

## Solving SQL Server Issues: The Long Road

You can easily resolve some issues related to SQL Server performance, but there are others that require more time and resources. Implementing these performance solutions may require some planning or perhaps a service outage. Once configurations are adjusted, they may need long-term monitoring or ongoing adjustments to succeed. That's why we call these solutions "The Long Road."

There are four issues under this umbrella: file locations, virtual log files, tempdb files, and missing index DMVs.

### File Location

File location is an important aspect of good database architecture. For example, locating system and user database files on the same disk creates problems in case data needs to be recovered. If that single disk goes down, all the information may be lost, but if the files are on separate disks, the situation is isolated. Also, From an organizational standpoint, it makes good sense to separate these types of files.

You may want to carefully consider using autogrow with SQL Server. If you're considering using the default values for database size and growth in Autogrow, you may want to rethink this. These types of initial configurations typically persist well after the database has grown into full production mode. If a database is active, using the feature may result in continual growth commands being issued for the file. Proper sizing and instituting growth procedures as needed — in off-hours — is a better option.

### Virtual Log Files

Misconfigured Autogrow settings may also cause an excess of virtual log files in the transaction log. When a log file grows, it grows in virtual log file "chunks". Similarly, when a log file is truncated or shrinks, this also affect the virtual log files. An excessive number of virtual log files can create unnecessary overhead, which translates into longer recovery times. The ideal approach, as with Autogrow, is to manually size the transaction log files. Of course, resizing may be part of a larger initiative to reduce virtual log file counts.

### Tempdb Files

Another issue is the number of files needed for tempdb. You may read several suggestions about how to implement tempdb, but the ideal is to have as many files as needed to eliminate page allocation contention. Every time tempdb is used, the thread accessing tempdb must first find adequate free space to build the temp object. Because tempdb is a shared resource, a bottleneck can develop when many threads need tempdb space.

A good solution is to distribute the load by adding more tempdb files. Proceed carefully. If there are too many tempdb files, this can create memory issues.

## Missing Index DMVs

Missing index Dynamic Management Views is one of the more underused features in SQL Server. These types of DMVs — `sys.dm_db_missing_index_group_stats`, `sys.dm_db_missing_index_groups`, `sys.dm_db_missing_index_details`, and `sys.dm_db_missing_index_columns` — allow a database administrator to discover the indexes the SQL query engine would have used had they been available. The missing index information can be reviewed later for possible implementation.

These missing index reports do have some limitations. Database administrators need to carefully proceed to prevent the production of needlessly large or duplicate indexes. A best practice is evaluating each suggested missing index.

These and other solutions designed to help you more fully leverage the time and resources invested in your SQL Server are discussed in detail in a white paper written by Andy McDermid, Datavail's SQL Server principal database administrator.

You may also wish to explore how a comprehensive health check can help your organization. The Datavail Health Check is a tool you can readily use to increase your system's operating efficiencies.

To learn more about long road performance tuning or our health checks, please contact Datavail and be sure to read our frequently updated blog for more DBA answers.

**About the Author:** Andy McDermid  
*SQL Server DBA, Datavail*

Andy is a MCITP certified MS SQL DBA who delivers and manages delivery of DBA services to many diverse clients. He enjoys helping his clients by finding and deploying pragmatic and practical solutions for their database issues. Andy is always working to improve and expand his DBA skills and he likes to share the experience via writing.