

Avoiding Costs During Downtime

Data warehouse platforms often support auto-suspension after certain inactivity time. This helps reducing the costs. Within this document, this is referred to as 'downtime'.

The activity is coming from queries by BI clients, and also from system queries by the AtScale system. While outside of business hours the former usually stop, the latter might continue, which would prevent the auto-suspension. If you wish to increase the downtime, you can follow the recommendations below to reduce the amount of system queries.

Disabling Warm-Up Queries

Connection-pool warm-up queries can be disabled by updating the following AtScale engine settings:

1. Follow the procedure in [Changing Engine Settings](#) to access the settings screen.
2. Locate the `connection.pool.minIdleConnections` setting and set its value to 0 (default).
3. Choose Show Custom Settings, enter `connection.pool.user.minIdleConnections`, and set the value to 0 (default).
4. Restart the AtScale engine.

As a result, your AtScale system would stop refreshing the connection-pool every 30 minutes, and would not interrupt the downtime. Consider that there might be a minimal performance impact with maintenance or user queries for initial connection-pool warm-up after a period of inactivity on the cluster, but cloud cost savings outweighs that.

Optimizing Aggregate Maintenance

Aggregate maintenance - which is time-based - can be scheduled not to interrupt the downtime. By default, aggregate maintenance is scheduled for 04:00 (am) UTC every day of the week. For details, see [Configuring Aggregate Maintenance](#)

Scheduling Batch Builds

Aggregate batch builds are also time-based, and can be scheduled not to interrupt the downtime. For details, see [Scheduler for Aggregate Creation](#).