

Aggregate Data Security

You can configure AtScale so that queries may be simultaneously accessing data from across multiple physical locations. By default, AtScale may create copies or aggregations of data across data warehouses to improve query performance.

You can also separately configure whether data from a dataset may be aggregated in the local data warehouse, in peer data warehouses, or in the [Preferred Aggregate Storage](#). If you need to disable local aggregations on specific datasets, or limit data to its original physical location for security reasons, you will edit a dataset's **Aggregate Locality** settings.

Before You Begin

Before setting aggregate locality, you must have:

- ▶ A cube draft that already includes modeled fact and dimension datasets
- ▶ Permission to update a cube

Editing Aggregate Locality

To edit aggregate locality for a dimensional dataset:

1. In the Cube Designer, double click the header for any dimension on the cube canvas. The dimension canvas opens.
2. On the dimension canvas, click the gear icon for that dataset and select "Edit Dataset".
3. Scroll down to the **Aggregate Preferences** section.

The aggregate locality options will only be displayed if the "Enable aggregate production from this table" option is enabled. Toggle the settings as desired:

- ▶ **Allow aggregates to be built locally:** Enables local aggregation on the dataset selected.
- ▶ **Allow aggregates to be built in peer datasources:** Enables aggregation on data derived from the dataset on data warehouses that are different from the source dataset.
- ▶ **Allow aggregates to be built in preferred storage:** Enables AtScale to promote aggregates from this dataset to the Preferred Aggregate Storage location.

To edit aggregate locality for a **fact dataset**

1. In the Cube Designer, double click the header for any fact dataset and select "Edit Dataset."
2. Scroll down to the **Aggregate Preferences** section.

The aggregate locality options will only show up if the "Enable aggregate production from this table" option is enabled. Toggle the settings as desired:

- ▶ **Allow aggregates to be built locally:** Enables local aggregation on the dataset selected.
- ▶ **Allow aggregates to be built in peer datasources:** Enables aggregation on data derived from the dataset on data warehouses that are different from the source dataset.
- ▶ **Allow aggregates to be built in preferred storage:** Enables AtScale to promote aggregates from this dataset to the Preferred Aggregate Storage location.

Figure 1. The Aggregate Preferences section of the Edit a Dataset dialog.

