

DigestQueries

A "digest" query is like a read query except that instead of the receiving node actually returning the data, it only returns a digest (hash) of the would-be data.

The intent of submitting a digest query is to discover whether two or more nodes agree on what the current data is, without sending the data over the network. In particular for large amounts of data, this is a significant saving of bandwidth cost relative to sending the full data response.

Keep in mind that the cost of potentially going down to disk, and most or all of the CPU cost, associated with a query will still be taken on nodes that receive digest queries. The optimization is only for bandwidth.

<https://c.statcounter.com/9397521/0/fe557aad/1/> | stats