

Package ‘sqltargets’

July 31, 2024

Type Package

Title 'Targets' Extension for 'SQL' Queries

Version 0.2.0

Maintainer David Ranzolin <daranzolin@gmail.com>

Description Provides an extension for 'SQL' queries as separate file within 'targets' pipelines. The shorthand creates two targets, the query file and the query result.

License MIT + file LICENSE

Encoding UTF-8

Imports cli, DBI, fs, glue, jinjar, purrr, readr, rlang, stringr, tarchetypes, targets, withr

URL <https://github.com/daranzolin/sqltargets>

BugReports <https://github.com/daranzolin/sqltargets/issues>

RoxygenNote 7.2.1

Suggests testthat (>= 3.0.0), RSQLite (>= 2.2.0)

Config/testthat/edition 3

NeedsCompilation no

Author David Ranzolin [aut, cre, cph]

Repository CRAN

Date/Publication 2024-07-31 02:30:02 UTC

Contents

sqltargets	2
sqltargets_option_get	2
tar_sql	3
tar_sql_deps	7
tar_sql_raw	7

Index	11
--------------	-----------

sqltargets	sqltargets <i>package</i>
------------	---------------------------

Description

targets extension for SQL files

Details

See the README on [GitHub](#)

sqltargets_option_get	<i>Get or Set sqltargets Options</i>
-----------------------	--------------------------------------

Description

Get or Set sqltargets Options

Usage

```
sqltargets_option_get(option_name)
```

```
sqltargets_option_set(option_name, option_value)
```

Arguments

option_name	Character. Option name. See Details.
option_value	Value to assign to option 'x'.

Details

Available Options

- "sqltargets.template_engine" - Either 'glue' or 'jinjar'. Determines how the query file should be parsed.
- "sqltargets.glue_sql_opening_delimiter" - character. Length 1. The opening delimiter passed to 'glue::glue_sql()'.
- "sqltargets.glue_sql_closing_delimiter" - character. Length 1. The closing delimiter passed to 'glue::glue_sql()'.
- "sqltargets.jinja_block_open" - character. Length 1. The opening delimiter passed to 'jinjar::jinjar_config()'.
- "sqltargets.jinja_block_close" - character. Length 1. The closing delimiter passed to 'jinjar::jinjar_config()'.
- "sqltargets.jinja_variable_open" - character. Length 1. The closing delimiter passed to 'jinjar::jinjar_config()'.

- "sqltargets.jinja_variable_close" - character. Length 1. The closing delimiter passed to 'jinjar::jinjar_config()'.
- "sqltargets.jinja_comment_open" - character. Length 1. The closing delimiter passed to 'jinjar::jinjar_config()'.
- "sqltargets.jinja_comment_close" - character. Length 1. The closing delimiter passed to 'jinjar::jinjar_config()'.

Value

No return value, called for side effects

tar_sql	<i>Target with a SQL query.</i>
---------	---------------------------------

Description

Shorthand to include a SQL query in a 'targets' pipeline.

Usage

```
tar_sql(
  name,
  path,
  params = list(),
  format = targets::tar_option_get("format"),
  tidy_eval = targets::tar_option_get("tidy_eval"),
  repository = targets::tar_option_get("repository"),
  iteration = targets::tar_option_get("iteration"),
  error = targets::tar_option_get("error"),
  memory = targets::tar_option_get("memory"),
  garbage_collection = targets::tar_option_get("garbage_collection"),
  deployment = targets::tar_option_get("deployment"),
  priority = targets::tar_option_get("priority"),
  resources = targets::tar_option_get("resources"),
  storage = targets::tar_option_get("storage"),
  retrieval = targets::tar_option_get("retrieval"),
  cue = targets::tar_option_get("cue")
)
```

Arguments

name	Symbol, name of the target. A target name must be a valid name for a symbol in R, and it must not start with a dot. Subsequent targets can refer to this name symbolically to induce a dependency relationship: e.g. <code>tar_target(downstream_target, f(upstream_target))</code> is a target named <code>downstream_target</code> which depends on a target <code>upstream_target</code> and a function <code>f()</code> . In addition, a target's name determines its random number generator seed. In this way, each target runs with
------	--

a reproducible seed so someone else running the same pipeline should get the same results, and no two targets in the same pipeline share the same seed. (Even dynamic branches have different names and thus different seeds.) You can recover the seed of a completed target with `tar_meta(your_target, seed)` and run `set.seed()` on the result to locally recreate the target's initial RNG state.

path	Character of length 1 to the single <code>*.sql</code> source file to be executed. Defaults to the working directory of the <code>'targets'</code> pipeline.
params	Code, can be <code>'NULL'</code> . <code>'params'</code> evaluates to a named list of parameters that are passed to <code>'jinja::render()'</code> . The list is quoted (not evaluated until the target runs) so that upstream targets can serve as parameter values.
format	Optional storage format for the target's return value. With the exception of <code>format = "file"</code> , each target gets a file in <code>_targets/objects</code> , and each format is a different way to save and load this file. See the "Storage formats" section for a detailed list of possible data storage formats.
tidy_eval	Logical, whether to enable tidy evaluation when interpreting command and pattern. If <code>TRUE</code> , you can use the "bang-bang" operator <code>!!</code> to programmatically insert the values of global objects.
repository	Character of length 1, remote repository for target storage. Choices: <ul style="list-style-type: none"> <code>"local"</code>: file system of the local machine. <code>"aws"</code>: Amazon Web Services (AWS) S3 bucket. Can be configured with a non-AWS S3 bucket using the <code>endpoint</code> argument of <code>tar_resources_aws()</code>, but versioning capabilities may be lost in doing so. See the cloud storage section of https://books.ropensci.org/targets/data.html for details for instructions. <code>"gcp"</code>: Google Cloud Platform storage bucket. See the cloud storage section of https://books.ropensci.org/targets/data.html for details for instructions. <p>Note: if <code>repository</code> is not <code>"local"</code> and <code>format</code> is <code>"file"</code> then the target should create a single output file. That output file is uploaded to the cloud and tracked for changes where it exists in the cloud. The local file is deleted after the target runs.</p>
iteration	Character of length 1, name of the iteration mode of the target. Choices: <ul style="list-style-type: none"> <code>"vector"</code>: branching happens with <code>vctrs::vec_slice()</code> and aggregation happens with <code>vctrs::vec_c()</code>. <code>"list"</code>, branching happens with <code>[[]]</code> and aggregation happens with <code>list()</code>. <code>"group"</code>: <code>dplyr::group_by()</code>-like functionality to branch over subsets of a data frame. The target's return value must be a data frame with a special <code>tar_group</code> column of consecutive integers from 1 through the number of groups. Each integer designates a group, and a branch is created for each collection of rows in a group. See the <code>tar_group()</code> function to see how you can create the special <code>tar_group</code> column with <code>dplyr::group_by()</code>.
error	Character of length 1, what to do if the target stops and throws an error. Options: <ul style="list-style-type: none"> <code>"stop"</code>: the whole pipeline stops and throws an error. <code>"continue"</code>: the whole pipeline keeps going.

- "abridge": any currently running targets keep running, but no new targets launch after that. (Visit <https://books.ropensci.org/targets/debugging.html> to learn how to debug targets using saved workspaces.)
- "null": The errored target continues and returns NULL. The data hash is deliberately wrong so the target is not up to date for the next run of the pipeline.

memory Character of length 1, memory strategy. If "persistent", the target stays in memory until the end of the pipeline (unless storage is "worker", in which case targets unloads the value from memory right after storing it in order to avoid sending copious data over a network). If "transient", the target gets unloaded after every new target completes. Either way, the target gets automatically loaded into memory whenever another target needs the value. For cloud-based dynamic files (e.g. format = "file" with repository = "aws"), this memory strategy applies to the temporary local copy of the file: "persistent" means it remains until the end of the pipeline and is then deleted, and "transient" means it gets deleted as soon as possible. The former conserves bandwidth, and the latter conserves local storage.

garbage_collection Logical, whether to run `base::gc()` just before the target runs.

deployment Character of length 1, only relevant to `tar_make_clustermq()` and `tar_make_future()`. If "worker", the target builds on a parallel worker. If "main", the target builds on the host machine / process managing the pipeline.

priority Numeric of length 1 between 0 and 1. Controls which targets get deployed first when multiple competing targets are ready simultaneously. Targets with priorities closer to 1 get built earlier (and polled earlier in `tar_make_future()`).

resources Object returned by `tar_resources()` with optional settings for high-performance computing functionality, alternative data storage formats, and other optional capabilities of targets. See `tar_resources()` for details.

storage Character of length 1, only relevant to `tar_make_clustermq()` and `tar_make_future()`. Must be one of the following values:

- "main": the target's return value is sent back to the host machine and saved/uploaded locally.
- "worker": the worker saves/uploads the value.
- "none": almost never recommended. It is only for niche situations, e.g. the data needs to be loaded explicitly from another language. If you do use it, then the return value of the target is totally ignored when the target ends, but each downstream target still attempts to load the data file (except when `retrieval = "none"`).

If you select `storage = "none"`, then the return value of the target's command is ignored, and the data is not saved automatically. As with dynamic files (`format = "file"`) it is the responsibility of the user to write to the data store from inside the target.

The distinguishing feature of `storage = "none"` (as opposed to `format = "file"`) is that in the general case, downstream targets will automatically try to load the data from the data store as a dependency. As a corollary, `storage = "none"` is completely unnecessary if `format` is "file".

retrieval	Character of length 1, only relevant to <code>tar_make_clustermq()</code> and <code>tar_make_future()</code> . Must be one of the following values: <ul style="list-style-type: none"> "main": the target's dependencies are loaded on the host machine and sent to the worker before the target builds. "worker": the worker loads the targets dependencies. "none": the dependencies are not loaded at all. This choice is almost never recommended. It is only for niche situations, e.g. the data needs to be loaded explicitly from another language.
cue	An optional object from <code>tar_cue()</code> to customize the rules that decide whether the target is up to date.

Details

'tar_sql()' is an alternative to 'tar_target()' for SQL queries that depend on upstream targets. The SQL source files (*.sql files) should mention dependency targets with 'tar_load()' within SQL comments ('--'). (Do not use 'tar_load_raw()' or 'tar_read_raw()' for this.) Then, 'tar_sql()' defines a special kind of target. It 1. Finds all the 'tar_load()'/ 'tar_read()' dependencies in the query and inserts them into the target's command. This enforces the proper dependency relationships. (Do not use 'tar_load_raw()' or 'tar_read_raw()' for this.) 2. Sets 'format = "file"' (see 'tar_target()') so 'targets' watches the files at the returned paths and reruns the query if those files change. 3. Creates another upstream target to watch the query file for changes '<target name>sqltargets_option_get("sqltargets.target_file_suffix")'.

Value

A data frame

Examples

```
targets::tar_dir({ # tar_dir() runs code from a temporary directory.
  # Unparameterized SQL query:
  lines <- c(
    "-- !preview conn=DBI::dbConnect(RSQLite::SQLite())",
    "-- targets::tar_load(data1)",
    "-- targets::tar_load(data2)",
    "select 1 AS my_col",
    ""
  )
  # In tar_dir(), not part of the user's file space:
  writeLines(lines, "query.sql")
  # Include the query in a pipeline as follows.
  targets::tar_script({
    library(tarchetypes)
    library(sqltargets)
    list(
      tar_sql(query, path = "query.sql")
    )
  }, ask = FALSE)
})
```

tar_sql_deps	<i>List SQL query dependencies.</i>
--------------	-------------------------------------

Description

List the target dependencies of one or more SQL queries.

Usage

```
tar_sql_deps(path)
```

Arguments

path Character vector, path to one or more SQL queries.

Value

Character vector of the names of targets that are dependencies of the SQL query.

Examples

```
lines <- c(
  "-- !preview conn=DBI::dbConnect(RSQLite::SQLite())",
  "-- targets::tar_load(data1)",
  "-- targets::tar_read(data2)",
  "select 1 as my_col",
  ""
)
query <- tempfile()
writeLines(lines, query)
tar_sql_deps(query)
```

tar_sql_raw	<i>Target with a SQL query.</i>
-------------	---------------------------------

Description

Shorthand to include a SQL query in a 'targets' pipeline.

Usage

```
tar_sql_raw(
  name,
  path = ".",
  params = params,
  format = format,
  error = targets::tar_option_get("error"),
  memory = targets::tar_option_get("memory"),
  garbage_collection = targets::tar_option_get("garbage_collection"),
  deployment = "main",
  priority = targets::tar_option_get("priority"),
  resources = targets::tar_option_get("resources"),
  retrieval = targets::tar_option_get("retrieval"),
  cue = targets::tar_option_get("cue"),
  params_nm = NULL
)
```

Arguments

name	Character of length 1, name of the target. A target name must be a valid name for a symbol in R, and it must not start with a dot. Subsequent targets can refer to this name symbolically to induce a dependency relationship: e.g. <code>tar_target(downstream_target, f(upstream_target))</code> is a target named <code>downstream_target</code> which depends on a target <code>upstream_target</code> and a function <code>f()</code> . In addition, a target's name determines its random number generator seed. In this way, each target runs with a reproducible seed so someone else running the same pipeline should get the same results, and no two targets in the same pipeline share the same seed. (Even dynamic branches have different names and thus different seeds.) You can recover the seed of a completed target with <code>tar_meta(your_target, seed)</code> and run <code>set.seed()</code> on the result to locally recreate the target's initial RNG state.
path	Character of length 1 to the single <code>*.sql</code> source file to be executed. Defaults to the working directory of the <code>'targets'</code> pipeline.
params	Code, can be <code>'NULL'</code> . <code>'params'</code> evaluates to a named list of parameters that are passed to <code>'jinja::render()'</code> . The list is quoted (not evaluated until the target runs) so that upstream targets can serve as parameter values.
format	Optional storage format for the target's return value. With the exception of <code>format = "file"</code> , each target gets a file in <code>_targets/objects</code> , and each format is a different way to save and load this file. See the "Storage formats" section for a detailed list of possible data storage formats.
error	Character of length 1, what to do if the target stops and throws an error. Options: <ul style="list-style-type: none"> <code>"stop"</code>: the whole pipeline stops and throws an error. <code>"continue"</code>: the whole pipeline keeps going. <code>"abridge"</code>: any currently running targets keep running, but no new targets launch after that. (Visit https://books.ropensci.org/targets/debugging.html to learn how to debug targets using saved workspaces.)

- "null": The errored target continues and returns NULL. The data hash is deliberately wrong so the target is not up to date for the next run of the pipeline.

memory	Character of length 1, memory strategy. If "persistent", the target stays in memory until the end of the pipeline (unless storage is "worker", in which case targets unloads the value from memory right after storing it in order to avoid sending copious data over a network). If "transient", the target gets unloaded after every new target completes. Either way, the target gets automatically loaded into memory whenever another target needs the value. For cloud-based dynamic files (e.g. format = "file" with repository = "aws"), this memory strategy applies to the temporary local copy of the file: "persistent" means it remains until the end of the pipeline and is then deleted, and "transient" means it gets deleted as soon as possible. The former conserves bandwidth, and the latter conserves local storage.
garbage_collection	Logical, whether to run <code>base::gc()</code> just before the target runs.
deployment	Character of length 1, only relevant to <code>tar_make_clustermq()</code> and <code>tar_make_future()</code> . If "worker", the target builds on a parallel worker. If "main", the target builds on the host machine / process managing the pipeline.
priority	Numeric of length 1 between 0 and 1. Controls which targets get deployed first when multiple competing targets are ready simultaneously. Targets with priorities closer to 1 get built earlier (and polled earlier in <code>tar_make_future()</code>).
resources	Object returned by <code>tar_resources()</code> with optional settings for high-performance computing functionality, alternative data storage formats, and other optional capabilities of targets. See <code>tar_resources()</code> for details.
retrieval	Character of length 1, only relevant to <code>tar_make_clustermq()</code> and <code>tar_make_future()</code> . Must be one of the following values: <ul style="list-style-type: none"> • "main": the target's dependencies are loaded on the host machine and sent to the worker before the target builds. • "worker": the worker loads the targets dependencies. • "none": the dependencies are not loaded at all. This choice is almost never recommended. It is only for niche situations, e.g. the data needs to be loaded explicitly from another language.
cue	An optional object from <code>tar_cue()</code> to customize the rules that decide whether the target is up to date.
params_nm	Character of length 1, name of object passed to 'params'.

Details

'tar_sql()' is an alternative to 'tar_target()' for SQL queries that depend on upstream targets. The SQL source files (*.sql files) should mention dependency targets with 'tar_load()' within SQL comments ('-'). (Do not use 'tar_load_raw()' or 'tar_read_raw()' for this.) Then, 'tar_sql()' defines a special kind of target. It 1. Finds all the 'tar_load()'/ 'tar_read()' dependencies in the query and inserts them into the target's command. This enforces the proper dependency relationships. (Do not use 'tar_load_raw()' or 'tar_read_raw()' for this.) 2. Sets 'format = "file"' (see 'tar_target()') so 'targets' watches the files at the returned paths and reruns the query if those files

change. 3. Creates another upstream target to watch the query file for changes '`<target name>`' `'sqltargets_option_get("sqltargets.target_file_suffix")'`.

Value

A data frame

Examples

```
targets::tar_dir({ # tar_dir() runs code from a temporary directory.
  # Unparameterized SQL query:
  lines <- c(
    "-- !preview conn=DBI::dbConnect(RSQLite::SQLite())",
    "-- targets::tar_load(data1)",
    "-- targets::tar_load(data2)",
    "select 1 AS my_col",
    ""
  )
  # In tar_dir(), not part of the user's file space:
  writeLines(lines, "query.sql")
  # Include the query in a pipeline as follows.
  targets::tar_script({
    library(tarchetypes)
    library(sqltargets)
    list(
      tar_sql(query, path = "query.sql")
    )
  }, ask = FALSE)
})
```

Index

* SQL query utilities

tar_sql_deps, 7

sqltargets, 2

sqltargets_option_get, 2

sqltargets_option_set
(sqltargets_option_get), 2

tar_group(), 4

tar_make_clustermq(), 5, 6, 9

tar_make_future(), 5, 6, 9

tar_resources_aws(), 4

tar_sql, 3

tar_sql_deps, 7

tar_sql_raw, 7