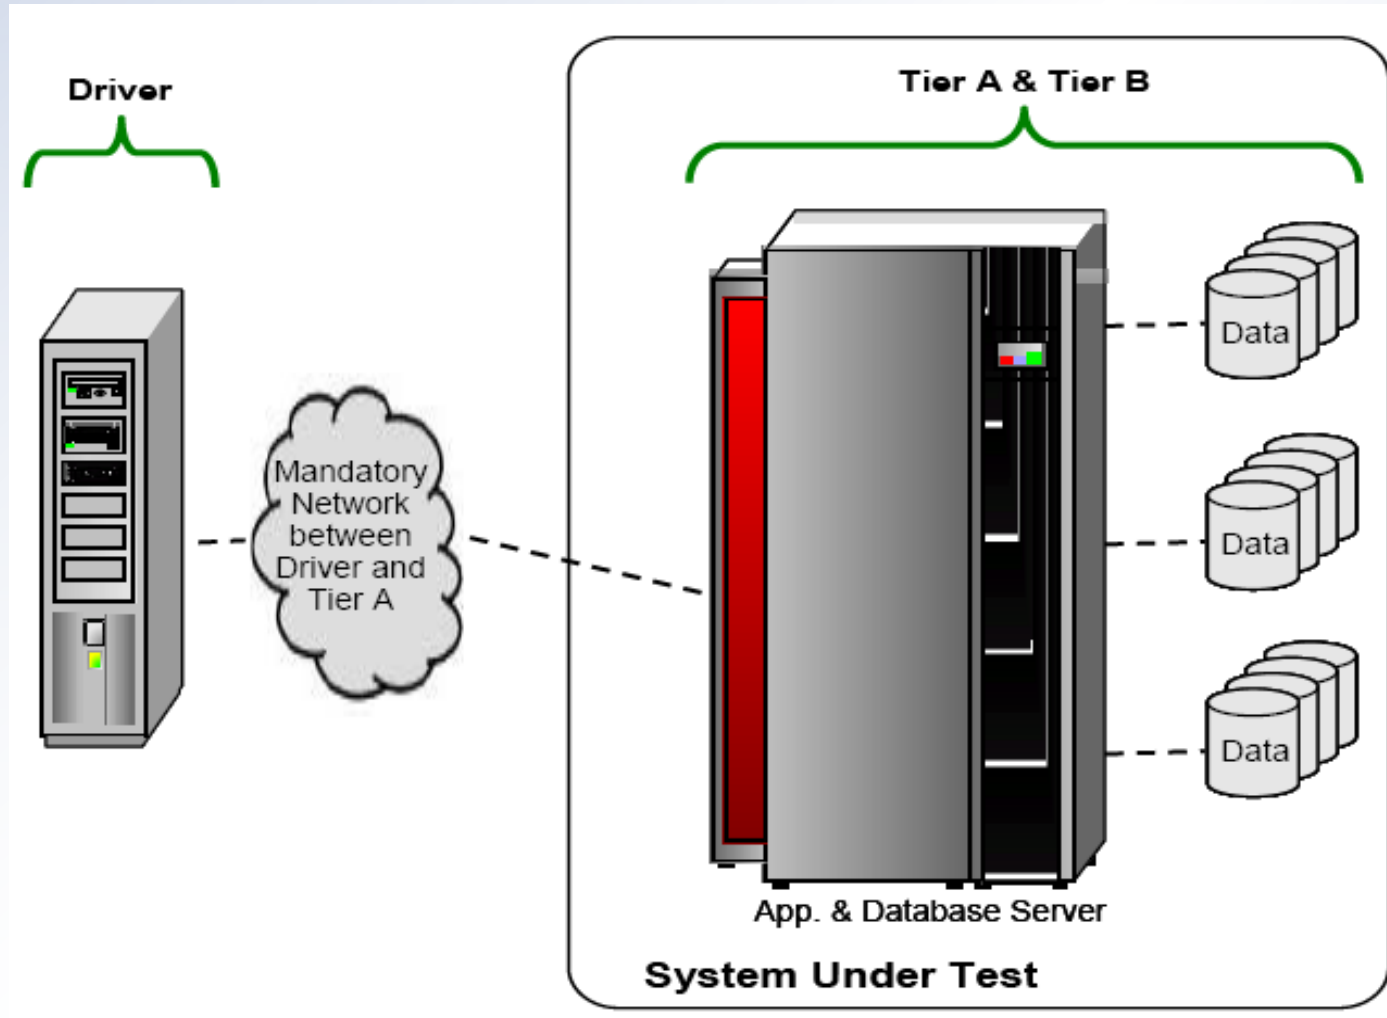
- Sample Test Configuration






# TPC-E Overview (11)

- Sample Test Configuration, Variation I



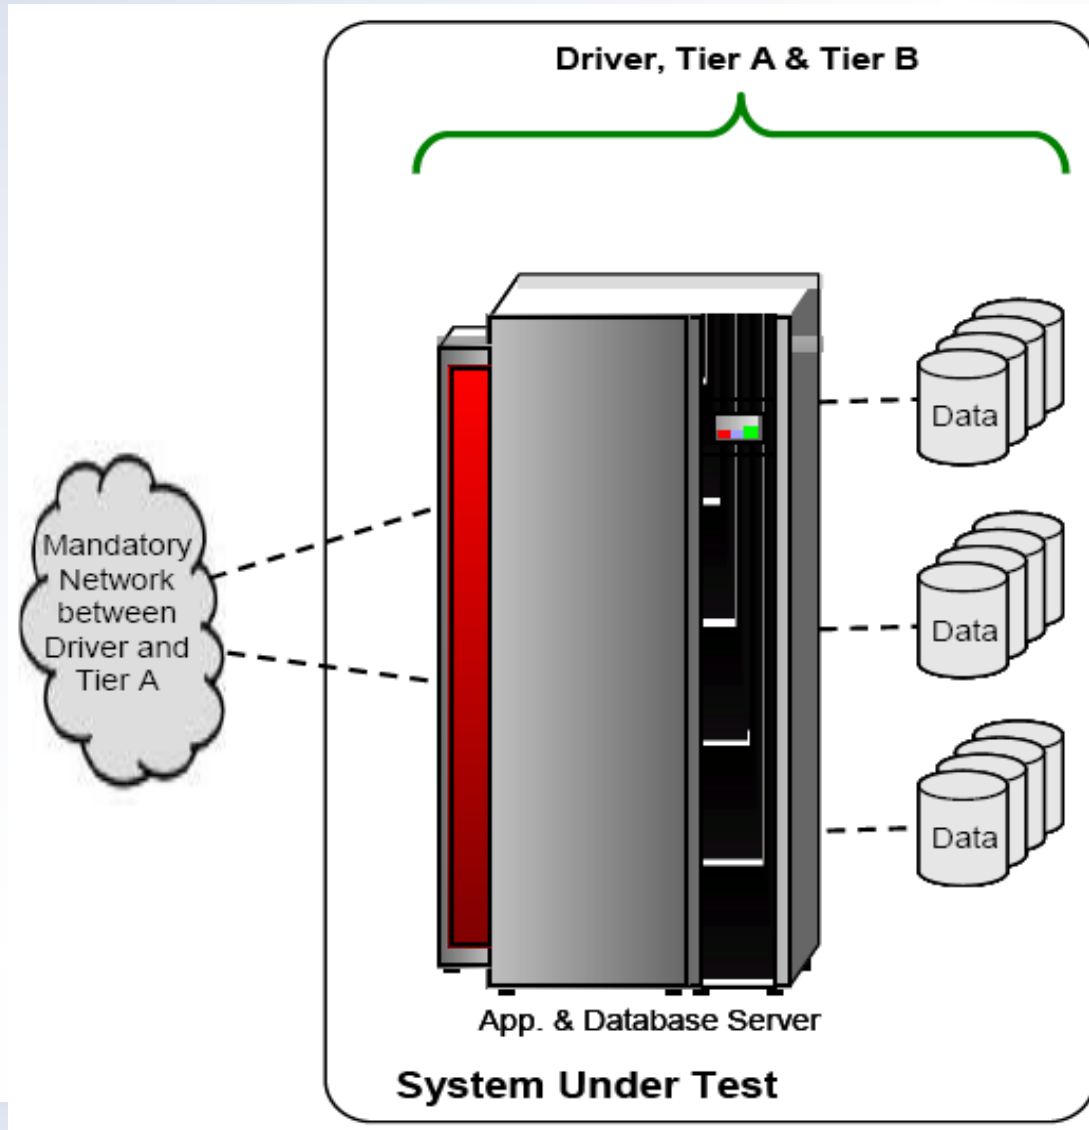
# TPC-E Overview (12)

- Sample Test Configuration, Variation II



# TPC-E Overview (13)

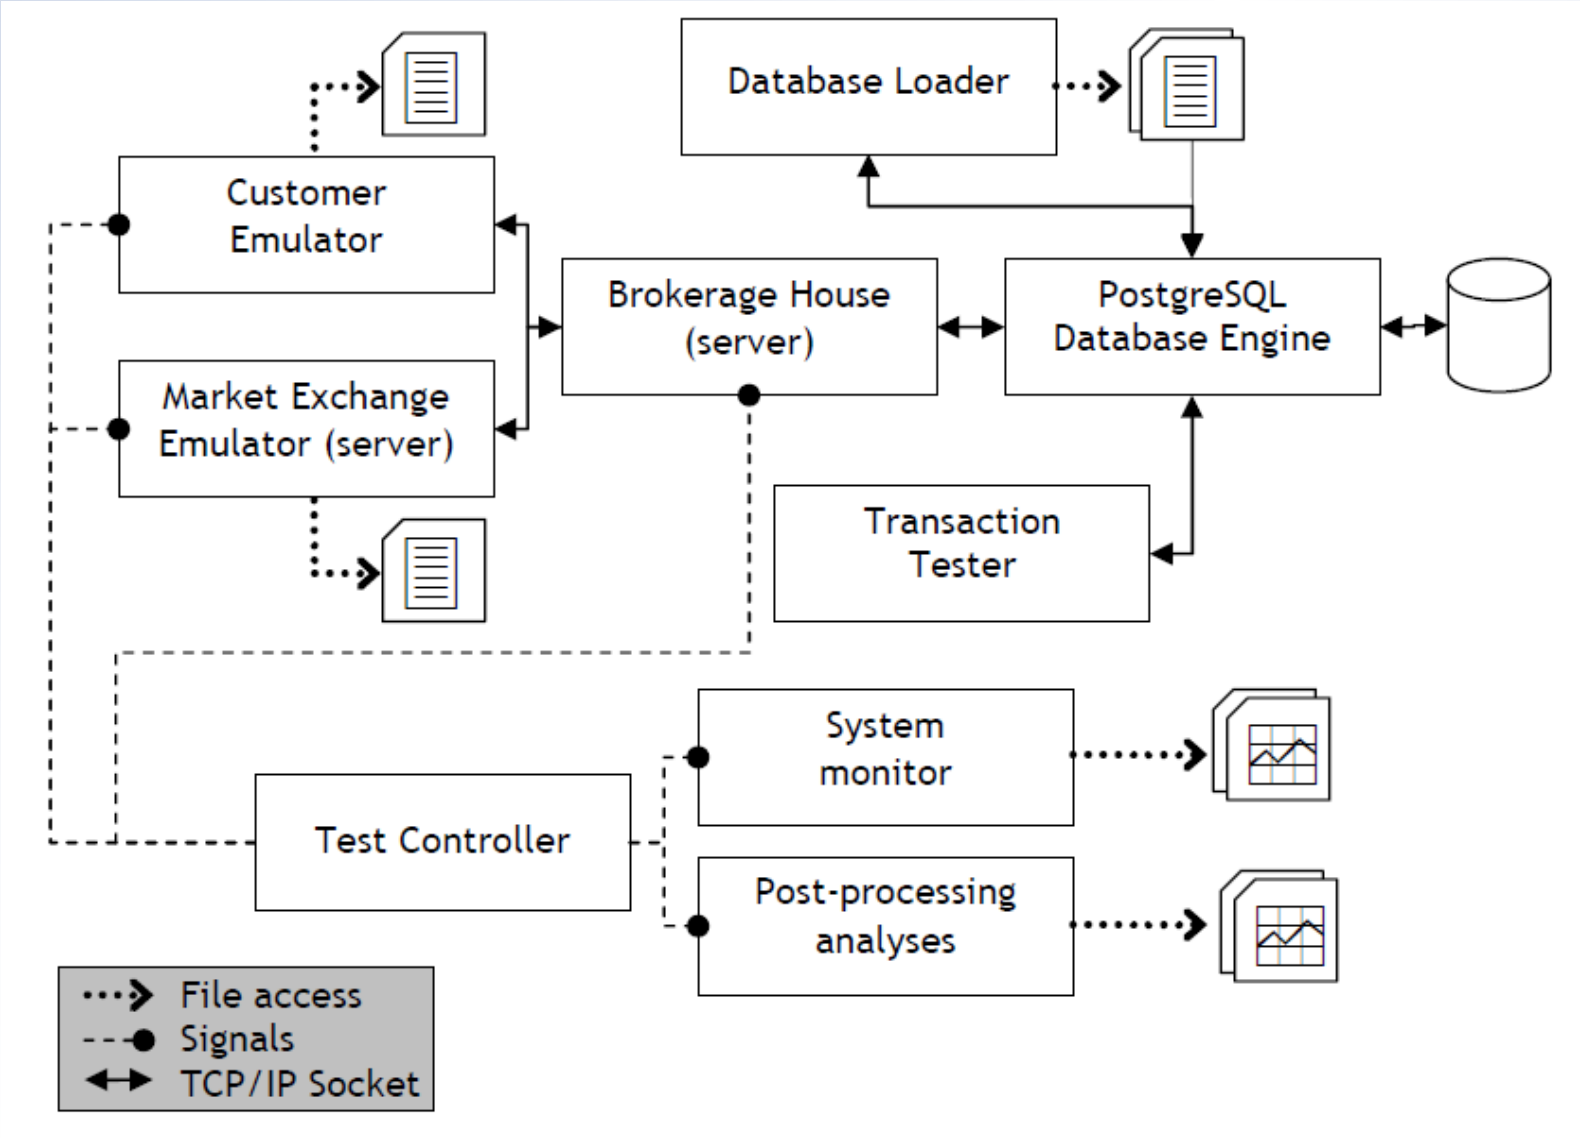
- Sample Test Configuration, Variation III



# TPC-E Overview – Highlights and Benefits

- Financial business model
- Rich transaction set
- Diverse, realistic schema
  
- Server-centric workload with DB focus
- Realistic application model
- Rebalanced hardware configuration
- Specification provides code where sponsor creativity is not being tested

# DBT-5 Architecture




## DBT-5 Architecture - Loader

- Database creation is vendor specific, but...
- Database population can be vendor neutral
- TPC-E includes a data generator for database loading
  - C++ code to generate data
  - Flat file generation is provided
  - Sponsor is free to create to customize interface
- Libpqxx: C++ API for PostgreSQL

# Experimental Results

- Test bed
  - Processor: Intel(R) Xeon(TM) CPU 2.80GHz w/HT
  - Memory: 3 GB
  - Disk: 14 disks in hw RAID 0 (data)
  - Operating System: Linux 2.6.20-gentoo-r4
  - Database Engine: PostgreSQL 8.2.3
  - Database Size: 2806 MB (1000 customers, 50 ITD)

# Experimental Results – DBT-5 Report



## Database Test 5 Report

Fri May 11 11:11:06 BRT 2007

### Results Summary


Trade Result Transactions per Seconds (trtps): 2.02

Scale Factor: 1000

Test Duration (min.): 30.00

Ramp-up Time (min.): 0.02

Total Unknown Errors: 20




### Transaction Summary

Transaction			Response Time		Rollbacks		Charts	
Name	Mix %	Total	Average (s)	90th %	Total	%	Response Time	Time Distribution
Trade Order	10.11	3831	0.02	0.04	38	0.99		
Trade Result	9.60	3635	0.04	0.08	0	0.00		
Trade Lookup	8.09	3063	0.35	0.94	0	0.00		



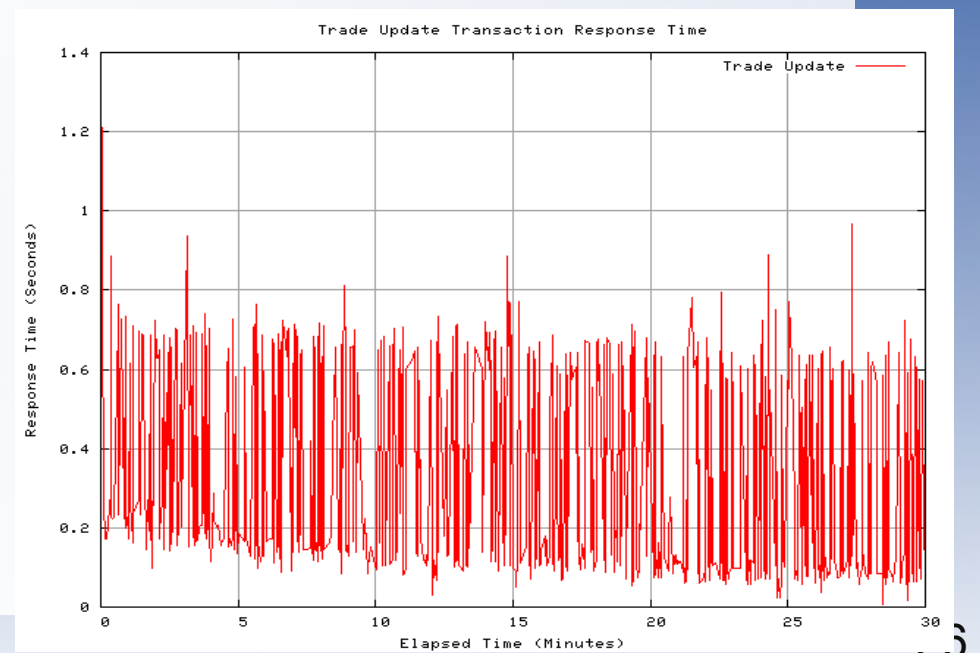
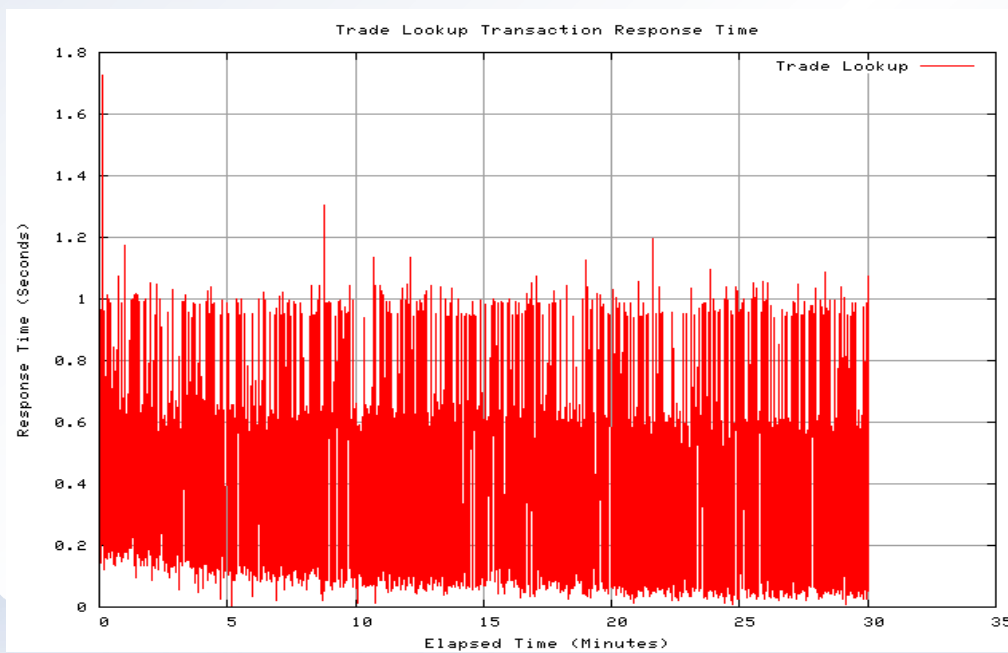
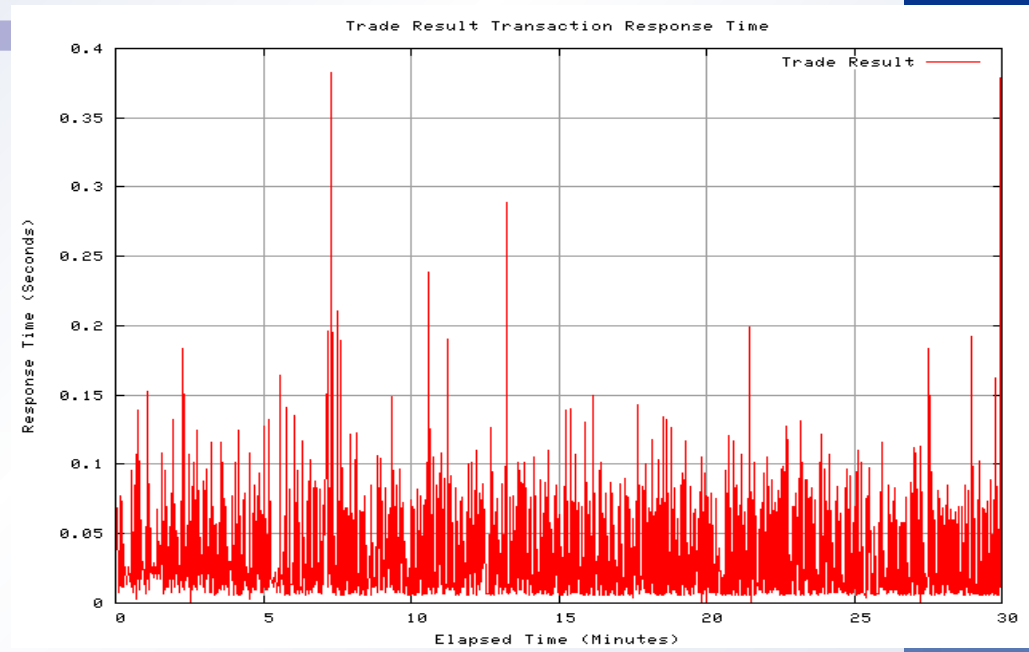

# Experimental Results



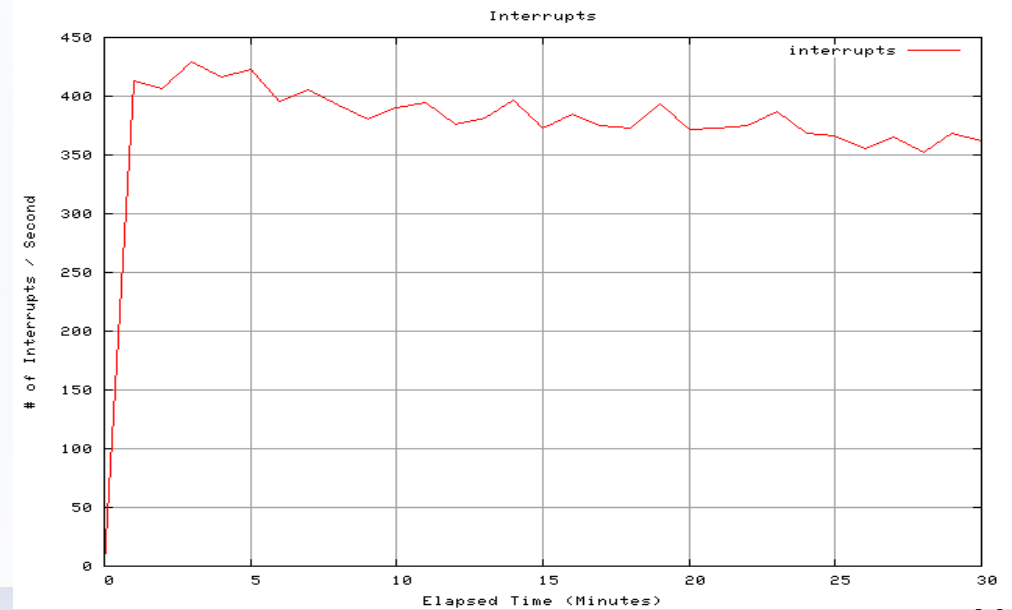
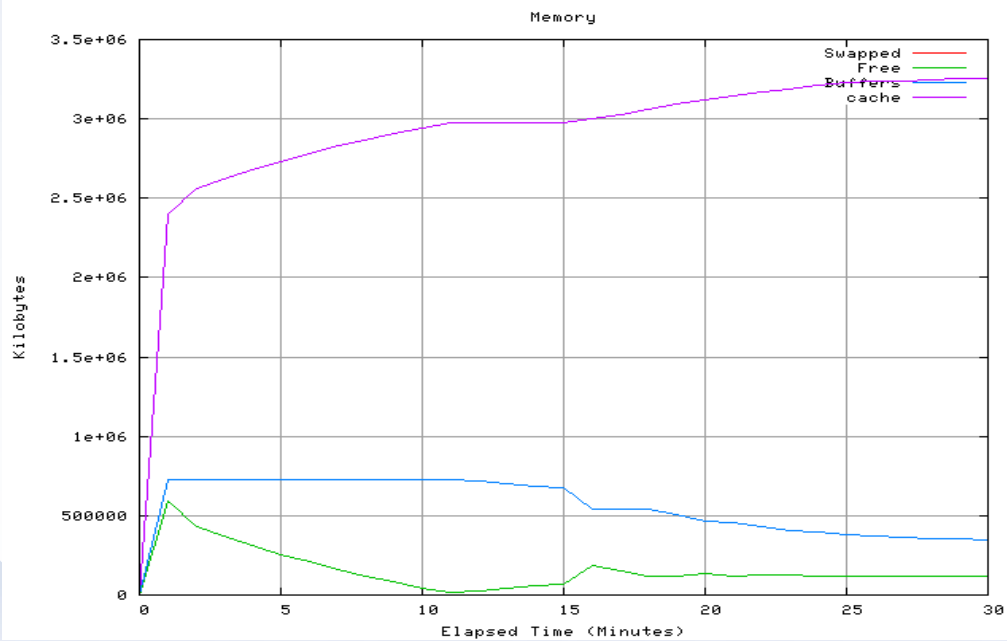

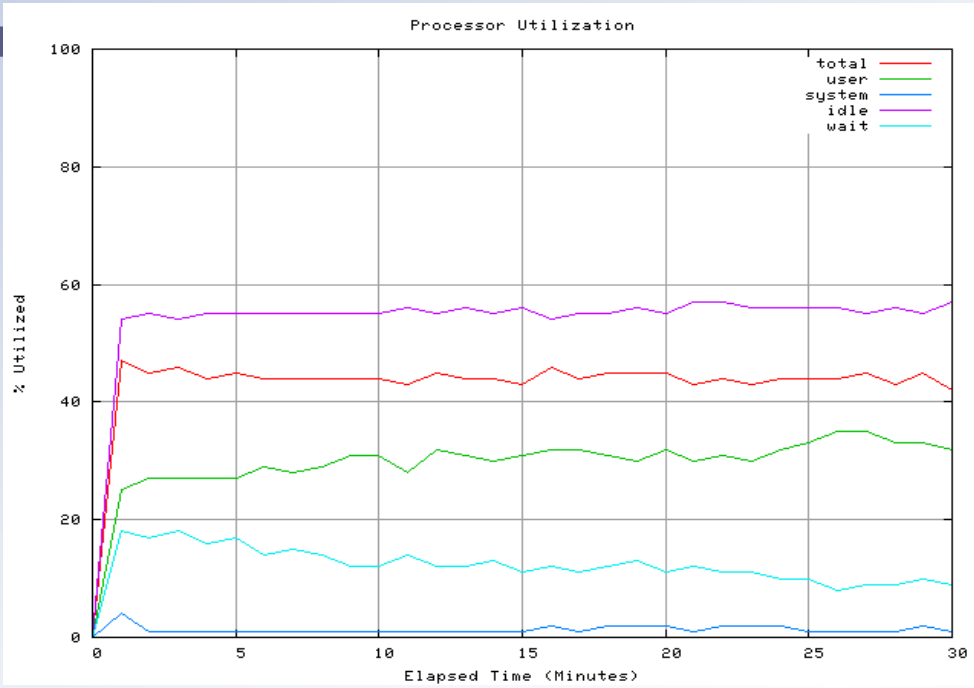
Transaction	%	Response Time (s)		Total	Rollbacks	%
		Average	90th %			
Trade Order	10.21	0.021	0.043	3869	38	0.99
Trade Result	9.59	0.040	0.083	3635	0	0.00
Trade Lookup	8.08	0.346	0.945	3063	0	0.00
Trade Update	1.96	0.313	0.677	743	0	0.00
Trade Status	19.03	0.006	0.010	7212	0	0.00
Customer Position	12.89	0.005	0.009	4885	0	0.00
Broker Volume	4.84	0.003	0.004	1835	0	0.00
Security Detail	14.43	0.014	0.018	5467	0	0.00
Market Feed	0.96	0.055	0.090	363	0	0.00
Market Watch	17.99	0.009	0.017	6817	0	0.00
Data Maintenance	n/a	0.036	0.122	9	0	0.00

2.02 trade-result transactions per second (TRTPS)  
 30.0 minute duration  
 20 total unknown errors  
 1 second(s) ramping up

# Experimental Results – Response Time Plots



# Experimental Results – System Statistics



## Research

- A scientific article portraying DBT-5 was accepted in the 27th Brazilian Computer Society Conference, that will be held in Rio de Janeiro in June 2007
  - <http://www.sbc.de9.ime.eb.br/en/index.php?view=wperformance&from=eventos&lang=en>
- Rilson's Master's dissertation is employing DBT-5: Synthesizing Representative I/O Workloads for TPC-E

## Future Work

- Update the workload to the latest TPC-E specification
  - EGen
  - Functions
- Write the Functions in C (in progress)
- Tune indexes/functions
- Support other databases

# Resources

- DBT-5  
[svn co https://osdl.dbt.svn.sourceforge.net/svnroot/osdl/dbt/trunk/dbt5](https://osdl.dbt.svn.sourceforge.net/svnroot/osdl/dbt/trunk/dbt5) dbt5
- libpqxx: C++ API for PostgreSQL  
<http://pqxx.org/>
- TPC-E Specification  
<http://www.tpc.org/tpce/spec/TPCE-v0.32.2g.pdf> (PDF)  
<http://www.tpc.org/tpce/spec/TPCE-v0.32.2g.doc> (DOC)

## Bibliography

- TPC BENCHMARK™ E Standard Specification      Version 1.0.0
- TPC Site - [www.tpc.org](http://www.tpc.org)
- Do Nascimento, R. O., M. Wong and P. R. M. Maciel. *DBT-5: A Fair Usage Open Source TPC-E Implementation for Performance Evaluation of Computer Systems*. XXVII Brazilian Computer Society Conference, 2007. (to appear)

**Thank you! :)**

markwkm@gmail.com  
rilson.nascimento@gmail.com  
PGCon 2007, Ottawa