

# **LDBC**

*The graph & RDF  
benchmark reference*

***Social Network Benchmark  
Interactive Workload  
Full Disclosure Report  
Moritz Kaufmann  
20/04/2015***

## ***Table of Contents***

Table of Contents	2
General Terms	3
1. System Description	4
1.1 Database System	4
1.2 Database Engine Configuration	4
1.3 Platform Description	4
1.4 Network Infrastructure Information	5
2. Data Generation & Loading	6
2.1 Dataset Information	6
2.1.1 Description	6
2.1.2 Data Generator Parameters	6
2.3. Bulk Loading	6
3. Benchmark Test Driver	7
3.1. Basic test driver configuration details	7
3.2. Configuration Parameters for Driver Warmup	7
3.3. Configuration Parameters for Driver Execution	7
4. Performance Metrics	8
5. Recovery	10
6. Pricing Summary	11
7. Attachment's CheckList	12

This is the Full Disclosure Report for the LDBC SNB interactive audit results for Virtuoso 07.50.3213 (v7fasttrack github) by OpenLink, Scale Factor SF100 on a single server configuration.

## *General Terms*

The Virtuoso SNB Interactive implementation is made in SQL and SQL stored procedures. It uses a relational schema physically represented as column-wise compressed tables and indices. The transitive graph operations are expressed using the Virtuoso transitive derived table SQL extension. The implementation materializes the connection weight between any two people who have replied to each other's posts or comments and the materialization is transactionally kept up to date by the stored procedures implementing the update workload. The schema represents posts and comments as a single table. The posts+comments table is indexed so as to duplicate some non-key columns in the indices, e.g. the creation date of a post together with the foreign key to the author. The test driver uses JDBC to connect to the SUT.

This report contains an audited LDBC benchmark run. The results have been gathered by an independent and impartial auditor, who has validated the implementation of the queries and the overall system's configuration conform to the description of the benchmark and his strict requirements.

## 1. System Description

### 1.1 Database System

<b>Vendor Name</b>	OpenLink
<b>Database Name</b>	Virtuoso
<b>Version No</b>	07.50.3213 (v7fasttrack/featureanalytics github)

**Table 1: DBMS Characteristics**

### 1.2 Database Engine Configuration

<b>Cache Configuration</b>	Default
<b>Transaction Isolation Level/Model</b>	Read committed
<b>Special settings</b>	See attached virtuoso.ini

**Table 2: Database Engine Configuration**

### 1.3 Platform Description

<b>Model</b>	
<b>Processors</b>	2 x INTEL XEON E5-2630
<b>Memory</b>	24x 8GB DDR3-1333
<b>No Disks/Type/Storage Configuration</b>	6 x SEAGATE CONSTELLATION ES 2TB, SATA 3GB/s 2 x CRUCIAL MX100 512GB, SATA 6Gb/s
<b>Network Adapters</b>	N/A
<b>Operating System</b>	CENTOS RELEASE 6.2 (FINAL)
<b>File System</b>	EXT2
<b>No Threads</b>	24
<b>No Cores</b>	12
<b>Memory</b>	192 GB
<b>Total Disks Capacity</b>	13 TB

**Table 3: System Configuration**

<b><i>Storage</i></b>	Intel Corporation Patsburg 6-Port SATA AHCI Controller (rev 06)
<b><i>RAID/HBA controller</i></b>	Not installed

***Table 4: No Disks/Type/Storage/Configuration***

#### ***1.4 Network Infrastructure Information***

<b><i>Model</i></b>	Not applicable
<b><i>Network Switches</i></b>	Not applicable
<b><i>Wiring Information</i></b>	Not applicable

***Table 5: Network Infrastructure Information***

## 2. Data Generation & Loading

### 2.1 Dataset Information

#### 2.1.1 Description

<i>Scale Factor</i>	30
<i>Data Format</i>	CSVMergedForeign
<i>Data Generator Version</i>	v.0.2.0
<i>Time Compression ratio</i>	0.16

*Table 6: Dataset characteristics*

#### 2.1.2 Data Generator Parameters

```
ldbc.snb.datagen.generator.scaleFactor: 100
ldbc.snb.datagen.serializer.compressed:false
ldbc.snb.datagen.serializer.personSerializer:
    ldbc.snb.datagen.serializer.snb.interactive.CSVMergeForeignPersonSerializer
ldbc.snb.datagen.serializer.invariantSerializer:
    ldbc.snb.datagen.serializer.snb.interactive.CSVMergeForeignInvariantSerializer
ldbc.snb.datagen.serializer.personActivitySerializer:
    ldbc.snb.datagen.serializer.snb.interactive.CSVMergeForeignPersonActivitySerializer
ldbc.snb.datagen.generator.numThreads:TODO:XXX
ldbc.snb.datagen.serializer.updateStreams:true
```

#### 2.3. Bulk Loading

<i>Loading Time</i>	43min 16s
---------------------	-----------

*Table 7: Bulk Loading Time*

### ***3. Benchmark Test Driver***

#### ***3.1. Basic test driver configuration details***

<b><i>No of worker threads</i></b>	25
<b><i>No of write streams</i></b>	12

***Table 8: Basic Test Driver Configuration Details***

#### ***3.2. Configuration Parameters for Driver Warmup***

Warmup using the first 150,000 operations

#### ***3.3. Configuration Parameters for Driver Execution***

See user\_parameters.properties, startup with 25 worker threads and TCR = 0.16

## 4. Performance Metrics

<i>Duration</i>	<i>Operations</i>	<i>Throughput</i>	<i>% above threshold</i>
<b>16.2min</b>	<b>1,200,000</b>	<b>1,234.38</b>	<b>2.20%</b>
<i>Query</i>	<i>Count</i>	<i>Mean</i>	<i>% Execution</i>
<i>Complex 1</i>	8,048.0	87.7	2.61%
<i>Complex 2</i>	5,656.0	36.6	0.76%
<i>Complex 3</i>	1,701.0	145.9	0.92%
<i>Complex 4</i>	5,814.0	66.5	1.43%
<i>Complex 5</i>	2,684.0	893.2	8.85%
<i>Complex 6</i>	483.0	1,515.7	2.70%
<i>Complex 7</i>	5,506.0	21.2	0.43%
<i>Complex 8</i>	41,849.0	54.2	8.38%
<i>Complex 9</i>	397.0	194.9	0.29%
<i>Complex 10</i>	5,232.0	120.7	2.33%
<i>Complex 11</i>	9,511.0	49.0	1.72%
<i>Complex 12</i>	4,756.0	91.7	1.61%
<i>Complex 13</i>	11,013.0	21.9	0.89%
<i>Complex 14</i>	4,270.0	114.7	1.81%
<i>Short 1</i>	133,969.0	16.2	8.01%
<i>Short 2</i>	133,970.0	36.8	18.21%
<i>Short 3</i>	133,972.0	33.0	16.32%
<i>Short 4</i>	133,919.0	1.7	0.86%
<i>Short 5</i>	133,917.0	15.5	7.67%
<i>Short 6</i>	133,917.0	15.4	7.64%
<i>Short 7</i>	133,918.0	10.3	5.09%
<i>Update 1</i>	34.0	4.3	0.00%
<i>Update 2</i>	24,332.0	2.3	0.21%
<i>Update 3</i>	44,281.0	2.2	0.37%
<i>Update 4</i>	528.0	2.8	0.01%
<i>Update 5</i>	60,467.0	1.9	0.42%
<i>Update 6</i>	6,623.0	4.3	0.11%
<i>Update 7</i>	16,901.0	3.7	0.23%
<i>Update 8</i>	2,332.0	16.3	0.14%

Table 10: Execution Summary



<i>Query</i>	<i>Min</i>	<i>Max</i>	<i>Mean</i>	<i>50th</i>	<i>90th</i>	<i>95th</i>	<i>99th</i>
<i>Complex 1</i>	10.0	2,204.0	91.6	64.0	178.0	247.0	358.7
<i>Complex 2</i>	2.0	1,872.0	40.8	45.0	56.0	68.0	120.0
<i>Complex 3</i>	4.0	3,172.0	198.9	111.0	225.0	660.1	1,864.6
<i>Complex 4</i>	5.0	3,408.0	139.6	20.0	73.0	1,637.0	1,859.7
<i>Complex 5</i>	3.0	2,855.0	853.8	849.0	996.0	1,046.3	2,276.4
<i>Complex 6</i>	10.0	3,470.0	1,563.0	1,540.0	2,223.7	2,520.9	3,021.9
<i>Complex 7</i>	7.0	1,891.0	27.8	16.0	56.0	62.0	98.0
<i>Complex 8</i>	9.0	2,028.0	58.2	57.0	71.0	83.0	134.0
<i>Complex 9</i>	12.0	1,844.0	198.9	190.0	222.6	236.6	264.0
<i>Complex 10</i>	7.0	1,953.0	126.9	120.0	161.0	176.0	217.0
<i>Complex 11</i>	3.0	1,933.0	53.4	51.0	64.0	80.0	134.0
<i>Complex 12</i>	9.0	1,932.0	91.4	84.0	116.0	130.0	162.5
<i>Complex 13</i>	1.0	1,877.0	25.6	25.0	42.0	56.0	97.7
<i>Complex 14</i>	1.0	1,980.0	119.3	135.0	174.0	189.0	230.2

*Table 11: Complex reads detail*

<i>Query</i>	<i>Min</i>	<i>Max</i>	<i>Mean</i>	<i>50th</i>	<i>90th</i>	<i>95th</i>	<i>99th</i>
<i>Short 1</i>	0.0	1,959.0	16.2	1.0	41.0	43.0	48.0
<i>Short 2</i>	0.0	1,972.0	36.8	42.0	46.0	48.0	53.0
<i>Short 3</i>	0.0	1,992.0	33.0	41.0	46.0	49.0	56.0
<i>Short 4</i>	0.0	1,945.0	1.7	0.0	3.0	6.0	40.0
<i>Short 5</i>	0.0	1,969.0	15.5	1.0	41.0	43.0	48.0
<i>Short 6</i>	0.0	1,979.0	15.4	1.0	41.0	43.0	48.0
<i>Short 7</i>	0.0	1,971.0	10.3	1.0	41.0	42.0	47.0

*Table 12: Short Reads detail*

<i>Query</i>	<i>Min</i>	<i>Max</i>	<i>Mean</i>	<i>50th</i>	<i>90th</i>	<i>95th</i>	<i>99th</i>
<i>Update 1</i>	1.0	13.0	4.3	3.0	8.4	10.1	12.7
<i>Update 2</i>	0.0	1,891.0	2.3	1.0	5.0	8.0	17.0
<i>Update 3</i>	0.0	1,929.0	2.2	1.0	5.0	7.0	16.0
<i>Update 4</i>	0.0	22.0	2.8	2.0	6.0	7.6	12.0
<i>Update 5</i>	0.0	1,921.0	1.9	1.0	5.0	7.0	15.0
<i>Update 6</i>	0.0	1,927.0	4.3	2.0	8.0	11.0	23.0
<i>Update 7</i>	0.0	1,823.0	3.7	2.0	8.0	11.0	22.0
<i>Update 8</i>	1.0	91.0	16.3	11.0	39.0	48.0	69.7

*Table 13: Updates detail*

## ***5. Recovery***

<b><i>Time to Recover</i></b>	151s
-------------------------------	------

***Table 14: Recovery***

## 6. Pricing Summary

	<i>Item</i>	<i>Price</i>
	BASIC VIRTUOSO OPEN SOURCE SUPPORT 3Y	15,000.00 €
	CUSTOM BUILT SERVER	5,212.59 €
<b>TOTAL</b>		20,212.59 €

**Table 15: Pricing Information**

## ***7. Attachment's CheckList***

DATAGEN PARAMS.INI	Attached
DRIVER'S LDBC-CONFIGURATION.PROPERTIES	Attached
DRIVER'S LDBC-RESULTS.JSON	Attached
VENDOR SPECIFIC CONFIGURATION FILES	virtuoso.ini

***Table 9: Checklist***