NXP-9494 performance review

Benoit Delbosc

2012-07-04

1. Context

Because VCS don't use a SQL Cursor [1] the full query result are transmited at the JDBC level.

NXP-9494 introduces an SQL LIMIT in queries to reduce the size of the results.

There are 3 potentials improvements:

- There is less network traffic
- The Java memory footprint is smaller
- The SQL execution can be faster

This evaluation try to answer to the question when an how it is efficient.

2. Benchmark

It sounds obvious but only a network bound benchmark can reveal the impact of traffic reduction.

The daily bench which is mostly CPU bound is not sensible to traffic improvement (Figure 1) furthermore as a mono server the loopback interface is not going to be a bottleneck.

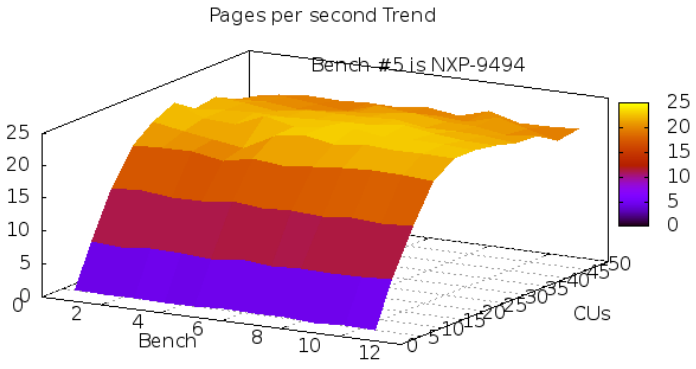

Figure 1: Continuous integration: Build #5 does not show any improvement

[1] VCS don't use a cursor because when possible we want the total number of results and it requires a scrollable result type which is not compatible with a cursor usage.

Simulating a network latency on a single host does not work in our case because of the TCP segmentation offload. When using a tool like `tc` [2] the tax is added before segmentation. This prevent to see any difference when reducing the size of a response.

For this reason the benchmark must be done on two physical servers that have enough CPU and memory. The network connection is only Fast Ethernet 100 Mbits/s to emphasis a network congestion.

Also NXP-9494 can only be seen as an improvment if the limitation is effective, like when browsing folders with few thousand of documents.

The scenario:

1. login among 50 different users
2. view a random folder among 100, each folder contains 5000 documents.
3. logout

Step 2 is done 5 time.

3. Results

The benchmark is network bound, there are no CPU, memory, disk saturation or Java concurrency contention.

The benchmark results indicate: **35% more throughput** (Fig 2) and responses times are **from 10% to 70% faster** (Fig 3).

This is because there are up to **5 times less traffic** from the database (Fig 5).

The memory impact and SQL execution are not visible may be because the bench with the limited results was able to process more requests.

4. Conclusion

The NXP-9494 optimisation will help on large-scale storage with a remote database. Especially when there are lots of hits and/or when the network is slow.

[2] Simulate latency with tc

A. Appendices

The full benchmark reports are available on-line:

- the differential report
- unlimited: bench report, monitoring, misc logs
- default limit to 200: bench report, monitoring, misc logs

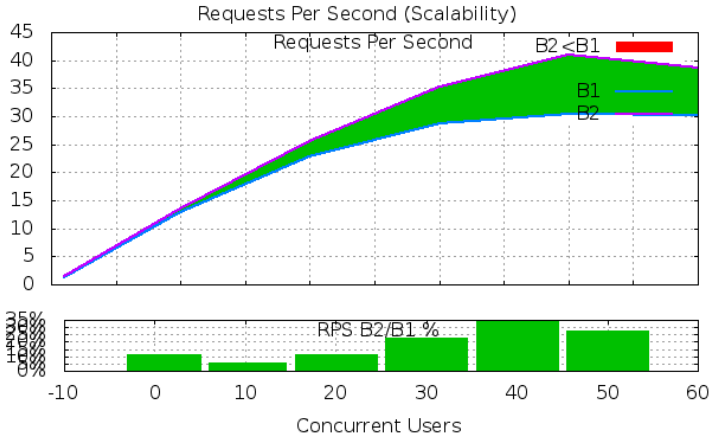

Figure 2: Throughput in request per second, B2 using LIMIT 200, B1 is unlimited

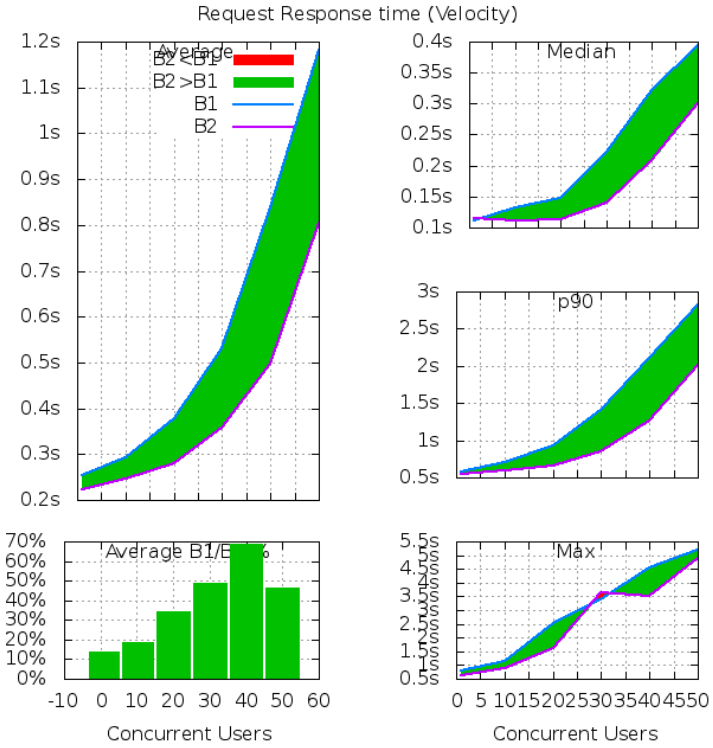

Figure 3: Request resonse time, B2 using LIMIT 200, B1 is unlimited

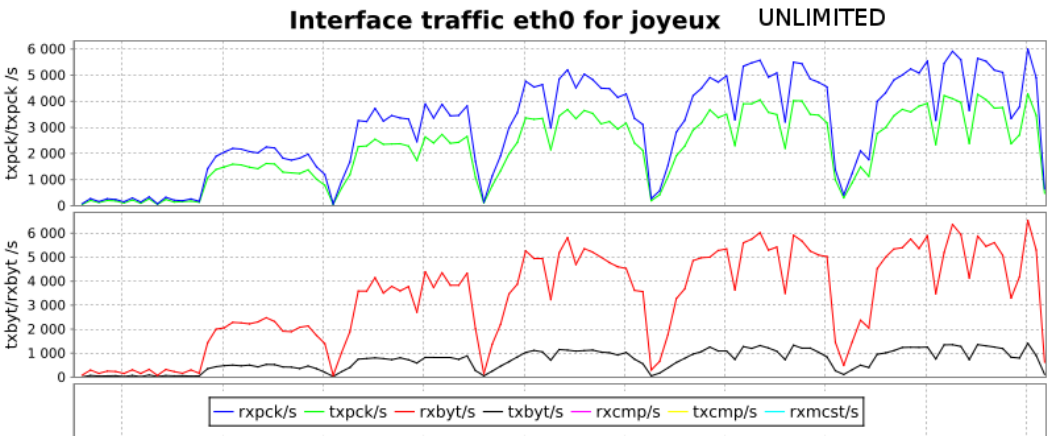

Figure 4: Network traffic unlimited results

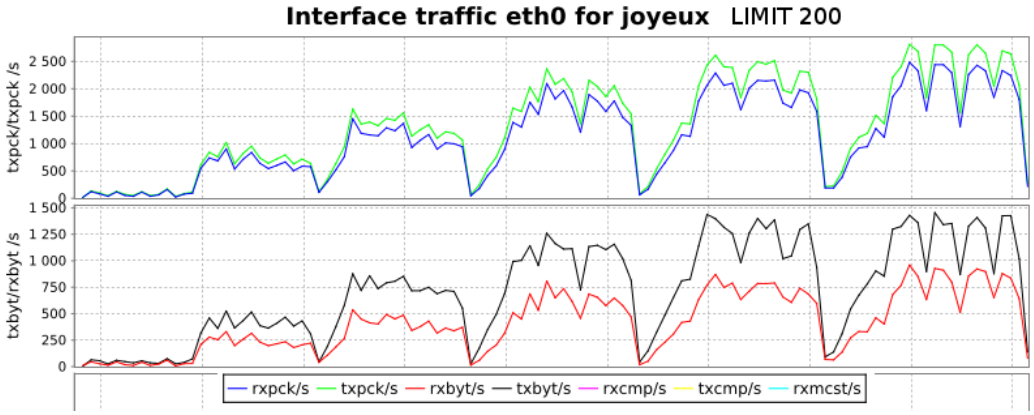

Figure 5: Network traffic 200 results