Orchestration (17.1)

Workflow management is an integral built-in part of Agile Data Engine and no administration or configuration effort is required for the end users.

Agile Data Engine uses an embedded Airflow instance to execute the workflows (per environment).

Reusable workflows definitions are configurable in the UI, workflows are built on the go and code generation is fully automated and the definitions form part of the internal CI/CD pipeline.

ADE determines the sequence of the individual processing steps based on the upstream entity loads in the chain, taking into account the interdependencies between them. It is easy to dynamically switch between various execution paths, i.e. execute a chain of dependent loads for a source object, isolate executions by layer or flexibly create multi-dependent loads.

Orchestration (17.2)

Workflow generation modes

ADE controls the workflow load processes by using slots (threads) where the level of parallelism and maximum database workload can be managed.

ADE provides three workflow generation modes to optimize the execution overhead which might suit best an individual data load:

- Load-oriented - created as load dependencies are defined in ADE
- Entity-oriented - run within an entity node in a sequential manner, i.e. multiple loads are executed sequentially for the same target
- Sequential – all loads are executed sequentially. It minimizes scheduling overhead for near real-time cycles but limits parallelism and monitoring

Orchestration (17.3) - DDVUG data

- Two workflows were defined for the test datasets on staging level - DDVUG_WEBSHOP_DB and DDVUG_ROADSHOW_DB
- All the downstream loads were generated automatically by ADE, following the data load dependencies logic:

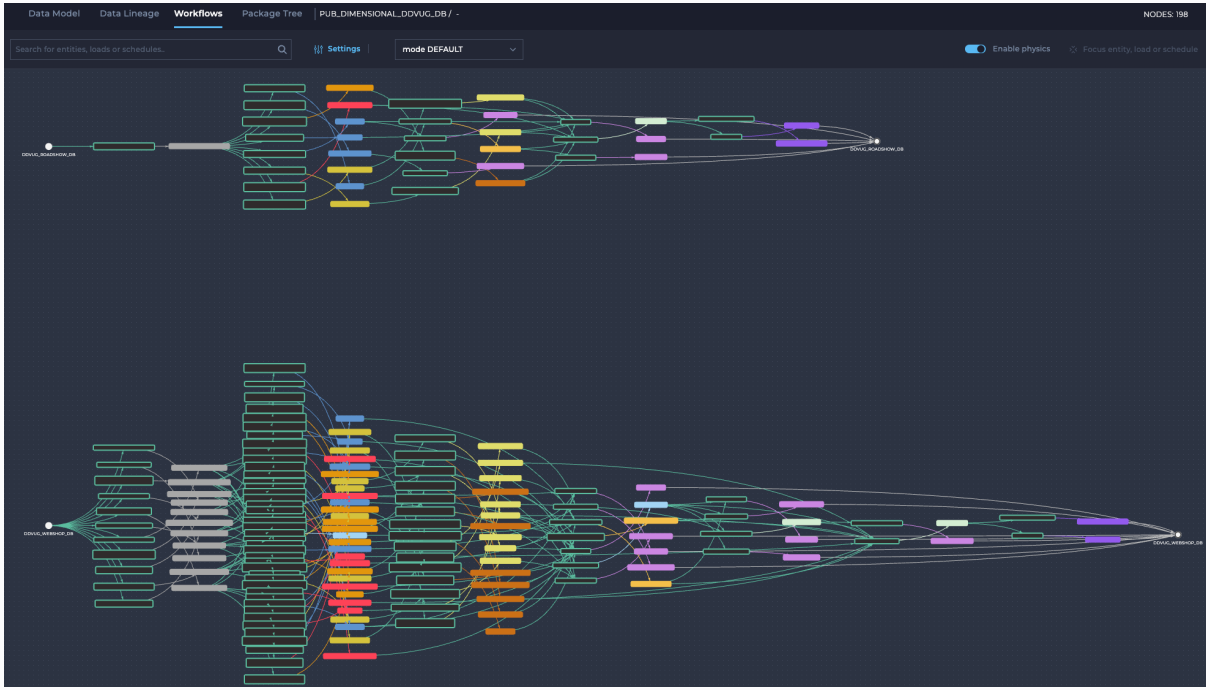

Orchestration (17.4) - Flexibility

- Let’s say for example we want to define WEBSHOP as a driving schedule for F_ORDER_LINE fact table. It’s as easy as assigning a schedule on a load level:

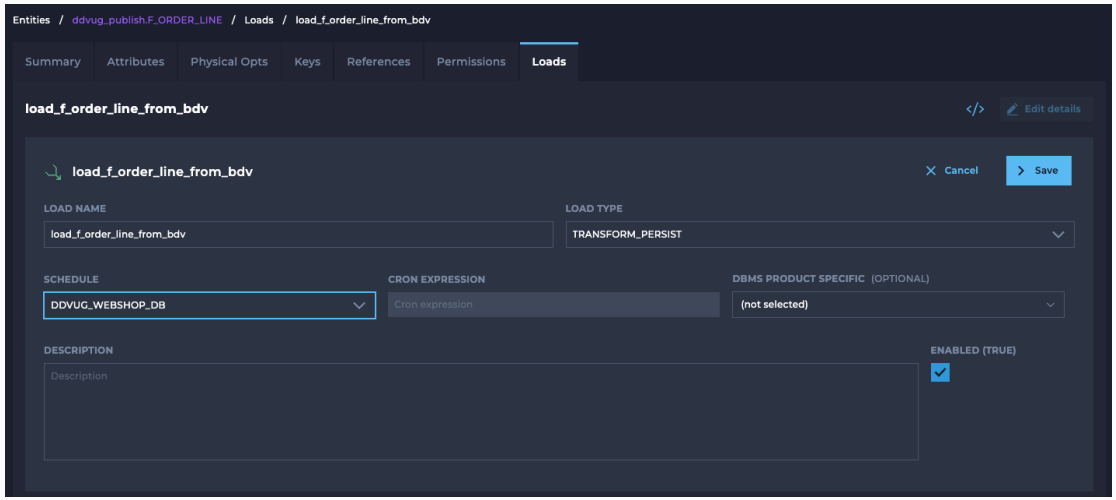

- In this case the F_ORDER_LINE fact table will be refereshed only if new WEBSHOP data arrived
- Can be easily and flexibly modified on the fly