

# **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	5
1.1 Database System	5
1.2 Database Engine Configuration	5
1.3 Platform Description	5
1.4 Network Infrastructure Information	6
2. Data Generation & Loading	7
2.1 Dataset Information	7
2.1.1 Description	7
2.1.2 Data Generator Parameters	7
2.3. Bulk Loading	7
3. Benchmark Test Driver	8
3.1. Basic test driver configuration details	8
3.2. Configuration Parameters for Driver Warmup	8
3.3. Configuration Parameters for Driver Execution	8
4. Performance Metrics	9
5. Recovery	11
6. Pricing Summary	12
7. Attachment's CheckList	13

This is the Full Disclosure Report for the LDBC SNB interactive audit results for Sparksee 5.1.1 by Sparsity Technologies, Scale Factor SF100 on a single server configuration.

## *General Terms*

Test sponsor of this benchmark is Sparsity Technologies, Spain. Sparsity technologies is a leading provider of graph database and applications that use the graph to add value to our customers. Our high end technologies, based on top research, combine the experience of University researchers and Industrial professionals. We have a vast Industrial experience to push your development to a more competitive customer solution. Sparsity was created based on the technologies developed at DAMA-UPC, and have successfully provided solutions to clients such as Barclays, Acceso, CA Technologies, MPG and contributed to leading European Projects such as LDBC and CoherentPaaS.

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>	Sparsity Technologies
<b>Database Name</b>	Sparksee
<b>Version No</b>	5.1.1

**Table 1: DBMS Characteristics**

### 1.2 Database Engine Configuration

<b>Cache Configuration</b>	Default
<b>Transaction Isolation Level/Model</b>	Serializable
<b>Special settings</b>	sparksee.io.rollback=false sparksee.io.recovery=true sparksee.io.recovery.checkpointTime=0

**Table 2: Database Engine Configuration**

### 1.3 Platform Description

<b>Model</b>	
<b>Processors</b>	2X INTEL XEON E5-2630 v3
<b>Memory</b>	8X 32GB LRDIMM, 2133 MT/s
<b>No Disks/Type/Storage Configuration</b>	2X80GB SSD SATA MLC 6GB/s IN HOT-PLUG DRIVE 1X2TB SATA 6GB/s 7,2k RPM 3,5" HDD
<b>Network Adapters</b>	BROADCOM 5720
<b>Operating System</b>	DEBIAN GNU/LINUX 7.8 (WHEEZY)
<b>File System</b>	EXT4
<b>No Threads</b>	32
<b>No Cores</b>	16
<b>Memory</b>	256 GB
<b>Total Disks Capacity</b>	2.16 TB

**Table 3: System Configuration**

<b><i>Storage</i></b>	Intel Corporation Wellsburg 6-Port SATA Controller
<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>	100
<i>Data Format</i>	CSV
<i>Data Generator Version</i>	v.0.2.0
<i>Time Compression ratio</i>	2.4

*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:
    ldc.snb.datagen.serializer.snb.interactive.CSVPersonSerializer
ldbc.snb.datagen.serializer.invariantSerializer:
    ldc.snb.datagen.serializer.snb.interactive.CSVInvariantSerializer
ldbc.snb.datagen.serializer.personActivitySerializer:
    ldc.snb.datagen.serializer.snb.interactive.CSVPersonActivitySerializer
ldbc.snb.datagen.generator.numThreads:4
ldbc.snb.datagen.serializer.updateStreams:true
```

### 2.3. Bulk Loading

<i>Loading Time</i>	698 min
---------------------	---------

*Table 7: Bulk Loading Time*

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

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

<b><i>No of worker threads</i></b>	16
<b><i>No of write streams</i></b>	8

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

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

Warmup using the first 50,000 operations

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

See user\_parameters.properties, startup with 16 worker threads and TCR = 2.4



## 4. Performance Metrics

<i>Duration</i>	<i>Operations</i>	<i>Throughput</i>	<i>% above threshold</i>
<b>224min</b>	<b>1,100,000</b>	<b>81.7</b>	<b>4.17%</b>
<i>Query</i>	<i>Count</i>	<i>Mean</i>	<i>% Execution</i>
<i>Complex 1</i>	7,437.0	412.9	2.39%
<i>Complex 2</i>	5,226.0	439.7	1.79%
<i>Complex 3</i>	1,572.0	1,002.2	1.23%
<i>Complex 4</i>	5,371.0	414.6	1.73%
<i>Complex 5</i>	2,479.0	512.8	0.99%
<i>Complex 6</i>	446.0	2,261.9	0.79%
<i>Complex 7</i>	4,028.0	403.8	1.27%
<i>Complex 8</i>	38,673.0	435.7	13.12%
<i>Complex 9</i>	367.0	4,638.1	1.33%
<i>Complex 10</i>	4,835.0	571.6	2.15%
<i>Complex 11</i>	8,790.0	437.8	3.00%
<i>Complex 12</i>	4,395.0	486.7	1.67%
<i>Complex 13</i>	10,177.0	1,305.0	10.35%
<i>Complex 14</i>	3,947.0	2,392.7	7.36%
<i>Short 1</i>	122,720.0	53.1	5.07%
<i>Short 2</i>	122,722.0	34.9	3.34%
<i>Short 3</i>	122,722.0	42.5	4.07%
<i>Short 4</i>	122,613.0	58.3	5.56%
<i>Short 5</i>	122,616.0	26.0	2.48%
<i>Short 6</i>	122,616.0	26.1	2.50%
<i>Short 7</i>	122,616.0	23.7	2.26%
<i>Update 1</i>	32.0	520.8	0.01%
<i>Update 2</i>	22,430.0	230.4	4.03%
<i>Update 3</i>	40,912.0	224.4	7.15%
<i>Update 4</i>	471.0	233.8	0.09%
<i>Update 5</i>	56,205.0	228.3	10.00%
<i>Update 6</i>	5,790.0	230.9	1.04%
<i>Update 7</i>	15,615.0	233.9	2.84%
<i>Update 8</i>	2,177.0	240.0	0.41%

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>	2.0	5,980.0	412.9	75.0	991.0	1,839.2	2,939.8
<i>Complex 2</i>	1.0	7,001.0	439.7	109.0	1,024.0	1,776.0	2,947.0
<i>Complex 3</i>	1.0	5,864.0	1,002.2	782.5	1,767.9	2,593.4	3,737.1
<i>Complex 4</i>	1.0	6,047.0	414.6	70.0	994.0	1,841.0	2,911.2
<i>Complex 5</i>	1.0	5,598.0	512.8	195.0	1,299.2	1,996.0	3,061.7
<i>Complex 6</i>	1.0	5,441.0	2,261.9	2,354.0	3,251.5	3,731.8	5,085.1
<i>Complex 7</i>	1.0	6,005.0	403.8	66.5	978.3	1,793.0	2,919.3
<i>Complex 8</i>	13.0	6,850.0	435.7	102.0	1,019.8	1,836.8	2,960.0
<i>Complex 9</i>	758.0	7,932.0	4,638.1	4,590.0	5,660.0	5,985.7	6,884.0
<i>Complex 10</i>	1.0	6,427.0	571.6	265.0	1,189.0	1,952.3	3,078.0
<i>Complex 11</i>	1.0	6,740.0	437.8	112.0	1,021.0	1,831.6	2,975.4
<i>Complex 12</i>	2.0	5,746.0	486.7	165.0	1,081.6	1,899.3	3,006.5
<i>Complex 13</i>	5.0	7,564.0	1,305.0	1,285.0	2,119.0	2,875.0	3,896.0
<i>Complex 14</i>	12.0	8,679.0	2,392.7	2,222.0	4,021.0	4,685.6	6,094.1

*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>	1.0	6,996.0	53.1	14.0	65.0	113.0	1,022.0
<i>Short 2</i>	1.0	7,052.0	34.9	21.0	52.0	68.0	134.0
<i>Short 3</i>	1.0	5,256.0	42.5	24.0	50.0	63.0	1,007.0
<i>Short 4</i>	1.0	7,040.0	58.3	14.0	67.0	123.0	1,038.0
<i>Short 5</i>	1.0	5,930.0	26.0	13.0	38.0	55.0	125.0
<i>Short 6</i>	1.0	5,239.0	26.1	12.0	37.0	50.0	108.0
<i>Short 7</i>	1.0	5,238.0	23.7	12.0	33.0	44.0	97.9

*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>	15.0	4,767.0	520.8	254.5	935.4	1,479.9	3,951.4
<i>Update 2</i>	6.0	7,628.0	230.4	26.0	872.1	1,269.1	3,105.7
<i>Update 3</i>	5.0	9,439.0	224.4	25.0	858.0	1,194.5	3,079.9
<i>Update 4</i>	8.0	4,070.0	233.8	29.0	851.0	1,306.5	3,069.1
<i>Update 5</i>	6.0	9,461.0	228.3	25.0	872.0	1,251.0	3,129.8
<i>Update 6</i>	9.0	7,775.0	230.9	33.0	858.0	1,229.0	3,158.2
<i>Update 7</i>	9.0	9,747.0	233.9	33.0	863.0	1,244.6	3,081.2
<i>Update 8</i>	6.0	5,259.0	240.0	26.0	903.2	1,381.0	3,060.9

*Table 13: Updates detail*

## ***5. Recovery***

<b><i>Time to Recover</i></b>	4.224s
-------------------------------	--------

***Table 14: Recovery***

## 6. Pricing Summary

	<i>Item</i>	<i>Price</i>
	SPARKSEE 5.1.1 SME LARGE 3Y	21,600.00 €
	DELL R730	8,827.94 €
<b>TOTAL</b>		30,427.94 €

**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	sparksee.cfg

***Table 9: Checklist***