

LDBC

*The graph & RDF
benchmark reference*

***Social Network Benchmark
Interactive Workload
Full Disclosure Report
Moritz Kaufmann
24/04/2015***

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 SF30 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

Vendor Name	OpenLink
Database Name	Virtuoso
Version No	07.50.3213 (v7fasttrack github)

Table 1: DBMS Characteristics

1.2 Database Engine Configuration

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

Table 2: Database Engine Configuration

1.3 Platform Description

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

Table 3: System Configuration

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

Table 4: No Disks/Type/Storage/Configuration

1.4 Network Infrastructure Information

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

Table 5: Network Infrastructure Information

2. Data Generation & Loading

2.1 Dataset Information

2.1.1 Description

Scale Factor	30
Data Format	CSVMergedForeign
Data Generator Version	v.0.2.0
Time Compression ratio	0.041

Table 6: Dataset characteristics

2.1.2 Data Generator Parameters

```
ldbc.snb.datagen.generator.scaleFactor: 30
ldbc.snb.datagen.serializer.compressed:false
ldbc.snb.datagen.serializer.personSerializer:
    ldc.snb.datagen.serializer.snb.interactive.CSVMergeForeignPersonSerializer
ldbc.snb.datagen.serializer.invariantSerializer:
    ldc.snb.datagen.serializer.snb.interactive.CSVMergeForeignInvariantSerializer
ldbc.snb.datagen.serializer.personActivitySerializer:
    ldc.snb.datagen.serializer.snb.interactive.CSVMergeForeignPersonActivitySerializer
ldbc.snb.datagen.generator.numThreads:4
ldbc.snb.datagen.serializer.updateStreams:true
```

2.3. Bulk Loading

Loading Time	12.3 min
---------------------	----------

Table 7: Bulk Loading Time

3. Benchmark Test Driver

3.1. Basic test driver configuration details

<i>No of worker threads</i>	25
<i>No of write streams</i>	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 = 0.041

4. Performance Metrics

<i>Duration</i>	<i>Operations</i>	<i>Throughput</i>	<i>% above threshold</i>
4.97min	384,000	1,287.17	4.98%
<i>Query</i>	<i>Count</i>	<i>Mean</i>	<i>% Execution</i>
<i>Complex 1</i>	3,009.0	57.3	2.65%
<i>Complex 2</i>	2,115.0	36.9	1.20%
<i>Complex 3</i>	738.0	75.9	0.86%
<i>Complex 4</i>	2,172.0	94.0	3.14%
<i>Complex 5</i>	1,086.0	388.6	6.49%
<i>Complex 6</i>	248.0	12.2	0.05%
<i>Complex 7</i>	1,629.0	17.5	0.44%
<i>Complex 8</i>	8,692.0	43.6	5.83%
<i>Complex 9</i>	204.0	116.6	0.37%
<i>Complex 10</i>	2,114.0	85.9	2.79%
<i>Complex 11</i>	3,912.0	44.4	2.67%
<i>Complex 12</i>	1,778.0	81.1	2.22%
<i>Complex 13</i>	4,118.0	13.0	0.82%
<i>Complex 14</i>	1,596.0	63.5	1.56%
<i>Short 1</i>	41,848.0	12.5	8.02%
<i>Short 2</i>	41,842.0	32.0	20.55%
<i>Short 3</i>	41,837.0	27.3	17.57%
<i>Short 4</i>	41,890.0	1.0	0.62%
<i>Short 5</i>	41,887.0	12.3	7.94%
<i>Short 6</i>	41,887.0	12.6	8.11%
<i>Short 7</i>	41,887.0	7.8	5.01%
<i>Update 1</i>	17.0	3.1	0.00%
<i>Update 2</i>	8,928.0	1.0	0.13%
<i>Update 3</i>	15,710.0	1.0	0.23%
<i>Update 4</i>	253.0	1.9	0.01%
<i>Update 5</i>	22,324.0	0.8	0.28%
<i>Update 6</i>	3,126.0	1.9	0.09%
<i>Update 7</i>	6,282.0	2.1	0.21%
<i>Update 8</i>	871.0	11.8	0.16%

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>	9.0	8,615.0	57.3	52.0	87.2	130.0	187.0
<i>Complex 2</i>	2.0	8,602.0	36.9	42.0	48.0	51.0	56.0
<i>Complex 3</i>	5.0	4,377.0	75.9	62.0	106.0	150.5	347.2
<i>Complex 4</i>	5.0	8,607.0	94.0	21.0	277.0	296.0	354.3
<i>Complex 5</i>	2.0	8,881.0	388.6	383.5	437.0	452.8	487.0
<i>Complex 6</i>	7.0	53.0	12.2	10.0	14.0	17.7	50.0
<i>Complex 7</i>	4.0	60.0	17.5	9.0	49.0	50.0	54.0
<i>Complex 8</i>	7.0	8,642.0	43.6	49.0	53.0	54.0	58.0
<i>Complex 9</i>	45.0	192.0	116.6	114.0	137.1	150.7	166.0
<i>Complex 10</i>	7.0	8,596.0	85.9	80.0	106.0	116.0	131.0
<i>Complex 11</i>	3.0	8,622.0	44.4	48.0	51.0	53.0	55.0
<i>Complex 12</i>	9.0	8,637.0	81.1	75.0	96.0	103.0	124.0
<i>Complex 13</i>	1.0	8,622.0	13.0	13.0	20.0	22.0	29.0
<i>Complex 14</i>	2.0	8,627.0	63.5	65.0	94.0	100.0	113.0

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	8,623.0	12.5	0.0	40.0	41.0	42.0
<i>Short 2</i>	0.0	8,617.0	32.0	41.0	44.0	45.0	47.0
<i>Short 3</i>	0.0	8,637.0	27.3	40.0	42.0	43.0	46.0
<i>Short 4</i>	0.0	46.0	1.0	0.0	1.0	2.0	40.0
<i>Short 5</i>	0.0	4,385.0	12.3	0.0	40.0	41.0	42.0
<i>Short 6</i>	0.0	8,607.0	12.6	1.0	41.0	41.0	43.0
<i>Short 7</i>	0.0	8,609.0	7.8	0.0	40.0	41.0	42.0

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	9.0	3.1	3.0	4.8	6.6	8.5
<i>Update 2</i>	0.0	18.0	1.0	1.0	2.0	3.0	6.0
<i>Update 3</i>	0.0	18.0	1.0	1.0	2.0	3.0	7.0
<i>Update 4</i>	0.0	20.0	1.9	1.0	3.0	5.0	10.0
<i>Update 5</i>	0.0	23.0	0.8	1.0	2.0	3.0	7.0
<i>Update 6</i>	0.0	39.0	1.9	1.0	4.0	6.0	12.0
<i>Update 7</i>	0.0	47.0	2.1	1.0	4.0	6.0	11.0
<i>Update 8</i>	1.0	58.0	11.8	8.0	27.0	34.0	45.0

Table 13: Updates detail

5. Recovery

<i>Time to Recover</i>	55s
------------------------	-----

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 €
TOTAL		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