

Rebuilding Aggregates Using The REST API

You can use the `aggregate-batch` endpoint of the AtScale engine REST API to trigger an initial build or rebuild for all aggregates of a published cube. The Design Center provides a sample Bash script that runs curl commands to authenticate with and post a request to the `aggregate-batch` endpoint for the cube.

One use case for this feature is the need to rebuild aggregates after ETL (Extract, Transform, Load) jobs. Many companies have ETL jobs that they routinely run in Hadoop to cleanse raw data and load that data into Hive. You can use AtScale's REST API to programmatically rebuild aggregates for a cube as a part of your ETL processing.

Before You Begin

- ▶ The user connecting to the API to trigger the aggregate build **must** have the "Manage Aggregates" permission. Assign the user the "Manage Aggregates" permission by creating a role with this permission present, or assigning the user an existing role that has this permission enabled.
- ▶ The user accessing the Build tab should have the `Login`, `View Aggregates`, and `Publish Projects` permissions. For more information, see [Creating and Editing Roles](#).

Procedure

1. At the top of the Design Center, click **Projects**.
2. In the list of published projects, select the project containing the cube you want to rebuild aggregates for.
3. Select the published version of the cube associated with the project, and select the **Build** tab. From the drop-down menu, select **Triggers**. The sample Bash script appears in the API Trigger Command dialog.

Detail	Description
Endpoint	<code>aggregate-batch/org_id/project_id</code>
URL Parameters	<p><code>org_id</code> - The AtScale organization ID</p> <p><code>project_id</code> - The AtScale project ID</p>
Query Parameters	<code>cubeID</code> - the AtScale cube ID
HTTP Method	POST
	Each API call to the AtScale engine must pass in the <code>client-api-key</code> field in the HTTP request header.

Authentication	The client API key value for your installation is shown in the example request.
Return Responses	201 - Created / Success 400 - Bad Request 401 - Unauthorized 404 - Not Found 500 - Server Error 503 - Server Unavailable

What To Do Next

After an initial build or rebuild is triggered, you can go to the **Aggregates > Build History** page to view the results. You can also open this page from the menu item **View History** in the menu for a cube.

This page shows whether builds are queued, running, or successful. It also shows whether builds failed.

If a build is running, you can cancel it, if you need to, by clicking **Cancel Build** in the entry for the build.

In an entry for a build, you can click the **View Instances** link to find out information about the instances that were or were not built.

Related concepts

[How Aggregate Tables Are Populated With Data](#)