SQL SERVER Performance Tuning Notes

http://msbiskills.com/

Why we should update statistics? When one should update statistics?

Why to update statistics?

If we do not update statistics then we will get bad query execution plan. Statistics are necessary for best query performance. Statistics essentially "describe" the data what is there in the index. SQL Server Query Optimizer uses statistics to estimate the distribution of values. Statistics helps query optimizer in making better decision making.

When to update statistics?

The best answer for this is we should update statistics regularly.

We should keep them updated using Auto Update Stats = true. If auto stats update is enabled then updates happen based on number of changes to the data. This is sometimes not enough.

SQL Server has index and column stats. We can use Update statistics to update stats on an index and on columns whenever you want using Update Statistics command.

On larger tables we should run update statistics on a schedule to refresh stats regularly. Otherwise we will land up in problem as your query is running slow.

Rebuilding indexes will also updates statistics. This process normally happens as a weekly maintenance task. Often when weekly index rebuild happens things works very fine on Monday and Tuesday. Now on Wednesday we will start getting issues like this query is

working slow and stuff like that. Now our query works fine on Monday and Tuesday because of updated statistics.

sp_updatestats - It runs Update Statistics against all User defined tables and internal tables in the current database. Well this is not recommended.

We can update statistics synchronously and asynchronously. Default is synchronously. Here when SQL Server tries to execute a query and before executing SQL Server will check the threshold value against the # of changes happened in the table and if the threshold has crossed than SQL Server first update the statistics and then the query will be executed. So depending on the table size and the other processes running on the server update statistics may take some time and slow down you query.

In asynchronous case the query will be executed first and then the update statistics will happen. In this case the benefit is that our query will not have to wait for update statistics to happen. The disadvantage of this case is that our query may get executed with less qualified execution plan as the statistics are not up to date. Most of the cases people don't ON this feature.

That's all folks, I hope you've enjoyed learning about why and when we should update statistics, and I'll see you soon with more "Performance Tuning" articles.

Thanks!

Pawan Kumar Khowal

MSBISKills.com