

PARALLEL_REDO_WORKER_WAIT_WORK

SQL Server uses a wait type called `PARALLEL_REDO_WORKER_WAIT_WORK` to track the amount of time that a worker is waiting for a redo operation to be assigned to it. This wait type is typically seen in databases that use the Always On Availability Groups feature, which allows multiple copies of a database to be maintained on different servers for high availability and disaster recovery purposes.

When a database is part of an Always On Availability Groups setup, the redo operation is responsible for applying changes made to the database on the primary server to the secondary databases on the other servers. This process, known as redo, is performed in parallel by multiple worker threads to improve performance and reduce the time it takes to apply changes to the secondary databases.

The `PARALLEL_REDO_WORKER_WAIT_WORK` wait type is used to track the time that a worker thread spends waiting for a redo operation to be assigned to it. If this wait type is seen frequently in the system, it can indicate that the redo operation is not being performed efficiently, which can impact the performance and availability of the database.

To address this issue, the performance of the redo operation can be optimized by increasing the number of worker threads, or by reducing the workload on the primary server to allow the redo operation to be performed more efficiently. Additionally, ensuring that the hardware and software used for the Always On Availability Groups setup are properly configured and optimized can also help improve the performance of the redo operation and reduce the occurrence of the `PARALLEL_REDO_WORKER_WAIT_WORK` wait type.