SQL Server Database Structure

Each SQL Server database consists of at least two files:

- A primary data file with an `.mdf` extension.
- A log file with an `.ldf` extension.

There are also secondary data files with `.ndf` extensions. A database can have only one primary data file, zero or more secondary data files, and one or more log files. Each database file can only be used by one database.

In a database, data files store persistent data. For ease of management, you can group one or more data files into logical tablespaces. The filegroup is the SQL Server equivalent of an Oracle tablespace. SQL Server filegroups fall under, and are associated with, the individual databases. The SQL Server data hierarchy is:

Instance / Database / FileGroup / Data file

Each data file can be a member of only one filegroup, but the log files are managed separately from one another. There are three types of filegroups:

- Primary
- User defined
- Default

When you configure your SQL Server databases, you can specify the maximum size of data files to prevent disk drives from running out of space. If you do not specify the size of data files, the database assumes that the size is unlimited.

> ⚠️ **Note**
>
> Unlimited data files can not initiate Uptime Infrastructure Monitor alerts. Because Uptime Infrastructure Monitor measures datafile and log file size as a percentage of their maximum size, if a datafile size has an infinite maximum size, the percent of maximum datafile size must be near zero. Therefore you should always set the maximum size of each datafile for effective monitoring.

The following diagram illustrates six data files in three filegroups in three databases across two instances of a system:

blocked URL

If you set a Critical threshold of 90% for SQL Server Instance_B and a Warning threshold of 70%, the SQL Server Tablespace Check monitor watches the size of all data files in that instance. The monitor will send an alert if any of the files reaches or exceeds the defined thresholds.