



**Full Disclosure Report
of the LDBC Social Network Benchmark**

Audit of the LDBC Social Network Benchmark's
Business Intelligence Workload over TigerGraph

April 6, 2023

GENERAL TERMS

Executive Summary

This report documents an audited execution of the LDBC SNB BI (Social Network Benchmark Business Intelligence) workload for TigerGraph.

TigerGraph is a massively parallel processing (MPP) graph database management system, designed for handling hybrid transaction/analytical processing (HTAP) query workloads. It is a distributed platform using a native graph storage format with an edge cut partitioning strategy. Within this, each segment (partition) of the graph holds a similar amount of vertices and processes requests in parallel. TigerGraph offers GSQL, a Turing-complete query language which provides both declarative features (e.g. graph patterns) as well as imperative ones (e.g. for expressing iterative graph algorithms with loops and accumulator primitives). TigerGraph previously passed the LDBC SNB BI benchmark with a single-machine on-premise setup using the SF1 000 dataset.

Declaration of Audit Success

This report contains details of a successful execution of the LDBC SNB BI benchmark. The results have been gathered by an independent auditor who has validated the implementation of the queries and verified the system's configuration conforms to the description of the benchmark and its strict requirements.

Sponsorship and Funding Disclaimer

TigerGraph, as an LDBC member, are the Test Sponsor of this audit. The audit and its associated execution costs (compute, storage) were funded by AMD. This arrangement was deemed acceptable by both parties, the LDBC Steering Committee and the Auditor.



DocuSigned by:

David Puroja

C7FE4F71E2554AE.....

4/6/2023

Mr. David Püroja
(Auditor)

Date

Pometry
(Audit Company)

DocuSigned by:

Gabor Szarnyas

FA33C2G8CB944A5.....

4/6/2023

Dr. Gábor Szárnyas
(Leader of LDBC SNB Task Force)

Date

DocuSigned by:

Songting Chen

BD86B2D9F5AA4G8.....

4/6/2023

Dr. Songting Chen
(Test Sponsor Representative)

Date



TABLE OF CONTENTS

1	BENCHMARK DESCRIPTION	5
2	SYSTEM DESCRIPTION AND PRICING SUMMARY	6
2.1	Details of machines driving and running the workload	6
2.1.1	Machine overview	6
2.1.2	CPU details	6
2.1.3	Memory details	6
2.1.4	Disk and storage details	6
2.1.5	Network details	7
2.1.6	Machine pricing	7
2.1.7	System version and availability	7
3	DATASET GENERATION	8
3.1	General information	8
3.2	Datagen configurations	8
3.2.1	SF10	8
3.2.2	SF1000	8
3.3	Data loading and data schema	8
4	IMPLEMENTATION DETAILS	10
4.1	Execution mode	10
4.2	Use of auxiliary data structures	10
4.3	Benchmark execution	11
5	PERFORMANCE RESULTS	12
5.1	TigerGraph performance results	12
6	VALIDATION OF THE RESULTS	14
7	SUPPLEMENTARY MATERIALS	15
A	CPU AND MEMORY DETAILS	16
B	IO PERFORMANCE	20
C	DATASET GENERATION INSTRUCTIONS	21
D	DATA SCHEMA	23



Benchmark Description

1 BENCHMARK DESCRIPTION

The audit was conducted in compliance with the Social Network Benchmark Business Intelligence workload's specification.

Table 1.1: Benchmark Overview

Artifact	Version	URL
Specification	2.2.0	https://arxiv.org/pdf/2001.02299v7.pdf
Data generator	0.5.0	https://github.com/ldbc/ldbc_snb_datagen_spark/releases/tag/v0.5.0
Driver and implementations	1.0.3	https://github.com/ldbc/ldbc_snb_bi/releases/tag/v1.0.3



System Description and Pricing Summary

2 SYSTEM DESCRIPTION AND PRICING SUMMARY

2.1 Details of machines driving and running the workload

2.1.1 Machine overview

The audit was conducted on 4 virtual machine instances in the Amazon Web Service (AWS) cloud, placed in an AWS Placement group with the cluster strategy to reduce network latency¹. The details below were obtained from the AWS console.

Table 2.1: Machine Type and Location

Cloud provider	Amazon Web Services
Machine region	N. Virginia (us-east-1)
Common name of the item	r6a.8xlarge
Operating system	Amazon Linux 2 (5.10.167-147.601.amzn2.x86_64)

2.1.2 CPU details

The details below were obtained using the command `cat /proc/cpuinfo` (Listing A.1) and `dmidecode -t processor` (Listing A.3) issued from the machine instance.

Table 2.2: CPU details summary

Type	AMD® AMD EPYC® 7R13 CPU
Total number	1
Cores per CPU	16
Threads per CPU core	2
CPU clock frequency	3.725 MHz
Total cache size per CPU	L1i cache: 32KiB L1d cache: 32KiB L2 cache: 512KiB L3 cache: 33MiB

2.1.3 Memory details

The total size of the memory installed is 256GiB and the type of memory is DDR4 with frequency 3200MHz. This information was obtained using the `sudo lshw -c memory` command (Listing A.5) issued from the virtual machine instance.

2.1.4 Disk and storage details

The virtual machine instance used a Amazon EBS General Purpose (GP) 3 storage, formatted in `xf`s, configured with properties shown in Table 2.3.

Disk controller or motherboard type was not obtainable from the virtual machine instance. Information on AWS General Storage can be found on the website <https://aws.amazon.com/ebs/general-purpose/> (accessed: March 16, 2023).

The 4KB QD1 write performance on the data disk was measured with the `fi`o command and the output (Listing B.1) showed an average of 1 551 IOPS.

¹<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html#placement-groups-cluster>



Table 2.3: Amazon EBS GP3 Configuration

Type	GP3
Size	1000GB
IOPS	16000
Throughput	1000

2.1.5 Network details

The benchmark was run using 4 r6a.8xlarge instances, both deployed in the same availability zone. All instances were assigned security groups to open the ports required for communication between the nodes². The r6a.8xlarge instances use the Elastic Network Adapter provided by Amazon. This information was obtained using the `lshw -class network` command (Listing B.2).

2.1.6 Machine pricing

The system pricing summary in US dollars (USD) is included in the table below. The pricing of the AWS machine instance is the price for a 3-year reserved dedicated instance machine without upfront payment using the EC2 Instance savings plan³.

Table 2.4: Pricing summary

Item	Price
r6a.8xlarge reserved instance machine in AWS (standard 3-year term)	95 166 USD
Software license (3 years)	1 097 280 USD
Maintenance fee (3 years)	109 728 USD
Total cost	1 302 174 USD

2.1.7 System version and availability

Table 2.5: System versions

System	Version	License
TigerGraph	3.7.0	Enterprise Licence provided by TigerGraph

²<https://docs.tigergraph.com/tigergraph-server/current/reference/ports>

³<https://aws.amazon.com/savingsplans/compute-pricing/>



3 DATASET GENERATION

3.1 General information

The data generation settings of the LDBC Datagen are described below.

Table 3.1: Datagen settings summary

Data format for TigerGraph	composite-projected-fk layout, compressed CSV files
Scale factors for TigerGraph	10 (validation), 1000 (benchmark)
Data format for Umbra	composite-merged-fk layout, compressed CSV files
Scale factors for Umbra	10 (validation)

3.2 Datagen configurations

The datasets and query substitution parameters used for the benchmark and the cross-validation runs were retrieved from the following URLs and copied to an AWS S3 Bucket, with the `tar.zst` files uncompressed. The URLs are served by LDBC's official data repository, available as a public bucket in the Cloudflare R2 object storage.¹

3.2.1 SF10

- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/parameters-2022-10-01.zip>
- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/bi-sf10-composite-merged-fk.tar.zst>
- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/bi-sf10-composite-projected-fk.tar.zst>

3.2.2 SF1000

- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/parameters-2022-10-01.zip>
- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/bi-sf1000-composite-projected-fk.tar.zst.000>
- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/bi-sf1000-composite-projected-fk.tar.zst.001>
- <https://pub-383410a98aef4cb686f0c7601eddd25f.r2.dev/bi-pre-audit/bi-sf1000-composite-projected-fk.tar.zst.002>

To re-generate these datasets from scratch, use the instructions provided in Appendix C.

3.3 Data loading and data schema

For SF1000, the data is partitioned evenly across each instance using a partitioning script. The partitioning is executed by calling the following command:

```

1 cd /data/ldbc_snb_bi/tigergraph
2 # change the following lines in k8s/vars.sh
3 # export NUM_NODES=4 # number of pods or nodes
4 # export SF=1000
5 # export TG_DATA_DIR=/home/tigergraph/sf${SF}
6 # export TG_PARAMETER=/home/tigergraph/parameters-sf${SF}
7
8 python3 download_all.py 1000 ./ip_list -t 20 --access_key_id $ACCESS_KEY_ID --secret_access_key
   $SECRET_ACCESS_KEY --region $BUCKET_REGION --bucket_name $BUCKET_NAME --provider $CLOUD_PROVIDER
9 tail -f ~/log.download

```

Listing 3.1: Script to execute the benchmark on TigerGraph for SF1000

¹<https://www.cloudflare.com/products/r2/>

The data preprocessing and loading times are reported below. Values were measured using the GNU Time tool (`/usr/bin/time`) with the `-v` flag, reading the *Elapsed (wall clock) time* from the output. The column **Data preprocessing time** shows how much time it took to preprocess the CSV files. For this benchmark execution, the preprocessing only consisted of decompressing the `.csv.gz` files, for which the timing of the first node is reported. The column **Data loading time** shows how long it took to create a graph from the input CSV files and perform the initial indexing, compilation of the queries and precomputation. The column **Total time** contains the sum of the data preprocessing and loading times.

The TigerGraph data schema is shown in Listing D.1.

Table 3.2: Data preprocessing and loading times for TigerGraph on scale factor 1000

Scale factor	Data preprocessing time (s)	Data loading time (s)	Total time (s)
1000	473	3 433	3 906

The decomposition of the **Data loading times** are shown in Table 3.3.

Table 3.3: Decomposition of data loading time

	Time (s)
Schema setup	43
Load data	2 782
Query install	81
Precompute	235
Rebuild	292
Total	3 433

4 IMPLEMENTATION DETAILS

4.1 Execution mode

Section 7.5.2.2 of the SNB specification defines two execution modes for the *throughput batches*. In *disjoint read-write mode*, the updates for each day of the benchmark's simulation are applied in bulk, separately from the read queries (i.e. there are no overlapping read and write operations). In *concurrent read-write mode*, the updates are applied concurrently with the reads. Systems opting for concurrent read-write mode are subject to the LDBC ACID test¹.

In the current audited run, TigerGraph was executed using the *disjoint read-write mode*. Therefore, no ACID tests were conducted.

4.2 Use of auxiliary data structures

The TigerGraph implementation precomputes the following auxiliary data structures. These are executed in each batch after the writes have been applied.

- **Root Post:** For each Message node (Comments and Posts), an edge to the corresponding Message thread's root Post is inserted. These derived edges are maintained incrementally, i.e. root Post edges are inserted for newly inserted Messages and removed for deleted Messages.
- **Q4:** For each Forum, the maximum number of members (for number of members per country) is pre-computed.
- **Q6:** For each Message, the popularityScore defined in the query is precomputed.
- **Q14:** The weight attributed on the knows edges are precomputed based on the number of interactions between the two Person nodes.
- **Q19:** The weight attributes on the knows edges are precomputed based on the number of interactions between the two Person nodes.
- **Q20:** The weight attributes on the knows edges are precomputed based on the classYear attributes on the studyAt edges that point to the same University from the endpoint Person nodes.

The precomputations for Q14 and Q19 are performed together using different scoring methods for establishing the edge weights. We display the runtime of this operation as “precomputation for Q14 and Q19” in Table 5.3.

¹https://github.com/ldbc/ldbc_acid

4.3 Benchmark execution

The benchmark is executed using the following commands.

Note: despite what the script's name suggests, this benchmark was executed on a AWS EC2 instance without Kubernetes running, with the TigerGraph instance running on a single server machine and not using any containerization technology.

```
1 cd /data/ldbc_snb_bi/tigergraph
2 # change the following lines in k8s/vars.sh
3 # export NUM_NODES=4 # number of pods or nodes
4 # export SF=1000
5 # export TG_DATA_DIR=/home/tigergraph/sf${SF}
6 # export TG_PARAMETER=/home/tigergraph/parameters-sf${SF}
7
8 nohup ./k8s/benchmark.sh > log.benchmark 2>&1 < /dev/null &
9 tail -f log.benchmark
```

Listing 4.1: Script to execute the benchmark on TigerGraph for SF1000

Performance Results

5 PERFORMANCE RESULTS

5.1 TigerGraph performance results

Table 5.1: Summary of results for TigerGraph on scale factor 1000

Benchmark duration	Power@SF	Power@SF/\$	Throughput@SF	Throughput@SF/\$
257.47 minutes	23 951.74	18.39	10 605.12	8.14

Table 5.2: Detailed **power test results** for TigerGraph on scale factor 1000. Execution times are reported in seconds.

Query	Count	Min.	Max.	Mean	P ₅₀	P ₉₀	P ₉₅	P ₉₉
1	30	5.417	6.430	5.786	5.699	6.101	6.158	6.430
2a	30	2.206	17.476	6.762	4.663	14.970	15.123	17.476
2b	30	1.143	2.945	1.854	1.682	2.725	2.776	2.945
3	30	2.762	9.300	4.355	3.680	6.846	7.457	9.300
4	30	1.527	1.978	1.738	1.765	1.825	1.872	1.978
5	30	1.288	1.336	1.312	1.311	1.329	1.335	1.336
6	30	1.201	1.503	1.296	1.284	1.327	1.495	1.503
7	30	2.724	3.060	2.855	2.850	2.981	3.013	3.060
8a	30	1.818	2.656	2.115	2.055	2.271	2.646	2.656
8b	30	1.041	1.588	1.220	1.248	1.309	1.459	1.588
9	30	8.025	9.447	8.986	9.014	9.396	9.436	9.447
10a	30	5.829	11.862	8.170	8.028	9.309	10.148	11.862
10b	30	3.226	4.697	4.087	4.105	4.575	4.685	4.697
11	30	4.383	5.823	5.241	5.166	5.613	5.656	5.823
12	30	5.055	9.695	7.420	8.505	9.579	9.683	9.695
13	30	28.938	33.204	31.187	31.319	32.184	32.404	33.204
14a	30	9.551	10.375	9.998	10.013	10.245	10.310	10.375
14b	30	4.013	4.863	4.464	4.442	4.718	4.843	4.863
15a	30	17.664	20.828	19.248	19.246	20.033	20.773	20.828
15b	30	20.972	72.847	48.963	48.317	68.036	69.391	72.847
16a	30	4.277	7.734	5.844	5.576	7.160	7.391	7.734
16b	30	1.735	2.980	2.147	1.935	2.766	2.848	2.980
17	30	4.162	5.087	4.704	4.756	4.951	5.052	5.087
18	30	16.715	21.693	18.137	17.860	19.289	20.438	21.693
19a	30	3.423	4.975	4.349	4.482	4.706	4.745	4.975
19b	30	3.394	4.969	4.632	4.714	4.864	4.912	4.969
20a	30	0.442	4.266	1.780	0.659	4.051	4.220	4.266
20b	30	1.060	1.922	1.436	1.294	1.903	1.911	1.922

Table 5.3: Operations in the **power test** for TigerGraph on scale factor 1000. Execution times are reported in seconds. Root Post precomputations are performed for each Comment insertion and deletion operation, therefore, they are reported as part of the writes.

Operation	Time
reads	6 602.802
writes	1 022.054
q4precomputation	86.371
q6precomputation	133.574
q19precomputation	492.757
q20precomputation	28.879

6 VALIDATION OF THE RESULTS

The results were cross-validated against the Umbra reference implementation¹ on scale factor 10, using Umbra version 664890d7f. Umbra is an in-memory relational database management system developed at the Technische Universität München. The queries of the BI workload are implemented using PostgreSQL-compatible SQL queries that use the `WITH RECURSIVE` clause to implement graph traversal operations.

Listing 6.1: Output of the Umbra–TigerGraph cross-validation command

```
1 $ export SF=10
2 $ scripts/cross-validate.sh umbra tigergraph
3
4 +++ Files "umbra/output/output-sf10/results.csv" and "tigergraph/output/output-sf10/results.csv" are equal
```

¹https://github.com/ldbc/ldbc_snb_bi/tree/a2d3ac18c946d6a698c6aa5e6cf5d8954296be63/umbra

Supplementary Materials

7 SUPPLEMENTARY MATERIALS

Table 7.1: Supplementary materials

File or Directory	Purpose
umbra/output/output-sf10	Output of the Umbra reference implementation
parameters/parameters-sf{10,1000}.tar	Query substitution parameters
scoring/calculate-scores.py	Python script to calculate the scores of the benchmark run
tigergraph/output/output-sf{10,1000}	Benchmark logs and outputs
ldbc_snb_bi-1.0.3.tar.gz	Benchmark driver and reference implementations
ldbc_snb_datagen_spark-0.5.1.tar.gz	Data generator
ldbc_snb_specification-2.2.0.pdf	Benchmark specification
tigergraph-3.7.0-offline.tar.gz	Installer of Tigergraph

The `ldbc_snb_bi_tigergraph_sf1000_attachments` folder's directory structure is as follows:

```
ldbc_snb_bi_tigergraph_sf1000_attachments
├── umbra
│   └── output
│       └── output-sf10
│           └── results.csv
├── parameters
│   ├── parameters-sf10.tar
│   └── parameters-sf1000.tar
├── scoring
│   └── calculate-scores.py
├── tigergraph
│   └── output
│       ├── output-sf10
│       │   ├── benchmark.csv
│       │   ├── load.csv
│       │   ├── results.csv
│       │   └── timings.csv
│       └── output-sf1000
│           ├── benchmark.csv
│           ├── load.csv
│           ├── results.csv
│           └── timings.csv
├── ldbc_snb_bi-1.0.3.tar.gz
├── ldbc_snb_datagen_spark-0.5.1.tar.gz
├── ldbc_snb_specification-2.2.0.pdf
└── tigergraph-3.7.0-offline.tar.gz
```

CPU and Memory details

A CPU AND MEMORY DETAILS

Listing A.1: Output of the `cat /proc/cpuinfo` command for a single CPU core

```

1  processor       : 0
2  vendor_id      : AuthenticAMD
3  cpu family     : 25
4  model         : 1
5  model name     : AMD EPYC 7R13 Processor
6  stepping      : 1
7  microcode     : 0xa001173
8  cpu MHz       : 1803.263
9  cache size    : 512 KB
10 physical id   : 0
11 siblings     : 32
12 core id      : 0
13 cpu cores    : 16
14 apicid      : 0
15 initial apicid : 0
16 fpu         : yes
17 fpu_exception : yes
18 cpuid level  : 16
19 wp         : yes
20 flags       : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr
sse sse2 ht syscall nx mmxext fxsr_opt pdpe1gb rdtscp lm constant_tsc rep_good nopl nonstop_tsc cpuid
extd_apicid aperfmperf tsc_known_freq pni pclmulqdq sse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt aes
xsave avx f16c rdrand hypervisor lahf_lm cmp_legacy cr8_legacy abm sse4a misalignsse 3dnowprefetch topoext
invpcid_single ssbd ibrs ibpb stibp vmmcall fsgsbase bmi1 avx2 smep bmi2 erms invpcid rdseed adx smap
clflushopt clwb sha_ni xsaveopt xsavec xgetbv1 clzero xsaveerptr rdpru wbnoinvd arat npt nrip_save vaes
vpclmulqdq rdpid
21 bugs       : sysret_ss_attrs null_seg spectre_v1 spectre_v2 spec_store_bypass
22 bogomips   : 5299.97
23 TLB size   : 2560 4K pages
24 clflush size : 64
25 cache_alignment : 64
26 address sizes : 48 bits physical, 48 bits virtual
27 power management:

```

Listing A.2: Output of the `lscpu` command

```

1  Architecture:      x86_64
2  CPU op-mode(s):   32-bit, 64-bit
3  Byte Order:       Little Endian
4  CPU(s):           32
5  On-line CPU(s) list: 0-31
6  Thread(s) per core: 2
7  Core(s) per socket: 16
8  Socket(s):        1
9  NUMA node(s):     2
10 Vendor ID:        AuthenticAMD
11 CPU family:       25
12 Model:           1
13 Model name:       AMD EPYC 7R13 Processor
14 Stepping:        1
15 CPU MHz:         1502.144
16 BogomIPS:        5299.97
17 Hypervisor vendor: KVM
18 Virtualization type: full

```



CPU and Memory details

```

19  L1d cache:          32K
20  L1i cache:          32K
21  L2 cache:           512K
22  L3 cache:           32768K
23  NUMA node0 CPU(s): 0-7,16-23
24  NUMA node1 CPU(s): 8-15,24-31
25  Flags:              fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr
                        sse sse2 ht syscall nx mmxext fxsr_opt pdpe1gb rdtscp lm constant_tsc rep_good nopl nonstop_tsc cpuid
                        extd_apicid aperfmperf tsc_known_freq pni pclmulqdq ssse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt aes
                        xsave avx f16c rdrand hypervisor lahf_lm cmp_legacy cr8_legacy abm sse4a misalignsse 3dnowprefetch topoext
                        invpcid_single ssbd ibrs ibpb stibp vmmcall fsgsbase bmi1 avx2 smep bmi2 erms invpcid rdseed adx smap
                        clflushopt clwb sha_ni xsaveopt xsavec xgetbv1 clzero xsaveerptr rdpru wbnoinvd arat npt nrip_save vaes
                        vpclmulqdq rdpid

```

Listing A.3: Output of the dmidecode -t processor command

```

1 # dmidecode 3.2
2 Getting SMBIOS data from sysfs.
3 SMBIOS 2.7 present.
4
5 Handle 0x0004, DMI type 4, 42 bytes
6 Processor Information
7   Socket Designation: CPU 0
8   Type: Central Processor
9   Family: Zen
10  Manufacturer: Advanced Micro Devices, Inc.
11  ID: 11 0F A0 00 FF FB 8B 17
12  Signature: Family 25, Model 1, Stepping 1
13  Flags:
14    FPU (Floating-point unit on-chip)
15    VME (Virtual mode extension)
16    DE (Debugging extension)
17    PSE (Page size extension)
18    TSC (Time stamp counter)
19    MSR (Model specific registers)
20    PAE (Physical address extension)
21    MCE (Machine check exception)
22    CX8 (CMPXCHG8 instruction supported)
23    APIC (On-chip APIC hardware supported)
24    SEP (Fast system call)
25    MTRR (Memory type range registers)
26    PGE (Page global enable)
27    MCA (Machine check architecture)
28    CMOV (Conditional move instruction supported)
29    PAT (Page attribute table)
30    PSE-36 (36-bit page size extension)
31    CLFSH (CLFLUSH instruction supported)
32    MMX (MMX technology supported)
33    FXSR (FXSAVE and FXSTOR instructions supported)
34    SSE (Streaming SIMD extensions)
35    SSE2 (Streaming SIMD extensions 2)
36    HTT (Multi-threading)
37  Version: AMD EPYC 7R13 Processor
38  Voltage: 1.1 V
39  External Clock: 100 MHz
40  Max Speed: 3725 MHz
41  Current Speed: 2650 MHz
42  Status: Populated, Enabled

```



CPU and Memory details

```

43 Upgrade: Socket SP3
44 L1 Cache Handle: 0x0005
45 L2 Cache Handle: 0x0006
46 L3 Cache Handle: 0x0007
47 Serial Number: Not Specified
48 Asset Tag: Not Specified
49 Part Number: Not Specified
50 Core Count: 16
51 Core Enabled: 16
52 Thread Count: 32
53 Characteristics:
54   64-bit capable
55   Multi-Core
56   Hardware Thread
57   Execute Protection

```

Listing A.4: Output of the `cat /proc/meminfo` command

```

1 MemTotal:      258586064 kB
2 MemFree:      256575764 kB
3 MemAvailable: 256563876 kB
4 Buffers:      2704 kB
5 Cached:       1407712 kB
6 SwapCached:   0 kB
7 Active:       269588 kB
8 Inactive:     1176892 kB
9 Active(anon): 256 kB
10 Inactive(anon): 36360 kB
11 Active(file): 269332 kB
12 Inactive(file): 1140532 kB
13 Unevictable: 0 kB
14 Mlocked:     0 kB
15 SwapTotal:   0 kB
16 SwapFree:    0 kB
17 Dirty:       628 kB
18 Writeback:   0 kB
19 AnonPages:   36372 kB
20 Mapped:      50336 kB
21 Shmem:       524 kB
22 KReclaimable: 85604 kB
23 Slab:        200248 kB
24 SReclaimable: 85604 kB
25 SUnreclaim: 114644 kB
26 KernelStack: 6400 kB
27 PageTables:  4276 kB
28 NFS_Unstable: 0 kB
29 Bounce:      0 kB
30 WritebackTmp: 0 kB
31 CommitLimit: 129293032 kB
32 Committed_AS: 282916 kB
33 VmallocTotal: 34359738367 kB
34 VmallocUsed:  228772 kB
35 VmallocChunk: 0 kB
36 Percpu:     9216 kB
37 HardwareCorrupted: 0 kB
38 AnonHugePages: 0 kB
39 ShmemHugePages: 0 kB
40 ShmemPmdMapped: 0 kB

```

CPU and Memory details

```

41 FileHugePages:          0 kB
42 FilePmdMapped:         0 kB
43 HugePages_Total:       0
44 HugePages_Free:        0
45 HugePages_Rsvd:        0
46 HugePages_Surp:        0
47 Hugepagesize:          2048 kB
48 Hugetlb:               0 kB
49 DirectMap4k:           145316 kB
50 DirectMap2M:           11003904 kB
51 DirectMap1G:           252706816 kB

```

Listing A.5: Output of the `lshw -C memory` command

```

1 *-memory
2   description: System Memory
3   physical id: 8
4   slot: System board or motherboard
5   size: 256GiB
6   *-bank:0
7     description: DIMM DDR4 Static column Pseudo-static Synchronous Window DRAM 3200 MHz (0.3 ns)
8     physical id: 0
9     size: 128GiB
10    width: 64 bits
11    clock: 3200MHz (0.3ns)
12   *-bank:1
13     description: DIMM DDR4 Static column Pseudo-static Synchronous Window DRAM 3200 MHz (0.3 ns)
14     physical id: 1
15     size: 128GiB
16     width: 64 bits
17     clock: 3200MHz (0.3ns)

```



IO performance

B IO PERFORMANCE

Listing B.1: Output of the `fiio` command

```

1  iotest: (groupid=0, jobs=1): err= 0: pid=8197: Mon Mar 18 15:14:08 2023
2  write: io=2048.0MB, bw=6207.7KB/s, iops=1551, runt=337834msec
3      clat (usec): min=312, max=11380, avg=642.29, stdev=257.29
4      lat (usec): min=312, max=11380, avg=642.40, stdev=257.29
5      clat percentiles (usec):
6      | 1.00th=[ 382], 5.00th=[ 418], 10.00th=[ 438], 20.00th=[ 478],
7      | 30.00th=[ 510], 40.00th=[ 548], 50.00th=[ 588], 60.00th=[ 628],
8      | 70.00th=[ 676], 80.00th=[ 748], 90.00th=[ 876], 95.00th=[ 1048],
9      | 99.00th=[ 1624], 99.50th=[ 1912], 99.90th=[ 2960], 99.95th=[ 3728],
10     | 99.99th=[ 5536]
11     lat (usec) : 500=26.96%, 750=53.21%, 1000=13.88%
12     lat (msec) : 2=5.54%, 4=0.37%, 10=0.04%, 20=0.01%
13     cpu        : usr=0.52%, sys=1.60%, ctx=524288, majf=0, minf=11
14     IO depths  : 1=100.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%, >=64=0.0%
15     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
16     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
17     issued    : total=r=0/w=524288/d=0, short=r=0/w=0/d=0, drop=r=0/w=0/d=0
18     latency   : target=0, window=0, percentile=100.00%, depth=1
19
20     Run status group 0 (all jobs):
21     WRITE: io=2048.0MB, aggrb=6207KB/s, minb=6207KB/s, maxb=6207KB/s, mint=337834msec, maxt=337834msec
22
23     Disk stats (read/write):
24     nvme0n1: ios=0/524249, merge=0/14, ticks=0/332462, in_queue=332461, util=100.00%
```

Listing B.2: Output of the `lshw -class network` command

```

1  *-network
2      description: Ethernet interface
3      product: Elastic Network Adapter (ENA)
4      vendor: Amazon.com, Inc.
5      physical id: 5
6      bus info: pci@0000:00:05.0
7      logical name: eth0
8      version: 00
9      serial: 0e:50:ee:9c:f1:af
10     width: 32 bits
11     clock: 33MHz
12     capabilities: pciexpress msix bus_master cap_list ethernet physical
13     configuration: broadcast=yes driver=ena driverversion=2.8.3g ip=10.0.1.48 latency=0 link=yes multicast=yes
14     resources: irq:0 memory:feb4000-feb5fff memory:feb6000-feb7fff memory:fe800000-fe8ffff
```

Dataset generation instructions

C DATASET GENERATION INSTRUCTIONS

The datasets can be generated using the LDBC SNB Datagen. To regenerate the data sets used in this benchmark, build the Datagen JAR as described in the project's README, configure the AWS EMR environment, upload the JAR to the S3 bucket (denoted as `${BUCKET_NAME}`) and run the following commands to generate the datasets used in this audit.

Note that while the datasets for TigerGraph were generated as gzip-compressed archives, they are decompressed during preprocessing. Decompressing the SF1000 data set took 2 minutes 45 seconds (wall clock) when performed by the following command: `time find /home/tigergraph/sf1000 -name "*.csv.gz" -print0 | parallel -q0 gunzip`

Listing C.1: Script to generate the SF10 dataset for TigerGraph in AWS EMR. This dataset is only used for cross-validation

```

1 export SCALE_FACTOR=10
2 export JOB_NAME=sf${SCALE_FACTOR}-projected-csv-gz
3
4 ./tools/emr/submit_datagen_job.py \
5     --use-spot \
6     --bucket ${BUCKET_NAME} \
7     --copy-all \
8     --az us-east-2c \
9     ${JOB_NAME} \
10    ${SCALE_FACTOR} \
11    csv \
12    bi \
13    -- \
14    --explode-edges \
15    --format-options compression=gzip \
16    --generate-factors

```

Listing C.2: Script to generate the SF1000 dataset for TigerGraph in AWS EMR. This dataset is used for the benchmark run

```

1 export SCALE_FACTOR=1000
2 export JOB_NAME=sf${SCALE_FACTOR}-projected-csv-gz
3
4 ./tools/emr/submit_datagen_job.py \
5     --use-spot \
6     --bucket ${BUCKET_NAME} \
7     --copy-all \
8     --az us-east-2c \
9     ${JOB_NAME} \
10    ${SCALE_FACTOR} \
11    csv \
12    bi \
13    -- \
14    --explode-edges \
15    --format-options compression=gzip \
16    --generate-factors

```

Listing C.3: Script to generate the SF10 dataset for Umbra locally. This dataset is only used for cross-validation

```

1 export SCALE_FACTOR=10
2 export LDBC_SNB_DATAGEN_MAX_MEM=60G
3 export LDBC_SNB_DATAGEN_JAR=$(sbt -batch -error 'print assembly / assemblyOutputPath')
4

```



Dataset generation instructions

```
5 tools/run.py \  
6   --cores $(nproc) \  
7   --memory ${LDBC_SNB_DATAGEN_MAX_MEM} \  
8   -- \  
9   --format csv \  
10  --scale-factor ${SCALE_FACTOR} \  
11  --explode-edges \  
12  --mode bi \  
13  --output-dir out-sf${SCALE_FACTOR}/ \  
14  --generate-factors \  
15  --format-options header=false,quoteAll=true,compression=gzip
```



Data schema

D DATA SCHEMA

Listing D.1: Content of the GSQL schema used by TigerGraph

```

1 ## Message
2 CREATE VERTEX Comment (PRIMARY_ID id UINT, creationDate INT, locationIP STRING, browserUsed STRING, content
   STRING, length UINT) WITH primary_id_as_attribute="TRUE"
3 CREATE VERTEX Post (PRIMARY_ID id UINT, imageFile STRING, creationDate INT, locationIP STRING, browserUsed STRING
   , language STRING, content STRING, length UINT) WITH primary_id_as_attribute="TRUE"
4 ## organisation
5 CREATE VERTEX Company (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
6 CREATE VERTEX University (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
7 ## place
8 CREATE VERTEX City (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
9 CREATE VERTEX Country (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
10 CREATE VERTEX Continent (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
11 ## etc
12 CREATE VERTEX Forum (PRIMARY_ID id UINT, title STRING, creationDate INT,
13   maxMember UINT) WITH primary_id_as_attribute="TRUE" // maxMember is for precompute in BI-4
14 CREATE VERTEX Person (PRIMARY_ID id UINT, firstName STRING, lastName STRING, gender STRING, birthday INT,
15   creationDate INT, locationIP STRING, browserUsed STRING, speaks SET<STRING>, email SET<STRING>,
16   popularityScore UINT) WITH primary_id_as_attribute="TRUE" // popularityScore is for precompute in BI-6
17 CREATE VERTEX Tag (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
18 CREATE VERTEX TagClass (PRIMARY_ID id UINT, name STRING, url STRING) WITH primary_id_as_attribute="TRUE"
19
20 # create edge
21 CREATE DIRECTED EDGE CONTAINER_OF (FROM Forum, TO Post) WITH REVERSE_EDGE="CONTAINER_OF_REVERSE"
22 CREATE DIRECTED EDGE HAS_CREATOR (FROM Comment|Post, TO Person) WITH REVERSE_EDGE="HAS_CREATOR_REVERSE"
23 CREATE DIRECTED EDGE HAS_INTEREST (FROM Person, TO Tag) WITH REVERSE_EDGE="HAS_INTEREST_REVERSE"
24 CREATE DIRECTED EDGE HAS_MEMBER (FROM Forum, TO Person, creationDate INT) WITH REVERSE_EDGE="HAS_MEMBER_REVERSE"
25 CREATE DIRECTED EDGE HAS_MODERATOR (FROM Forum, TO Person) WITH REVERSE_EDGE="HAS_MODERATOR_REVERSE"
26 CREATE DIRECTED EDGE HAS_TAG (FROM Comment|Post|Forum, TO Tag) WITH REVERSE_EDGE="HAS_TAG_REVERSE"
27 CREATE DIRECTED EDGE HAS_TYPE (FROM Tag, TO TagClass) WITH REVERSE_EDGE="HAS_TYPE_REVERSE"
28 CREATE DIRECTED EDGE IS_LOCATED_IN (FROM Company, TO Country | FROM Person, TO City | FROM University, TO City)
   WITH REVERSE_EDGE="IS_LOCATED_IN_REVERSE"
29 CREATE DIRECTED EDGE MSG_LOCATED_IN (FROM Comment, TO Country | FROM Post, TO Country) // Reverse edge of
   Comment/Post -IS_Located_IN-> Country will cause Country connected by too many edges, which makes loading
   slow
30 CREATE DIRECTED EDGE IS_PART_OF (FROM City, TO Country | FROM Country, TO Continent) WITH REVERSE_EDGE="
   IS_PART_OF_REVERSE"
31 CREATE DIRECTED EDGE IS_SUBCLASS_OF (FROM TagClass, TO TagClass) WITH REVERSE_EDGE="IS_SUBCLASS_OF_REVERSE"
32 CREATE UNDIRECTED EDGE KNOWS (FROM Person, TO Person, creationDate INT, weight19 UINT, weight20 UINT DEFAULT
   10000)
33 CREATE DIRECTED EDGE LIKES (FROM Person, TO Comment|Post, creationDate INT) WITH REVERSE_EDGE="LIKES_REVERSE"
34 CREATE DIRECTED EDGE REPLY_OF (FROM Comment, TO Comment|Post) WITH REVERSE_EDGE="REPLY_OF_REVERSE"
35 CREATE DIRECTED EDGE STUDY_AT (FROM Person, TO University, classYear INT) WITH REVERSE_EDGE="STUDY_AT_REVERSE"
36 CREATE DIRECTED EDGE WORK_AT (FROM Person, TO Company, workFrom INT) WITH REVERSE_EDGE="WORK_AT_REVERSE"
37
38 CREATE DIRECTED EDGE ROOT_POST (FROM Comment, TO Post) WITH REVERSE_EDGE="ROOT_POST_REVERSE" //FOR BI-3,9,17
39 CREATE DIRECTED EDGE REPLY_COUNT (FROM Person, TO Person, cnt UINT)
40
41 CREATE GLOBAL SCHEMA_CHANGE JOB addIndex {
42   ALTER VERTEX Country ADD INDEX country_name ON (name);
43   ALTER VERTEX Company ADD INDEX company_name ON (name);
44   ALTER VERTEX University ADD INDEX university_name ON (name);
45   ALTER VERTEX Tag ADD INDEX tag_name ON (name);
46   ALTER VERTEX TagClass ADD INDEX tagclass_name ON (name);

```



Data schema

```
47 | }  
48 |  
49 | RUN GLOBAL SCHEMA_CHANGE JOB addIndex  
50 | CREATE GRAPH ldbc_snb (*)
```