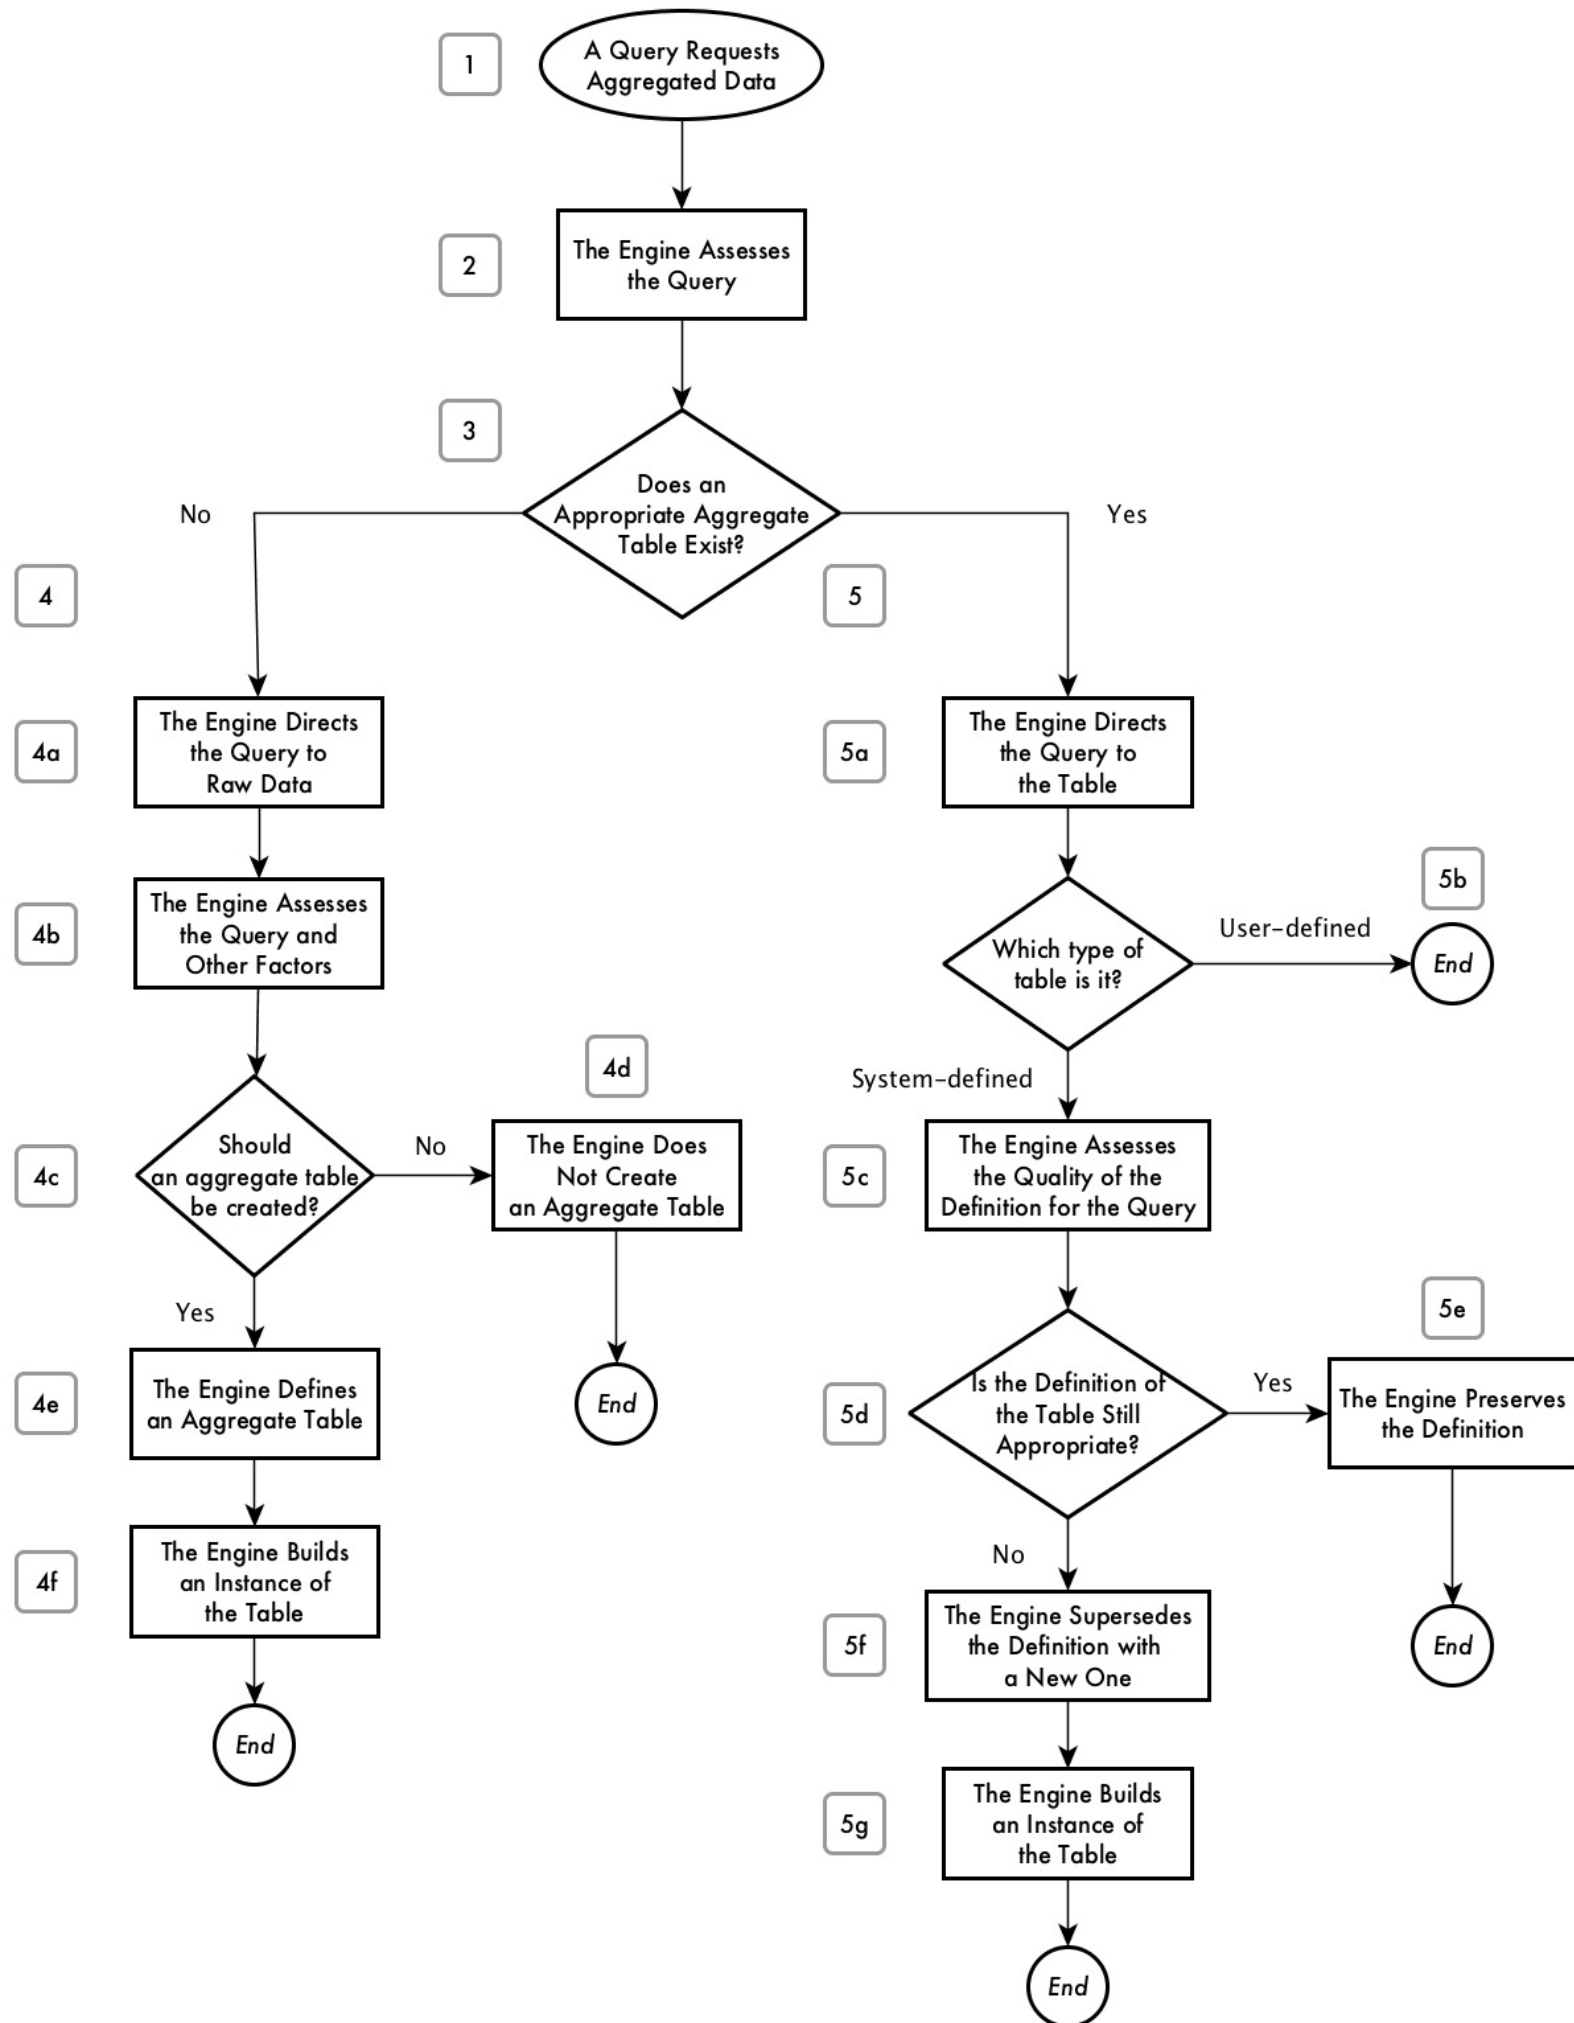


Flowchart Of The Actions Taken For Each Query

Each query that requests aggregated data causes the AtScale engine to find out whether an aggregate table exists that will satisfy the query. There are two possible outcomes of this search: an appropriate aggregate table does not exist and an appropriate aggregate table does exist. This flowchart shows the actions that the engine takes for both outcomes.



1. A BI client application issues a query on aggregated data.

2. The AtScale engine receives the query and examines it.
3. The engine determines whether an aggregate table (either system-defined or user-defined) that can satisfy the query exists.
4. If an appropriate aggregate table does not exist, the engine takes these actions:
 1. Directs the query against the data in the cluster.
 2. Assesses the need for an aggregate table to satisfy the query and similar queries in the future.
 3. Reaches a decision on whether to create an aggregate table.
 4. If the decision is no, the actions triggered by the query are complete and the workflow ends.
 5. If the decision is yes, the engine defines an aggregate table.
 6. The next build process for the cube builds an instance of the new aggregate table. The workflow ends.
5. If an appropriate aggregate table does exist, the engine takes these actions:
 1. Directs the query to the current instance of the aggregate table, if the instance is active. The actions that the engine takes to run the query are complete.
 2. If the aggregate table is user-defined, the workflow ends.
 3. If the aggregate table is system-defined, assesses the quality of the definition for the query and future similar queries, determining whether a better aggregate table could be created.
 4. Decides whether or not the definition is still sufficient.
 5. If the definition is sufficient, preserves the definition. The workflow ends.
 6. If the definition is not sufficient, supersedes the aggregate-table definition with a new one. Retains the previous definition for comparisons and that definitions instances for statistics. Eventually, both the definition and instances are deleted.
 7. The next build process for the cube builds an instance of the new aggregate table. The actions triggered by the query are complete.