



**Hewlett Packard
Enterprise**

HP Enterprise

The Experience On Running
Industry First Big Data Benchmark
TPCx-BB

Paul Cao

May 11, 2016



Big Data Benchmark


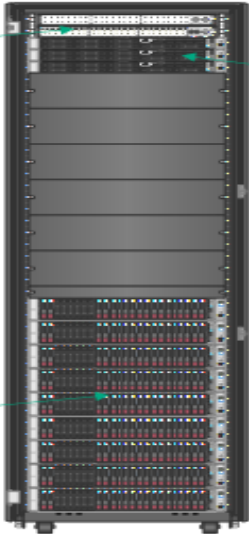
TPCx-BB Agenda

- ✓ What system you used.
- ✓ How easy/difficult it was to set the benchmark up.
- ✓ How much time you spent tuning.
- ✓ Did you run into any issues.
- ✓ Assessment of the benchmark.
- ✓ How well does it measure the performance of your system.


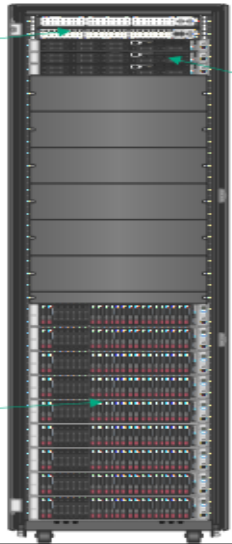
✓ Hardware Configuration

- ✓ Two TPC Official Publications

Configuration 1

 Hewlett Packard Enterprise		Hewlett Packard Enterprise ProLiant DL for Big Data			TPCx-BB Rev. 1.0.1 TPC-Pricing Rev. 1.7.0	
Total System Cost		TPCx-BB Performance Metric			Report Date: March 31, 2016	
371,977 USD		337.26 BBQpm@3000			1,102.94 USD \$/BBQpm@3000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams	
Cloudera for Apache Hadoop (CDH) 5.6	Red Hat Enterprise Linux Server 6.7	None	March 31, 2016	3000	2	
System Configuration						
Ethernet Switch: HPE 1620-24G Switch (ILO connection) HP5900AF-48XGT-4QSFP+ (main connection)					3 Management Nodes Each: HPE ProLiant DL360 Gen9 2x Intel E5-2640 v4 2.40GHz 1x HPE 800GB SSD 256GB Memory	
9 Worker Nodes Each: HPE ProLiant DL380 Gen9 2x Intel E5-2697A v4 2.60GHz 1x HPE 480GB SSD 16x HPE 1TB HDD 1x HPE 1.2TB SDD 192GB Memory					Software: Red Hat Enterprise Linux 6.7 Cloudera Enterprise 5.6	
Physical Storage/Scale Factor: 53.84			Scale Factor/Physical Memory: 1.20			
Servers:		9x HPE ProLiant DL 380 Gen 9, 3x HPE ProLiant DL360 Gen9				
Total Processors/Cores/Threads		24/348/696				
Server Configuration:	Per HPE ProLiant DL 380 Gen 9:		Per HPE ProLiant DL360 Gen9:			
Processors	2x Intel Xeon E5-2697A v4 @ 2.60GHz		2x Intel Xeon E5-2640 v4 @ 2.40GHz			
Memory	192GB		256GB			
Storage Controller	HPE Smart Array P840/4G		Smart HBA H240ar			
Storage Device	1x HPE 480GB SSD 16x HPE 1 TB HDD 1x HPE 1.2 TB SSD		1x HPE 800GB SSD			
Network	HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr		HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr			
Connectivity:	HPE 1620-24G Switch, HPE 5900AF-48XG-4QSFP+ Switch					

Configuration 2

 Hewlett Packard Enterprise		Hewlett Packard Enterprise ProLiant DL for Big Data			TPCx-BB Rev. 1.0.1 TPC-Pricing Rev. 1.7.0	
Total System Cost		TPCx-BB Performance Metric			Report Date: March 23, 2016	
325,156 USD		265.93 BBQpm@3000			1,222.72 USD \$/BBQpm@3000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams	
Cloudera for Apache Hadoop (CDH) 5.6	Red Hat Enterprise Linux Server 6.7	None	March 23, 2016	3000	2	
System Configuration						
<p>Ethernet Switch: HPE 1620-24G Switch (ILO connection) HP5900AF-48XGT-4QSFP+ (main connection)</p>					<p>3 Management Nodes Each: HPE ProLiant DL360 Gen9 2x Intel E5-2640 v3 2.60GHz 1x HPE 480GB SSD 1x HPE 800GB SSD (1 node) 128GB Memory</p>	
<p>9 Worker Nodes Each: HPE ProLiant DL380 Gen9 2x Intel E5-2697 v3 2.60GHz 1x HPE 480GB SSD 16x HPE 1TB HDD 128GB Memory</p>					<p>Software: Red Hat Enterprise Linux 6.7 Cloudera Enterprise 5.6</p>	
Physical Storage/Scale Factor: 50.19				Scale Factor/Physical Memory: 1.95		
Servers:		9x HPE ProLiant DL 380 Gen 9, 3x HPE ProLiant DL360 Gen9				
Total Processors/Cores/Threads		24/300/600				
Server Configuration:		Per HPE ProLiant DL 380 Gen 9:		Per HPE ProLiant DL360 Gen9:		
Processors		2x Intel Xeon E5-2697 v3 @ 2.60GHz		2x Intel Xeon E5-2640 v3 @ 2.60GHz		
Memory		128GB		128GB		
Storage Controller		HPE Smart Array P840/4G		Smart HBA H240ar		
Storage Device		1x HPE 480GB SSD 16x HPE 1 TB HDD		1x HPE 480GB SSD 1x HPE 800GB SSD (1 node)		
Network		HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr		HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr		
Connectivity:		HPE 1620-24G Switch, HPE 5900AF-48XG-4QSFP+ Switch				

✓ The Challenge

- ✓ Multiple nodes (12 nodes x 2)
- ✓ Hardware setup
- ✓ Software setup
- ✓ Runtime
- ✓ Tuning
- ✓ Debugging (hardware, software)
- ✓ Audit
- ✓ FDR

How To

1. Get your cluster ready with HW/SW

2. Setup Passwordless ssh for pssh to work.

3. Edit "userSettings.conf" and insert your intended variables.

```
export BIG_BENCH_DEFAULT_DATABASE="bigbench" */ Name of Hive Metastore */
```

```
export BIG_BENCH_DEFAULT_ENGINE="hive" */ Default framework Engine */
```

```
export BIG_BENCH_DEFAULT_MAP_TASKS="80" */ Number of map tasks to generate data, read extended
```

readme on selecting one */

```
export BIG_BENCH_DEFAULT_SCALE_FACTOR="10" */ Scale Factor you would like to test 1000=1TB,
```

```
3000=3TB, 10000=10TB, 30000=30TB, 100000=100TB, 300000=300TB, 1000000=1PB */
```

```
export BIG_BENCH_DEFAULT_NUMBER_OF_PARALLEL_STREAMS="2" */Number of concurrent streams to use
```

during throughput phase, default is 2 */

```
export BIG_BENCH_DEFAULT_BENCHMARK_PHASE="run_query"
```

```
export BIG_BENCH_HADOOP_CONF="/etc/hadoop/conf.cloudera.hdfs" */ Adjust this to whatever distrubution of
```

Hadoop you are using */

```
export BIG_BENCH_HADOOP_LIBS_NATIVE="/opt/cloudera/parcels/CDH/lib/hadoop/lib/native" */ Adjust this to
```

whatever distrubution of Hadoop you are using */

```
export BIG_BENCH_DATAGEN_DFS_REPLICATION="1" */ Use this setting to select number of replicas for your
```

generated data on your HDFS file system, "1" means no replica, "3" is HDFS default of 3 copies. WARNING:* This setting has no bearing on default HDFS replication for all other files, which is set to 3" */

```
export BIG_BENCH_STOP_AFTER_FAILURE="1" */ the default behaviour is to stop when a query error occurs, set this
```

to 0 to keep on running when an error occurs*/Challenges */

✓ How To

4. Run `$/bin/TPCxBB_Validation.sh */ Engine Validation Phase on SF1“ */`
5. Run `$/TPCxBB_Benchmarkrun.sh */` To execute the Benchmark, run twice to obtain "Performance" and "Repeatability" numbers. `*/`
6. Copy 3 set of logs from `$/../logs` 1. From Engine Validation Phase logs. 2. Performance run logs. 3. Repeatability run logs. E.g logfile from one run: `logs-20151029-135147-hive-sf1xxx.zip`
7. Engage the auditor and submit the report for publication.


✓ Assessment of the Benchmark

- ✓ Stability
- ✓ Scaling
- ✓ Challenges
- ✓ Audit

✓ System Response to the Workload

✓ Better than expected


HPE ProLiant DL Clusters Intel Haswell vs Broadwell on Red Hat RHEL 6




Q & A

TPC Big Bench CPU Utilization % 2 vs. 4 Streams

Total 67%




Total 83%

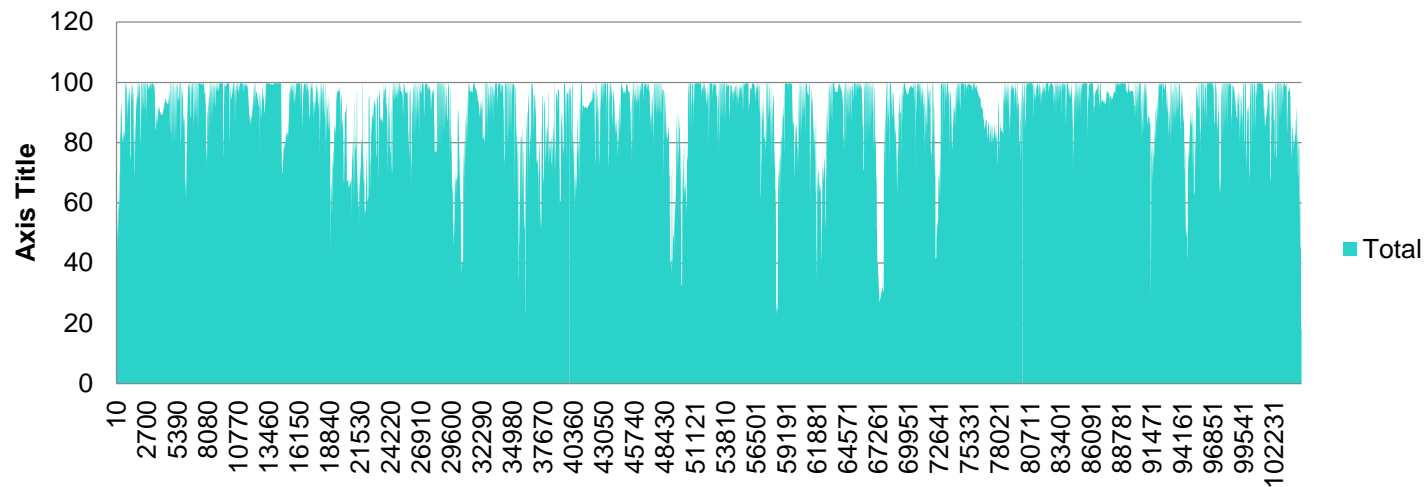


TPC Big Bench Memory Utilization

Memory Utilization




Memory Utilization




TPC Big Bench Disk Bandwidth

Disk Bandwidth




Disk Bandwidth




TPC Big Bench Network IO

Network IO




Network IO




TPC Big Bench IO Latencies

IO Latencies




IO Latencies

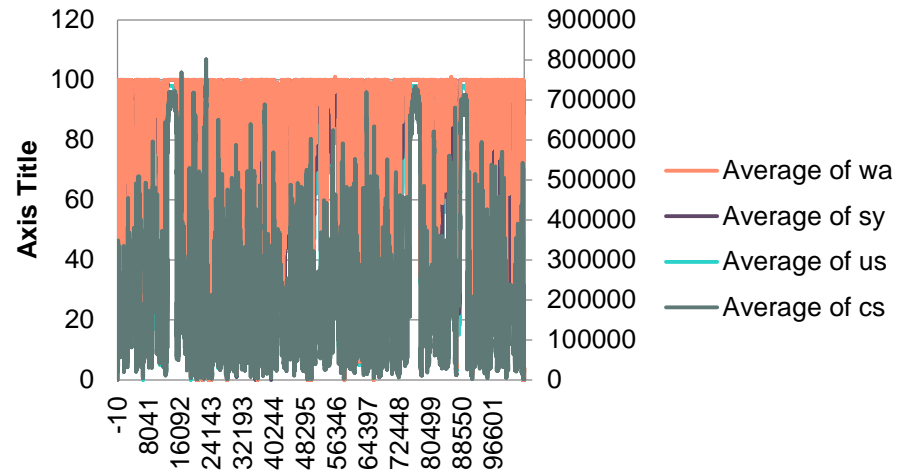


TPC Big Bench Context Switches

Context Switches




Context Switches



TPC Big Bench Free Memory

Free Memory



Free Memory

