

# SecondaryIndexes

## Articles/Blogs

- A good article by Jonathan Ellis: <http://www.datastax.com/dev/blog/whats-new-cassandra-07-secondary-indexes>.
- An article that discusses when and when not to use secondary indexes: <http://www.datastax.com/docs/1.1/ddl/indexes>.

## FAQ for Secondary Indexes

- Q: Are there any limitations beside the hash properties (no range queries), like size, memory, etc.?  
A: No.
- Q: Are they distributed? If so, how does that work? How are they stored on the nodes?  
A: Each node only indexes data that it holds locally.
- Q: When you write a new row, when/how does the index get updated? What I would like to know is the atomicity of the operation--is the "index write" part of the "row write"?  
A: The row and index updates are one, atomic operation.
- Q: Is there a difference between creating a secondary index vs creating an "index" CF manually such as "users\_by\_country"?  
A: Yes. First, when creating your own index, a node may index data held by another node. Second, updates to the index and data are not atomic.
- Q: Why is it necessary to always have at least one EQ comparison on secondary indices?  
A: Inequalities on secondary indices are always done in memory, so without at least one EQ on another secondary index you will be loading every row in the database, which with a massive database isn't a good idea. So by requiring at least one EQ on an index, you hopefully limit the set of rows that need to be read into memory to a manageable size. (Although obviously you can still get into trouble with that as well).
- Q: How does choice of Consistency Level affect cluster availability when using secondary indexes?  
A: Because secondary indexes are distributed, you must have CL nodes available for **all** token ranges in the cluster in order to complete a query. For example, with RF = 3, when two out of three consecutive nodes in the ring are unavailable, **all** secondary index queries at CL = QUORUM will fail, however secondary index queries at CL = ONE will succeed. This is true regardless of cluster size.

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