
Resolve

Release 38.3.2.dev2+gdd3d4dd


Genialis, Inc.

Jan 25, 2024

CONTENTS

1	Contents	3
1.1	Overview	3
1.2	Getting started	4
1.3	Writing processes	4
1.4	API	16
1.5	Type extension composition	18
1.6	Reference	19
1.7	Resolve Flow Design	61
1.8	Resolve Storage Framework	64
1.9	Change Log	66
1.10	Contributing	113
	Python Module Index	117
	Index	119


Resolve is an open source dataflow package for [Django framework](#). It offers a complete RESTful API to connect with external resources. A higher layer of convenience APIs for JavaScript, Python and R are in development. A collection of bioinformatics pipelines is available within the [Resolve Bioinformatics](#) project. We envision a similar toolkit for machine learning.



CONTENTS

1.1 Overview


Resolve consists of two major components: a RESTful API and the Flow Engine. The RESTful API is based on the [Django REST Framework](#) and offers complete control over the workflow, the data involved and the permissions on those data. The Resolve Flow engine, on the other hand, handles pipeline execution. It resolves dependencies between *processes* (jobs or tasks), and executes them on worker nodes. Results are saved to a PostgreSQL database and a clustered file system.



The Flow Engine has several layers of execution that can be configured either on the server or by the individual processes.



Processes can be executed on a server cluster. In this case the Executor, Runtime and Expression Engine layers span multiple worker nodes.



Resolve can be configured for lightweight desktop use (*e.g.*, by bioinformatics professionals) or deployed as a complex set-up of multiple servers and worker nodes. In addition to the components described above, customizing the configuration of the web server (*e.g.*, NGINX or Apache HTTP), workload manager, and the database offer high scaling potential.

Example of a lightweight configuration: synchronous workload manager that runs locally, Docker executor and runtime, Django web server, and local file system.

Example of a complex deploy: Slurm workload manager with a range of computational nodes, Docker executor and runtime on each worker node, NGINX web server, and a fast file system shared between worker nodes.

1.2 Getting started

TODO: Write about how to include Resolve in a Django project and explain settings parameters. Create an example Django project in the docs/example folder and give code references where a detailed explanation is needed.

1.3 Writing processes

Process is a central building block of the Resolve’s dataflow. Formally, a process is an algorithm that transforms inputs to outputs. For example, a *Word Count* process would take a text file as input and report the number of words on the output.




Fig. 1: *Word Count* process with input `doc` of type `basic:file` and output `words` of type `basic:integer`.

When you execute the process, Resolve creates a new `Data` object with information about the process instance. In this case the *document* and the *words* would be saved to the same `Data` object. What if you would like to execute another analysis on the same document, say count the number of lines? We could create a similar process *Number of Lines* that would also take the file and report the number of lines. However, when we would execute the process we would have 2 copies of the same *document* file stored on the platform. In most cases it makes sense to split the upload (data storage) from the analysis. For example, we could create 3 processes: *Upload Document*, *Word Count* and *Number of Lines*.




Fig. 2: Separate the data storage (*Upload Document*) and analysis (*Word Count*, *Number of Lines*). Notice that the *Word Count* and *Number of Lines* processes accept `Data` objects of type `data:doc`—the type of the *Upload Document* process.

Resolve handles the execution of the dataflow automatically. If you were to execute all three processes at the same time, Resolve would delay the execution of *Word Count* and *Number of Lines* until the completion of *Upload Document*. Resolve resolves dependencies between processes.

A process is defined by:

- Inputs
- Outputs
- Meta-data
- Algorithm

Processes are stored in the data base in the `Process` model. A process' algorithm runs automatically when you create a new `Data` object. The inputs and the process name are required at `Data` create, the outputs are saved by the algorithm, and users can update the meta-data at any time. The *Process syntax* chapter explains how to add a process definition to the `Process` data base model

Processes can be chained into a dataflow. Each process is assigned a type (e.g., `data:wc`). The `Data` object created by a process is implicitly assigned a type of that process. When you define a new process, you can specify which data types are required on the input. In the figure below, the *Word Count* process accepts `Data` objects of type `data:doc` on the input. Types are hierarchical with each level of the hierarchy separated by a colon. For instance, `data:doc:text` would be a sub-type of `data:doc`. A process that accepts `Data` objects of type `data:doc`, also accepts `Data` objects of type `data:doc:text`. However, a process that accepts `Data` objects of type `data:doc:text`, does not accept `Data` objects of type `data:doc`.




Fig. 3: Types are hierarchical. When you define the type on the input, keep in mind that the process should also handle all sub-types.

1.3.1 Process syntax

A process can be written in any syntax as long as you can save it to the Process model. The most straight-forward would be to write in Python, using the Django ORM:

```
p = Process(name='Word Count',
            slug='wc-basic',
            type='data:wc:',
            inputs = [{
                'name': 'document',
                'type': 'basic:file:'
            }],
            outputs = [{
                'name': 'words',
                'type': 'basic:integer:'
            }],
            run = {
                'bash': 'WORDS=`wc {{ document.file }}\n` +
                        'echo {"words": $WORDS}'
            })
p.save()
```

We suggest to write processes in the YAML syntax. Resolwe includes a `register Django` command that parses `.yaml` files in the `processes` directory and adds the discovered processes to the Process model:

```
./manage.py register
```

Do not forget to re-register the process after you make changes to the `.yaml` file. You have to increase the process version each time you register it. For development, you can use the `--force` option (or `-f` for short):

```
./manage.py register -f
```

This is an example of the `smallest` processor in YAML syntax:

```

1 - slug: mini
2   name: Minimalistic Process
3   requirements:
4     expression-engine: jinja
5   type: "data:mini"
6   run:
7     language: bash
8     program: |
9       echo 'Hello bioinformatician!'

```

This is the example of the basic Word Count implementation in the YAML syntax (with the document file as input):

```

1 - name: Word Count
2   slug: wc-basic
3   type: "data:wc"
4   inputs:
5     - name: document
6       type: basic:file
7   outputs:
8     - name: words
9       type: basic:integer
10  run:
11    language: bash
12    program: |
13      WORDS=$(wc {{ document.file }})
14      echo {"words": $WORDS}

```

If you would like to review the examples of the three processes mentioned above (*Upload Document*, *Word Count* and *Number of Lines*), follow [this link](#). Read more about the process options in [Process schema](#) below.

1.3.2 Process schema

Process is defined by a set of fields in the Process model. We will describe how to write the process schema in YAML syntax. Some fields in the YAML syntax have different name or values than the actual fields in the Process model. See an example of a process with all fields. Fields in a process schema:

Field	Short description	Required	Default
<i>slug</i>	unique id	required	
<i>name</i>	human readable name	required	
<i>description</i>	detailed description	optional	<empty string>
<i>version</i>	version numbering	optional	
<i>type</i>	data type	required	
<i>category</i>	menu category	optional	<empty string>
<i>entity</i>	automatic grouping	optional	
<i>persistence</i>	storage optimization	optional	RAW
<i>scheduling_class</i>	scheduling class	optional	batch
<i>input</i>	list of input fields	optional	<empty list>
<i>output</i>	list of result fields	optional	<empty list>
<i>run</i>	the algorithm	required	
<i>requirements</i>	requirements	optional	<empty dict>

Slug

TODO

Name

TODO

Description

TODO

Version

TODO

Type

TODO

Category

The category is used to arrange processes in a GUI. A category can be any string of lowercase letters, numbers, - and `:`. The colon is used to split categories into sub-categories (*e.g.*, `analyses:alignment`).

We have predefined three top categories: `upload`, `import` and `analyses`. Processes without this top category will not be displayed in the GenBoard interface, but will be available on the platform.

Entity

With defining the `entity` field in the process, new data objects will be automatically attached to a new or existing Entity, depending on it's parents and the definition of the field.

`entity` field has 3 subfields:

- `type` is required and defines the type of entity that the new Data object is attached to
- `input` limits the group of parents' entities to a single field (dot separated path to the field in the definition of input)
- `descriptor_schema` specifies the slug of the descriptor schema that is attached to newly created entity. It defaults to the value of `type`

Persistence

Use RAW for imports. CACHED or TMP processes should be idempotent.

Scheduling class

The scheduling class specifies how the process should be treated by the scheduler. There are two possible values:

- `batch` is for long running tasks, which require high throughput.
- `interactive` is for short running tasks, which require low latency. Processes in this scheduling class are given a limited amount of time to execute (default: 30 seconds).

The default value for processes is `batch`.

Input and Output

A list of *Resolve Fields* that define the inputs and outputs of a process. A *Resolve Field* is defined as a dictionary of the following properties:

Required *Resolve Field* properties:

- `name` - unique name of the field
- `label` - human readable name
- `type` - type of field (either `basic:<...>` or `data:<...>`)

Optional *Resolve Field* properties (except for `group`):

- `description` - displayed under titles or as a tooltip
- `required` - (choices: `true`, `false`)
- `disabled` - (choices: `true`, `false`)
- `hidden` - (choices: `true`, `false`)
- `default` - initial value
- `placeholder` - placeholder value displayed if nothing is specified
- `validate_regex` - client-side validation with regular expression
- `choices` - list of choices to select from (`label`, `value` pairs)

Optional *Resolve Field* properties for `group` fields:

- `description` - displayed under titles or as a tooltip
- `disabled` - (choices: `true`, `false`)
- `hidden` - (choices: `true`, `false`)
- `collapsed` - (choices: `true`, `false`)
- `group` - list of process fields

TODO: explain what is field schema. For field schema details see `fieldSchema.json`.

Run

The algorithm that transforms inputs into outputs. Bash and workflow languages are currently supported and we envision more language support in the future (*e.g.*, directly writing processes in Python or R). Commands should be written to a `program` subfield.

TODO: link a few lines from the `all_fields.yml` process

Requirements

A dictionary defining optional features that should be available in order for the process to run. There are several different types of requirements that may be specified:

- `expression-engine` defines the name of the engine that should be used to evaluate expressions embedded in the `run` section. Currently, only the `jinja` expression engine is supported. By default no expression engine is set, so expressions cannot be used and will be ignored.
- `executor` defines executor-specific options. The value should be a dictionary, where each key defines requirements for a specific executor. The following executor requirements are available:
 - `docker`:
 - * `image` defines the name of the Docker container image that the process should run under.
- `resources` define resources that should be made available to the process. The following resources may be requested:
 - `cores` defines the number of CPU cores available to the process. By default, this value is set to 1 core.
 - `memory` defines the amount of memory (in megabytes) that the process may use. By default, this value is set to 4096 MiB.
 - `network` should be a boolean value, specifying whether the process requires network access. By default this value is `false`.

1.3.3 Types

Types are defined for processes and *Resolwe Fields*. Data objects have implicitly defined types, based on the corresponding processor. Types define the type of objects that are passed as inputs to the process or saved as outputs of the process. Resolwe uses 2 kinds of types:

- `basic`:
- `data`:

`Basic`: types are defined by Resolwe and represent the data building blocks. `Data`: types are defined by processes. In terms of programming languages you could think of `basic`: as primitive types (like integer, float or boolean) and of `data`: types as classes.

Resolwe matches inputs based on the type. Types are hierarchical, so the same or more specific inputs are matched. For example:

- `data:genome:fasta`: will match the `data:genome`: input, but
- `data:genome`: will not match the `data:genome:fasta`: input.

Note: Types in a process schema do not have to end with a colon. The last colon can be omitted for readability and is added automatically by Resolwe.

Basic types

Basic types are entered by the user. Resolve implements the backend handling (storage and retrieval) of basic types and GenBoard supports the HTML5 controls.

The following basic types are supported:

- `basic:boolean:` - boolean
- `basic:date:` - date (format *yyyy-mm-dd*)
- `basic:datetime:` - date and time (format *yyyy-mm-dd hh:mm:ss*)
- `basic:decimal:` - decimal number (*e.g.*, *-123.345*)
- `basic:integer:` - whole number (*e.g.*, *-123*)
- `basic:string:` - short string
- `basic:text:` - multi-line string
- `basic:url:link:` - visit link
- `basic:url:download:` - download link
- `basic:url:view:` - view link (in a popup or iframe)
- `basic:file:` - a file, stored on shared file system
- `basic:dir:` - a directory, stored on shared file system
- `basic:json:` - a JSON object, stored in MongoDB collection
- `basic:group:` - list of form fields (default if nothing specified)

The values of basic data types are different for each type, for example: `basic:file:` data type is a JSON dictionary: {"file": "file name"} `basic:dir:` data type is a JSON dictionary: {"dir": "directory name"} `basic:string:` data type is just a JSON string

Resolve treats types differently. All but `basic:file:`, `basic:dir:` and `basic:json:` are treated as meta-data. `basic:file:` and `basic:dir:` objects are saved to the shared file storage, and `basic:json:` objects are stored in PostgreSQL bson field. Meta-data entries have references to `basic:file:`, `basic:dir:` and `basic:json:` objects.

Data types

Data types are defined by processes. Each process is itself a `data:` sub-type named with the `type` attribute. A `data:` sub-type is defined by a list process outputs. All processes of the same `type` should have the same outputs.

Data type name:

- `data:<type>[:<sub-type>[...]]:`

1.3.4 The algorithm

Algorithm is the key component of a process. The algorithm transforms process's inputs into outputs. It is written as a sequence of Bash commands in process's `run.program` field.

Note: In this section, we assume that the program is written using the `bash` language and having the `expression-engine` requirement set to `jinja`.

To write the algorithm in a different language (*e.g.*, Python), just put it in a file with an appropriate *shebang* at the top (*e.g.*, `#!/usr/bin/env python2` for Python2 programs) and add it to the *tools* directory. To run it simply call the script with appropriate arguments.

For example, to compute a Volcano plot of the baySeq data, use:

```
volcanoplot.py diffexp_bayseq.tab
```

Platform utilities

Resolwe provides some convenience utilities for writing processes:

- `re-import`

is a convenience utility that copies/downloads a file from the given temporary location, extracts/compresses it and moves it to the given final location. It takes six arguments:

1. file's temporary location or URL
2. file's final location
3. file's input format, which can have one of the following forms:
 - `ending1|ending2`: matches files that end with `ending1` or `ending2` or a combination of `(ending1|ending2).(gz|bz2|zip|rar|7z|tgz|tar.gz|tar.bz2)`
 - `ending1|ending2|compression`: matches files that end with `ending1` or `ending2` or a combination of `(ending1|ending2).(gz|bz2|zip|rar|7z|tgz|tar.gz|tar.bz2)` or just with a supported compression format line ending `(gz|bz2|zip|rar|7z)`
4. file's output format (*e.g.*, `fasta`)
5. maximum progress at the end of transfer (a number between 0.0 and 1.0)
6. file's output format, which can be one of the following:
 - `compress`: to produce a compressed file
 - `extract`: to produce an extracted file

If this argument is not given, both, the compressed and the extracted file are produced.

For storing the results to process's output fields, Resolwe provides a series of utilities. They are described in the *Outputs* section.

Runtime

TODO: Write about BioLinux and what is available in the Docker runtime.

Inputs

To access values stored in process's input fields, use [Jinja2's template language syntax for accessing variables](#). For example, to access the value of process's `fastq` input field, write `{{ fastq }}`.

In addition to all process's input fields, Resolwe provides the following system variables:

- `proc.case_ids`: ids of the corresponding cases
- `proc.data_id`: id of the data object
- `proc.slugs_path`: file system path to Resolwe's slugs

Resolve also provides some custom built-in filters to access the fields of the referenced data objects:

- `id`: returns the id of the referenced data object
- `type`: returns the type of the referenced data object
- `name`: returns the value of the `static.name` field if it exists

For example, to use these filters on the `reads` field, use `{{ reads|id }}`, `{{ reads|type }}` or `{{ reads|name }}`, respectively.

You can also use any [Jinja2's built in template tags and filters](#) in your algorithm.

Note: All input variables should be considered *unsafe* and will be automatically quoted when used in your scripts. For example, the following call:

```
volcanoplot.py {{ reads.fastq.0.file }}
```

will actually be transformed into something like (depending on the value):

```
volcanoplot.py '/path/to/reads with spaces.gz'
```

If you do not want this behaviour for a certain variable and you are sure that it is safe to do so, you can use the `safe` filter as follows:

```
volcanoplot.py {{ known_good_input | safe }}
```

Outputs

Processes have three options for storing the results:

- as files in data object's directory (i.e. `{{ proc.data_dir }}`)
- as constants in process's output fields
- as entries in the MongoDB data storage

Note: Files are stored on a shared file system that supports fast read and write accesss by the processes. Accessing MongoDB from a process requires more time and is suggested for interactive data retrieval from GenPackages only.

Saving status

There are two special fields that you should use:

- `proc.rc`: the return code of the process
- `proc.progress`: the process's progress

If you set the `proc.rc` field to a positive value, the process will fail and its status will be set to `ERROR`. All processes that depend on this process will subsequently fail and their status will be set to `ERROR` as well.

The `proc.progress` field can be used to report processing progress interactively. You can set it to a value between 0 and 1 that represents an estimate for process's progress.

To set them, use the `re-progress` and `re-checkrc` utilities described in the [Saving constants](#) section.

Resolwe provides some specialized utilities for reporting process status:

- **re-error**

takes one argument and stores it to `proc.error` field. For example:

```
re-error "Error! Something went wrong."
```

- **re-warning**

takes one argument and stores it to `proc.warning` field. For example:

```
re-warning "Be careful there might be a problem."
```

- **re-info**

takes one argument and stores it to `proc.info` field. For example:

```
re-info "Just say hello."
```

- **re-progress**

takes one argument and stores it to `proc.progress` field. The argument should be a float between 0 and 1 and represents an estimate for process's progress. For example, to estimate the progress to 42%, use:

```
re-progress 0.42
```

- **re-checkrc**

saves the return code of the previous command to `proc.rc` field. To use it, just call:

```
re-checkrc
```

As some programs exit with a non-zero return code, even though they finished successfully, you can pass additional return codes as arguments to the `re-checkrc` command and they will be translated to zero. For example:

```
re-checkrc 2 15
```

will set `proc.rc` to 0 if the return code is 0, 2 or 15, and to the actual return code otherwise.

It is also possible to set the `proc.error` field with this command in case the return code is not zero (or is not given as one of the acceptable return codes). To do that, just pass the error message as the last argument to the `re-checkrc` command. For example:

```
re-checkrc "Error ocurred."  
re-checkrc 2 "Return code was not 0 or 2."
```

Saving constants

To store a value in a process's output field, use the `re-save` utility. The `re-save` utility requires two arguments, a key (i.e. field's name) and a value (i.e. field's value).

For example, executing:

```
re-save quality_mean $QUALITY_MEAN
```

will store the value of the `QUALITY_MEAN` Bash variable in process's `quality_mean` field.

Note: To use the `re-save` utility, add `re-require common` to the beginning of the algorithm. For more details, see [Platform utilities](#).

You can pass any JSON object as the second argument to the `re-save` utility, *e.g.*:

```
re-save foo '{"extra_output": "output.txt"}'
```

Note: Make sure to put the second argument into quotes (*e.g.*, `""` or `'`) if you pass a JSON object containing a space to the `re-save` utility.

Saving files

A convenience function for saving files is:

```
re-save-file
```

It takes two arguments and stores the value of the second argument in the first argument's `file` subfield. For example:

```
re-save-file fastq $NAME.fastq.gz
```

stores `$NAME.fastq.gz` to the `fastq.file` field which has to be of type `basic:file:`.

To reference additional files/folders, pass them as extra arguments to the `re-save-file` utility. They will be saved to the `refs` subfield of type `basic:file:`. For example:

```
re-save-file fastq $NAME.fastq.gz fastqc/${NAME}_fastqc
```

stores `fastqc/${NAME}_fastqc` to the `fastq.refs` field in addition to storing `$NAME.fastq.gz` to the `fastq.file` field.

Note: Resolve will automatically add files' sizes to the files' `size` subfields.

Warning: After the process has finished, Resolve will automatically check if all the referenced files exist. If any file is missing, it will set the data object's status to `ERROR`. Files that are not referenced are automatically deleted by the platform, so make sure to reference all the files you want to keep!

Saving JSON blobs in MongoDB

To store a JSON blob to the MongoDB storage, simply create a field of type `data:json:` and use the `re-save` utility to store it. The platform will automatically detect that you are trying to store to a `data:json:` field and it will store the blob to a separate collection.

For example:

```
re-save etc { JSON blob }
```

will store the `{ JSON blob }` to the `etc` field.

Note: Printing a lot of data to standard output can cause problems when using the Docker executor due to its current implementation. Therefore, it is advised to save big JSON blobs to a file and only pass the file name to the `re-save` function.

For example:

```
command_that_generates_large_json > json.txt  
re-save etc json.txt
```

Warning: Do not store large JSON blobs into the data collection directly as this will slow down the retrieval of data objects.

1.4 API

The Resolwe framework provides a RESTful API through which most of its functionality is exposed.

TODO

1.4.1 Elasticsearch endpoints

Advanced lookups

All fields that can be filtered upon (as defined for each viewset) support specific lookup operators that can be used for some more advanced lookups.

Currently the supported lookup operators are:

- `lt` creates an ES range query with `lt` bound. Supported for number and date fields.
- `lte` creates an ES range query with `lte` bound. Supported for number and date fields.
- `gt` creates an ES range query with `gt` bound. Supported for number and date fields.
- `gte` creates an ES range query with `gte` bound. Supported for number and date fields.
- `in` creates an ES boolean query with all values passed as a should match. For GET requests, multiple values should be comma-separated.
- `exact` creates an ES query on the raw subfield of the given field, requiring the value to match exactly with the raw value that was supplied during indexing.

1.4.2 Limiting fields in responses

As responses from the Resolve API can contain a lot of data, especially with nested JSON outputs and schemas, the API provides a way of limiting what is returned with each response.

This is achieved through the use of a special `fields` GET parameter, which can specify one or multiple field projections. Each projection defines what should be returned. As a working example, let's assume we have the following API response when no field projections are applied:

```
[
  {
    "foo": {
      "name": "Foo",
      "bar": {
        "level3": 42,
        "another": "hello"
      }
    },
    "name": "Boo"
  },
  {
    "foo": {
      "name": "Different",
    },
    "name": "Another"
  }
]
```

A field projection may reference any of the top-level fields. For example, by using the `fields=name` projection, we get the following result:

```
[
  {
    "name": "Boo"
  },
  {
    "name": "Another"
  }
]
```

Basically all fields not matching the projection are gone. We can go further and also project deeply nested fields, e.g., `fields=foo__name`:

```
[
  {
    "foo": {
      "name": "Foo"
    }
  },
  {
    "foo": {
      "name": "Different"
    }
  }
]
```

(continues on next page)

(continued from previous page)

]

And at last, we can combine multiple projections by separating them with commas, e.g., `fields=name,foo__name`, giving us:

```
[
  {
    "foo": {
      "name": "Foo"
    },
    "name": "Boo"
  },
  {
    "foo": {
      "name": "Different"
    },
    "name": "Another"
  }
]
```

1.5 Type extension composition

Many types that are part of the core Resolwe framework contain logic that users of the framework may need to extend. To facilitate this in a controlled manner, the Resolwe framework provides a generic type extension composition system.

1.5.1 Making a type extendable

The composition system is very generic and as such can be used on any type. It provides a single method which allows you to retrieve a list of all registered extensions for a type or an instance of that type.

```
>>> composer.get_extensions(my_type_or_instance)
[<Extension1>, <Extension2>]
```

The type can then use this API to incorporate the registered extensions into its current instance however it chooses. Note that what these extensions are is entirely dependent upon the type that uses them.

For example, in the core Resolwe framework we make all index definitions extendable by using something like:

```
for extension in composer.get_extensions(attr):
    mapping = getattr(extension, 'mapping', {})
    index.mapping.update(mapping)
```

1.5.2 Writing an extension

On the other side, you can also define extensions for types that are using the above mentioned API. All extensions are automatically discovered during Django application registration if they are placed in a module called `extensions` in the given application.

Extensions can be registered using a simple API:

```
class MyExtension:
    pass

composer.add_extension('fully.qualified.type.Path', MyExtension)
```

Again, what the extension is depends on the type that is being extended. Now we describe some common extension types for types that are part of the Resolve core.

Data viewset

It is possible to extend the filters of the Data viewset by defining an extension as follows:

```
class ExtendedDataViewSet:
    """Data viewset extensions."""
    filtering_fields = ('source', 'species', 'build', 'feature_type')

    def text_filter(self, value):
        return [
            Q('match', species={'query': value, 'operator': 'and', 'boost': 2.0}),
            Q('match', source={'query': value, 'operator': 'and', 'boost': 2.0}),
            Q('match', build={'query': value, 'operator': 'and', 'boost': 2.0}),
            Q('match', feature_type={'query': value, 'operator': 'and', 'boost': 1.0}),
        ]

composer.add_extension('resolve.flow.views.data.DataViewSet', ExtendedDataViewSet)
```

1.6 Reference

1.6.1 Permissions shortcuts

`resolve.permissions.shortcuts.get_object_perms`(*obj*: Model, *user*: User | None = None, *mock_superuser_permissions*: bool = False) → List[Dict]

Return permissions for given object in Resolve specific format.

Function returns permissions for given object *obj* in the following format:

```
{
  "type": "group"/"user"/"public",
  "id": <group_or_user_id>,
  "name": <group_or_user_name>,
  "username": <username>,
  "permissions": [<first_permission>, <second_permission>, ...]
}
```

For public type id and name keys are omitted.

If user parameter is given, permissions are limited only to given user, groups he belongs to and public permissions.

This function should be only used from Resolwe views: since permissions for the current user (users when user has share permission on the given object) are prefetched, we only iterate through objects here and filter them in Python. Using filter method would result in a new database query.

Parameters

- **obj** – Resolwe’s DB model’s instance
- **user** – Django user
- **mock_superuser_permissions** – when True return all permissions for users that are superusers

Returns

list of permissions object in described format

1.6.2 Permissions utils

`resolwe.permissions.utils.copy_permissions(src_obj: Model, dest_obj: Model)`

Copy permissions from src_obj to dest_obj.

Warning: Existing permissions in dest_obj will be deleted.

1.6.3 Flow Managers

Workflow workload managers.

`resolwe.flow.managers.manager`

The global manager instance.

Type

Manager

Dispatcher

`class resolwe.flow.managers.dispatcher.Manager(*args, **kwargs)`

The manager handles process job dispatching.

Each *Data* object that’s still waiting to be resolved is dispatched to a concrete workload management system (such as Celery or SLURM). The specific manager for that system (descended from *BaseConnector*) then handles actual job setup and submission. The job itself is an executor invocation; the executor then in turn sets up a safe and well-defined environment within the workload manager’s task in which the process is finally run.

async communicate(*data_id=None, run_sync=False*)

Scan database for resolving Data objects and process them.

This is submitted as a task to the manager’s channel workers.

Parameters

- **data_id** – Optional id of Data object which (+ its children) should be processed. If it is not given, all resolving objects are processed.
- **run_sync** – If True, wait until all processes spawned from this point on have finished processing.

discover_engines()

Discover configured engines.

Parameters

executor – Optional executor module override

async drain_messages()

Drain Django Channel messages.

async execution_barrier()

Wait for executors to finish.

At least one must finish after this point to avoid a deadlock.

get_execution_engine(name: str)

Return an execution engine instance.

get_executor()

Return an executor instance.

get_expression_engine(name: str)

Return an expression engine instance.

async handle_control_event(message: dict)

Handle the control event.

The method is called from the channels layer when there is some change either in the state of the Data object or the executors have finished with processing.

When running in sync state check that all database objects are in final state before raising the execution_barrier.

Channels layer callback, do not call directly.

load_execution_engines(engines: List[dict | str])

Load execution engines.

load_executor(executor_name: str)

Load process executor.

load_expression_engines(engines: List[dict | str])

Load expression engines.

run(data: Data, argv: List)

Select a concrete connector and run the process through it.

Parameters

- **data** – The *Data* object that is to be run.
- **argv** – The argument vector used to spawn the executor.

```
class resolve.flow.managers.dispatcher.SettingsJSONifier(*, skipkeys=False, ensure_ascii=True,
check_circular=True, allow_nan=True,
sort_keys=False, indent=None,
separators=None, default=None)
```

Customized JSON encoder, coercing all unknown types into strings.

Needed due to the class hierarchy coming out of the database, which can't be serialized using the vanilla json encoder.

default(*o*)

Try default; otherwise, coerce the object into a string.

Workload Connectors

The workload management system connectors are used as glue between the Resolwe Manager and various concrete workload management systems that might be used by it. Since the only functional requirement is job submission, they can be simple and nearly contextless.

Base Class

class `resolwe.flow.managers.workload_connectors.base.BaseConnector`

The abstract base class for workload manager connectors.

The main *Manager* instance in *manager* uses connectors to handle communication with concrete backend workload management systems, such as Celery and SLURM. The connectors need not worry about how jobs are discovered or how they're prepared for execution; this is all done by the manager.

cleanup(*data_id: int*)

Perform final cleanup after data object is finished processing.

submit(*data: Data, argv*)

Submit the job to the workload management system.

Parameters

- **data** – The *Data* object that is to be run.
- **argv** – The argument vector used to spawn the executor.

Local Connector

class `resolwe.flow.managers.workload_connectors.local.Connector`

Local connector for job execution.

cleanup(*data_id: int*)

Cleanup.

submit(*data: Data, argv*)

Run process locally.

For details, see `submit()`.

Celery Connector

class `resolve.flow.managers.workload_connectors.celery.Connector`

Celery-based connector for job execution.

cleanup(*data_id: int*)

Cleanup.

submit(*data: Data, argv*)

Run process.

For details, see `submit()`.

Slurm Connector

class `resolve.flow.managers.workload_connectors.slurm.Connector`

Slurm-based connector for job execution.

cleanup(*data_id: int*)

Cleanup.

submit(*data: Data, argv*)

Run process with SLURM.

For details, see `submit()`.

Kubernetes Connector

class `resolve.flow.managers.workload_connectors.kubernetes.ConfigLocation`(*value*)

The enum specifying where to read the configuration from.

class `resolve.flow.managers.workload_connectors.kubernetes.Connector`

Kubernetes-based connector for job execution.

cleanup(*data_id: int*)

Remove the persistent volume claims created by the executor.

optimize_job_scheduling(*data: Data, job_description: dict*)

Optimize the scheduling by modifying the job description.

Moodify the job description to be a better fit on the scheduler. Currently this is used to schedule based on the process scheduling class.

start(*data: Data, listener_connection: Tuple[str, str, str]*)

Start process execution.

Construct kubernetes job description and pass it to the kubernetes.

submit(*data: Data, argv*)

Run process.

For details, see `submit()`.

`resolwe.flow.managers.workload_connectors.kubernetes.get_mountable_connectors()` → `Iterable[Tuple[str, BaseStorageConnector]]`

Iterate through all the storages and find mountable connectors.

Returns

list of tuples (storage_name, connector).

`resolwe.flow.managers.workload_connectors.kubernetes.get_upload_dir()` → `str`

Get the upload path.

: returns: the path of the first mountable connector for storage
‘upload’.

Raises

`RuntimeError` – if no applicable connector is found.

`resolwe.flow.managers.workload_connectors.kubernetes.sanitize_kubernetes_label`(*label: str*,
trim_end: bool = True)
→ `str`

Make sure kubernetes label complies with the rules.

See the URL bellow for details.

<https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/>

`resolwe.flow.managers.workload_connectors.kubernetes.unique_volume_name`(*base_name: str*,
data_id: int, *postfix: str*) → `str`

Get unique persistent volume claim name.

Listener

Consumer

Manager Channels consumer.

`class resolwe.flow.managers.consumer.HealthCheckConsumer`

Channels consumer for handling health-check events.

`check_database()` → `bool`

Perform a simple database check.

`async health_check`(*message: dict*)

Perform health check.

We are testing the channels layer and database layer. The channels layer is already functioning if this method is called so we have to perform database check.

If the check is successfull touch the file specified in the channels message.

`class resolwe.flow.managers.consumer.ManagerConsumer`(*args, **kwargs)

Channels consumer for handling manager events.

`async control_event`(*message*)

Forward control events to the manager dispatcher.

async `resolve.flow.managers.consumer.exit_consumer()`

Cause the synchronous consumer to exit cleanly.

async `resolve.flow.managers.consumer.run_consumer(timeout=None)`

Run the consumer until it finishes processing.

Parameters

timeout – Set maximum execution time before cancellation, or None (default) for unlimited.

async `resolve.flow.managers.consumer.send_event(message)`

Construct a Channels event packet with the given message.

Parameters

message – The message to send to the manager workers.

Utilities

Utilities for using global manager features.

`resolve.flow.managers.utils.disable_auto_calls()`

Decorator/context manager which stops automatic manager calls.

When entered, automatic `communicate()` calls from the Django transaction signal are not done.

1.6.4 Flow Executors

Base Class

class `resolve.flow.executors.run.BaseFlowExecutor`(*data_id: int, communicator: ZMQCommunicator, listener_connection: Tuple[str, str, str], *args, **kwargs*)

Represents a workflow executor.

get_tools_paths()

Get tools paths.

async run()

Execute the script and save results.

async start()

Start process execution.

Flow Executor Preparer

Framework for the manager-resident executor preparation facilities.

class `resolve.flow.executors.prepare.BaseFlowExecutorPreparer`

Represents the preparation functionality of the executor.

extend_settings(*data_id, files, secrets*)

Extend the settings the manager will serialize.

Parameters

- **data_id** – The *Data* object id being prepared for.

- **files** – The settings dictionary to be serialized. Keys are filenames, values are the objects that will be serialized into those files. Standard filenames are listed in `resolwe.flow.managers.protocol.ExecutorFiles`.
- **secrets** – Secret files dictionary describing additional secret file content that should be created and made available to processes with special permissions. Keys are filenames, values are the raw strings that should be written into those files.

get_environment_variables()

Return dict of environment variables that will be added to executor.

get_tools_paths(*from_applications=False*)

Get tools' paths.

post_register_hook(*verbosity=1*)

Run hook after the 'register' management command finishes.

Subclasses may implement this hook to e.g. pull Docker images at this point. By default, it does nothing.

prepare_for_execution(*data*)

Prepare the data object for the execution.

This is mostly needed for the null executor to change the status of the data and worker object to done.

resolve_data_path(*data=None, filename=None*)

Resolve data path for use with the executor.

Parameters

- **data** – Data object instance
- **filename** – Filename to resolve

Returns

Resolved filename, which can be used to access the given data file in programs executed using this executor

resolve_upload_path(*filename=None*)

Resolve upload path for use with the executor.

Parameters

filename – Filename to resolve

Returns

Resolved filename, which can be used to access the given uploaded file in programs executed using this executor

Raises

RuntimeError – when no storage connectors are configured for upload storage or path could not be resolved.

Docker Flow Executor

Preparation

class `resolve.flow.executors.docker.prepare.FlowExecutorPreparer`

Specialized manager assist for the docker executor.

get_environment_variables()

Return dict of environment variables that will be added to executor.

post_register_hook(*verbosity=1*)

Pull Docker images needed by processes after registering.

resolve_data_path(*data=None, filename=None*)

Resolve data path for use with the executor.

Parameters

- **data** – Data object instance
- **filename** – Filename to resolve

Returns

Resolved filename, which can be used to access the given data file in programs executed using this executor

Raises

RuntimeError – when data path can not be resolved.

resolve_upload_path(*filename=None*)

Resolve upload path for use with the executor.

Parameters

- **filename** – Filename to resolve

Returns

Resolved filename, which can be used to access the given uploaded file in programs executed using this executor

Local Flow Executor

class `resolve.flow.executors.local.run.FlowExecutor`(*args, **kwargs)

Local dataflow executor proxy.

Preparation

class `resolve.flow.executors.local.prepare.FlowExecutorPreparer`

Specialized manager assist for the local executor.

extend_settings(*data_id, files, secrets*)

Prevent processes requiring access to secrets from being run.

Null Flow Executor

```
class resolwe.flow.executors.null.run.FlowExecutor(data_id: int, communicator: ZMQCommunicator,
                                                    listener_connection: Tuple[str, str, str], *args,
                                                    **kwargs)
```

Null dataflow executor proxy.

This executor is intended to be used in tests where you want to save the object to the database but don't need to run it.

1.6.5 Flow Models

Base Model

Base model for all other models.

```
class resolwe.flow.models.base.BaseModel(*args, **kwargs)
```

Abstract model that includes common fields for other models.

```
class Meta
```

BaseModel Meta options.

```
contributor
```

user that created the entry

```
created
```

creation date and time

```
modified
```

modified date and time

```
name
```

object name

```
save(*args, **kwargs)
```

Save the model.

```
slug
```

URL slug

```
version
```

process version

Collection Model

Postgres ORM model for the organization of collections.

```
class resolwe.flow.models.collection.BaseCollection(*args, **kwargs)
```

Template for Postgres model for storing a collection.

```
class Meta
```

BaseCollection Meta options.

```
contributor
```

user that created the entry

created

creation date and time

description

detailed description

descriptor

collection descriptor

descriptor_dirtyindicate whether *descriptor* doesn't match *descriptor_schema* (is dirty)**descriptor_schema**

collection descriptor schema

modified

modified date and time

name

object name

save(*args, **kwargs)

Perform descriptor validation and save object.

search

field used for full-text search

slug

URL slug

tags

tags for categorizing objects

version

process version

class `resolve.flow.models.Collection(*args, **kwargs)`

Postgres model for storing a collection.

exception `DoesNotExist`**exception** `MultipleObjectsReturned`**annotation_fields**

annotation fields available to samples in this collection

contributor

user that created the entry

created

creation date and time

delete_background()

Delete the object in the background.

description

detailed description

descriptor

collection descriptor

descriptor_dirty

indicate whether *descriptor* doesn't match *descriptor_schema* (is dirty)

descriptor_schema

collection descriptor schema

duplicate(*contributor*) → BackgroundTask

Duplicate (make a copy) object in the background.

duplicated

duplication date and time

is_duplicate()

Return True if collection is a duplicate.

modified

modified date and time

name

object name

objects = <django.db.models.manager.ManagerFromCollectionQuerySet object>

manager

permission_group

permission group for the object

save(*args, **kwargs)

Add required annotation fields to the collection.

search

field used for full-text search

slug

URL slug

tags

tags for categorizing objects

version

process version

Data model

Postgres ORM model for keeping the data structured.

class resolwe.flow.models.Data(*args, **kwargs)

Postgres model for storing data.

exception DoesNotExist

exception MultipleObjectsReturned

STATUS_DIRTY = 'DR'

data object is in dirty state

STATUS_DONE = 'OK'

data object is done

STATUS_ERROR = 'ER'

data object is in error state

STATUS_PREPARING = 'PP'

data object is preparing

STATUS_PROCESSING = 'PR'

data object is processing

STATUS_RESOLVING = 'RE'

data object is being resolved

STATUS_UPLOADING = 'UP'

data object is uploading

STATUS_WAITING = 'WT'

data object is waiting

checksum

checksum field calculated on inputs

collection

collection

contributor

user that created the entry

created

creation date and time

delete(*args, **kwargs)

Delete the data model.

delete_background()

Delete the object in the background.

dependency_status() → str | None

Return abstracted status of instance IO dependencies.

Returns

- STATUS_ERROR .. one dependency has error status or was deleted
- STATUS_DONE .. all dependencies have done status
- None .. other

descriptor

actual descriptor

descriptor_dirty

indicate whether *descriptor* doesn't match *descriptor_schema* (is dirty)

descriptor_schema

data descriptor schema

duplicate(*contributor*) → BackgroundTask

Duplicate (make a copy) object in the background.

duplicated

duplication date and time

entity

entity

finished

process finished date date and time (set by `resolwe.flow.executors.run.BaseFlowExecutor.run` or its derivatives)

get_resource_limits()

Return the resource limits for this data.

input

actual inputs used by the process

is_duplicate()

Return True if data object is a duplicate.

location

data location

modified

modified date and time

name

object name

named_by_user

track if user set the data name explicitly

objects = <`django.db.models.manager.ManagerFromDataQuerySet` object>

manager

output

actual outputs of the process

parents

dependencies between data objects

permission_group

permission group for the object

process

process used to compute the data object

process_cores

actual allocated cores

process_error

error log message

process_info

info log message

process_memory

actual allocated memory

process_pid

process id

process_progress

progress

process_rc

return code

process_resources

process requirements overrides

process_warning

warning log message

resolve_secrets()

Retrieve handles for all basic:secret: fields on input.

The process must have the `secrets` resource requirement specified in order to access any secrets. Otherwise this method will raise a `PermissionDenied` exception.

Returns

A dictionary of secrets where key is the secret handle and value is the secret value.

restart(*resource_overrides: dict = {}*)

Restart the data object and all its children.

The status of the data object must be `ERROR` and status of its dependencies must be `DONE`.

Parameters

resource_overrides – dictionary mapping ids of data objects to resource overrides.

Raises

- `RuntimeError` – if the object is not in the right state.
- `RuntimeError` – when object dependencies are not in the status `DONE`.

save(*render_name=False, *args, **kwargs*)

Save the data model.

save_dependencies(*instance, schema*)

Save data: and list:data: references as parents.

scheduled

date and time when process was dispatched to the scheduling system (set by `resolve.flow.managers.dispatcher.Manager.run``)

search

field used for full-text search

size

total size of data's outputs in bytes

slug

URL slug

started

process started date and time (set by `resolwe.flow.executors.run.BaseFlowExecutor.run` or its derivatives)

status

Data status

It can be one of the following:

- *STATUS_UPLOADING*
- *STATUS_RESOLVING*
- *STATUS_WAITING*
- *STATUS_PROCESSING*
- *STATUS_DONE*
- *STATUS_ERROR*

tags

tags for categorizing objects

validate_change_collection(*collection*)

Raise validation error if data object cannot change collection.

version

process version

class `resolwe.flow.models.DataDependency`(*args, **kwargs)

Dependency relation between data objects.

exception `DoesNotExist`

`KIND_IO = 'io'`

child uses parent's output as its input

`KIND_SUBPROCESS = 'subprocess'`

child was spawned by the parent

exception `MultipleObjectsReturned`

child

child data object

kind

kind of dependency

parent

parent data object

Entity–relationship model

Postgres ORM to define the entity–relationship model that describes how data objects are related in a specific domain.

```

class resolve.flow.models.Entity(*args, **kwargs)
    Postgres model for storing entities.

    exception DoesNotExist

    exception MultipleObjectsReturned

    collection
        collection to which entity belongs

    contributor
        user that created the entry

    copy_annotations(destination: Entity) → List[AnnotationValue]
        Copy annotation from this entity to the destination.

        Raises
            ValidationError – when some of the annotation fields are missing on the destination entity.

    created
        creation date and time

    delete_background()
        Delete the object in the background.

    description
        detailed description

    descriptor
        collection descriptor

    descriptor_dirty
        indicate whether descriptor doesn't match descriptor_schema (is dirty)

    descriptor_schema
        collection descriptor schema

    duplicate(contributor) → BackgroundTask
        Duplicate (make a copy) object in the background.

    duplicated
        duplication date and time

    get_annotation(path: str, default: Any | None = None) → Any
        Get the annotation for the given path.

        Attr path
            the path to the annotation in the format 'group.field'.

        Attr default
            default value when annotation is not found.

    Returns
        value of the annotation or default if not found.

```

invalid_annotation_fields(*annotation_fields=None*)

Get the Queryset of invalid annotation fields.

The invalid annotation field is a field that has annotation but it is not allowed in the collection this entity belongs to.

Attr annotation_fields

the iterable containing annotations fields to be checked. When None is given the annotation fields belonging to the entity are checked.

is_duplicate()

Return True if entity is a duplicate.

modified

modified date and time

name

object name

objects = <django.db.models.manager.ManagerFromEntityQuerySet object>

manager

permission_group

permission group for the object

search

field used for full-text search

set_annotation(*path: str, value: Any*)

Get the annotation for the given path.

Attr path

the path to the annotation in the format 'group.field'.

Attr value

the annotation value.

slug

URL slug

tags

tags for categorizing objects

type

entity type

update_annotations(*annotations: dict[str, Any], update=True*)

Update annotations with the given values.

When annotation value is set no None it is deleted.

Attr annotations

the dictionary with annotation values. Keys are annotation paths.

validate_annotations()

Perform streamlined descriptor validation.

Raises

ValidationError – when annotations do not pass validation. All fields are validated and error messages aggregated into single exception.

version

process version

class `resolve.flow.models.Relation(*args, **kwargs)`

Relations between entities.

The Relation model defines the associations and dependencies between entities in a given collection:

```
{
  "collection": "<collection_id>",
  "type": "comparison",
  "category": "case-control study",
  "entities": [
    {"entity": "<entity1_id>", "label": "control"},
    {"entity": "<entity2_id>", "label": "case"},
    {"entity": "<entity3_id>", "label": "case"}
  ]
}
```

Relation type defines a specific set of associations among entities. It can be something like `group`, `comparison` or `series`. The relation type is an instance of `RelationType` and should be defined in any Django app that uses relations (e.g., as a fixture). Multiple relations of the same type are allowed on the collection.

Relation category defines a specific use case. The relation category must be unique in a collection, so that users can distinguish between different relations. In the example above, we could add another `comparison` relation of category, say `Case-case study` to compare `<entity2>` with `<entity3>`.

Relation is linked to `resolve.flow.models.Collection` to enable defining different relations structures in different collections. This also greatly speed up retrieving of relations, as they are envisioned to be mainly used on a collection level.

`unit` defines units used in partitions where it is applicable, e.g. in relations of type `series`.

exception `DoesNotExist`**exception** `MultipleObjectsReturned`**category**

category of the relation

collection

collection to which relation belongs

contributor

user that created the entry

created

creation date and time

descriptor

relation descriptor

descriptor_dirtyindicate whether `descriptor` doesn't match `descriptor_schema` (is dirty)**descriptor_schema**

relation descriptor schema

entities

partitions of entities in the relation

modified

modified date and time

name

object name

objects = <django.db.models.manager.ManagerFromPermissionQuerySet object>

custom manager with permission filtering methods

permission_group

permission group for the object

save(*args, **kwargs)

Perform descriptor validation and save object.

slug

URL slug

type

type of the relation

unit

unit used in the partitions' positions (where applicable, e.g. for serieses)

version

process version

class resolwe.flow.models.**RelationType**(*args, **kwargs)

Model for storing relation types.

exception **DoesNotExist**

exception **MultipleObjectsReturned**

name

relation type name

ordered

indicates if order of entities in relation is important or not

DescriptorSchema model

Postgres ORM model for storing descriptors.

class resolwe.flow.models.**DescriptorSchema**(*args, **kwargs)

Postgres model for storing descriptors.

exception **DoesNotExist**

exception **MultipleObjectsReturned**

contributor

user that created the entry

created
creation date and time

description
detailed description

modified
modified date and time

name
object name

permission_group
permission group for the object

schema
user descriptor schema represented as a JSON object

slug
URL slug

version
process version

Process model

Postgres ORM model for storing processes.

class `resolve.flow.models.Process(*args, **kwargs)`

Postgres model for storing processes.

exception `DoesNotExist`

exception `MultipleObjectsReturned`

PERSISTENCE_CACHED = 'CAC'

cached persistence

PERSISTENCE_RAW = 'RAW'

raw persistence

PERSISTENCE_TEMP = 'TMP'

temp persistence

category

category

contributor

user that created the entry

created

creation date and time

data_name

template for name of Data object created with Process

description

detailed description

entity_always_create

Create new entity, regardless of `entity_input` or `entity_descriptor_schema` fields.

entity_descriptor_schema

Slug of the descriptor schema assigned to the Entity created with `entity_type`.

entity_input

Limit the entity selection in `entity_type` to a single input.

entity_type

Automatically add `Data` object created with this process to an `Entity` object representing a data-flow. If all input `Data` objects belong to the same entity, add newly created `Data` object to it, otherwise create a new one.

get_resource_limits(*data*: `Data` | `None` = `None`)

Get the core count and memory usage limits for this process.

Returns

A dictionary with the resource limits, containing the following keys:

- `memory`: Memory usage limit, in MB. Defaults to 4096 if not otherwise specified in the resource requirements.
- `cores`: Core count limit. Defaults to 1.
- `storage`: Size (in gibibytes) of temporary volume used for processing in kubernetes. Defaults to 200.

Return type

`dict`

input_schema

process input schema (describes input parameters, form layout “**Inputs**” for `Data.input`)

Handling:

- schema defined by: `dev`
- default by: `user`
- changable by: `none`

is_active

designates whether this process should be treated as active

modified

modified date and time

name

object name

output_schema

process output schema (describes output JSON, form layout “**Results**” for `Data.output`)

Handling:

- schema defined by: `dev`
- default by: `dev`

- changable by: *dev*

Implicitly defined fields (by `resolve.flow.management.commands.register()` or `resolve.flow.executors.run.BaseFlowExecutor.run` or its derivatives):

- progress of type `basic:float` (from 0.0 to 1.0)
- proc of type `basic:group` containing:
 - stdout of type `basic:text`
 - rc of type `basic:integer`
 - task of type `basic:string` (celery task id)
 - worker of type `basic:string` (celery worker hostname)
 - runtime of type `basic:string` (runtime instance hostname)
 - pid of type `basic:integer` (process ID)

permission_group

permission group for the object

persistence

Persistence of *Data* objects created with this process. It can be one of the following:

- *PERSISTENCE_RAW*
- *PERSISTENCE_CACHED*
- *PERSISTENCE_TEMP*

Note: If persistence is set to *PERSISTENCE_CACHED* or *PERSISTENCE_TEMP*, the process must be idempotent.

requirements

process requirements

run

process command and environment description for internal use

Handling:

- schema defined by: *dev*
- default by: *dev*
- changable by: *dev*

scheduling_class

process scheduling class

slug

URL slug

type

data type

version

process version

Storage model

Postgres ORM model for storing JSON.

```
class resolwe.flow.models.Storage(*args, **kwargs)
```

Postgres model for storing storages.

exception DoesNotExist

exception MultipleObjectsReturned

contributor

user that created the entry

created

creation date and time

data

corresponding data objects

json

actual JSON stored

modified

modified date and time

name

object name

objects = <django.db.models.manager.StorageManagerFromPermissionQuerySet object>

storage manager

slug

URL slug

version

process version

Secret model

Postgres ORM model for storing secrets.

```
class resolwe.flow.models.Secret(*args, **kwargs)
```

Postgres model for storing secrets.

ProcessMigrationHistory model

Postgres ORM model for storing proces migration history.

```
class resolwe.flow.models.ProcessMigrationHistory(*args, **kwargs)
```

Model for storing process migration history.

exception DoesNotExist

exception MultipleObjectsReturned

created
creation date and time

metadata
migration-specific metadata

migration
migration identifier

DataMigrationHistory model

Postgres ORM model for storing data migration history.

```
class resolve.flow.models.DataMigrationHistory(*args, **kwargs)
    Model for storing data migration history.
exception DoesNotExist
exception MultipleObjectsReturned
created
    creation date and time
metadata
    migration-specific metadata
migration
    migration identifier
```

AnnotationGroup model

Postgres ORM model for storing annotation group data.

```
class resolve.flow.models.annotations.AnnotationGroup(*args, **kwargs)
    Group of annotation fields.
exception DoesNotExist
exception MultipleObjectsReturned
label
    the label of the annotation group
name
    the name of the annotation group
sort_order
    the sorting order among annotation groups
```

AnnotationField model

Postgres ORM model for storing annotation field data.

class `resolwe.flow.models.annotations.AnnotationField(*args, **kwargs)`

Annotation field.

exception `DoesNotExist`

exception `MultipleObjectsReturned`

static `add_to_collection(source: Collection, destination: Collection)`

Add fields from the source to the destination collection.

property `annotation_type: AnnotationType`

Get the field type as enum.

description

user visible field description

classmethod `field_from_path(path: str) → AnnotationField`

Get the field id from the field path.

Raises

ValidationError – when field does not exist.

group

the annotation group this field belongs to

static `group_field_from_path(path: str) → List[str]`

Return the group and field name from path.

classmethod `id_from_path(path: str) → int | None`

Get the field id from the field path.

label

user visible field name

label_by_value(`label: str`) → `str`

Get the value by label.

When no value is found the label is returned.

name

the name of the annotation fields

required

is this field required

revalidate_values()

Revalidate all annotation values.

Raises

ValidationError – when validation fails.

save(*args, **kwargs)

Recompute the labels for annotation values if vocabulary changes.

Raises

ValidationError – when vocabulary changes so that annotation values are no longer valid.

sort_order

the sorting order among annotation fields

type

the type of the annotation field

validator_regex

optional regular expression for validation

vocabulary

optional map of valid values to labels

AnnotationPreset model

Postgres ORM model for storing annotation presets data.

class `resolwe.flow.models.annotations.AnnotationPreset(*args, **kwargs)`

The named set of annotation fields.

The presets have permissions.

exception DoesNotExist**exception MultipleObjectsReturned****contributor**

user that created the entry

created

creation date and time

fields

the fields belonging to this preset

modified

modified date and time

name

object name

permission_group

permission group for the object

slug

URL slug

version

process version

AnnotationValue model

Postgres ORM model for storing annotation values.

class `resolwe.flow.models.annotations.AnnotationValue(*args, **kwargs)`

The value of the annotation.

exception `DoesNotExist`

exception `MultipleObjectsReturned`

entity: `Entity`

the entity this field belongs to

field: `AnnotationField`

the field this field belongs to

static `from_path(entity_id: int, path: str) → AnnotationValue | None`

Get the annotation value from the path.

has_permission(permission: Permission, user) → bool

Return if user permission on this object.

The permission is checked on the entity.

property label: `Any`

Return the cached label.

modified

the date when field was last modified

recompute_label()

Recompute label from value and set it to the model instance.

save(*args, **kwargs)

Save the annotation value after validation has passed.

validate()

Validate the given value.

The validation is always done in full, all errors are gathered and sent in the response.

Raises

ValidationError – when the validation fails.

property value: `str | int | float | date`

Get the actual value.

History model

Basic model for history tracking.

DataHistory model

Model for tracking changes in Data objects.

```
class resolve.flow.models.history.DataHistory(*args, **kwargs)
```

Track Data history.

exception DoesNotExist

exception MultipleObjectsReturned

datum

Pointer to the data object, can be null

deleted

When the object was deleted

valid

When this object is valid for

CollectionHistory model

Model for tracking changes in Collection objects.

```
class resolve.flow.models.history.CollectionHistory(*args, **kwargs)
```

Track Collection history.

exception DoesNotExist

exception MultipleObjectsReturned

datum

Pointer to the collection object, can be null

deleted

When the object was deleted

valid

When this object is valid for

TrackChange model

Basic model to track changes to particular model fields.

```
class resolve.flow.models.history.TrackChange(*args, **kwargs)
```

Base class to track the changes in the model.

class Meta

TrackChange meta options.

get_value(*field_name: str*) → Any

Return the value used to track changes.

objects: Manager

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

static preprocess_value(*instance: Model, field_name: str*) → *Any*

Preprocess the value before saving.

classmethod process_change(*instance: HistoryMixin, field_name: str*)

Create new track change object.

timestamp

Timestamp of the change

DataSlugChange model

Track data slug changes.

class `resolwe.flow.models.history.DataSlugChange(*args, **kwargs)`

Track data slug changes.

exception DoesNotExist

exception MultipleObjectsReturned

history

Reference to the history object

objects: Manager = <django.db.models.manager.Manager object>

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

timestamp

Timestamp of the change

value: Any

New value.

CollectionSlugChange model

Track collection slug changes.

class `resolwe.flow.models.history.CollectionSlugChange(*args, **kwargs)`

Track collection slug changes.

exception DoesNotExist

exception MultipleObjectsReturned

history

Reference to the history object

objects: Manager = <django.db.models.manager.Manager object>

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

timestamp

Timestamp of the change

value: Any

New value.

DataNameChange model

Track data name changes.

```
class resolve.flow.models.history.DataNameChange(*args, **kwargs)
    Track data name changes.
    exception DoesNotExist
    exception MultipleObjectsReturned
    history
        Reference to the history object
    objects: Manager = <django.db.models.manager.Manager object>
        Defined to keep mypy from reporting errors (objects do not exist since class is abstract).
    timestamp
        Timestamp of the change
    value: Any
        New value.
```

CollectionNameChange model

Track collection name changes.

```
class resolve.flow.models.history.CollectionNameChange(*args, **kwargs)
    Track collection name changes.
    exception DoesNotExist
    exception MultipleObjectsReturned
    history
        Reference to the history object
    objects: Manager = <django.db.models.manager.Manager object>
        Defined to keep mypy from reporting errors (objects do not exist since class is abstract).
    timestamp
        Timestamp of the change
    value: Any
        New value.
```

CollectionChange model

Track collection changes.

```
class resolve.flow.models.history.CollectionChange(*args, **kwargs)
    Track collection changes.
    exception DoesNotExist
    exception MultipleObjectsReturned
```

history

Reference to the history object

objects: `Manager = <django.db.models.manager.Manager object>`

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

static preprocess_value(*instance: Model, field_name: str*) → *Any*

Preprocess the value before saving.

timestamp

Timestamp of the change

value: `Any`

New value.

SizeChange model

Track size changes.

class `resolwe.flow.models.history.SizeChange(*args, **kwargs)`

Track size changes.

exception `DoesNotExist`

exception `MultipleObjectsReturned`

history

Reference to the history object

objects: `Manager = <django.db.models.manager.Manager object>`

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

timestamp

Timestamp of the change

value: `Any`

New value

ProcessingHistory model

Track processing data objects.

class `resolwe.flow.models.history.ProcessingHistory(*args, **kwargs)`

Track processing changes.

exception `DoesNotExist`

exception `MultipleObjectsReturned`

allocated_cpu

Number of CPU allocated to the process

allocated_memory

Amount of memory (in bytes) allocated to the process

get_value(*field_name: str*) → *Any*

Return the value used to track changes.

history

Reference to the history object

interval

Processing interval

node_cpu

Number of CPU on machine running the process

node_memory

Amount of memory on machine running the process

objects: Manager = <django.db.models.manager.Manager object>

Defined to keep mypy from reporting errors (objects do not exist since class is abstract).

classmethod process_change(*instance: HistoryMixin, field_name: str*)

Preprocess the value before saving.

timestamp

Timestamp of the change

Utility functions

1.6.6 Flow Utilities

Resolwe Exceptions Utils

Utils functions for working with exceptions.

`resolwe.flow.utils.exceptions.resolwe_exception_handler`(*exc, context*)

Handle exceptions raised in API and make them nicer.

To enable this, you have to add it to the settings:

```
REST_FRAMEWORK = {
    'EXCEPTION_HANDLER': 'resolwe.flow.utils.exceptions.resolwe_exception_
↪ handler',
}
```

Statistics

Various statistical utilities, used mostly for manager load tracking.

class `resolwe.flow.utils.stats.NumberSeriesShape`

Helper class for computing characteristics for numerical data.

Given a series of numerical data, the class will keep a record of the extremes seen, arithmetic mean and standard deviation.

to_dict()

Pack the stats computed into a dictionary.

update(*num*)

Update metrics with the new number.

class `resolwe.flow.utils.stats.SimpleLoadAvg`(*intervals*)

Helper class for a sort of load average based on event times.

Given a series of queue depth events, it will compute the average number of events for three different window lengths, emulating a form of ‘load average’. The calculation itself is modelled after the Linux scheduler, with a 5-second sampling rate. Because we don’t get consistent (time-wise) samples, the sample taken is the average of a simple moving window for the last 5 seconds; this is to avoid numerical errors if actual time deltas were used to compute the scaled decay.

add(*count*, *timestamp=None*)

Add a value at the specified time to the series.

Parameters

- **count** – The number of work items ready at the specified time.
- **timestamp** – The timestamp to add. Defaults to None, meaning current time. It should be strictly greater (newer) than the last added timestamp.

to_dict()

Pack the load averages into a nicely-keyed dictionary.

1.6.7 Flow Management

Register Processes

class `resolwe.flow.management.commands.register.Command`(*stdout=None*, *stderr=None*,
no_color=False, *force_color=False*)

Register processes.

add_arguments(*parser*)

Command arguments.

find_descriptor_schemas(*schema_file*)

Find descriptor schemas in given path.

find_schemas(*schema_path*, *schema_type='process'*, *verbosity=1*)

Find schemas in packages that match filters.

handle(**args*, ***options*)

Register processes.

register_descriptors(*descriptor_schemas*, *user*, *force=False*, *verbosity=1*)

Read and register descriptors.

register_processes(*process_schemas*, *user*, *force=False*, *verbosity=1*)

Read and register processors.

retire(*process_schemas*)

Retire obsolete processes.

Remove old process versions without data. Find processes that have been registered but do not exist in the code anymore, then:

- If they do not have data: remove them

- If they have data: flag them not active (`is_active=False`)

`valid(instance, schema)`

Validate schema.

1.6.8 Resolwe Test Framework

Resolwe Test Cases

`class resolwe.test.TestCaseHelpers(methodName='runTest')`

Mixin for test case helpers.

`assertAlmostEqualGeneric(actual, expected, msg=None)`

Assert almost equality for common types of objects.

This is the same as `assertEqual()`, but using `assertAlmostEqual()` when floats are encountered inside common containers (currently this includes `dict`, `list` and `tuple` types).

Parameters

- **actual** – object to compare
- **expected** – object to compare against
- **msg** – optional message printed on failures

`assertAnnotation(entity, path, value)`

Compare the entity annotation with the given value.

`keep_data(mock_purge=True)`

Do not delete output files after tests.

`setUp()`

Prepare environment for test.

`class resolwe.test.TransactionTestCase(methodName='runTest')`

Base class for writing Resolwe tests not enclosed in a transaction.

It is based on Django's `TransactionTestCase`. Use it if you need to access the test's database from another thread/process.

`setUp()`

Initialize test data.

`class resolwe.test.TestCase(methodName='runTest')`

Base class for writing Resolwe tests.

It is based on `TransactionTestCase` and Django's `TestCase`. The latter encloses the test code in a database transaction that is rolled back at the end of the test.

`class resolwe.test.ProcessTestCase(methodName='runTest')`

Base class for writing process tests.

It is a subclass of `TransactionTestCase` with some specific functions used for testing processes.

To write a process test use standard Django's syntax for writing tests and follow the next steps:

1. Put input files (if any) in `tests/files` directory of a Django application.
2. Run the process using `run_process()`.

3. Check if the process has the expected status using `assertStatus()`.
4. Check process's output using `assertFields()`, `assertFile()`, `assertFileExists()`, `assertFiles()` and `assertJSON()`.

Note: When creating a test case for a custom Django application, subclass this class and over-ride the `self.files_path` with:

```
self.files_path = os.path.join(os.path.dirname(os.path.abspath(__file__)), 'files')
```

Danger: If output files don't exist in `tests/files` directory of a Django application, they are created automatically. But you have to check that they are correct before using them for further runs.

assertDir(*obj, field_path, fn*)

Compare process output directory to correct compressed directory.

Parameters

- **obj** (*Data*) – object that includes the directory to compare
- **field_path** (*str*) – path to *Data* object's field with the file name
- **fn** (*str*) – file name (and relative path) of the correct compressed directory to compare against. Path should be relative to the `tests/files` directory of a Django application. Compressed directory needs to be in `tar.gz` format.

assertDirExists(*obj, field_path*)

Assert that a directory in the output field of the given object exists.

Parameters

- **obj** – object that includes the file for which to check if it exists
- **field_path** – directory name/path

assertDirStructure(*obj, field_path, dir_struct, exact=True*)

Assert correct tree structure in output field of given object.

Only names of directories and files are asserted. Content of files is not compared.

Parameters

- **obj** (*Data*) – object that includes the directory to compare
- **dir_path** (*str*) – path to the directory to compare
- **dir_struct** (*dict*) – correct tree structure of the directory. Dictionary keys are directory and file names with the correct nested structure. Dictionary value associated with each directory is a new dictionary which lists the content of the directory. Dictionary value associated with each file name is `None`
- **exact** (*bool*) – if `True` tested directory structure must exactly match *dir_struct*. If `False` *dir_struct* must be a partial structure of the directory to compare

assertFields(*obj, path, value*)

Compare object's field to the given value.

The file size is ignored. Use `assertFile` to validate file contents.

Parameters

- **obj** (*Data*) – object with the field to compare
- **path** (*str*) – path to *Data* object's field
- **value** (*str*) – desired value of *Data* object's field

assertFile(*obj, field_path, fn, **kwargs*)

Compare a process's output file to the given correct file.

Parameters

- **obj** (*Data*) – object that includes the file to compare
- **field_path** (*str*) – path to *Data* object's field with the file name
- **fn** (*str*) – file name (and relative path) of the correct file to compare against. Path should be relative to the `tests/files` directory of a Django application.
- **compression** (*str*) – if not `None`, files will be uncompressed with the appropriate compression library before comparison. Currently supported compression formats are *gzip* and *zip*.
- **filter** (*FunctionType*) – function for filtering the contents of output files. It is used in `itertools.filterfalse()` function and takes one parameter, a line of the output file. If it returns `True`, the line is excluded from comparison of the two files.
- **sort** (*bool*) – if set to `True`, basic sort will be performed on file contents before computing hash value.

assertFileExists(*obj, field_path*)

Ensure a file in the given object's field exists.

Parameters

- **obj** (*Data*) – object that includes the file for which to check if it exists
- **field_path** (*str*) – path to *Data* object's field with the file name/path

assertFiles(*obj, field_path, fn_list, **kwargs*)

Compare a process's output file to the given correct file.

Parameters

- **obj** (*Data*) – object which includes the files to compare
- **field_path** (*str*) – path to *Data* object's field with the list of file names
- **fn_list** (*list*) – list of file names (and relative paths) of files to compare against. Paths should be relative to the `tests/files` directory of a Django application.
- **compression** (*str*) – if not `None`, files will be uncompressed with the appropriate compression library before comparison. Currently supported compression formats are *gzip* and *zip*.
- **filter** (*FunctionType*) – Function for filtering the contents of output files. It is used in `itertools.filterfalse` function and takes one parameter, a line of the output file. If it returns `True`, the line is excluded from comparison of the two files.
- **sort** (*bool*) – if set to `True`, basic sort will be performed on file contents before computing hash value.

assertFilesExist(*obj, field_path*)

Ensure files in the given object's field exists.

Parameters

- **obj** (*Data*) – object that includes list of files for which to check existence
- **field_path** (*str*) – path to *Data* object's field with the file name/path

assertJSON(*obj, storage, field_path, file_name*)

Compare JSON in Storage object to the given correct JSON.

Parameters

- **obj** (*Data*) – object to which the *Storage* object belongs
- **storage** (*Storage* or *str*) – object or id which contains JSON to compare
- **field_path** (*str*) – path to JSON subset in the *Storage*'s object to compare against. If it is empty, the entire object will be compared.
- **file_name** (*str*) – file name (and relative path) of the file with the correct JSON to compare against. Path should be relative to the `tests/files` directory of a Django application.

Note: The given JSON file should be compressed with *gzip* and have the `.gz` extension.

assertStatus(*obj, status*)

Check if object's status is equal to the given status.

Parameters

- **obj** (*Data*) – object for which to check the status
- **status** (*str*) – desired value of object's *status* attribute

property files_path

Path to test files.

get_json(*file_name, storage*)

Return JSON saved in file and test JSON to compare it to.

The method returns a tuple of the saved JSON and the test JSON. In your test you should then compare the test JSON to the saved JSON that is committed to the repository.

The storage argument could be a Storage object, Storage ID or a Python dictionary. The test JSON is assigned a json field of the Storage object or the complete Python dictionary (if a dict is given).

If the file does not exist it is created, the test JSON is written to the new file and an exception is raised.

Parameters

- **file_name** (*str*) – file name (and relative path) of a JSON file. Path should be relative to the `tests/files` directory of a Django app. The file name must have a `.gz` extension.
- **storage** (*Storage, str* or *dict*) – Storage object, Storage ID or a dict.

Returns

(reference JSON, test JSON)

Return type

tuple

preparation_stage()

Context manager to mark input preparation stage.

run_process(*process_slug*, *input_*={}, *assert_status*='OK', *descriptor*=None, *descriptor_schema*=None, *verbosity*=0, *tags*=None, *contributor*=None, *collection*=None, *process_resources*=None)

Run the specified process with the given inputs.

If input is a file, file path should be given relative to the `tests/files` directory of a Django application. If `assert_status` is given, check if `Data` object's status matches it after the process has finished.

Note: If you need to delay calling the manager, you must put the desired code in a `with transaction.atomic()` block.

Parameters

- **process_slug** (*str*) – slug of the `Process` to run
- **input_** (*dict*) – `Process`'s input parameters

Note: You don't have to specify parameters with defined default values.

- **assert_status** (*str*) – desired status of the `Data` object
- **descriptor** (*dict*) – descriptor to set on the `Data` object
- **descriptor_schema** (*dict*) – descriptor schema to set on the `Data` object
- **tags** (*list*) – list of tags that will be added to the created `Data` object

Returns

object created by `Process`

Return type

`Data`

run_processor(*args, **kwargs)

Run process.

Deprecated method: use `run_process`.

setUp()

Initialize test data.

tearDown()

Clean up after the test.

class `resolve.test.TransactionResolveAPITestCase`(*methodName*='runTest')

Base class for testing Resolve REST API.

This class is derived from Django REST Framework's `APITransactionTestCase` class and has implemented some basic features that makes testing Resolve API easier. These features includes following functions:

_get_list(*user*=None, *query_params*={})

Make GET request to `self.list_view` view.

If `user` is not `None`, the given user is authenticated before making the request.

Parameters

user (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

`_get_detail(pk, user=None, query_params={})`

Make GET request to `self.detail_view` view.

If `user` is not `None`, the given user is authenticated before making the request.

Parameters

- **pk** (*int*) – Primary key of the corresponding object
- **user** (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

`_post(data={}, user=None, query_params={})`

Make POST request to `self.list_view` view.

If `user` is not `None`, the given user is authenticated before making the request.

Parameters

- **data** (*dict*) – data for posting in request's body
- **user** (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

`_patch(pk, data={}, user=None, query_params={})`

Make PATCH request to `self.detail_view` view.

If `user` is not `None`, the given user is authenticated before making the request.

Parameters

- **pk** (*int*) – Primary key of the corresponding object
- **data** (*dict*) – data for posting in request's body
- **user** (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

`_delete(pk, user=None, query_params={})`

Make DELETE request to `self.detail_view` view.

If `user` is not `None`, the given user is authenticated before making the request.

Parameters

- **pk** (*int*) – Primary key of the corresponding object
- **user** (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

`_detail_permissions(pk, data={}, user=None)`

Make POST request to `self.detail_view` view.

If user is not *None*, the given user is authenticated before making the request.

Parameters

- **pk** (*int*) – Primary key of the corresponding object
- **data** (*dict*) – data for posting in request's body
- **user** (*User* or *None*) – User to authenticate in request

Returns

API response object

Return type

Response

It also has included 2 views made from referenced DRF's *ViewSet*. First mimic list view and has following links between request's methods and *ViewSet*'s methods:

- GET -> `list`
- POST -> `create`

Second mimic detail view and has following links between request's methods and *ViewSet*'s methods:

- GET -> `retrieve`
- PUT -> `update`
- PATCH -> `partial_update`
- DELETE -> `destroy`
- POST -> `permissions`

If any of the listed methods is not defined in the *ViewSet*, corresponding link is omitted.

Note: `self.viewset` (instance of DRF's *ViewSet*) and `self.resource_name` (string) must be defined before calling super `setUp` method to work properly.

`self.factory` is instance of DRF's *APIRequestFactory*.

`assertKeys(data, wanted)`

Assert dictionary keys.

`detail_permissions(pk)`

Get detail permissions url.

detail_url(pk)

Get detail url.

property list_url

Get list url.

setUp()

Prepare data.

class `resolwe.test.ResolweAPITestCase`(*methodName='runTest'*)

Base class for writing Resolwe API tests.

It is based on *TransactionResolweAPITestCase* and Django's *TestCase*. The latter encloses the test code in a database transaction that is rolled back at the end of the test.

Resolwe Test Helpers and Decorators

`resolwe.test.utils.check_docker()`

Check if Docker is installed and working.

Returns

(indicator of the availability of Docker, reason for unavailability)

Return type

`tuple(bool, str)`

`resolwe.test.utils.check_installed(command)`

Check if the given command is installed.

Parameters

command (*str*) – name of the command

Returns

(indicator of the availability of the command, message saying command is not available)

Return type

`tuple(bool, str)`

`resolwe.test.utils.create_data_location(subpath=None)`

Create equivalent of old `DataLocation` object.

When argument is `None`, store the ID of the file storage object in the subpath.

`resolwe.test.utils.is_testing()`

Return current testing status.

This assumes that the Resolwe test runner is being used.

`resolwe.test.utils.with_custom_executor(wrapped=None, **custom_executor_settings)`

Decorate unit test to run processes with a custom executor.

Parameters

custom_executor_settings (*dict*) – custom `FLOW_EXECUTOR` settings with which you wish to override the current settings

`resolwe.test.utils.with_docker_executor(wrapped=None)`

Decorate unit test to run processes with the Docker executor.

`resolwe.test.utils.with_null_executor(wrapper=None, enabled=None, adapter=None, proxy=<class 'FunctionWrapper'>)`

Decorate unit test to run processes with the Null executor.

`resolwe.test.utils.with_resolwe_host(wrapper=None, enabled=None, adapter=None, proxy=<class 'FunctionWrapper'>)`

Decorate unit test to give it access to a live Resolwe host.

Set `RESOLWE_HOST_URL` setting to the address where the testing live Resolwe host listens to.

Note: This decorator must be used with a (sub)class of `LiveServerTestCase` which starts a live Django server in the background.

1.6.9 Resolwe Utilities

`class resolwe.utils.BraceMessage(fmt, *args, **kwargs)`

Log messages with the new `{}`-string formatting syntax.

Note: When using this helper class, one pays no significant performance penalty since the actual formatting only happens when (and if) the logged message is actually outputted to a log by a handler.

Example of usage:

```
from resolwe.utils import BraceMessage as __
logger.error(__("Message with {0} {name}", 2, name="placeholders"))
```

Source: <https://docs.python.org/3/howto/logging-cookbook.html#use-of-alternative-formatting-styles>.

1.7 Resolwe Flow Design

The Resolwe Flow workflow engine comprises the execution framework and other layers which make up the internal data model and facilitate dependency resolution, permissions enforcement and process filtering.

1.7.1 Overview

Execution Framework

Flow consists of a number of active services which need to be running before job execution can proceed.

The core message transport and coordination facility, as currently used, is `Redis`. It serves as the central status hub for keeping track of shared dynamic information used by parts of the framework, and as a contact point for those parts of the framework that run remotely. These connect to well-known ‘channels’ (specially named Redis list objects), into which they can deposit JSON-formatted messages and commands.

Flow’s execution manager, or just the ‘manager’, is an independent service which runs as a `Django Channels` event consumer. When objects are added to the database to be executed, they will trigger events for the appropriate channels. These will be processed by the manager, which will carry out all the preparatory tasks necessary to start execution and

then communicate with a concrete workload management system so that the job can eventually be scheduled and run on a worker node.

Finally, the jobs are executed by the aptly named ‘executors’. These are run on worker nodes and act as local execution managers: preparing a controlled execution environment, running the job’s code, collecting results and communicating them back to the manager which stores them in the database.

Utility Layers

Django’s facilities are used for interfacing with the database, thus all models used in Flow are actually Django Model objects. The most important two models are the *Data* model and the *Process* model.

A Data object represents a single instance of data to be processed, i.e. a node in the flow graph being executed. It contains properties which mainly concern execution, such as various process and task IDs, output statuses and the results produced by executors.

A Process object represents the way in which its Data object will be ‘executed’, i.e. the type of node in the flow graph and the associated code. It contains properties defining its relationship to other nodes in the graph currently being executed, the permissions that control access rights for users and other processes, and the actual code that is run by the executors.

The code in the process object can be either final code that is already ready for execution, or it can be a form of template, for which an ‘expression engine’ is needed. An expression engine (the only one currently in use is *Jinja*) pre-processes the process’ code to produce text that can then be executed by an ‘execution engine’.

An execution engine is, simply put, the interpreter that will run the processed code, just after an additional metadata discovery step. It is done by the execution engine because the encoding might be language-dependent. The properties to be discovered include process resource limits, secret requirements, etc. These properties are passed on to the executor, so that it can properly set up the target environment. The only currently supported execution engine is Bash.

1.7.2 Technicalities

The Manager

Being a Django Channels consumer application, the Flow Manager is entirely event-driven and mostly contextless. The main input events are data object creation, processing termination and intermediate executor messages. Once run, it consists of two distinct servers and a modularized connection framework used to interface with workload managers used by the deployment site.

Dispatcher

This is the central job scheduler. On receipt of an appropriate event through Django Channels (in this service, only data object creation and processing termination), the dispatcher will scan the database for outstanding data objects. For each object found to still not be processed, dependencies will be calculated and scanned for completion. If all the requirements are satisfied, its execution cycle will commence. The manager-side of this cycle entails job pre-processing and a part of the environment preparation steps:

- The data object’s process is loaded, its code preprocessed with the configured expression engine and the result of that fed into the selected execution engine to discover further details about the process’ environmental requirements (resource limits).
- The runtime directories on the global shared file system are prepared: file dependencies are copied out from the database, the process’ processed code (as output by the expression engine) is stored into a file that the executor will run.

- The executor platform is created by copying the Flow Executor source code to the destination (per-data) directories on the shared file system, along with serialized (JSON) settings and support metadata (file lists, directory paths, Docker configuration and other information the configured executor will need for execution).
- After all this is done, control is handed over to the configured ‘workload connector’, see below for a description.

Listener

As the name might imply to some, the purpose of the listener is to listen for status updates and distressing thoughts sent by executors. The service itself is an independent (*i.e.* not Django Channels-based) process which waits for events to arrive on the executor contact point channels in Redis.

The events are JSON-formatted messages and include:

- processing status updates, such as execution progress and any computed output values,
- spawn commands, with which a process can request the creation of new data objects,
- execution termination, upon which the listener will finalize the Data object in question: delete temporary files from the global shared file system, update process exit code fields in the database, store the process’ standard output and standard error sent by the executor and notify the dispatcher about the termination, so that any state internal to it may be updated properly,
- ancillary status updates from the executor, such as logging. Because executors are running remotely with respect to the manager’s host machine, they may not have access to any centralized logging infrastructure, so the listener is used as a proxy.

Workload Connectors

Workload connectors are thin glue libraries which communicate with the concrete workload managers used on the deployment site. The dispatcher only contains logic to prepare execution environments and generate the command line necessary to kick off an executor instance. The purpose of the workload connector is to submit that command line to the workload manager which will then execute it on one of its worker nodes. The currently supported workload managers are [Celery](#), [SLURM](#) and a simple local dummy for test environments.

The Executor

The Flow Executor is the program that controls Process execution on a worker node managed by the site workload manager, for which it is a job. Depending on the configured executor, it further prepares an execution environment, configures runtime limitations enforced by the system and spawns the code in the Process object. The currently supported executor types are a simple local executor for testing deployments and a [Docker](#)-based one.

Once started, the executor will carry out any additional preparation based on its type (*e.g.* the Docker executor constructs a command line to create an instance of a pre-prepared Docker container, with all necessary file system mappings and communication conduits). After that, it executes the Process code as prepared by the manager, by running a command to start it (this need not be anything more complicated than a simple *subprocess.Popen*).

Following this, the executor acts as a proxy between the process and the database by relaying messages generated by the process to the manager-side listener. When the process is finished (or when it terminates abnormally), the executor will send a final cleanup message and terminate, thus finishing the job from the point of view of the workload manager.

1.7.3 Example Execution, from Start to Finish

- Flow services are started: the dispatcher Django Channels application and the listener process.
- The user, through any applicable intricacy, creates a Data object.
- Django signals will fire on creation and submit a data scan event to the dispatcher through Django Channels.
- The dispatcher will scan the database for outstanding data objects (alternatively, only for a specific one, given an ID). The following steps are then performed for each discovered data object whose dependencies are all processed:
 - The runtime directory is populated with data files, executor source and configuration files.
 - The process code template is run through an expression engine to transform it into executable text. This is also scanned with an execution engine to discover runtime resource limits and other process-local configuration.
 - A command line is generated which can be run on a processing node to start an executor.
 - The command line is handed over to a workload connector, which submits it as a job to the workload manager installed on the site.
 - At this point, the dispatcher's job for this data object is done. Eventually, the workload manager will start processing the submitted job, thereby spawning an executor.
 - The executor will prepare a safe runtime context, such as a Docker container, configure it with appropriate communication channels (stdin/out redirection or sockets) and run the command to execute the process code.
 - The code executes, periodically generating status update messages. These are received by the executor and re-sent to the listener. The listener responds appropriately, updating database fields for the data object, notifying the dispatcher about lifetime events or forwarding log messages to any configured infrastructure.
 - Once the process is done, the executor will send a finalizing command to the listener and terminate.
 - The listener will notify the dispatcher about the termination and finalize the database status of this data object (processing exit code, outputs).
 - The dispatcher will update processing states and statistics, and re-scan the database for data objects which might have dependencies on the one that just finished and could therefore potentially be started up.

1.8 Resolwe Storage Framework

Resolwe Storage Framework is storage management system for Resolwe Flow Design. Currently it supports storing data to local filesystem, Google Cloud Storage and Amazon Simple Storage Service.

1.8.1 Testing connectors

When testing connectors make sure there are credentials available (see sample settings below for details). These tests start with `storage_credentials_test_` and are skipped by default when running `tox`. Since `tox` does not allow decrypting files on pull requests these tests would fail on all pull requests blocking merge.

1.8.2 Example settings

```

# Storage connector sample settings.
STORAGE_LOCAL_CONNECTOR = 'local'
STORAGE_CONNECTORS = {
    # Code assumes that connector named 'local' exists and points
    # to the location FLOW_EXECUTOR["DATA_DIR"].
    # If this is not true BAD THING WILL HAPPEN.
    "local": {
        "connector": "resolwe.storage.connectors.localconnector.LocalFilesystemConnector
→",
        "config": {
            "priority": 0, # If ommited, default 100 is used
            "path": FLOW_EXECUTOR["DATA_DIR"],
            # Public URL from where data is served
            "public_url": "/local_data",
            # Delete from here after delay days from last access to this storage
            # location and when min_other_copies of data exist on other
            # locations.
            "delete": {
                "delay": 2, # in days
                "min_other_copies": 2,
            },
            "copy": {
                # Override default settings for this data_slug.
                "data_slug": {
                    "00hrfastqgz-66": {
                        "delay": 0
                    }
                },
                "delay": 10,
            }
        }
    },
    "S3": {
        "connector": "resolwe.storage.connectors.s3connector.AwsS3Connector",
        "config": {
            "priority": 10,
            "bucket": "genialis-test-storage",
            "copy": { # copy here from delay days from creation of filestorage object
                "delay": 5, # in days
            },
            # Region name is needed to generate valid pre-signed urls.
            "region_name": "eu-central-1",
            "credentials": os.path.join(
                PROJECT_ROOT, "testing_credentials_s3.json"
            ),
        },
    },
    "GCS": {
        "connector": "resolwe.storage.connectors.googleconnector.GoogleConnector",
        "config": {
            "bucket": "genialis-test-storage",

```

(continues on next page)

(continued from previous page)

```
        "credentials": os.path.join(
            PROJECT_ROOT, "testing_credentials_gcs.json"
        ),
        "copy": {
            "delay": 2, # in days
        },
    },
},
}
```

1.9 Change Log

All notable changes to this project are documented in this file. This project adheres to [Semantic Versioning](#).

1.9.1 38.3.2 - 2024-01-25

Fixed

- Correctly decorate class based views

1.9.2 38.3.1 - 2024-01-25

Changed

- Do not require CSRF token to access uri resolver endpoint

1.9.3 38.3.0 - 2024-01-11

Fix

- Correctly clear Redis cache on data restart

Add

- Add `clear_redis_cache` management command
- Add modified field to the `AnnotationValue` model and expose it in API

1.9.4 38.2.0 - 2023-12-15

Fix

- Add default value for `FLOW_PROCESSES_ALLOW_LIST` and `FLOW_PROCESSES_IGNORE_LIST` in case of missing settings.
- User defined slug must not be changed
- Raise `ValidationError` when trying to get non-existent `AnnotationField`
- Propagate arguments to parent save method in `BaseCollection` model

Added

- Add `extend_lock` method to redis cache in listener
- Add extending processing lock task to the listener
- Add `FLOW_KUBERNETES_PULL_POLICY` setting to set the image pull policy in Kubernetes instead of using `Always`

Changed

- Wrap listener command processing in a transaction
- Return `False` when `has_perm` is called with `obj` set to `None`, since object level permissions are granted using Django `ModelBackend` authentication backend
- Use `keys` instead of `scan_iter` in `clear` method of redis cache, since `scan_iter` is much slower
- Do not allow deleting annotation field if annotation values are present

1.9.5 38.1.0 - 2023-11-21

Changed

- Redis cache in listener is updated when data fields are retrieved from the database

Added

- Add `processes allow` and `ignore list` to dispatcher, controlled by environmental variables `FLOW_PROCESSES_ALLOW_LIST` and `FLOW_PROCESSES_IGNORE_LIST`
- Allow overriding the maximal number of commands listener can process concurrently
- Support bulk create/update/delete on `AnnotationValues` endpoint

Fixed

- Store random postfix to redis for use at cleanup time
- Fix `_save_data` method in listener that caused changes to started and progress fields to be ignored
- Already processed messages in listener are ignored for one day so messages are not processed twice

Changed

- Bulk annotations on entity endpoint now accept field path instead of id

1.9.6 38.0.0 - 2023-11-13

Added

- Add `restart` method to the Data model
- Add `restart` endpoint to `DataViewSet`
- Add history tracking for Data and Collection objects
- Add `bulk_delete` to data, entity and collection endpoints

Changed

- **BACKWARD INCOMPATIBLE:** Remove arguments `inherit_collection` and `inherit_entity` from `duplicate` endpoints
- The method `run_process` in Python processes returns the created Data object
- The docker container names have random postfix to support restarting data object
- The kubernetes resources have random postfix to support restarting data object
- Remove `inherit_collection` and `inherit_entity` parameters from `duplicate` methods on Entity and Data models and corresponding query sets
- Add filtering and searching to annotation field admin page
- Annotation values are re-validated if annotation field type is changed

Fixed

- Add default ordering on `AnnotationValue` model
- Create history object for `Collection` if none exists
- Use the same postfix for all kubernetes volumes for the given data object
- Fix `handle_progress` handler in the listener
- Fix moving data to entity in collection also sets collection on the data object

1.9.7 37.0.2 - 2023-10-26

Fixed

- Set value to AnnotationValue object on duplication when it is created
- Send post_duplicate signal only on successful duplication

Changed

- Simplify permission checks on AnontationValue endpoint

1.9.8 37.0.1 - 2023-10-23

Changed

- Removed descriptor and DescriptorSchema fields from Entity endpoint

1.9.9 37.0.0 - 2023-10-23

Added

- Allow sorting entities by collection name

Changed

- **BACKWARD INCOMPATIBLE:** Command re-annotate-entity creates annotation objects instead of updating entity descriptor
- **BACKWARD INCOMPATIBLE:** Method update_entity_descriptor in Python processes creates annotation objects instead of updating the entity descriptor
- Add annotations property to the Entity object in Python processes
- Allow filtering AnnotationField by ids
- Allow filtering AnnotationField by entity
- Allow filtering AnnotationValue by group name
- Allow create/update/delete REST API calls for AnnotationValue objects
- Remove collection field from AnnotationField serializer

1.9.10 36.1.0 - 2023-09-14

Changed

- Relation slug can be updated
- Anonymous user can access the annotation field endpoint
- Annotations contain the field label
- Filtering entities by annotations now filters by label and value fields

Added

- Add `clean_observers` management command to clean old subscriptions
- Allow filtering annotation values endpoint by label

1.9.11 36.0.0 - 2023-07-17

Added

- Add `slug` to the relation serializer
- Add support for Python 3.11

Changed

- **BACKWARD INCOMPATIBLE:** Bump Django requirement to version 4.2
- **BACKWARD INCOMPATIBLE:** Bump channels requirement to version 4.0
- Bump other requirements to the latest versions
- Replace `django-fernet-fields` with `django-fernet-fields-v2` to support Django 4.2
- Bump PostgreSQL version to 14 since 10 is no longer supported
- Bump dependencies in the communication container image to the latest versions
- Bump `psycopy` version to 3.1.9
- Upgrade Fedora base image to version 37
- Remove Django Priority Batch dependency
- Set filename with *Content-Disposition* header when downloading data from S3 to avoid issues with some browsers renaming the file due to media type

1.9.12 35.1.0 - 2023-06-19

Added

- Method `groups_with_permissions` was added to the `PermissionObject`
- Register annotation models with Django admin
- Enable editing `AnnotationPreset` permissions in the Django admin
- Extend `AnnotationPreset` from `BaseModel`
- Enable filtering `Entity` objects by annotation values
- Add required field to the `AnnotationField` model
- Add required field to the `AnnotationFieldSerializer` and `AnnotationFieldFilter`
- Enable sorting `Entity` by annotation values

Changed

- Annotation fields endpoint serializes the annotation group instead of only returning its id
- Do not allow removing required `AnnotationField` objects from the collection
- Add required annotation fields to collection when it is created
- Replace `add_fields_to_collection` and `remove_fields_from_collection` with `set_annotation_fields`

Fixed

- Filtering annotation values by entity ids
- Add `PluginManager` and `Plugin` classes to unify plugin management
- Add Redis cache management plugin manager and basic plugins
- Add the option to skip the response to the already processed command
- Add the ability to add locks to Redis and wait for them
- Lock the message in listener when processing it
- Duplicate filter loopup lists
- Fix `AnnotationPreset` contributor to point to the user model specified in the settings

1.9.13 35.0.0 - 2023-05-15

Added

- Add `delete` method to the background consumer
- Add `delete_background` method to data, entity and collection model and queryset
- Allow IP protocol sockets to be opened in the processing container by the seccomp security profile

Changed

- **BACKWARD INCOMPATIBLE:** Delete objects in the background in the data, entity and collection endpoints
- Bump Docker SDK version to fix requests/urllib incompatibility issue (<https://github.com/docker/docker-py/issues/3113>)
- Always use Docker default seccomp profile
- Increase resend timeout in processing container to 60 seconds

Fixed

- Close websocket requests without session id

1.9.14 34.1.0 - 2023-05-02

Fixed

- The contributor is assigned VIEW permission to the BackgroundTask

Added

- Add endpoints for new annotation models
- Send observer create signals on object duplication
- Add filtering by id to the background task endpoint

1.9.15 34.0.0 - 2023-04-17

Added

- Add BackgroundTask model to enable support for observable long running backend tasks
- Add BackgroundTaskConsumer class to handle starting background jobs
- Make Kubernetes overcommitment factors configurable via settings
- Support filtering by tags on relation API endpoint

Changed

- **BACKWARD INCOMPATIBLE** Remove the duplicated collection subfield from the entity field when serializing Data objects
- **BACKWARD INCOMPATIBLE:** Replace duplicate methods on Data, Entity and Collection objects with the call to the corresponding background duplicate job.
- **BACKWARD INCOMPATIBLE:** The duplicate endpoints on Data, Entity and Collection return the details of the BackgroundTask instance that is used to track the duplication status
- When subscribing to an object for notifications the CREATE notifications are also sent along. This is necessary to receive CREATE notifications triggered from the data inside the container on the container.
- Remove the redundant prefetch in the BaseCollectionViewSet
- Annotate the queryset in the BaseCollectionViewSet with data_statuses, entity_count and data_count attributes to reduce the number of queries
- Remove the redundant prefetch in the EntityViewSet
- Annotate the queryset in the EntityViewSet with data_statuses and data_count attributes to reduce the number of queries
- Modify the DataViewSet to remove redundant queries

- Add `duplicate_data`, `duplicate_entity` and `duplicate_collection` methods to `BackgroundTaskConsumer` class
- Set `is_active` to `True` for all available processes every time `register` management command is run even if the same version of the process already exists in the database
- The `FLOW_KUBERNETES_AFFINITY` variable is now a JSON dictionary with information about affinity for interactive and batch processes
- Fix values in observer `ChangeType` enum to prevent accidental value change
- Always start `BackgroundTaskConsumer` when running tests

Fixed

- Send observer notifications when data object is automatically assigned to the entity via `entity_type` in the process
- Remove observer from subscription if object is deleted
- Handle empty iterable in `observe_instance_container` method
- Remove call to `close` in observers consumer that caused `RuntimeError`
- Instances of `Observer` class could be duplicated when `object_id` was `None` due to PostgreSQL handling of `NULL` values
- Delete the observer subscription when object is deleted
- Check response status when bootstrapping executor and print out sensible error message on failure
- Do not run the test twice when it is tagged and contained in the changed test file

1.9.16 33.0.0 - 2023-03-13

Added

- Add `suppress_observer_notifications` and `move_to_container` decorators to observers
- Add types `ChannelsMessage` and `WebsocketMessage` describing type of observers messages
- Add property `containers` on the `PermissionObject` that returns all the containers for the given instance
- Add `observe_instance_containers` method to observer models

Changed

- **BACKWARD INCOMPATIBLE:** Drop support for Python 3.6, Python 3.7, Python 3.8 and Python 3.9. The processing container continues to support Python ≥ 3.6 .
- Add `source` attribute to observer notifications
- When subscribing to observer endpoint without explicit ids subscribe to `CREATE` and `DELETE` notifications by default
- Send `CREATE` and `DELETE` notifications to the containers of the given object when permissions to it are gained/lost or it is moved between containers. The notifications also contain the `source` and `source_id` context containing content type and id of the object that triggered the change.

1.9.17 32.1.0 - 2023-02-13

Added

- Add new entity annotation framework
- Add set_permission method on PermissionQuerySet
- Add notify_create method on Subscription model in observers to enable sending notification when object is created
- Allow ordering data endpoint by entity name
- Add with_superuser argument to users_with_permission method

Changed

- When slug collision occurs in the listener when creating new objects retry up to ten times before raising the exception
- Enable overriding user and group id of the processing container with environmental variable

Fixed

- Notify user when object (data, entity) is created in the container
- Do not delete observers in use by other subscriptions when unsubscribing
- Notify subscribers to the collection if object inside them is modified
- Notify superusers without explicit permissions when object is created

1.9.18 32.0.0 - 2022-11-14

Changed

- **BACKWARD INCOMPATIBLE:** Drop support for Python 3.4 in the processing container
- **BACKWARD INCOMPATIBLE:** Rewrite the listener to allow more than one of them to run at the same time
- Rewrite the processing container code to make it more stable
- Rewrite the communication container code to make it more stable
- Add resolve_url call to the re_import.sh bash script

Changed

- Migrate base docker images to Ubuntu 22.04 and Fedora 36

1.9.19 31.3.1 - 2022-10-19

Fix

- Make sure signals are triggered for data objects when their sample is moved to a new collection

1.9.20 31.3.0 - 2022-10-17

Added

- Add handler `resolve_url` to listener
- Add support for resolving urls in Python processes
- Add `url` method to the `BaseConnector` class and override it in `S3` and `local` connectors

1.9.21 31.2.1 - 2022-09-20

Changed

- Add `cleanup_callback` to `retry` decorator
- Clean `kubernetes` temporary files with credentials on config load error

1.9.22 31.2.0 - 2022-09-19

Added

- Add `descriptor` and `descriptor_schema` fields to the `Relation` model
- Setting `KUBERNETES_DISPATCHER_CONFIG_LOCATION` specifying the location of the `kubernetes` config to load in the `Kubernetes` workload connector

Changed

- Retry loading `Kubernetes` configuration couple of times before giving up in the `kubernetes` workload connector
- Add model observers that notify clients about model changes via a `websocket` connection

1.9.23 31.1.0 - 2022-08-23

Fixed

- Clear `custom_messages` array on `auditlog` reset call

1.9.24 31.0.0 - 2022-07-18

Changed

- **BACKWARD INCOMPATIBLE:** Always try to load kubernetes configuration via `load_kube_config` before falling back to `load_incluster_config`
- Rename `docker-compose.yml` to `compose.yml`
- Listener can always modify the data object that it is processing
- Add original objects to `post_duplicate` call
- Install `asgiref` version based on the version of the installed Python interpreter

Fixed

- Fix a typo in listener permission handling when creating data model: a check was performed on the wrong object type

Added

- Make requests and limits for the `communication` container configurable

1.9.25 30.3.0 - 2022-06-13

Added

- Support custom user model in Python processes

1.9.26 30.2.0 - 2022-05-16

Added

- Send custom signal `post_duplicate` when objects are duplicated since regular Django signals are not sent
- Add `auditlog` application to log user actions

1.9.27 30.1.0 - 2022-04-15

Added

- Custom signal `resolwe.flow.signals.before_processing` is sent when data object is ready for processing
- Support setting `descriptor` and `DescriptorSchema` on Data during process runtime
- Support filtering Data, Entity and Collections based on permissions (`view`, `edit`, `owner`, `group`, `shared_with_me`)
- Support filtering Data and Entity objects by relation id
- Create `upload_config` API endpoint that specifies upload connector type and credentials client can use for optimized upload

Fixed

- Do not return multiple version of the same process while checking for permissions in Python processes
- Change misleading error message when importing file if the response with status code indicating error was received from the upstream server

Changed

- Use `data_id` instead of `data.id` when notifying dispatcher to avoid potential database query inside async context

1.9.28 30.0.0 - 2022-03-14

Added

- Add support for Python 3.10
- Add health checks support for deploy in Kubernetes
- Add `collecttools_kubernetes` management command
- Add `COMMUNICATION_CONTAINER_LISTENER_CONNECTION` to separate settings for listener (where to bind to) and containers (where to connect to)
- Support `docker_volume` setting in connector config
- Support use of named volumes as processing or input volume in Docker executor
- Support SSL connection to Redis

Changed

- **BACKWARD INCOMPATIBLE:** Require Django 3.2
- **BACKWARD INCOMPATIBLE:** Require Django Priority Batch version 4
- Do not prepare tools configmaps in Kubernetes workload connector
- Enable Docker containers to connect to the custom network
- Auto-delete completed jobs in Kubernetes after 5 minutes
- Optionally add affinity to the Kubernetes job
- Remove support for setting permissions using old syntax

1.9.29 29.3.0 - 2022-02-15

Added

- Add MD5 checksum to `DataBrowseView` view

1.9.30 29.2.0 - 2022-01-17

Added

- Support range parameter in fields of Python proces

Changed

- Flush stdout/stderr on Python processes on every write

Fixed

- Add attribute as a field to a `resolwe.process.fields.GroupField` in Python process only if it is an instance of `resolwe.process.fields.Field`.

1.9.31 29.1.0 - 2021-12-12

Changed

- Do not fail in case of missing files in `UriResolverView`

Fixed

- Remove references to temporary export files from the database and make sure they are not created anymore
- Wrap `move_to_collection` in transaction and only call method if collection has changed

1.9.32 29.0.0 - 2021-11-11

Added

- **BACKWARD INCOMPATIBLE:** New permission architecture: it is not based on Guardian anymore. The main benefits of new architecture are speed gains in common operations, such as setting a permission and retrieving objects with the given permission.
- Allow overriding process resources in data object

Changed

- Allow mounting connectors into pods as persistent volume claim instead of volume of type `hostPath`

Fixed

- use the same connector inside pod to handle files and directories
- When data object was deleted listener did not receive the terminate message and pod did not terminate immediately

1.9.33 28.5.0 - 2021-09-13

Added

- Add `compare_models_and_csv` management script to check if all `ReferencedPath`s` point to a valid file in the aws database
- Add method `get_latest` to `Process` class in `Python Processes` returning the latest version of the process with the given slug
- Support assuming role in S3 connector

Changed

- Set hashes during upload to avoid creating multiple versions of the object in S3 bucket with enabled versioning

1.9.34 28.4.0 - 2021-08-16

Changed

- Remove dependency on EFS/NFS when running on Kubernetes
- When running on Kubernetes the runtime volume configuration can be omitted

1.9.35 28.3.0 - 2021-07-20

Fixed

- Prepare release 28.3.0 due to preexisting 28.3.0a1 pre-release

1.9.36 28.2.1 - 2021-07-13

Fixed

- Speed up deleting storage locations by considering only referenced paths belonging to the given storage location
- Temporary pin `asteval` to version `0.9.23` due to compatibility issues with Python 3.6

Changed

- Improve logging in cleanup manager

1.9.37 28.2.0 - 2021-06-15

Fixed

- Create ReferencedPath objects during transfer only when needed

Changed

- Retry data transfer if `botocore.exceptions.ClientError` is raised during transfer

Added

- Add `FLOW_PROCESS_MAX_MEM` Django setting to limit the amount of memory used by a process
- Support disabled fields in Python processes
- Add method `get_latest` to the `Process` class in Python processes which retrieves latest version of the process with the given slug

1.9.38 28.1.0 - 2021-05-17

Fixed

- Do not raise exception when terminating `runlistener` management command
- Change concurrency issue in the listener causing processes to sometimes get incorrect value for `RUNTIME_VOLUME_MAPS` settings

Changed

- Make S3 connectors use system credentials when they are not explicitly given in settings

Added

- Make it possible to rewrite container image names in kubernetes workload connector

1.9.39 28.0.4 - 2021-05-04

Fixed

- Use per process storage overrides

1.9.40 28.0.3 - 2021-05-04

Changed

- Make gp2 default EBS storage class

1.9.41 28.0.2 - 2021-05-03

Fixed

- Log peer activity on every received message to avoid declaring otherwise healthy node as failed
- Fix possible data loss caused by parallel command processing when uploading log files interfered with processing command from a script
- Fix deadlock when uploading empty files

1.9.42 28.0.1 - 2021-04-28

Changed

- Make logger level inside init and communication containers configurable via environmental variable
- Change default logger level inside init and communication containers for AWS S3 and Google Cloud Storage components to WARNING

Fixed

- Stop timer that uploads log files in the processing container immediately after the processing is finished to avoid timing issues that could cause the data object to be marked as failed

1.9.43 28.0.0 - 2021-04-19

Changed

- **BACKWARD INCOMPATIBLE:** Require Django 3.1.x
- **BACKWARD INCOMPATIBLE:** Require Django Channels version 3.0.x
- **BACKWARD INCOMPATIBLE:** Require asgiref version 3.3.x
- **BACKWARD INCOMPATIBLE:** Require jsonschema version 3.2.x
- **BACKWARD INCOMPATIBLE:** Require Sphinx version 3.5.x
- **BACKWARD INCOMPATIBLE:** Require django-guardian version 2.3.x
- Refresh versions of the other dependencies
- Replace django-versionfield2 with django-versionfield
- Overhaul of the storage configuration
- Remove per-process runtime directory
- Increase socket timeouts in the processing and the communication container

Added

- Add multipart upload capability to LocalFilesystemConnector and AwsS3Connector
- Support uploading files to LocalFilesystemConnector or AwsS3Connector
- Add support for Python 3.9

1.9.44 27.1.2 - 2021-03-22

Fixed

- Bump version of upload-dir process to use the default version of processing image instead of the previous one
- Use Signature Version 4 when generating presigned URLs in S3 connector
- Fix possible socket timeout when uploading files in the processing container
- Remove static rnaseq image from list of docker images

Changed

- Use tagged base image in upload-file process
- Allow to change Entity descriptor in Python process.

1.9.45 27.1.1 - 2021-03-21

Fixed

- Fix connection timeout in communication container when sending initial message to the listener

1.9.46 27.1.0 - 2021-03-15

Fixed

- Account for file system overhead when processing Data objects with large inputs

Changed

- Improve storage manager to only process applicable storage locations instead of iterating through all of them
- Skip hash computation when connector itself provides data integrity check
- Remove job prefix from kubernetes job name
- Make error messages in Python processes more useful

Added

- Add label `job_type` to Kubernetes job to separate interactive jobs from batch jobs

1.9.47 27.0.0 - 2021-02-22

Fixed

- Fixed progress reporting in Python processes
- Do not override content-type of S3 object when storing hashes
- Support upload of files larger than 80G to AWS S3

Changed

- Download input data in init container
- Storage objects are sent to the listener over socket instead of using files on the shared filesystem
- Make it possible to run the platform without shared filesystem. All inputs for processed data object are prepared in input container and all outputs are uploaded to the chosen storage connector when they are referenced.
- Overcommit CPU in kubertenes processing container by 20%
- Move docker images from Docker Hub to Amazon ECR

Added

- Make automatic removal of Docker containers configurable
- Terminate processing immediately when data object is deleted
- Make default processing image configurable

1.9.48 26.0.0 - 2021-01-20

Changed

- **BACKWARD INCOMPATIBLE:** Remove `resolwe/upload-tab-file` and `resolwe/archiver` Docker images
- **BACKWARD INCOMPATIBLE:** Remove obsolete processes: `archiver`, `upload-tab-file` and `upload-image-file`
- **BACKWARD INCOMPATIBLE:** Python process syntax has changed: all the attributes of Data object are now available in Python process and therefore accessing outputs using syntax `data_object.output_name` is no longer valid since `output_name` could be the name of the attribute. The new syntax is `data_object.output.output_name`.
- **BACKWARD INCOMPATIBLE:** Communication between the processing script and listener has changed from printing to stdout to sending messages over sockets. Messages printed to stdout or sent using old version of the `resolwe-runtime-utils` (YAML processes) are no longer processed. YAML processes need new version of `resolwe-runtime-utils` while Python processes require a rewrite to the new syntax and Python version 3.4 or higher in the container (`resolwe-runtime-utils` package is no longer needed).

- Use Github Actions to run the tests
- Listener communicates with containers through ZeroMQ instead of Redis
- Start two containers for each process instead of one: the second one is used to communicate with the listener
- Move settings for Python processes from files to environmental variables

Added

- Add Kubernetes workload connector
- Support running process instant termination
- Support registering custom command handlers in listener and exposing data objects (possibly defined in other Django applications) to Python processes
- Support Django-like syntax in Python processes to create, filter or access attributes of the exposed data objects
- Support creating new base classes for Python processes

1.9.49 25.2.0 - 2020-12-15

Fixed

- Allow retrieval of Storage object that was linked to more than one Data object

Changed

- Migrate docker images to Fedora 33 and Ubuntu 20.04

1.9.50 25.1.0 - 2020-11-16

Added

- Support Python processes in Sphinx `autoprocess*::` directive

1.9.51 25.0.0 - 2020-10-16

Added

- **BACKWARD INCOMPATIBLE:** Only copy parent relations when duplicating Data objects
- Add duplicate data dependency to indicate from which object the Data object was duplicated
- Support accessing Data name in Python processes through `self.name`
- Add permission filter to `collection`, `entity` and `data` that returns only objects on which current user has given permission

Changed

- Make relations in collection visible to public user if he has view permissions on the collection

1.9.52 24.0.0 - 2020-09-14

Changed

- **BACKWARD INCOMPATIBLE:** Terminate Python process immediately after `self.error` method is called inside the process

Fixed

- Make sure to terminate Docker container before executor exits
- Speed-up duplication of Data objects, Entities, and Collections
- Lock inputs' storage locations while the process is waiting and processing to make sure that they are not deleted
- Don't validate input objects when Data object is marked as done as they may already be deleted at that point

1.9.53 23.0.0 - 2020-08-17

Fixed

- Fix ordering options in Data viewset to enable ordering by `process__name` and `process__type`
- Handle exception when processing deleted Data object

Changed

- **BACKWARD INCOMPATIBLE:** Delete `elastic` application
- Don't pass undefined values to steps of workflows

Added

- Add relations property to `data:` and `list:data:` fields to support relations on client
- Add `entity_id` property to `DataField` in Python processes
- Add relations in Python processes

1.9.54 22.1.3 - 2020-07-13

Fixed

- When deciding which StorageLocation objects will be deleted consider only completed StorageLocation objects.
- Add `google.resumable_media.common.DataCorruption` exception to `transfer_exceptions` tuple.

1.9.55 22.1.2 - 2020-06-30

Fixed

- Celery sometimes starts more than one worker for a given Data object. In such case the download and purge part of the worker must be skipped or errors processing Data objects might occur.

1.9.56 22.1.1 - 2020-06-16

Changed

- Remove `asgiref` version pin due to new release that fixed previous regression.

1.9.57 22.1.0 - 2020-06-15

Changed

- Rename `transfer_rec` to `transfer_objects` and change its signature to accept dictionary objects with information about name, size and hashes of objects to transfer
- Move part of Data object validation to listener
- Improve loading time of collection endpoint

Added

- Add `move_to_collection` method to Data viewset
- Report registration failure in `ProcessTestCase`
- Add a pseudo Python process to serve as a template
- Add `validate_urls` method to storage `BaseConnector` class
- Validate storage connector settings on registration
- Add `transfer_data` method to `StorageLocation` class
- Remove data when `StorageLocation` object is deleted
- Store file hashes inside `ReferencedPath` model and connect it to `StorageLocation` model
- Add `get_hashes` method to storage connectors
- Add `open_stream` method to storage connectors
- Add `compute_hashes` function to `storage.connectors.hasher` module
- Use threads when transferring files with `AwsS3Connector`

- Add `duplicate` method to storage connectors
- Add pre-processing and post-processing hooks to storage connectors
- Use multiple threads for file transfer

Fixed

- Add missing decorator `validate_url` to `AwsS3Connector`
- Always import exceptions from `requests` library
- Fix bug that sometimes caused objects inside workflow to fail with `Failed to transfer data.`
- Fix dependency handling bug in listener: when checking for missing data listener must only consider dependencies with kind `KIND_IO` instead of all dependencies.
- Raise exception when data transfer failed.

1.9.58 22.0.0 - 2020-05-18

Changed

- **BACKWARD INCOMPATIBLE:** Move purge code inside worker, remove old purge code
- Various code fixes to make code work with the new storage model
- Use storage connectors in workers to download data not available locally

Added

- Add `resolve.storage` application, a framework for storage management
- Add storage connectors for Google Cloud Storage, Amazon Simple Storage service and local filesystem.
- Add migrations to move from old storage model to the new one
- Add storage manager
- Add management command to start storage manager
- Add cleanup manager for removing unreferenced data
- Add `isnull` related lookup filter
- Add `entity_count` to the `Collection` serializer
- Add `inherit_collection` to `Data` viewset
- Add `entity_always_create` in `Process` serializer

1.9.59 21.1.0 - 2020-04-14

Added

- Add support for the `allow_custom_choice` field property in Python processes
- Add ordering by contributor's first and last name to Collection and Data viewsets
- Add `data_count` and `status` fields to the Collection serializer

Fixed

- Enable all top-level class definitions in Python processes
- Make filtering via foreign key more 'Django like': when foreign key does not exist return empty set instead of raising validation exception. Also when filtering using list of foreign keys do not raise validation exception if some foreign keys in the list do not exist.
- Reduce number of database queries in API viewsets by prefetching all required data

1.9.60 21.0.0 - 2020-03-16

Changed

- **BACKWARD INCOMPATIBLE:** Use Postgres filtering instead of Elasticsearch on API endpoints
- **BACKWARD INCOMPATIBLE:** Remove filtering by year, month, day, hour, minute and second on API endpoints
- Migrate docker images to Fedora 31
- Use `DictRelatedField` for collection field in `RelationSerializer`. In practice this fixes inconsistency comparing with how other serializers handle collections field.

Added

- Add a custom database connector to optimize queries and enable them to use database indexes
- Add database indexes to improve search performance
- Add database fields and triggers for full-text search in Postgres
- Add support for annotating entities in processes
- Add support for Python 3.8

1.9.61 20.2.0 - 2020-02-17

Added

- Support workflows as inputs to Python processes
- Support retrieval of `Data.name` in Python process
- Add `name_contains`, `contributor_name`, and `owners_name` collection and data filtering fields on API
- Add `username` to `current_user_permissions` field of objects on API

- Add `delete_chunked` method to Collection, Entity and Storage managers

Fixed

- Delete orphaned Storage objects in chunks in purge command to prevent running out of memory

1.9.62 20.1.0 - 2019-12-16

Added

- Add description field to Collection full-text search

1.9.63 20.0.0 - 2019-11-18

Changed

- **BACKWARD INCOMPATIBLE:** Remove download permission from Data objects, samples and collections and add permission from samples and collections
- **BACKWARD INCOMPATIBLE:** Remove `Entity.descriptor_completed` field

Fixed

- Fix Docker executor command with `--cpus` limit option. This solves the issue where process is killed before the timeout 30s is reached

1.9.64 19.1.0 - 2019-09-17

Added

- Support filtering by `process_slug` in `DataViewSet`

Fixed

- Fix `DictRelatedField` so it can be used in browsable-API
- Fix access to subfields of empty `GroupField` in Python processes

1.9.65 19.0.0 - 2019-08-20

Changed

- **BACKWARD INCOMPATIBLE:** Change relations between Data, Entity and Collection from ManyToMany to ManyToOne. In practice this means that `Data.entity`, `Data.collection` and `Entity.collection` are now `ForeignKey`-s. This also implies the following changes:
 - `CollectionViewSet` methods `add_data` and `remove_data` are removed
 - `EntityViewSet` methods `add_data`, `remove_data`, `add_to_collection` and `remove_from_collection` are removed

- EntityQuerySet and Entity method duplicate argument inherit_collections is renamed to inherit_collection.
- EntityFilter FilterSet field collections is renamed to collection.
- **BACKWARD INCOMPATIBLE:** Change following fields in DataSerializer:
 - process_slug, process_name, process_type, process_input_schema, process_output_schema are removed and moved in process field which is now DictRelatedField that uses ProcessSerializer for representation
 - Remove entity_names and collection_names fields
 - add entity and collection fields which are DictRelatedField-s that use corresponding serializers for representation
 - Remove support for hydrate_entities and hydrate_collections query parameters
- **BACKWARD INCOMPATIBLE:** Remove data field in EntitySerializer and CollectionSerializer. This implies that parameter hydrate_data is no longer supported.
- **BACKWARD INCOMPATIBLE:** Remove delete_content parameter in delete method of EntityViewSet and CollectionViewSet. From now on, when Entity/Collection is deleted, all it's objects are removed as well
- Gather all Data creation logic into DataQuerySet.create method

Added

- Enable sharing based on user email
- Support running tests with live Resolwe host on non-linux platforms
- Add inherit_entity and inherit_collection arguments to Data.duplicate and DataQuerySet.duplicate method
- Implement DictRelatedField

1.9.66 18.0.0 - 2019-07-15

Changed

- **BACKWARD INCOMPATIBLE:** Remove parents and children query filters from Data API endpoint

Added

- /api/data/:id/parents and /api/data/:id/children API endpoints for listing parents and children Data objects of the object with given id
- Add entity_always_create field to Process model

Fixed

- Make sure that Elasticsearch index exists before executing a search query

1.9.67 17.0.0 - 2019-06-17

Changed

- **BACKWARD INCOMPATIBLE:** Use Elasticsearch version 6.x
- **BACKWARD INCOMPATIBLE:** Bump Django requirement to version 2.2
- **BACKWARD INCOMPATIBLE:** Remove not used `django-mathfilters` requirement

Added

- Support Python 3.7
- Support forward and reverse many-to-one relations in Elasticsearch
- Add `collection_names` field to `DataSerializer`
- Add test methods to `ProcessTestCase` that assert directory structure and content: `assertDirExists`, `assertDir`, and `assertDirStructure`
- Add `upload-dir` process

1.9.68 16.0.1 - 2019-04-29

Fixed

- Pin `django-priority-batch` to version 1.1 to fix compatibility issues

1.9.69 16.0.0 - 2019-04-16

Changed

- **BACKWARD INCOMPATIBLE:** Access to `DataField` members (in Python process input) changed from dict to Python objects. For example, `input_field.file_field['name']` changed to `input_field.file_field.path`.
- **BACKWARD INCOMPATIBLE:** Filters that are based on `django-filter FilterSet` now use dict-declaring-syntax. This requires that subclasses of respective filters modify their syntax too.
- Interactively save results in Python processes

Added

- Add `get_data_id_by_slug` method to Python processes' `Process` class
- Python process syntax enhancements:
 - Support `.entity_name` in data inputs
 - Easy access to process resources through `self.resources`
- Raise error if `ViewSet` receives invalid filter parameter(s)
- Report process error for exceptions in Python processes
- Report process error if spawning fails
- Automatically export files for spawned processes (in Python process syntax)
- Import files of Python process `FileField` inputs (usage: `inputs.src.import_file()`)

Fixed

- Interactively write to standard output within Python processes
- Fix writing to integer and float output fields
- Allow non-required `DataField` as Python process input

1.9.70 15.0.1 - 2019-03-19

Fixed

- Fix storage migration to use less memory

1.9.71 15.0.0 - 2019-03-19

Changed

- Log plumbum commands to standard output
- Change storage data relation from many-to-one to many-to-many
- Moved `purged` field from `Data` to `DataLocation` model

Added

- Add `run_process` method to `Process` to support triggering of a new process from the running Python process
- Add `DataLocation` model and pair it with `Data` model to handle data location
- Add `entity_names` field to `DataSerializer`
- Support duplication of `Data`, `Entity` and `Collection`
- Support moving entities between collections
- Support relations requirement in process syntax

1.9.72 14.4.0 - 2019-03-07

Changed

- Purge processes only not jet purged Data objects

Fixed

- Allow references to missing Data objects in the output of finished Data objects, as we don't have the control over what (and when) is deleted

1.9.73 14.3.0 - 2019-02-19

Added

- Add `scheduled` field to Data objects to store the date when object was dispatched to the scheduling system
- Add `purge` field to Data model that indicates whether Data object was processed by purge

Fixed

- Make Elasticsearch build arguments cache thread-safe and namespace cache keys to make sure they don't interfere
- Trigger the purge outside of the transaction, to make sure the Data object is committed in the database when purger grabs it

1.9.74 14.2.0 - 2019-01-28

Added

- Add input Jinja filter to access input fields

1.9.75 14.1.0 - 2019-01-17

Added

- Add `assertFilesExist` method to `ProcessTestCase`
- Add `clean_test_dir` management command that removes files created during testing

Fixed

- Support registration of Python processes inherited from `process.Process`
- Skip docker image pull if image exists locally. This solves the issue where pull would fail if process uses an image that is only used locally.

1.9.76 14.0.1 - 2018-12-17

Fixed

- Make sure that tmp dir exists in Docker executor

1.9.77 14.0.0 - 2018-12-17

Changed

- **BACKWARD INCOMPATIBLE:** Run data purge in a separate worker to make sure that listener replies to the executor within 60 seconds
- Use batcher for spawned processes in listener
- Increase Docker's memory limit for 100MB to make sure processes are not killed when using all available memory and tune Docker memory limits to avoid OOM.

Added

- Raise an exception in Docker executor if container doesn't start for 60 seconds
- Set TMPDIR environment variable in Docker executor to .tmp dir in data directory to prevent filling up container's local storage

Fixed

- Process SIGTERM signal in executor as expected - set the Data status to error and set the process_error field
- Clear cached Django settings from the manager's shared state on startup

1.9.78 13.3.0 - 2018-11-20

Changed

- Switch channels_redis dependency to upstream version

Added

- Python execution engine
- Support multiple entity types
- Support extending viewsets with custom filter methods
- Add tags attribute to ProcessTestCase.run_process method which adds listed tag to the created Data object
- Copy Data objects tags from parent objects for spawned Data objects and Data objects created by workflows

Fixed

- Fix manager shutdown in the test runner. If an unrecoverable exception occurred while running a test, and never got caught (e.g. an unpicklable exception in a parallel test worker), the listener would not get terminated properly, leading to a hang.
- Data and collection name API filters were fixed to work as expected (ngrams was switched to raw).

1.9.79 13.2.0 - 2018-10-23

Added

- Use prioritized batcher in listener

1.9.80 13.1.0 - 2018-10-19

Added

- Use batching for ES index builds

Fixed

- Fix handling of M2M dependencies in ES indexer

1.9.81 13.0.0 - 2018-10-10

Changed

- **BACKWARD INCOMPATIBLE:** Remove Data descriptors from Entity Elasticsearch index
- Support searching by `slug` and `descriptor_data` in entity viewset text search

Added

- Add tags to collections

1.9.82 12.0.0 - 2018-09-18

Changed

- **BACKWARD INCOMPATIBLE:** Switch Collection and Entity API viewsets to use Elasticsearch
- **BACKWARD INCOMPATIBLE:** Refactor Relation model, which includes:
 - renaming `position` to `partition`
 - renaming `label` to `category` and making it required
 - adding `unit`
 - making `collection` field required

- requiring unique combination of `collection` and `category`
- renaming partition's `position` to `label`
- adding (integer) `position` to partition (used for sorting)
- deleting `Relation` when the last `Entity` is removed
- **BACKWARD INCOMPATIBLE:** Remove rarely used parameters of the `register` command `--path` and `--schemas`.
- Omit `current_user_permissions` field in serialization if only a subset of fields is requested
- Allow slug to be null on update to enable slug auto-generation
- Retire obsolete processes. We have added the `is_active` field to the `Process` model. The field is read-only on the API and can only be changed through Django ORM. Inactive processes can not be executed. The `register` command was extended with the `--retire` flag that removes old process versions which do not have associated data. Then it finds the processes that have been registered but do not exist in the code anymore, and:
 - If they do not have data: removes them
 - If they have data: flags them not active (`is_active=False`)

Added

- Add support for URLs in `basic:file:` fields in Django tests
- Add `collections` and `entities` fields to Data serializer, with optional hydration using `hydrate_collections` and/or `hydrate_entities`
- Support importing large files from Google Drive in re-import
- Add `python3-plumbum` package to `resolwe/base:ubuntu-18.04` image

Fixed

- Prevent mutation of `input_` parameter in `ProcessTestCase.run_process`
- Return 400 instead of 500 error when slug already exists
- Add trailing colon to process category default
- Increase stdout buffer size in the Docker executor

1.9.83 11.0.0 - 2018-08-13

Changed

- **BACKWARD INCOMPATIBLE:** Remove option to list all objects on Storage API endpoint
- Make the main executor non-blocking by using Python `asyncio`
- Debug logs are not send from executors to the listener anymore to limit the amount of traffic on Redis

Added

- Add size to Data serializer
- Implement ResolveSlugRelatedField. As a result, DescriptorSchema objects can only be referenced by slug (instead of id)
- Add options to filter by type and scheduling_class on Process API endpoint

Fixed

- Inherit collections from Entity when adding Data object to it

1.9.84 10.1.0 - 2018-07-16

Changed

- Lower the level of all INFO logs in elastic app to DEBUG

Added

- Add load tracking to the listener with log messages on overload
- Add job partition selection in the SLURM workload connector
- Add slug Jinja filter
- Set Data status to ERROR if executor is killed by the scheduling system

Fixed

- Include the manager in the documentation, make sure all references work and tidy the content up a bit

1.9.85 10.0.1 - 2018-07-07

Changed

- Convert the listener to use asyncio
- Switched to channels_redis_persist temporarily to mitigate connection storms

Fixed

- Attempt to reconnect to Redis in the listener in case of connection errors

1.9.86 10.0.0 - 2018-06-19

Changed

- **BACKWARD INCOMPATIBLE:** Drop support for Python 3.4 and 3.5
- **BACKWARD INCOMPATIBLE:** Start using Channels 2.x

Added

- Add the options to skip creating of fresh mapping after dropping ES indices with `elastic_purge` management command
- Add `dirname` and `relative_path` Jinja filters

1.9.87 9.0.0 - 2018-05-15

Changed

- Make sorting by contributor case insensitive in Elasticsearch endpoints
- Delete ES documents in post delete signal instead of pre delete one

Added

- **BACKWARD INCOMPATIBLE:** Add on-register validation of default values in process and schemas
- **BACKWARD INCOMPATIBLE:** Validate that field names in processes and schemas start with a letter and only contain alpha-numeric characters
- Support Python 3.6
- Add `range` parameter and related validation to fields of type `basic:integer:`, `basic:decimal`, `list:basic:integer:` and `list:basic:decimal`
- Support filtering and sorting by `process_type` parameter on Data API endpoint
- Add `dirname` Jinja filter
- Add `relative_path` Jinja filter

Fixed

- Add missing `list:basic:decimal` type to JSON schema
- Don't crash on empty `in` lookup
- Fix `{{ requirements.resources.* }}` variables in processes to take in to the account overrides specified in Django settings
- Create Elasticsearch mapping even if there is no document to push

1.9.88 8.0.0 - 2018-04-11

Changed

- **BACKWARD INCOMPATIBLE:** Use Elasticsearch version 5.x
- **BACKWARD INCOMPATIBLE:** Raise an error if an invalid query argument is used in Elasticsearch viewsets
- **BACKWARD INCOMPATIBLE:** Switch Data API viewset to use Elasticsearch
- Terminate the executor if listener response with error message
- verbosity setting is no longer propagated to the executor
- Only create Elasticsearch mappings on first push

Added

- Add sort argument to assertFile and assertFiles methods in ProcessTestCase to sort file lines before asserting the content
- Add process_slug field to DataSerializer
- Improve log messages in executor and workload connectors
- Add process_memory and process_cores fields to Data model and DataSerializer
- Support lookup expressions (lt, lte, gt, gte, in, exact) in ES viewsets
- Support for easier dynamic composition of type extensions
- Add elastic_mapping management command

Fixed

- Fix Elasticsearch index rebuilding after a dependant object is deleted
- Send response to executor even if data object was already deleted
- Correctly handle reverse m2m relations when processing ES index dependencies

1.9.89 7.0.0 - 2018-03-12

Changed

- **BACKWARD INCOMPATIBLE:** Remove Ubuntu 17.04 base Docker image due to end of lifetime
- **BACKWARD INCOMPATIBLE:** Remove support for Jinja in DescriptorSchema's default values
- **BACKWARD INCOMPATIBLE:** Remove CONTAINER_IMAGE configuration option from the Docker executor; if no container image is specified for a process using the Docker executor, the same pre-defined default image is used (currently this is resolve/base:ubuntu-16.04)
- Add mechanism to change test database name from the environment, appending a _test suffix to it; this replaces the static name used before

Added

- Add Ubuntu 17.10 and Ubuntu 18.04 base Docker images
- Add database migration operations for process schema migrations
- Add `delete_chunked` method to `Data` objects queryset which is needed due to Django's extreme memory usage when deleting a large count of `Data` objects
- Add `validate_process_types` utility function, which checks that all registered processes conform to their supertypes
- Add `FLOW_CONTAINER_VALIDATE_IMAGE` setting which can be used to validate container image names
- Only pull Docker images at most once per process in `list_docker_images`
- Add `FLOW_PROCESS_MAX_CORES` Django setting to limit the number of CPU cores used by a process

Fixed

- Make parallel test suite worker threads clean up after initialization failures
- Add mechanism to override the manager's control channel prefix from the environment
- Fix deletion of a `Data` objects which belongs to more than one `Entity`
- Hydrate paths in refs of `basic:file:`, `list:basic:file:`, `basic:dir:` and `list:basic:dir:` fields before processing `Data` object

1.9.90 6.1.0 - 2018-02-21

Changed

- Remove runtime directory during general data purge instead of immediately after each process finishes
- Only process the `Data` object (and its children) for which the dispatcher's `communicate()` was triggered
- Propagate logging messages from executors to the listener
- Use process' slug instead of data id when logging errors in listener
- Improve log messages in dispatcher

Added

- Add `descriptor_completed` field to the `Entity` filter
- Validate manager semaphors after each test case, to ease debugging of tests which execute processes

Fixed

- Don't set Data object's status to error if executor is run multiple times to mitigate the Celery issue of tasks being run multiple times
- Make management commands respect the set verbosity level

1.9.91 6.0.1 - 2018-01-29**Fixed**

- Make manager more robust to ORM/database failures during data object processing
- Rebuild the Elasticsearch index after permission is removed from an object
- Trim `Data.process_error`, `Data.process_warning` and `Data.process_info` fields before saving them
- Make sure values in `Data.process_error`, `Data.process_warning` and `Data.process_info` cannot be overwritten
- Handle missing Data objects in `hydrate_input_references` function
- Make executor fail early when executed twice on the same data directory

1.9.92 6.0.0 - 2018-01-17**Changed**

- **BACKWARD INCOMPATIBLE:** `FLOW_DOCKER_LIMIT_DEFAULTS` has been renamed to `FLOW_PROCESS_RESOURCE_DEFAULTS` and `FLOW_DOCKER_LIMIT_OVERRIDES` has been renamed to `FLOW_PROCESS_RESOURCE_OVERRIDES`
- **BACKWARD INCOMPATIBLE:** `Process.PERSISTENCE_TEMP` is not used for execution priority anymore
- **BACKWARD INCOMPATIBLE:** There is only one available manager class, which includes dispatch logic; custom manager support has been removed and their role subsumed into the new connector system
- **BACKWARD INCOMPATIBLE:** Removed `FLOW_DOCKER_MAPPINGS` in favor of new `FLOW_DOCKER_VOLUME_EXTRA_OPTIONS` and `FLOW_DOCKER_EXTRA_VOLUMES`
- Parent relations are kept even after the parent is deleted and are deleted when the child is deleted
- Dependency resolver in manager is sped up by use of parent relations
- Validation of Data inputs is performed on save instead of on create

Added

- Support for the SLURM workload manager
- Support for dispatching Data objects to different managers
- Support for passing secrets to processes in a controlled way using a newly defined `basic:secret` input type
- `is_testing` test helper function, which returns `True` when invoked in tests and `False` otherwise
- Add `collecttools` Django command for collecting tools' files in single location defined in `FLOW_TOOLS_ROOT` Django setting which is used for mapping tools in executor when `DEBUG` is set to `False` (but not in tests)

Fixed

- Fix Data object preparation race condition in `communicate()`
- Set correct executor in flow manager
- Make executors more robust to unhandled failures
- Calculate `Data.size` by summing `total_size` of all file-type outputs
- Don't change slug explicitly defined by user - raise an error instead
- Objects are locked while updated over API, so concurrent operations don't override each other
- Make manager more robust to unhandled failures during data object processing
- Fix manager deadlock during tests
- Fix ctypes cache clear during tests
- Don't raise `ChannelFull` error in manager's `communicate` call
- Don't trim predefined slugs in `ResolweSlugField`

1.9.93 5.1.0 - 2017-12-12

Added

- Database-side JSON projections for `Storage` models
- Compute total size (including refs size) for file-type outputs
- Add `size` field to `Data` model and migrate all existing objects

Change

- Change Test Runner's test directory creation so it always creates a subdirectory in `FLOW_EXECUTOR`'s `DATA_DIR`, `UPLOAD_DIR` and `RUNTIME_DIR` directories

Fixed

- Do not report additional failure when testing a tagged process errors or fails
- Fix Test Runner's `changes-only` mode when used together with a Git repository in detached `HEAD` state
- Fix handling of tags and test labels together in Test Runner's `changes-only` mode
- Fix parallel test execution where more test processes than databases were created during tests

1.9.94 5.0.0 - 2017-11-28

Changed

- **BACKWARD INCOMPATIBLE:** The `keep_data()` method in `TransactionTestCase` is no longer supported. Use the `--keep-data` option to the test runner instead.
- **BACKWARD INCOMPATIBLE:** Convert the manager to Django Channels
- **BACKWARD INCOMPATIBLE:** Refactor executors into standalone programs
- **BACKWARD INCOMPATIBLE:** Drop Python 2 support, require Python 3.4+
- Move common test environment preparation to `TestCaseHelpers` mixin

Fixed

- Fix parents/children filter on `Data` objects
- Correctly handle removed processes in the changes-only mode of the Resolwe test runner

1.9.95 4.0.0 - 2017-10-25

Added

- **BACKWARD INCOMPATIBLE:** Add option to build only subset of specified queryset in Elasticsearch index builder
- `--pull` option to the `list_docker_images` management command
- Test profiling and process tagging
- New test runner, which supports partial test suite execution based on changed files
- Add `all` and `any` Jinja filters

Changed

- **BACKWARD INCOMPATIBLE:** Bump Django requirement to version 1.11.x
- **BACKWARD INCOMPATIBLE:** Make `ProcessTestCase` non-transactional
- Pull Docker images after process registration is complete
- Generalize Jinja filters to accept lists of `Data` objects
- When new `Data` object is created, permissions are copied from collections and entity to which it belongs

Fixed

- Close schema (YAML) files after `register` command stops using them
- Close schema files used for validating JSON schemas after they are no longer used
- Close stdout used to retrieve process results in executor after the process is finished
- Remove unrelated permissions occasionally listed among group permissions on `permissions` endpoint
- Fix `ResolwePermissionsMixin` to work correctly with multi-words model names, i.e. `DescriptorSchema`
- Fix incorrect handling of offset/limit in Elasticsearch viewsets

1.9.96 3.1.0 - 2017-10-05

Added

- `resolwe/base` Docker image based on Ubuntu 17.04
- Support different dependency kinds between data objects

Fixed

- Serialize `current_user_permissions` field in a way that is compatible with DRF 3.6.4+
- API requests on single object endpoints are allowed to all users if object has appropriate public permissions

1.9.97 3.0.1 - 2017-09-15

Fixed

- Correctly relabel SELinux contexts on user/group files

1.9.98 3.0.0 - 2017-09-13

Added

- Add filtering by id on `descriptor_schema` API endpoint
- Support assigning descriptor schema by id (if set value is of type int) on `Collection`, `Data` and `Entity` endpoints
- `assertAlmostEqualGeneric` test case helper, which enables recursive comparison for almost equality of floats in nested containers

Changed

- **BACKWARD INCOMPATIBLE:** Run Docker containers as non-root user

Fixed

- Use per-process upload dir in tests to avoid race conditions

1.9.99 2.0.0 - 2017-08-24

Added

- `descriptor` jinja filter to get the descriptor (or part of it) in processes
- Ubuntu 14.04/16.04 based Docker images for Resolwe
- Add `list_docker_images` management command that lists all Docker images required by registered processes in either plain text or YAML
- Data status is set to `ERROR` and error message is appended to `process_error` if value of `basic:storage:` field is set to a file with invalid JSON

Changed

- **BACKWARD INCOMPATIBLE:** Quote all unsafe strings when evaluating expressions in Bash execution engine
- **BACKWARD INCOMPATIBLE:** Rename `permissions` attribute on API endpoints to `current_user_permissions`
- API `permissions` endpoint raises error if no owner is assigned to the object after applied changes
- `owner` permission cannot be assigned to a group
- Objects with public permissions are included in list API views for logged-in users
- Owner permission is assigned to the contributor of the processes and descriptor schemas in the `register` management command
- The base image Dockerfile is renamed to `Dockerfile.fedora-26`

Fixed

- Add `basic:url:link` field to the JSON schema
- Return more descriptive error if non-existing permission is sent to the `permissions` endpoint
- Handle errors occurred while processing Elasticsearch indices and log them
- Return 400 error with a descriptive message if permissions on API are assigned to a non-existing user/group

1.9.100 1.5.1 - 2017-07-20

Changed

- Add more descriptive message if user has no permission to add Data object to the collection when the object is created

Fixed

- Set contributor of Data object to public user if it is created by not authenticated user
- Remove remaining references to calling pip with `--process-dependency-links` argument

1.9.101 1.5.0 - 2017-07-04

Added

- Add Resolwe test framework
- Add `with_custom_executor` and `with_resolwe_host` test decorators
- Add `isort` linter to check order of imports
- Support basic test case based on Django's `TransactionTestCase`
- Support ES test case based on Django's `TransactionTestCase`
- Support process test case based on Resolwe's `TransactionTestCase`
- Add ability to set a custom command for the Docker executor via the `FLOW_DOCKER_COMMAND` setting.
- `get_url` jinja filter
- When running `register` management command, permissions are automatically granted based on the permissions of previous latest version of the process or descriptor schema.
- Set `parent` relation in spawned Data objects and workflows
- Relations between entities
- Resolwe toolkit Docker images
- Archive file process
- File upload processes
- Resolwe process tests
- Add `SET_ENV` setting to set environment variables in executor
- Support ordering by version for descriptor schema
- Add `NullExecutor`
- If `choices` are defined in JSON schema, value of field is validated with them
- Add `cpu core`, `memory` and `network resource limits`
- Add scheduling class for processes (`interactive`, `batch`), which replaces the previously unused `process priority` field
- Add `share_content` flag to the collection and entity permissions endpoint to also share the content of the corresponding object

- Add `delete_content` flag to the collection and entity destroy method on API to also delete the content of the corresponding object

Changed

- Support running tests in parallel
- Split `flow.models` module to multiple files
- Remove ability to set a custom executor command for any executor via the `FLOW_EXECUTOR['COMMAND']` setting.
- Rename `RESOLWE_API_HOST` setting and environment variable in executor to `RESOLWE_HOST_URL`
- Remove `keep_failed` function in tests.
- Rename `keep_all` function to `keep_data`.
- Manager is automatically run when new `Data` object is created
- Outputs of `Data` objects with status `Error` are not validated
- Superusers are no longer included in response in `permissions` endpoint of resources
- Remove `public_processes` field from the `Collection` model as it is never used
- Public users can create new `Data` objects with processes and descriptor schemas on which they have appropriate permissions
- Add custom `ResolweSlugField` and use it instead of `django-autoslug`

Fixed

- **SECURITY:** Prevent normal users from creating new `Processes` over API
- Configure parallel tests
- Isolate Elasticsearch indices for parallel tests
- Fix Docker container name for parallel tests
- Generate temporary names for upload files in tests
- Fix permissions in Elasticsearch tests
- Do not purge data in tests
- Remove primary keys before using cached schemas' in process tests
- Set appropriate SELinux labels when mounting tools in Docker containers
- `Data` objects created by the workflow inherit its permissions
- If user doesn't have permissions on the latest versions of processes and descriptor schemas, older ones are used or error is returned
- Support `data:` and `list:data:` types
- Set `Data` object status to error if worker cannot update the object in the database
- `Data` objects returned in `CollectionViewset` and `EntityViewSet` are filtered by permissions of the user in request
- Public permissions are taken into account in elastic app
- Treat `None` field value as if the field is missing

- Copy parent's permissions to spawned Data objects

1.9.102 1.4.1 - 2017-01-27

Fixed

- Update instructions on preparing a release to no longer build the wheel distribution which currently fails to install Resolwe's dependency links

1.9.103 1.4.0 - 2017-01-26

Added

- Auto-process style, type tree and category index
- Support loading JSON from a file if the string passed to the `basic:json:` field is a file.
- `list:basic:integer:` field
- Data object's checksum is automatically calculated on save
- `get_or_create` end point for Data objects
- `basic:file:html:` field for HTML files
- Helper function for comparing JSON fields in tests
- Purge directories not belonging to any data objects
- Ordering options to API endpoints
- Workflow execution engine
- `data_by_slug` filter for jinja expression engine
- Export `RESOLWE_API_HOST` environment variable in executor
- Add `check_installed()` test utility function
- Add support for configuring the network mode of Docker executor
- Add `with_docker_executor` test utility decorator
- Support for Docker image requirements
- Support version in descriptor schema YAML files
- Add `Entity` model that allows grouping of Data objects
- Introduce priority of Data objects
- Data objects created with processes with temporary persistence are given high priority.
- Add `resolwe.elastic` application, a framework for advanced indexing of Django models with ElasticSearch

Changed

- Refactor linters, check PEP 8 and PEP 257
- Split expression engines into expression engines and execution engines
- Use Jinja2 instead of Django Template syntax
- Expression engine must be declared in requirements
- Set Docker Compose's project name to `resolwe` to avoid name clashes
- Expose `check_docker()` test utility function
- Update versionfield to 0.5.0
- Support Django 1.10 and update filters
- Executor is no longer serialized
- Put Data objects with high priority into `hipri` Celery queue.

Fixed

- Fix pylint warnings (PEP 8)
- Fix pydocstyle warnings (PEP 257)
- Take last version of process for spawned objects
- Use default values for descriptor fields that are not given
- Improve handling of validation errors
- Ignore file size in `assertFields`
- Order data objects in `CollectionViewSet`
- Fix tests for Django 1.10
- Add quotes to paths in a test process `test-save-file`

1.9.104 1.3.1 - 2016-07-27

Added

- Sphinx extension `autoprocess` for automatic process documentation

1.9.105 1.3.0 - 2016-07-27

Added

- Ability to pass certain information to the process running in the container via environment variables (currently, user's uid and gid)
- Explicitly set working directory inside the container to the mapped directory of the current Data's directory
- Allow overriding any `FLOW_EXECUTOR` setting for testing
- Support GET request on `/api/<model>/<id>/permissions/` url
- Add OWNER permissions

- Validate JSON fields before saving Data object
- Add basic:dir field
- RESOLWE_CUSTOM_TOOLS_PATHS setting to support custom paths for tools directories
- Add test coverage and track it with Codecov
- Implement data purge
- Add process_fields.name custom tamplate tag
- Return contributor information together with objects
- Added permissions filter to determine Storage permissions based on referenced Data object

Changed

- Move filters to separate file and systemize them
- Unify file loading in tests
- Simplify ProcessTestCase by removing the logic for handling different uid/gid of the user running inside the Docker container
- Upgrade to django-guardian 1.4.2
- Rename FLOW_EXECUTOR['DATA_PATH'] setting to FLOW_EXECUTOR['DATA_DIR']
- Rename FLOW_EXECUTOR['UPLOAD_PATH'] setting to FLOW_EXECUTOR['UPLOAD_DIR']
- Rename proc.data_path system variable to proc.data_dir
- Rename test project's data and upload directories to .test_data and .test_upload
- Serve permissions in new format
- Rename assertFiles method in ProcessTestCase to assertFile and add new assertFiles method to check list:basic:file field
- Make flow.tests.run_process function also handle file paths
- Use Travis CI to run the tests
- Include all necessary files for running the tests in source distribution
- Exclude tests from built/installed version of the package
- Put packaging tests in a separate Tox testing environment
- Put linters (pylint, pep8) into a separate Tox testing environment
- Drop django-jenkins package since we no longer use Jenkins for CI
- Move testing utilities from resolwe.flow.tests to resolwe.flow.utils.test and from resolwe.permissions.tests.base to resolwe.permissions.utils.test
- Add Tox testing environment for building documentation
- Extend Reference documentation

Fixed

- Spawn processors (add data to current collection)
- Set collection name to avoid warnings in test output
- Improve Python 3 compatibility
- Fix setting descriptor schema on create

1.9.106 1.2.1 - 2016-05-15

Added

- Add docker-compose configuration for PostgreSQL
- Processes can be created on API
- Enable spawned processes

Changed

- Move logic from `Collection` model to the `BaseCollection` abstract model and make it its parent
- Remove all logic for handling `flow_collection`
- Change default database user and port in test project's settings
- Keep track of upload files created during tests and purge them afterwards

Fixed

- Test processes location agnostic
- Test ignore timezone support

1.9.107 1.2.0 - 2016-05-06

Changed

- Rename `assertFileExist` to `assertFileExists`
- Drop `--process-dependency-links` from Tox's pip configuration
- Improve documentation on preparing a new release

Added

- Ability to use a custom executor command by specifying the `FLOW_EXECUTOR['COMMAND']` setting
- Make workload manager configurable in settings

Fixed

- Make Resolwe work with Python 3 again
- Fix tests
- Render data name again after inputs are resolved
- Ensure Tox installs the package from sdist
- Pass all Resolwe's environment variables to Tox's testing environment
- Ensure tests gracefully handle unavailability of Docker

1.9.108 1.1.0 - 2016-04-18

Changed

- Rename `process_register` manage.py command to `register`
- Reference process by slug when creating new Data object
- Run manager when new Data object is created through API
- Include full DescriptorSchema object when hydrating Data and Collection objects
- Add `djangoestframework-filters` package instead of `django-filters`

Added

- Tox tests for ensuring high-quality Python packaging
- Timezone support in executors
- Generating slugs with `django-autoslug` package
- Auto-generate Data name on creation based on template defined in Process
- Added endpoint for adding/removing Data objects to/from Collection

Fixed

- Pass all Resolwe's environment variables to Tox's testing environment
- Include all source files and supplementary package data in sdist
- Make Celery engine work
- Add all permissions to creator of `flow_collection` Collection
- Set DescriptorSchema on creating Data objects and Collections
- Loading DescriptorSchema in tests
- Handle Exceptions if input field doesn't match input schema

- Trigger ORM signals on Data status updates
- Don't set status of Data object to error status if return code of tool is 0

1.9.109 1.0.0 - 2016-03-31

Changed

- Renamed Project to Collection
- Register processes from packages and custom paths
- Removed support for Python 3.3

Added

- Permissions
- API for flow
- Docker executor
- Expression engine support
- Celery engine
- Purge command
- Framework for testing processors
- Processor finders
- Support for Django 1.9
- Support for Python 3.5
- Initial migrations
- Introductory documentation

1.9.110 0.9.0 - 2015-04-09

Added

Initial release.

1.10 Contributing

1.10.1 Installing prerequisites

Resolve runs on [Python 3.6](#) or later. If you don't have it yet, follow [these instructions](#).

It's easiest to run other required services in Docker containers, which is assumed in this tutorial. If you don't have it yet, you can follow the [official Docker tutorial](#) for Mac and for Windows or install it as a distribution's package in most of standard Linux distributions (Fedora, Ubuntu, ...).

1.10.2 Preparing environment

Fork the main Resolwe's git repository.

If you don't have Git installed on your system, follow [these instructions](#).

Clone your fork (replace <username> with your GitHub account name) and change directory:

```
git clone https://github.com/<username>/resolwe.git
cd resolwe
```

Prepare Resolwe for development:

```
pip install -e .[docs,package,test]
```

Note: We recommend using [pyvenv](#) to create an isolated Python environment for Resolwe.

1.10.3 Preparing database

Start Docker containers:

```
cd tests
docker-compose up --detach
```

Set-up database:

```
./manage.py migrate
./manage.py createsuperuser --username admin --email admin@genialis.com
```

1.10.4 Registering processes

```
cd tests
./manage.py register
```

1.10.5 Running tests

To run the tests, use:

```
cd tests
./manage.py test resolwe --parallel=2
```

To run the tests with [Tox](#), use:

```
tox -r
```

1.10.6 Building documentation

```
python setup.py build_sphinx
```

1.10.7 Submitting changes upstream

Signed commits are required in the Resolwe upstream repository. Generate your personal GPG key and [configure Git to use it automatically](#).

1.10.8 Preparing release

Checkout the latest code and create a release branch:

```
git checkout master
git pull
git checkout -b release-<new-version>
```

Replace the *Unreleased* heading in docs/CHANGELOG.rst with the new version, followed by release's date (e.g. *13.2.0 - 2018-10-23*).

Note: Use [Semantic versioning](#).

Commit changes to git:

```
git commit -a -m "Prepare release <new-version>"
```

Push changes to your fork and open a pull request:

```
git push --set-upstream <resolwe-fork-name> release-<new-version>
```

Wait for the tests to pass and the pull request to be approved. Merge the code to master:

```
git checkout master
git merge --ff-only release-<new-version>
git push <resolwe-upstream-name> master <new-version>
```

Tag the new release from the latest commit:

```
git checkout master
git tag -sm "Version <new-version>" <new-version>
```

Note: Project's version will be automatically inferred from the git tag using [setuptools_scm](#).

Push the tag to the main [Resolwe's git repository](#):

```
git push <resolwe-upstream-name> master <new-version>
```

The tagged code will be released to PyPI automatically. Inspect Travis logs of the Release step if errors occur.

Preparing pre-release

When preparing a pre-release (i.e. an alpha release), one can skip the “release” commit that updates the change log and just tag the desired commit with a pre-release tag (e.g. *13.3.0a1*). By pushing it to GitHub, the tagged code will be automatically tested by Travis CI and then released to PyPI.

PYTHON MODULE INDEX

r

- resolwe.flow.executors, 25
- resolwe.flow.executors.docker, 26
- resolwe.flow.executors.docker.prepare, 27
- resolwe.flow.executors.local, 27
- resolwe.flow.executors.local.prepare, 27
- resolwe.flow.executors.local.run, 27
- resolwe.flow.executors.null, 27
- resolwe.flow.executors.null.run, 28
- resolwe.flow.executors.prepare, 25
- resolwe.flow.executors.run, 25
- resolwe.flow.management, 52
- resolwe.flow.management.commands.register, 52
- resolwe.flow.managers, 20
- resolwe.flow.managers.consumer, 24
- resolwe.flow.managers.dispatcher, 20
- resolwe.flow.managers.listener, 24
- resolwe.flow.managers.utils, 25
- resolwe.flow.managers.workload_connectors, 22
- resolwe.flow.managers.workload_connectors.base, 22
- resolwe.flow.managers.workload_connectors.celery, 22
- resolwe.flow.managers.workload_connectors.kubernetes, 23
- resolwe.flow.managers.workload_connectors.local, 22
- resolwe.flow.managers.workload_connectors.slurm, 23
- resolwe.flow.models, 28
- resolwe.flow.utils, 51
- resolwe.flow.utils.exceptions, 51
- resolwe.flow.utils.stats, 51
- resolwe.permissions.shortcuts, 19
- resolwe.permissions.utils, 20
- resolwe.test, 53
- resolwe.test.testcases, 53
- resolwe.test.testcases.api, 57
- resolwe.test.testcases.process, 53
- resolwe.test.utils, 60
- resolwe.utils, 61

Symbols

- `_delete()` (*resolve.test.TransactionResolveAPITestCase* method), 58
 - `_detail_permissions()` (*resolve.test.TransactionResolveAPITestCase* method), 59
 - `_get_detail()` (*resolve.test.TransactionResolveAPITestCase* method), 58
 - `_get_list()` (*resolve.test.TransactionResolveAPITestCase* method), 57
 - `_patch()` (*resolve.test.TransactionResolveAPITestCase* method), 58
 - `_post()` (*resolve.test.TransactionResolveAPITestCase* method), 58
- ## A
- `add()` (*resolve.flow.utils.stats.SimpleLoadAvg* method), 52
 - `add_arguments()` (*resolve.flow.management.commands.register.Command* method), 52
 - `add_to_collection()` (*resolve.flow.models.annotations.AnnotationField* static method), 44
 - `allocated_cpu` (*resolve.flow.models.history.ProcessingHistory* attribute), 50
 - `allocated_memory` (*resolve.flow.models.history.ProcessingHistory* attribute), 50
 - `annotation_fields` (*resolve.flow.models.Collection* attribute), 29
 - `annotation_type` (*resolve.flow.models.annotations.AnnotationField* property), 44
 - `AnnotationField` (class in *resolve.flow.models.annotations*), 44
 - `AnnotationField.DoesNotExist`, 44
 - `AnnotationField.MultipleObjectsReturned`, 44
 - `AnnotationGroup` (class in *resolve.flow.models.annotations*), 43
 - `AnnotationGroup.DoesNotExist`, 43
 - `AnnotationGroup.MultipleObjectsReturned`, 43
 - `AnnotationPreset` (class in *resolve.flow.models.annotations*), 45
 - `AnnotationPreset.DoesNotExist`, 45
 - `AnnotationPreset.MultipleObjectsReturned`, 45
 - `AnnotationValue` (class in *resolve.flow.models.annotations*), 46
 - `AnnotationValue.DoesNotExist`, 46
 - `AnnotationValue.MultipleObjectsReturned`, 46
 - `assertAlmostEqualGeneric()` (*resolve.test.TestCaseHelpers* method), 53
 - `assertAnnotation()` (*resolve.test.TestCaseHelpers* method), 53
 - `assertDir()` (*resolve.test.ProcessTestCase* method), 54
 - `assertDirExists()` (*resolve.test.ProcessTestCase* method), 54
 - `assertDirStructure()` (*resolve.test.ProcessTestCase* method), 54
 - `assertFields()` (*resolve.test.ProcessTestCase* method), 54
 - `assertFile()` (*resolve.test.ProcessTestCase* method), 55
 - `assertFileExists()` (*resolve.test.ProcessTestCase* method), 55
 - `assertFiles()` (*resolve.test.ProcessTestCase* method), 55
 - `assertFilesExist()` (*resolve.test.ProcessTestCase* method), 55
 - `assertJSON()` (*resolve.test.ProcessTestCase* method), 56
 - `assertKeys()` (*resolve.test.TransactionResolveAPITestCase* method), 59
 - `assertStatus()` (*resolve.test.ProcessTestCase* method), 56
- ## B
- `BaseCollection` (class in *resolve.flow.models.collection*), 28
 - `BaseCollection.Meta` (class in *resolve.flow.models.collection*), 28
 - `BaseConnector` (class in *resolve.flow.managers.workload_connectors.base*), 22
 - `BaseFlowExecutor` (class in *resolve.flow.executors.run*), 25
 - `BaseFlowExecutorPreparer` (class in *resolve.flow.executors.run*), 25

- resolve.flow.executors.prepare*), 25
 - BaseModel (class in *resolve.flow.models.base*), 28
 - BaseModel.Meta (class in *resolve.flow.models.base*), 28
 - BraceMessage (class in *resolve.util*s), 61
- C**
- category (*resolve.flow.models.Process* attribute), 39
 - category (*resolve.flow.models.Relation* attribute), 37
 - check_database() (*resolve.flow.managers.consumer.HealthCheckConsumer* method), 24
 - check_docker() (in module *resolve.test.util*s), 60
 - check_installed() (in module *resolve.test.util*s), 60
 - checksum (*resolve.flow.models.Data* attribute), 31
 - child (*resolve.flow.models.DataDependency* attribute), 34
 - cleanup() (*resolve.flow.managers.workload_connectors.base.BaseConnector* method), 22
 - cleanup() (*resolve.flow.managers.workload_connectors.celery.Connector* method), 23
 - cleanup() (*resolve.flow.managers.workload_connectors.kubernetes.Connector* method), 23
 - cleanup() (*resolve.flow.managers.workload_connectors.local.Connector* method), 22
 - cleanup() (*resolve.flow.managers.workload_connectors.slurm.Connector* method), 23
 - Collection (class in *resolve.flow.models*), 29
 - collection (*resolve.flow.models.Data* attribute), 31
 - collection (*resolve.flow.models.Entity* attribute), 35
 - collection (*resolve.flow.models.Relation* attribute), 37
 - Collection.DoesNotExist, 29
 - Collection.MultipleObjectsReturned, 29
 - CollectionChange (class in *resolve.flow.models.history*), 49
 - CollectionChange.DoesNotExist, 49
 - CollectionChange.MultipleObjectsReturned, 49
 - CollectionHistory (class in *resolve.flow.models.history*), 47
 - CollectionHistory.DoesNotExist, 47
 - CollectionHistory.MultipleObjectsReturned, 47
 - CollectionNameChange (class in *resolve.flow.models.history*), 49
 - CollectionNameChange.DoesNotExist, 49
 - CollectionNameChange.MultipleObjectsReturned, 49
 - CollectionSlugChange (class in *resolve.flow.models.history*), 48
 - CollectionSlugChange.DoesNotExist, 48
 - CollectionSlugChange.MultipleObjectsReturned, 48
 - Command (class in *resolve.flow.management.commands.register*), 52
 - communicate() (*resolve.flow.managers.dispatcher.Manager* method), 20
 - ConfigLocation (class in *resolve.flow.managers.workload_connectors.kubernetes*), 23
 - Connector (class in *resolve.flow.managers.workload_connectors.celery*), 23
 - Connector (class in *resolve.flow.managers.workload_connectors.kubernetes*), 23
 - Connector (class in *resolve.flow.managers.workload_connectors.local*), 23
 - Connector (class in *resolve.flow.managers.workload_connectors.slurm*), 23
 - contributor (*resolve.flow.models.annotations.AnnotationPreset* attribute), 45
 - contributor (*resolve.flow.models.base.BaseModel* attribute), 28
 - contributor (*resolve.flow.models.Collection* attribute), 29
 - contributor (*resolve.flow.models.collection.BaseCollection* attribute), 28
 - contributor (*resolve.flow.models.Data* attribute), 31
 - contributor (*resolve.flow.models.DescriptorSchema* attribute), 38
 - contributor (*resolve.flow.models.Entity* attribute), 35
 - contributor (*resolve.flow.models.Process* attribute), 39
 - contributor (*resolve.flow.models.Relation* attribute), 37
 - contributor (*resolve.flow.models.Storage* attribute), 42
 - control_event() (*resolve.flow.managers.consumer.ManagerConsumer* method), 24
 - copy_annotations() (*resolve.flow.models.Entity* method), 35
 - copy_permissions() (in module *resolve.permissions.util*s), 20
 - create_data_location() (in module *resolve.test.util*s), 60
 - created (*resolve.flow.models.annotations.AnnotationPreset* attribute), 45
 - created (*resolve.flow.models.base.BaseModel* attribute), 28
 - created (*resolve.flow.models.Collection* attribute), 29
 - created (*resolve.flow.models.collection.BaseCollection* attribute), 28
 - created (*resolve.flow.models.Data* attribute), 31
 - created (*resolve.flow.models.DataMigrationHistory* attribute), 43
 - created (*resolve.flow.models.DescriptorSchema* attribute), 38
 - created (*resolve.flow.models.Entity* attribute), 35
 - created (*resolve.flow.models.Process* attribute), 39
 - created (*resolve.flow.models.ProcessMigrationHistory* attribute), 42
 - created (*resolve.flow.models.Relation* attribute), 37
 - created (*resolve.flow.models.Storage* attribute), 42

D

- Data (class in *resolwe.flow.models*), 30
- data (*resolwe.flow.models.Storage* attribute), 42
- Data.DoesNotExist, 30
- Data.MultipleObjectsReturned, 30
- data_name (*resolwe.flow.models.Process* attribute), 39
- DataDependency (class in *resolwe.flow.models*), 34
- DataDependency.DoesNotExist, 34
- DataDependency.MultipleObjectsReturned, 34
- DataHistory (class in *resolwe.flow.models.history*), 47
- DataHistory.DoesNotExist, 47
- DataHistory.MultipleObjectsReturned, 47
- DataMigrationHistory (class in *resolwe.flow.models*), 43
- DataMigrationHistory.DoesNotExist, 43
- DataMigrationHistory.MultipleObjectsReturned, 43
- DataNameChange (class in *resolwe.flow.models.history*), 49
- DataNameChange.DoesNotExist, 49
- DataNameChange.MultipleObjectsReturned, 49
- DataSlugChange (class in *resolwe.flow.models.history*), 48
- DataSlugChange.DoesNotExist, 48
- DataSlugChange.MultipleObjectsReturned, 48
- datum (*resolwe.flow.models.history.CollectionHistory* attribute), 47
- datum (*resolwe.flow.models.history.DataHistory* attribute), 47
- default() (*resolwe.flow.managers.dispatcher.SettingsJSONFile* method), 22
- delete() (*resolwe.flow.models.Data* method), 31
- delete_background() (*resolwe.flow.models.Collection* method), 29
- delete_background() (*resolwe.flow.models.Data* method), 31
- delete_background() (*resolwe.flow.models.Entity* method), 35
- deleted (*resolwe.flow.models.history.CollectionHistory* attribute), 47
- deleted (*resolwe.flow.models.history.DataHistory* attribute), 47
- dependency_status() (*resolwe.flow.models.Data* method), 31
- description (*resolwe.flow.models.annotations.AnnotationField* attribute), 44
- description (*resolwe.flow.models.Collection* attribute), 29
- description (*resolwe.flow.models.collection.BaseCollection* attribute), 29
- description (*resolwe.flow.models.DescriptorSchema* attribute), 39
- description (*resolwe.flow.models.Entity* attribute), 35
- description (*resolwe.flow.models.Process* attribute), 39
- descriptor (*resolwe.flow.models.Collection* attribute), 29
- descriptor (*resolwe.flow.models.collection.BaseCollection* attribute), 29
- descriptor (*resolwe.flow.models.Data* attribute), 31
- descriptor (*resolwe.flow.models.Entity* attribute), 35
- descriptor (*resolwe.flow.models.Relation* attribute), 37
- descriptor_dirty (*resolwe.flow.models.Collection* attribute), 30
- descriptor_dirty (*resolwe.flow.models.collection.BaseCollection* attribute), 29
- descriptor_dirty (*resolwe.flow.models.Data* attribute), 31
- descriptor_dirty (*resolwe.flow.models.Entity* attribute), 35
- descriptor_dirty (*resolwe.flow.models.Relation* attribute), 37
- descriptor_schema (*resolwe.flow.models.Collection* attribute), 30
- descriptor_schema (*resolwe.flow.models.collection.BaseCollection* attribute), 29
- descriptor_schema (*resolwe.flow.models.Data* attribute), 31
- descriptor_schema (*resolwe.flow.models.Entity* attribute), 35
- descriptor_schema (*resolwe.flow.models.Relation* attribute), 37
- DescriptorSchema (class in *resolwe.flow.models*), 38
- DescriptorSchema.DoesNotExist, 38
- DescriptorSchema.MultipleObjectsReturned, 38
- detail_permissions() (*resolwe.test.TransactionResolweAPITestCase* method), 59
- detail_url() (*resolwe.test.TransactionResolweAPITestCase* method), 59
- disable_auto_calls() (in module *resolwe.flow.managers.utils*), 25
- discover_engines() (*resolwe.flow.managers.dispatcher.Manager* method), 21
- drain_messages() (*resolwe.flow.managers.dispatcher.Manager* method), 21
- duplicate() (*resolwe.flow.models.Collection* method), 30
- duplicate() (*resolwe.flow.models.Data* method), 32
- duplicate() (*resolwe.flow.models.Entity* method), 35
- duplicated (*resolwe.flow.models.Collection* attribute), 30
- duplicated (*resolwe.flow.models.Data* attribute), 32
- duplicated (*resolwe.flow.models.Entity* attribute), 35

E

- entities (*resolwe.flow.models.Relation* attribute), 37

Entity (class in *resolve.flow.models*), 35
 entity (*resolve.flow.models.annotations.AnnotationValue* attribute), 46
 entity (*resolve.flow.models.Data* attribute), 32
 Entity.DoesNotExist, 35
 Entity.MultipleObjectsReturned, 35
 entity_always_create (*resolve.flow.models.Process* attribute), 40
 entity_descriptor_schema (*resolve.flow.models.Process* attribute), 40
 entity_input (*resolve.flow.models.Process* attribute), 40
 entity_type (*resolve.flow.models.Process* attribute), 40
 execution_barrier() (*resolve.flow.managers.dispatcher.Manager* method), 21
 exit_consumer() (in module *resolve.flow.managers.consumer*), 24
 extend_settings() (*resolve.flow.executors.local.prepare.BaseFlowExecutorPreparer* method), 27
 extend_settings() (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer* method), 25

F

field (*resolve.flow.models.annotations.AnnotationValue* attribute), 46
 field_from_path() (*resolve.flow.models.annotations.AnnotationField* class method), 44
 fields (*resolve.flow.models.annotations.AnnotationPreset* attribute), 45
 files_path (*resolve.test.ProcessTestCase* property), 56
 find_descriptor_schemas() (*resolve.flow.management.commands.register.Command* method), 52
 find_schemas() (*resolve.flow.management.commands.register.Command* method), 52
 finished (*resolve.flow.models.Data* attribute), 32
 FlowExecutor (class in *resolve.flow.executors.local.run*), 27
 FlowExecutor (class in *resolve.flow.executors.null.run*), 28
 FlowExecutorPreparer (class in *resolve.flow.executors.docker.prepare*), 27
 FlowExecutorPreparer (class in *resolve.flow.executors.local.prepare*), 27
 from_path() (*resolve.flow.models.annotations.AnnotationField* static method), 46

G

get_annotation() (*resolve.flow.models.Entity* method), 35
 get_environment_variables() (*resolve.flow.executors.docker.prepare.FlowExecutorPreparer* method), 27
 get_environment_variables() (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer* method), 26
 get_execution_engine() (*resolve.flow.managers.dispatcher.Manager* method), 21
 get_executor() (*resolve.flow.managers.dispatcher.Manager* method), 21
 get_expression_engine() (*resolve.flow.managers.dispatcher.Manager* method), 21
 get_json() (*resolve.test.ProcessTestCase* method), 56
 get_mountable_connectors() (in module *resolve.flow.managers.workload_connectors.kubernetes*), 23
 get_object_perms() (in module *resolve.permissions.shortcuts*), 19
 get_resource_primitives() (*resolve.flow.models.Data* method), 32
 get_resource_primitives() (*resolve.flow.models.Process* method), 40
 get_tools_paths() (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer* method), 26
 get_tools_paths() (*resolve.flow.executors.run.BaseFlowExecutorPreparer* method), 25
 get_upload_dir() (in module *resolve.flow.managers.workload_connectors.kubernetes*), 24
 get_value() (*resolve.flow.models.history.ProcessingHistory* method), 50
 get_value() (*resolve.flow.models.history.TrackChange* method), 47
 group (*resolve.flow.models.annotations.AnnotationField* attribute), 44
 group_field_from_path() (*resolve.flow.models.annotations.AnnotationField* static method), 44

H

handle() (*resolve.flow.management.commands.register.Command* method), 52
 handle_control_event() (*resolve.flow.managers.dispatcher.Manager* method), 21
 has_permission() (*resolve.flow.models.annotations.AnnotationValue* method), 46
 HealthCheckConsumer (class in *resolve.flow.managers.consumer*), 24
 health_check() (*resolve.flow.managers.consumer.HealthCheckConsumer* method), 24

history (*resolwe.flow.models.history.CollectionChange attribute*), 49
 history (*resolwe.flow.models.history.CollectionNameChange attribute*), 49
 history (*resolwe.flow.models.history.CollectionSlugChange attribute*), 48
 history (*resolwe.flow.models.history.DataNameChange attribute*), 49
 history (*resolwe.flow.models.history.DataSlugChange attribute*), 48
 history (*resolwe.flow.models.history.ProcessingHistory attribute*), 51
 history (*resolwe.flow.models.history.SizeChange attribute*), 50

I

id_from_path() (*resolwe.flow.models.annotations.AnnotationField class method*), 44
 input (*resolwe.flow.models.Data attribute*), 32
 input_schema (*resolwe.flow.models.Process attribute*), 40
 interval (*resolwe.flow.models.history.ProcessingHistory attribute*), 51
 invalid_annotation_fields() (*resolwe.flow.models.Entity method*), 35
 is_active (*resolwe.flow.models.Process attribute*), 40
 is_duplicate() (*resolwe.flow.models.Collection method*), 30
 is_duplicate() (*resolwe.flow.models.Data method*), 32
 is_duplicate() (*resolwe.flow.models.Entity method*), 36
 is_testing() (*in module resolwe.test.utils*), 60

J

json (*resolwe.flow.models.Storage attribute*), 42

K

keep_data() (*resolwe.test.TestCaseHelpers method*), 53
 kind (*resolwe.flow.models.DataDependency attribute*), 34
 KIND_IO (*resolwe.flow.models.DataDependency attribute*), 34
 KIND_SUBPROCESS (*resolwe.flow.models.DataDependency attribute*), 34

L

label (*resolwe.flow.models.annotations.AnnotationField attribute*), 44
 label (*resolwe.flow.models.annotations.AnnotationGroup attribute*), 43
 label (*resolwe.flow.models.annotations.AnnotationValue property*), 46

label_by_value() (*resolwe.flow.models.annotations.AnnotationField method*), 44
 list_url (*resolwe.test.TransactionResolweAPITestCase property*), 60
 load_execution_engines() (*resolwe.flow.managers.dispatcher.Manager method*), 21
 load_executor() (*resolwe.flow.managers.dispatcher.Manager method*), 21
 load_expression_engines() (*resolwe.flow.managers.dispatcher.Manager method*), 21
 location (*resolwe.flow.models.Data attribute*), 32

M

Manager (*class in resolwe.flow.managers.dispatcher*), 20
 ManagerField (*in module resolwe.flow.managers*), 20
 ManagerConsumer (*class in resolwe.flow.managers.consumer*), 24
 metadata (*resolwe.flow.models.DataMigrationHistory attribute*), 43
 metadata (*resolwe.flow.models.ProcessMigrationHistory attribute*), 43
 migration (*resolwe.flow.models.DataMigrationHistory attribute*), 43
 migration (*resolwe.flow.models.ProcessMigrationHistory attribute*), 43
 modified (*resolwe.flow.models.annotations.AnnotationPreset attribute*), 45
 modified (*resolwe.flow.models.annotations.AnnotationValue attribute*), 46
 modified (*resolwe.flow.models.base.BaseModel attribute*), 28
 modified (*resolwe.flow.models.Collection attribute*), 30
 modified (*resolwe.flow.models.collection.BaseCollection attribute*), 29
 modified (*resolwe.flow.models.Data attribute*), 32
 modified (*resolwe.flow.models.DescriptorSchema attribute*), 39
 modified (*resolwe.flow.models.Entity attribute*), 36
 modified (*resolwe.flow.models.Process attribute*), 40
 modified (*resolwe.flow.models.Relation attribute*), 38
 modified (*resolwe.flow.models.Storage attribute*), 42
 module
 resolwe.flow.executors, 25
 resolwe.flow.executors.docker, 26
 resolwe.flow.executors.docker.prepare, 27
 resolwe.flow.executors.local, 27
 resolwe.flow.executors.local.prepare, 27
 resolwe.flow.executors.local.run, 27
 resolwe.flow.executors.null, 27
 resolwe.flow.executors.null.run, 28
 resolwe.flow.executors.prepare, 25
 resolwe.flow.executors.run, 25

resolwe.flow.management, 52
 resolwe.flow.management.commands.register, 52
 resolwe.flow.managers, 20
 resolwe.flow.managers.consumer, 24
 resolwe.flow.managers.dispatcher, 20
 resolwe.flow.managers.listener, 24
 resolwe.flow.managers.utils, 25
 resolwe.flow.managers.workload_connectors, 22
 resolwe.flow.managers.workload_connectors.base, 22
 resolwe.flow.managers.workload_connectors.celery, 22
 resolwe.flow.managers.workload_connectors.kubernetes, 23
 resolwe.flow.managers.workload_connectors.local, 22
 resolwe.flow.managers.workload_connectors.slurm, 23
 resolwe.flow.models, 28
 resolwe.flow.utils, 51
 resolwe.flow.utils.exceptions, 51
 resolwe.flow.utils.stats, 51
 resolwe.permissions.shortcuts, 19
 resolwe.permissions.utils, 20
 resolwe.test, 53
 resolwe.test.testcases, 53
 resolwe.test.testcases.api, 57
 resolwe.test.testcases.process, 53
 resolwe.test.utils, 60
 resolwe.utils, 61
 node_cpu (*resolwe.flow.models.history.ProcessingHistory* attribute), 51
 node_memory (*resolwe.flow.models.history.ProcessingHistory* attribute), 51
 NumberSeriesShape (*class in resolwe.flow.utils.stats*), 51
O
 objects (*resolwe.flow.models.Collection* attribute), 30
 objects (*resolwe.flow.models.Data* attribute), 32
 objects (*resolwe.flow.models.Entity* attribute), 36
 objects (*resolwe.flow.models.history.CollectionChange* attribute), 50
 objects (*resolwe.flow.models.history.CollectionNameChange* attribute), 49
 objects (*resolwe.flow.models.history.CollectionSlugChange* attribute), 48
 objects (*resolwe.flow.models.history.DataNameChange* attribute), 49
 objects (*resolwe.flow.models.history.DataSlugChange* attribute), 48
 objects (*resolwe.flow.models.history.ProcessingHistory* attribute), 51
 objects (*resolwe.flow.models.history.SizeChange* attribute), 50
 objects (*resolwe.flow.models.history.TrackChange* attribute), 47
 objects (*resolwe.flow.models.Relation* attribute), 38
 objects (*resolwe.flow.models.Storage* attribute), 42
 optimize_job_scheduling()
 (*resolwe.flow.managers.workload_connectors.kubernetes.Connection* method), 23
 ordered (*resolwe.flow.models.RelationType* attribute), 38
 output (*resolwe.flow.models.Data* attribute), 32
 output_schema (*resolwe.flow.models.Process* attribute), 40

N

name (*resolwe.flow.models.annotations.AnnotationField* attribute), 44
 name (*resolwe.flow.models.annotations.AnnotationGroup* attribute), 43
 name (*resolwe.flow.models.annotations.AnnotationPreset* attribute), 45
 name (*resolwe.flow.models.base.BaseModel* attribute), 28
 name (*resolwe.flow.models.Collection* attribute), 30
 name (*resolwe.flow.models.collection.BaseCollection* attribute), 29
 name (*resolwe.flow.models.Data* attribute), 32
 name (*resolwe.flow.models.DescriptorSchema* attribute), 39
 name (*resolwe.flow.models.Entity* attribute), 36
 name (*resolwe.flow.models.Process* attribute), 40
 name (*resolwe.flow.models.Relation* attribute), 38
 name (*resolwe.flow.models.RelationType* attribute), 38
 name (*resolwe.flow.models.Storage* attribute), 42
 named_by_user (*resolwe.flow.models.Data* attribute), 32

P

parent (*resolwe.flow.models.DataDependency* attribute), 34
 parents (*resolwe.flow.models.Data* attribute), 32
 permission_group (*resolwe.flow.models.annotations.AnnotationPreset* attribute), 45
 permission_group (*resolwe.flow.models.Collection* attribute), 30
 permission_group (*resolwe.flow.models.Data* attribute), 32
 permission_group (*resolwe.flow.models.DescriptorSchema* attribute), 39
 permission_group (*resolwe.flow.models.Entity* attribute), 36
 permission_group (*resolwe.flow.models.Process* attribute), 41

permission_group (*resolve.flow.models.Relation attribute*), 38
 persistence (*resolve.flow.models.Process attribute*), 41
 PERSISTENCE_CACHED (*resolve.flow.models.Process attribute*), 39
 PERSISTENCE_RAW (*resolve.flow.models.Process attribute*), 39
 PERSISTENCE_TEMP (*resolve.flow.models.Process attribute*), 39
 post_register_hook()
 (*resolve.flow.executors.docker.prepare.FlowExecutorPreparer method*), 27
 post_register_hook()
 (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer method*), 26
 preparation_stage() (*resolve.test.ProcessTestCase method*), 56
 prepare_for_execution()
 (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer method*), 26
 preprocess_value() (*resolve.flow.models.history.CollectionChange static method*), 50
 preprocess_value() (*resolve.flow.models.history.TrackChange static method*), 47
 Process (*class in resolve.flow.models*), 39
 process (*resolve.flow.models.Data attribute*), 32
 Process.DoesNotExist, 39
 Process.MultipleObjectsReturned, 39
 process_change() (*resolve.flow.models.history.ProcessingHistory class method*), 51
 process_change() (*resolve.flow.models.history.TrackChange class method*), 48
 process_cores (*resolve.flow.models.Data attribute*), 32
 process_error (*resolve.flow.models.Data attribute*), 32
 process_info (*resolve.flow.models.Data attribute*), 32
 process_memory (*resolve.flow.models.Data attribute*), 33
 process_pid (*resolve.flow.models.Data attribute*), 33
 process_progress (*resolve.flow.models.Data attribute*), 33
 process_rc (*resolve.flow.models.Data attribute*), 33
 process_resources (*resolve.flow.models.Data attribute*), 33
 process_warning (*resolve.flow.models.Data attribute*), 33
 ProcessingHistory (*class in resolve.flow.models.history*), 50
 ProcessingHistory.DoesNotExist, 50
 ProcessingHistory.MultipleObjectsReturned, 50
 ProcessMigrationHistory (*class in resolve.flow.models*), 42
 ProcessMigrationHistory.DoesNotExist, 42
 ProcessMigrationHistory.MultipleObjectsReturned, 42
 ProcessTestCase (*class in resolve.test*), 53
R
 recompute_label() (*resolve.flow.models.annotations.AnnotationValue method*), 46
 register_descriptors()
 (*resolve.flow.management.commands.register.Command method*), 52
 register_processes()
 (*resolve.flow.management.commands.register.Command method*), 52
 Relation (*class in resolve.flow.models*), 37
 Relation.DoesNotExist, 37
 Relation.MultipleObjectsReturned, 37
 RelationType (*class in resolve.flow.models*), 38
 RelationType.DoesNotExist, 38
 RelationType.MultipleObjectsReturned, 38
 required (*resolve.flow.models.annotations.AnnotationField attribute*), 44
 requirements (*resolve.flow.models.Process attribute*), 41
 resolve_data_path()
 (*resolve.flow.executors.docker.prepare.FlowExecutorPreparer method*), 27
 resolve_data_path()
 (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer method*), 26
 resolve_secrets() (*resolve.flow.models.Data method*), 33
 resolve_upload_path()
 (*resolve.flow.executors.docker.prepare.FlowExecutorPreparer method*), 27
 resolve_upload_path()
 (*resolve.flow.executors.prepare.BaseFlowExecutorPreparer method*), 26
 resolve.flow.executors
 module, 25
 resolve.flow.executors.docker
 module, 26
 resolve.flow.executors.docker.prepare
 module, 27
 resolve.flow.executors.local
 module, 27
 resolve.flow.executors.local.prepare
 module, 27
 resolve.flow.executors.local.run
 module, 27
 resolve.flow.executors.null
 module, 27
 resolve.flow.executors.null.run
 module, 28
 resolve.flow.executors.prepare
 module, 25

resolwe.flow.executors.run
 module, 25
 resolwe.flow.management
 module, 52
 resolwe.flow.management.commands.register
 module, 52
 resolwe.flow.managers
 module, 20
 resolwe.flow.managers.consumer
 module, 24
 resolwe.flow.managers.dispatcher
 module, 20
 resolwe.flow.managers.listener
 module, 24
 resolwe.flow.managers.utils
 module, 25
 resolwe.flow.managers.workload_connectors
 module, 22
 resolwe.flow.managers.workload_connectors.base
 module, 22
 resolwe.flow.managers.workload_connectors.celery
 module, 22
 resolwe.flow.managers.workload_connectors.kubernetes
 module, 23
 resolwe.flow.managers.workload_connectors.local
 module, 22
 resolwe.flow.managers.workload_connectors.slurm
 module, 23
 resolwe.flow.models
 module, 28
 resolwe.flow.utils
 module, 51
 resolwe.flow.utils.exceptions
 module, 51
 resolwe.flow.utils.stats
 module, 51
 resolwe.permissions.shortcuts
 module, 19
 resolwe.permissions.utils
 module, 20
 resolwe.test
 module, 53
 resolwe.test.testcases
 module, 53
 resolwe.test.testcases.api
 module, 57
 resolwe.test.testcases.process
 module, 53
 resolwe.test.utils
 module, 60
 resolwe.utils
 module, 61
 resolwe_exception_handler() (in module
 resolwe.flow.utils.exceptions), 51
 ResolweAPITestCase (class in resolwe.test), 60
 restart() (resolwe.flow.models.Data method), 33
 retire() (resolwe.flow.management.commands.register.Command
 method), 52
 revalidate_values()
 (resolwe.flow.models.annotations.AnnotationField
 method), 44
 run (resolwe.flow.models.Process attribute), 41
 run() (resolwe.flow.executors.run.BaseFlowExecutor
 method), 25
 run() (resolwe.flow.managers.dispatcher.Manager
 method), 21
 run_consumer() (in module
 resolwe.flow.managers.consumer), 25
 run_process() (resolwe.test.ProcessTestCase method),
 57
 run_processor() (resolwe.test.ProcessTestCase
 method), 57
S
 sanitize_kubernetes_label() (in module
 resolwe.flow.managers.workload_connectors.kubernetes),
 24
 save() (resolwe.flow.models.annotations.AnnotationField
 method), 44
 save() (resolwe.flow.models.annotations.AnnotationValue
 method), 46
 save() (resolwe.flow.models.base.BaseModel method),
 28
 save() (resolwe.flow.models.Collection method), 30
 save() (resolwe.flow.models.collection.BaseCollection
 method), 29
 save() (resolwe.flow.models.Data method), 33
 save() (resolwe.flow.models.Relation method), 38
 save_dependencies() (resolwe.flow.models.Data
 method), 33
 scheduled (resolwe.flow.models.Data attribute), 33
 scheduling_class (resolwe.flow.models.Process
 attribute), 41
 schema (resolwe.flow.models.DescriptorSchema at-
 tribute), 39
 search (resolwe.flow.models.Collection attribute), 30
 search (resolwe.flow.models.collection.BaseCollection
 attribute), 29
 search (resolwe.flow.models.Data attribute), 33
 search (resolwe.flow.models.Entity attribute), 36
 Secret (class in resolwe.flow.models), 42
 send_event() (in module
 resolwe.flow.managers.consumer), 25
 set_annotation() (resolwe.flow.models.Entity
 method), 36
 SettingsJSONifier (class in
 resolwe.flow.managers.dispatcher), 21
 setUp() (resolwe.test.ProcessTestCase method), 57

- setUp() (*resolwe.test.TestCaseHelpers* method), 53
 - setUp() (*resolwe.test.TransactionResolweAPITestCase* method), 60
 - setUp() (*resolwe.test.TransactionTestCase* method), 53
 - SimpleLoadAvg (class in *resolwe.flow.utils.stats*), 52
 - size (*resolwe.flow.models.Data* attribute), 33
 - SizeChange (class in *resolwe.flow.models.history*), 50
 - SizeChange.DoesNotExist, 50
 - SizeChange.MultipleObjectsReturned, 50
 - slug (*resolwe.flow.models.annotations.AnnotationPreset* attribute), 45
 - slug (*resolwe.flow.models.base.BaseModel* attribute), 28
 - slug (*resolwe.flow.models.Collection* attribute), 30
 - slug (*resolwe.flow.models.collection.BaseCollection* attribute), 29
 - slug (*resolwe.flow.models.Data* attribute), 33
 - slug (*resolwe.flow.models.DescriptorSchema* attribute), 39
 - slug (*resolwe.flow.models.Entity* attribute), 36
 - slug (*resolwe.flow.models.Process* attribute), 41
 - slug (*resolwe.flow.models.Relation* attribute), 38
 - slug (*resolwe.flow.models.Storage* attribute), 42
 - sort_order (*resolwe.flow.models.annotations.AnnotationField* attribute), 44
 - sort_order (*resolwe.flow.models.annotations.AnnotationGroup* attribute), 43
 - start() (*resolwe.flow.executors.run.BaseFlowExecutor* method), 25
 - start() (*resolwe.flow.managers.workload_connectors.kubernetes.Connector* method), 23
 - started (*resolwe.flow.models.Data* attribute), 34
 - status (*resolwe.flow.models.Data* attribute), 34
 - STATUS_DIRTY (*resolwe.flow.models.Data* attribute), 30
 - STATUS_DONE (*resolwe.flow.models.Data* attribute), 31
 - STATUS_ERROR (*resolwe.flow.models.Data* attribute), 31
 - STATUS_PREPARING (*resolwe.flow.models.Data* attribute), 31
 - STATUS_PROCESSING (*resolwe.flow.models.Data* attribute), 31
 - STATUS_RESOLVING (*resolwe.flow.models.Data* attribute), 31
 - STATUS_UPLOADING (*resolwe.flow.models.Data* attribute), 31
 - STATUS_WAITING (*resolwe.flow.models.Data* attribute), 31
 - Storage (class in *resolwe.flow.models*), 42
 - Storage.DoesNotExist, 42
 - Storage.MultipleObjectsReturned, 42
 - submit() (*resolwe.flow.managers.workload_connectors.base.BaseConnector* method), 22
 - submit() (*resolwe.flow.managers.workload_connectors.celery.Connector* method), 23
 - submit() (*resolwe.flow.managers.workload_connectors.kubernetes.Connector* method), 23
 - submit() (*resolwe.flow.managers.workload_connectors.local.Connector* method), 22
 - submit() (*resolwe.flow.managers.workload_connectors.slurm.Connector* method), 23
- ## T
- tags (*resolwe.flow.models.Collection* attribute), 30
 - tags (*resolwe.flow.models.collection.BaseCollection* attribute), 29
 - tags (*resolwe.flow.models.Data* attribute), 34
 - tags (*resolwe.flow.models.Entity* attribute), 36
 - tearDown() (*resolwe.test.ProcessTestCase* method), 57
 - TestCase (class in *resolwe.test*), 53
 - TestCaseHelpers (class in *resolwe.test*), 53
 - timestamp (*resolwe.flow.models.history.CollectionChange* attribute), 50
 - timestamp (*resolwe.flow.models.history.CollectionNameChange* attribute), 49
 - timestamp (*resolwe.flow.models.history.CollectionSlugChange* attribute), 48
 - timestamp (*resolwe.flow.models.history.DataNameChange* attribute), 49
 - timestamp (*resolwe.flow.models.history.DataSlugChange* attribute), 48
 - timestamp (*resolwe.flow.models.history.DataTimestampChange* attribute), 48
 - timestamp (*resolwe.flow.models.history.ProcessingHistory* attribute), 51
 - timestamp (*resolwe.flow.models.history.SizeChange* attribute), 50
 - timestamp (*resolwe.flow.models.history.TrackChange* attribute), 48
 - to_dict() (*resolwe.flow.utils.stats.NumberSeriesShape* method), 51
 - to_dict() (*resolwe.flow.utils.stats.SimpleLoadAvg* method), 52
 - TrackChange (class in *resolwe.flow.models.history*), 47
 - TrackChange.Meta (class in *resolwe.flow.models.history*), 47
 - TransactionResolweAPITestCase (class in *resolwe.test*), 57
 - TransactionTestCase (class in *resolwe.test*), 53
 - type (*resolwe.flow.models.annotations.AnnotationField* attribute), 45
 - type (*resolwe.flow.models.Entity* attribute), 36
 - type (*resolwe.flow.models.Process* attribute), 41
 - type (*resolwe.flow.models.Relation* attribute), 38
- ## U
- unique_volume_name() (in module *resolwe.flow.managers.workload_connectors.kubernetes*), 24
 - unit (*resolwe.flow.models.Relation* attribute), 38
 - update() (*resolwe.flow.utils.stats.NumberSeriesShape* method), 51

update_annotations() (*resolwe.flow.models.Entity* method), 36

with_null_executor() (*in module resolwe.test.utils*), 60

with_resolwe_host() (*in module resolwe.test.utils*), 61

V

valid (*resolwe.flow.models.history.CollectionHistory* attribute), 47

valid (*resolwe.flow.models.history.DataHistory* attribute), 47

valid() (*resolwe.flow.management.commands.register.Command* method), 53

validate() (*resolwe.flow.models.annotations.AnnotationValue* method), 46

validate_annotations() (*resolwe.flow.models.Entity* method), 36

validate_change_collection() (*resolwe.flow.models.Data* method), 34

validator_regex (*resolwe.flow.models.annotations.AnnotationField* attribute), 45

value (*resolwe.flow.models.annotations.AnnotationValue* property), 46

value (*resolwe.flow.models.history.CollectionChange* attribute), 50

value (*resolwe.flow.models.history.CollectionNameChange* attribute), 49

value (*resolwe.flow.models.history.CollectionSlugChange* attribute), 48

value (*resolwe.flow.models.history.DataNameChange* attribute), 49

value (*resolwe.flow.models.history.DataSlugChange* attribute), 48

value (*resolwe.flow.models.history.SizeChange* attribute), 50

version (*resolwe.flow.models.annotations.AnnotationPreset* attribute), 45

version (*resolwe.flow.models.base.BaseModel* attribute), 28

version (*resolwe.flow.models.Collection* attribute), 30

version (*resolwe.flow.models.collection.BaseCollection* attribute), 29

version (*resolwe.flow.models.Data* attribute), 34

version (*resolwe.flow.models.DescriptorSchema* attribute), 39

version (*resolwe.flow.models.Entity* attribute), 36

version (*resolwe.flow.models.Process* attribute), 41

version (*resolwe.flow.models.Relation* attribute), 38

version (*resolwe.flow.models.Storage* attribute), 42

vocabulary (*resolwe.flow.models.annotations.AnnotationField* attribute), 45

W

with_custom_executor() (*in module resolwe.test.utils*), 60

with_docker_executor() (*in module resolwe.test.utils*), 60