

Duplicate Indexes

Duplicate indexes are indexes in a SQL Server database that have the same key columns and the same included columns, and are created on the same set of data. These indexes serve the same purpose as each other and provide no additional benefit to the database.

Duplicate indexes can be bad for a few reasons. First, they take up space in the database, which can increase the size of the database and consume storage resources. This can be particularly problematic for large databases, where the amount of space consumed by duplicate indexes can be significant.

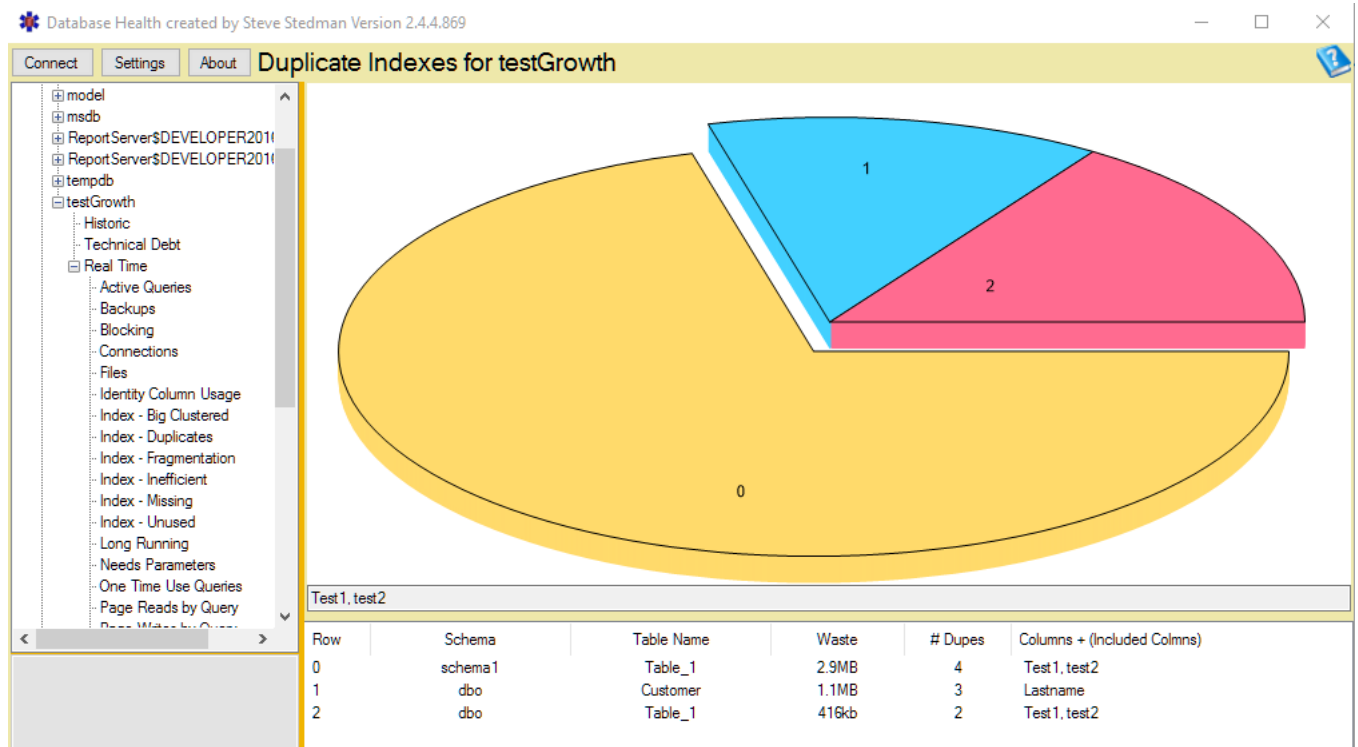
Second, duplicate indexes can slow down performance. When the database engine accesses the data in the database, it has to consider all of the available indexes, including the duplicate ones. This can add overhead to the query and slow down performance.

Third, duplicate indexes can make it more difficult to manage the database. Because there are multiple indexes that serve the same purpose, it can be confusing for users to determine which index to use in a given situation. This can lead to poor indexing choices and suboptimal performance.

Duplicate Indexes on your database can be very expensive from a number of different perspectives.

- Increases database size with no value.
- Increase database backup size with no value.
- Slows inserts, updates and deletes because there are duplicate indexes that need to be updated.
- Can slow queries as the query optimizer needs to consider both indexes when analyzing query performance and determining a plan.
- Increase Deadlocks.
- Bloated transaction log size.
- Add to the amount of data that replicated databases need to transfer.

With the Duplicate Index Report, you can track down the waste and eliminate it.



From the duplicate indexes page, you can jump to to the [Duplicate Index Advisor](#) by double clicking on the grid or on the chart.

On SQL Server 2008 and newer, filtered indexes are excluded from the Duplicate Indexes report.