

# [HBASE-13071] Hbase Streaming Scan Feature: Evaluation Results for Multiple Clients

March 16, 2015

Eshcar Hillel ([eshcar@yahoo-inc.com](mailto:eshcar@yahoo-inc.com))

## Experimental Settings

The experiments are described in details in the [Hbase Streaming Scan Evaluation](#) document.

A table with 1M records, each record with 10 fields of 100b, total of 1GB table.

Executed 1000 big scans, each scanning 20K records.

Average delay per next operation 20us

Tested prefetch size 100 and 1000 with 1, 2, and 10 clients on a single machine, sync vs async scanners.



## Evaluation results


The results with 2 and 10 parallel clients show similar trends to these of a single client, namely reduced latency with async scanner.

As the prefetch size increases from 100 to 1000, the client-server-client prefetch round trip takes longer time, and the benefit from running the prefetch in the background increases.

The figures below depicts the accumulated latency of the x-th step at each experiment.

## Prefetch size 100





## Prefetch size 1000

