

Catalog

`atscale.yml` (the catalog file) is the control file for an AtScale repository. It contains all repository-level definitions, such as the repository name and settings for building aggregates. Each repository must contain a file called `atscale.yml` at the root level.



Note: Some properties can appear in both `atscale.yml` and [model files](#). Those defined in model files override their counterparts in `atscale.yml`.

`atscale.yml` supports the following properties.

Unique_name

- ▲ **Type:** string
- ▲ **Required:** Y

The name of the repository. This must be unique across all repositories and subrepositories.

Object_type

- ▲ **Type:** const
- ▲ **Required:** Y

The type of object defined by the file. For `atscale.yml`, this must be `catalog`.

Label

- ▲ **Type:** string
- ▲ **Required:** Y

The name of the repository, as it appears in AtScale. This value does not need to be unique.

As_version

- ▲ **Type:** number
- ▲ **Required:** Y

The version of SML being used.

Aggressive_agg_promotion

- ▲ **Type:** boolean
- ▲ **Required:** Y

Enables/disables aggressive aggregate promotion for the repository. When enabled, all aggregates referenced by a query are considered for promotion, regardless of whether a join to other non-preferred or non-aggregate datasets was required.

Supported values:

- ▲ true
- ▲ false

Build_speculative_aggs

- ▲ **Type:** boolean
- ▲ **Required:** Y

Enables/disables speculative aggregates for the repository.

When enabled, the AtScale engine automatically creates aggregate tables that it anticipates being useful based on your models. These are intended to improve the performance of queries from client BI tools faster than with demand-defined aggregates alone.



Note: Speculative aggregates are also called prediction-defined aggregates.

Supported values:

- ▲ true
- ▲ false

For more information on speculative aggregates, see [Types of Aggregate Tables in AtScale](#).

Dataset_properties

- ▲ **Type:** object
- ▲ **Required:** N

Defines dataset properties to use within the repository.

Supported properties:

- ▶ `allow_aggregates` : Boolean, optional. Enables the AtScale engine to create aggregates for datasets in the repository.
- ▶ `allow_local_aggs` : Boolean, optional. Enables local aggregation for datasets in the repository.
- ▶ `allow_peer_aggs` : Boolean, optional. Enables aggregation on data derived from datasets in data warehouses that are different from the source dataset.

Specify the `unique_name` of the dataset followed by the properties and values you want to set for it at the repository level. For example:

```
dataset1:  
  create_hinted_aggregate: true
```



Note: Datasets are typically defined at the repository level, in `atscale.yml`; however, datasets used by a specific model (typically fact datasets) can be defined within the model itself. For more information, see [Models](#).