Performance tests

TBLignite 0.1.0 performance tests for apache Ignite 2.7.0 on Google Compute using the NOAA NCDC weather dataset. If you want to reproduce these tests, download the dataset from ftp://ftp.ncdc.noaa.gov/pub/data/noaa.

The source code of these tests is available from http://tblcore.com/download/.

Results

Machine	Google Compute n1-standard-4
Memory	15GB
Cores	4
OS	Debian 9.8
Rows	400mln
Size on disk	164GB

Clients	14GB Ignite	9GB Ignite 5GB TBLignite plugin	Improvement
1	01:00:46	00:32:04	+89%
4	01:07:30	01:37:26	-31%

Machine	Google Compute n1-standard-16
Memory	60GB
Cores	16
OS	Debian 9.8
Rows	690mln
Size on disk	283GB

Clients	56GB Ignite	38GB Ignite 18GB TBLignite plugin	Improvement
1	01:18:23	00:47:14	+65%
12	01:37:37	01:07:56	+44%

Conclusion

The TBLignite 0.1.0 plugin can almost double the speed of queries in Ignite provided you meet these conditions:

- Your dataset is fit for columnar compression, ea. repeating values and/or a timeseries-like dataset.
- Your active dataset does not entirely fit in memory.
- The Ignite nodes are not CPU starved.

In some cases, the use of TBLignite 0.1.0 can lead to a decrease in performance.