

With many thanks to Vadim at Percona for his analysis of different capabilities of different columnar dbms. Definitely good information, and well documented. Queries and results available at:

<http://www.mysqlperformanceblog.com/2009/11/02/air-traffic-queries-in-infinidb-early-alpha/>

Of course InfiniDB, does offer multi-threaded processing for all offerings and distributed processing (enabled with Enterprise Edition additional functionality) that was beyond the scope of his test platform.

So, using the same airline on-time queries, we can show the relative performance of the queries as we scale both threads and servers.

System under test is a Dell server with dual Xeon CPU's, with quad cores (8 total cores).

Intel(R) Xeon(R) CPU E5335 @ 2.00GHz

To gain an approximation of different server capabilities, individual cores were then taken offline using Linux hotplug capabilities and the queries were run:

```
echo 0 > /sys/devices/system/cpu/cpu7/online
```

```
...
```

```
echo 0 > /sys/devices/system/cpu/cpu2/online
```

The timings are captured below, with some analysis to follow. The 2 core timings measured here appear to correspond well with Vadim's timings (total of 178.6 seconds reported by Vadim, and 174.56 seconds measured here). Columns 3-5 below show benefits of additional cores, columns 6 and 7 show benefits of additional servers.

Query

InfiniDB
metrics -Vadim at Percona

Community InfiniDB Repeated w/ 2 cores enabled

Community InfiniDB Repeated w/ 4 cores enabled

Community InfiniDB Repeated w/ 8 cores enabled

InfiniDB Two Performance Modules (Enterprise Edition)

InfiniDB Four Performance Modules (Enterprise Edition)

Q0

NA

Q1

6.79

6.97

3.43

1.91

1.29

0.87

Q2

4.59

4.83

2.48

1.71

1.34

0.69

Q3

4.96

5.56

3.19

1.93

1.7

0.96

Q4

0.75

0.8

0.43

0.36

0.32

0.32

Q5

NA

Q6

NA

Q7

NA

Q8 (1y)

8.13

8.03

4.17

2.54

2.23

2.02

Q8 (2y)

16.54

16.03

8.75

5.13

3.26

3.79

Q8 (3y)

24.46

23.62

12.66

7.72

5.15

3.3

Q8 (4y)

32.49

31.21

16.61

10.06

5.77

3.44

Q8 (10y)

70.35

68.3

37.06

23.08

14.09

7.16

Q9

9.54

9.21

4.55

2.4

1.39

0.86

Q0, Q5, Q6, Q7 – on roadmap

Q1, Q2, Q3, Q4, Q9 - show good scalability overall, however queries under about 1/2 second show either smaller benefits from scaling, or no benefit (Q4 shows the close to the same time for 4 cores and above).

Q8 variations - show overall good scalability, but we think there may be some performance opportunities here.

Look to this site for additional scalability measurements on Scan operations, Aggregation, and (Hash) Join operations.