

JDBC Client Connections

Java Database Connectivity (JDBC) is an API standard that allows Java client applications to connect to a relational database and issue SQL queries using a standard interface. This section explains how to connect to AtScale from JDBC-compliant client applications.

To applications that send SQL queries, AtScale uses the same protocols and drivers as a remote HiveServer2 instance. See the [Apache Hive Wiki](#) for general information on Hive JDBC client connections.

Connecting to AtScale from a Java client is the same as connecting to Hive, however you connect to the AtScale engine. The JDBC connection URL points to a particular AtScale project, for example:

```
jdbc:hive2://atscale_host:11111/project_name
```

To issue an SQL query to AtScale, the SQL table name is the name of the cube and the column names are the cube attribute's query names. For example (notice the back-ticks to escape table and column names containing spaces):

```
SELECT `Internet Sales Cube`.`Order Month`,  
       SUM(Internet Sales Cube.orderquantity) AS `Order Quantity`  
FROM `Sales Insights`.`Internet Sales Cube`  
GROUP BY `Internet Sales Cube`.`Order Month`
```