

About Perspectives

Perspectives are publishable subsets of the data model for an underlying AtScale cube. They are meant to make it easier for analysts to query only the subset of data in a virtual cube that is relevant to their purposes or responsibilities. Rather than provide analysts with the entire data model for a cube, you can make dimensions, hierarchies, levels, secondary attributes, measures and calculated measures in that cube invisible to them.

For example, suppose that you worked for a retailer as a data modeler. Four business analysts need access to company data to create reports: one report related to inventory, one related to marketing, and the last related to orders. Rather than create three separate data models in three separate cubes, you can create one data model and then create three perspectives on that cube, one perspective for each of the analysts. So, you create your cube and then create the perspectives that are described in the next two tables.

Table 1. Fact-table columns that are visible in the different perspectives that are based on your cube

Fact-Table Columns	Perspective: Inventory	Perspective: Marketing	Perspective: Orders
Clicks		✓	
Quantity	✓		✓
Sales Amount			✓
Spend		✓	
Tax			✓
Views	✓		
Wholesale Price	✓		

Table 2. Dimensions that are visible in the different perspectives that are based on your cube

Dimensions	Perspective: Inventory	Perspective: Marketing	Perspective: Orders
Color	✓		
Customer		✓	✓
Date	✓	✓	✓
Product	✓	✓	✓
Size	✓		

You can create perspectives in the **Perspectives** panel in the Cube Designer canvas, represented in the canvas toolbar by an icon of a pair of binoculars. Click the binoculars icon to open the panel. In the panel, click the menu icon and choose **Add Perspective** to create a new perspective. The panel also lists perspectives that have already been created on the cube.

Then, rather than grant permission to the three analysts to query the entire cube, you grant permission to each analyst to query only the perspective that the analyst needs to build a report. You can add users to a perspective from the cube's overview page. The perspectives are listed under the **Perspectives** tab on that page. To access the security dialog for that perspective, click the ellipsis icon next to its name and select **Security**.



Note: Perspectives are not secure against users querying entities in the cube that are not visible. If someone who has permission to query a perspective knows the names of other entities that are in the underlying cube but hidden from the perspective, that person can write an MDX query that runs against the perspective but references those other entities. Perspectives control visibility of items in a cube and are not meant to be used as a security measure.

When the perspectives are ready, you publish the project to make them available to the business analysts. After you do that, the perspectives are listed under their base cubes in the **Publish Status** section. As with cubes, the business analysts can get the .tds file or other connection information for a perspective by clicking the down arrow next to the perspective.

As the analysts query the perspectives, the AtScale engine creates demand-defined aggregates for the base cube. Demand-defined aggregates that are created based on one or more queries on a perspective can also be used to satisfy queries against the perspective's underlying cube. Likewise, demand-defined aggregates that are created based on one or more queries on a cube can be used to satisfy queries against perspectives on that cube. Moreover, when you published your project, the AtScale engine created prediction-defined aggregates, as usual, that are based on the cube. These aggregates can also be used by the engine to satisfy queries against the perspectives. In fact, it is also possible for any user-defined aggregates defined for the cube satisfy queries against perspectives.

If an analyst requests a drill-through that uses attributes or measures that are not part of the perspective that the analyst is using, the AtScale engine removes those attributes and measures from the result set.

When you monitor aggregate tables on the **Definitions** and **Instances**, you notice that the entries always refer to the base cube for the perspectives. When you monitor queries on the **Queries** page, you notice that only inbound queries always refer to the perspective that is being queried. Therefore, on these pages, it isn't possible to search for queries, aggregate definitions, or aggregate instances by perspective.

Eventually, you realize that you need to modify one of the perspectives. You edit it by opening the **Perspectives** panel in the Cube Designer canvas and double-clicking the name of the perspective. However, changes made to a perspective are not available to people querying the perspective until the following steps are taken:

1. The corresponding AtScale project is republished.
2. Users who are currently connected to the perspective close the workbook, sheets, or other workspaces in which they are querying the perspective and then reconnect to the perspective. For instance, users of Tableau Desktop cannot just download a new .tds file and then refresh their connection; they must close their workbook and then reopen it.

Related tasks

[Adding Perspectives](#)

[Granting Runtime Permissions on Perspectives](#)

[Editing Perspectives](#)

[Deleting Perspectives](#)