
Spynl Documentation

Release 5.50.rc1

Softwear BV

February 09, 2017

1	Installation	3
2	Short developer tutorial	5
2.1	Development tutorial	5
3	Short operations tutorial	11
3.1	Building & Deploying Tutorial	11
4	In-depth documentation	15
4.1	ini-settings	15
4.2	Your production environment(s)	15
4.3	Parameter Handling	16
4.4	Serialisation	16
4.5	Translations	16
4.6	Routing	16
4.7	HTML emails	17
4.8	Custom get_user_info	17
4.9	Validation per JSON Schema	18
4.10	Aggregation of Logs	18

Spynl is a Python web framework which extends the [Pyramid](#) web framework.

Spynl helps you to manage your web application:

- **Build** your app (via [Jenkins](#))
- **Deploy** your app (via [Docker](#))
- Serve **endpoint documentation** for frontend devs (via [Swagger](#))
- **Inspect** settings and meta-data of running instances in the browser
- **Aggregate** performance indicators and error messages in [NewRelic](#) and/or [Sentry](#)

Spynl also has a few other in-built utilities which are often necessary in a modern professional web application but easily take a few days to get right:

- Manage translations (via [Babel](#))
- Send templatable, translatable HTML emails
- validate JSON input and output with Schemas

Installation

Here is a (very) quick How-To for installing Spynl:

```
$ pip install spynl
$ spynl dev.serve
```

Now you can visit Spynl's in-built /about endpoint:

```
$ curl http://localhost:6543/about
```

And you should get a response like this:

```
{
  "status": "ok",
  "language": "en",
  "time": "2017-02-03T10:13:41+0000",
  "plugins": {},
  "message": "This is a Spynl web application. You can get more information at about/endpoints, about/ini, and about/faq.",
  "spynl_version": "6.0.1"
}
```

So you see there are a few endpoints with more specific information. Try visiting <http://localhost:6543/about/endpoints> to see the documentation for in-built endpoints and <http://localhost:6543/about/ini> to see possible settings.

The *plugins* part is empty because we haven't written any code of our own yet. See the next section for that.

Short developer tutorial

A small tutorial where we build up a small app step by step.

2.1 Development tutorial

In this minimalistic tutorial we will make a Spynl plugin called *my-package* with one endpoint. It will be served in a dev environment, documented, tested and translated. We will build the code part by part, but you can browse [my-package-index](#).

The *spynl* CLI has support for common tasks in the development cycle, which we'll use as they become relevant.

2.1.1 A plugin with one endpoint

First, you probably want to [make a virtual environment and activate it](#). Remember to use Python3, preferably ≥ 3.5 .

Then we install Spynl:

```
$ pip install spynl
```

Now we'll create a folder for our package

```
$ cd $VIRTUAL_ENV/src
$ mkdir my-package
$ touch my-package/setup.py my-package/hello_world.py
```

Here is the content of our (very) simple *setup.py*, where we state that *my-package* is a Spynl plugin:

```
from setuptools import setup

setup(name='my-package',
      entry_points={
          'spynl.plugins': [
              'hello_world = 1',
          ]
      }
)
```

And here is *hello_world.py*. We write one endpoint and the registration for it:

```
def hello(request):
    return dict(message="Hello, world!")
```

```
def includeme(config):
    config.add_endpoint(hello, 'hello')
```

Finally, we develop our package so Spynl knows about it and serve the application:

```
$ python setup.py develop
$ spynl dev.serve
```

(As you see, the *spynl* CLI command works from anywhere when you have your virtual environment activated).

The endpoint <http://localhost:6543/hello> answers:

```
{
  "status": "ok",
  "message": "Hello, world!"
}
```

2.1.2 Adding documentation for the endpoint

Now let's document the endpoint for frontend developers:

```
def hello(request):
    """
    Say hello to the world.

    ---
    get:
    description: >

        #####Response

        JSON keys | Content Type | Description\n
        ----- | ----- | -----\n
        status    | string | 'ok' or 'error'\n
        message   | string | Hello, world!\n

    tags:
        - my-package
    show-try: true
    """
    return dict(message="Hello, world!")

def includeme(config):
    config.add_endpoint(hello, 'hello')
```

Then, the Swagger doc at <http://localhost/about/endpoints> actually lists our endpoint:

Spynl Endpoints

A list of all endpoints on this Spynl instance, with a short description on how to use them.

All endpoints usually return application/json, unless otherwise specified here or requested differently by the request. They will have a "status" (ok|error) field and all error responses also will have a "message" field.

about	Show/Hide	List Operations	Expand Operations
contact	Show/Hide	List Operations	Expand Operations
my-package	Show/Hide	List Operations	Expand Operations
GET /hello	Say hello to the world.		

Click on the endpoint to see details or try it out:

my-package	Show/Hide	List Operations	Expand Operations
GET /hello	Say hello to the world.		
Implementation Notes			
This is a simple endpoint used for the basic Spynl tutorial.			
Response			
JSON keys	Content Type	Description	
status	string	'ok' or 'error'	
message	string	Hello, world!	
<input type="button" value="Try it out!"/>			

We are not using Swagger to its full potential here w.r.t. to its schema capabilities, we know. We chose not to, you can choose otherwise.

2.1.3 Serve on localhost

You already saw how to serve the app:

```
$ spynl dev.serve
```

2.1.4 Getting help about spynl CLI tasks

Now that we begin using the *sspynl* CLI, we should note that for each CLI task, you can get help:

```
$ spynl --help dev.serve
Usage: spynl [--core-opts] dev.serve [other tasks here ...]

Docstring:
Run a local server. The ini-file development.ini is searched for in
installed Spynl plugins. If there is none, minimal.ini is used.

Options:
none
```

2.1.5 Testing the endpoint

Let's write a simple test in *my-package/test_hello.py*:

```
import pytest
from webtest import TestApp
from splyn.main import main

@pytest.fixture(scope="session")
def app():
    splyn_app = main(None)
    return TestApp(splyn_app)

def test_hello(app):
    response = app.get('/hello', status=200)
    assert response.json['message'] == "Hello, world!"
```

Then, we can run:

```
$ splyn dev.test
```

I hope you saw this (the dot says it succeeded):

```
[splyn dev.test] Testing package: my-package
===== test session starts =====
platform linux -- Python 3.5.2, pytest-3.0.5, py-1.4.32, pluggy-0.4.0
rootdir: /home/nicolas/workspace/splyn-git/venv/src/my-package, inifile:
plugins: sugar-0.8.0, cov-2.4.0, raisesregexp-2.1
collected 1 items

test_hello.py .
```

2.1.6 Adding translations

Then we have support for translating the app. Let us add a translatable string to the *hello_world* endpoint:

```
from splyn.main.locale import SplynTranslationString as _

def hello(request):
    return dict(message=_('hello-msg', default="Hello, world!"))
```

We can now refresh the translation catalogue of our package:

```
$ splyn dev.translate --packages my-package --languages nl --action refresh
[splyn dev.translate] Package: my-package ...
[splyn dev.translate] Creating locale folder ...
running extract_messages
extracting messages from hello_world.py
extracting messages from setup.py
extracting messages from test_hello.py
writing PO template file to ./locale/messages.pot
[splyn dev.translate] File ./locale/nl/LC_MESSAGES/my-package.po does not exist. Initializing.
running init_catalog
creating catalog ./locale/nl/LC_MESSAGES/my-package.po based on ./locale/messages.pot
[splyn dev.translate] Done with language nl.
-----
```

Splyn created all necessary folders and initialised a catalogue. Now a human needs to translate our string to Dutch. Make this change in *my-package/locale/nl/LC_MESSAGES/my-package.po*:

```
#: hello_world.py:23
msgid "hello-msg"
msgstr "Hallo, Wereld!"
```

Then we can compile the catalogue, so that Spynl will serve Dutch when it should:

```
$ spynl dev.translate --packages my-package --languages nl
[spynl dev.translate] Package: my-package ...
[spynl dev.translate] Located locale folder in /home/nicolas/workspace/spynl/venv/src/my-package ...
running compile_catalog
compiling catalog /home/nicolas/workspace/spynl/venv/src/my-package/locale/nl/LC_MESSAGES/my-package
[spynl dev.translate] Done with language nl.
-----
```

There are only two actions, *refresh* and *compile*. If the `-action` parameter is not given, `spynl dev.translate` compiles.

The compilation step is not necessary and you don't have to include the binary `.mo` file in your SCM. When we build a Docker image on Jenkins (see below), Jenkins runs the compile action.

we need to tell Pyramid that the new locale directory exists. Add this to the `include_me` function in `my-package/hello_world.py`:

```
config.add_translation_dirs('%s/src/my-package/locale'
                             % os.environ['VIRTUAL_ENV'])
```

Now we want to see our app serve Dutch. We still need to configure the list of languages we serve in our app. This is a great opportunity to start using our own `.ini` file. Copy Spynl's `minimal.ini` to `my-package/development.ini` and add the `spynl.languages` setting in the `[app:main]` section:

```
[app:main]
use = egg:spynl
spynl.pretty = 1
spynl.languages = nl,en
```

It is crucial which language is first in this list. Because `nl` is first, we'll get a Dutch reply from Spynl, e.g. by visiting <http://localhost:6543>:

FIXME: However, <http://localhost:6543/hello> still returns english ...

2.1.7 Tab completion for the spynl CLI

Now that we're `spynl` power users, it's time to reveal an important feature: There is tab completion for the `spynl` CLI. To activate it, run

```
$ source $VIRTUAL_ENV/lib/python3.5/site-packages/spynl/spynl/cli/zsh.completion
```

(you might need to adapt the path to `spynl`, it depends on your environment and method of installation)

You can list (a subset of) tasks by pressing TAB and if the task is complete also the available options. To see options, type a dash ("-") and the press TAB.

This is available for `bash` and `fish` as well, simply replace `zsh` in the command.

2.1.8 Installing the package from SCM

Of course, we will want to use Source Code Management (SCM) for our own code, e.g. on github or bitbucket. `spynl` provides a task called `dev.install` which makes it easy to get started in a new dev environment with developing your app further.

Let's assume your project uses git as SCM and lives in a bitbucket repo:

```
$ spynl dev.install --scm-url git@bitbucket.org:my-team/my-package.git
```

spynl dev.install will clone the code and develop it.

In general, Spynl also supports mercurial repositories.

There are some configuration options here (try *spynl -help dev.install* for all of them). For example, let's assume you want to work with a feature branch and you want/need to specify in which directory the code should be installed:

```
$ spynl dev.install --scm-url git@bitbucket.org:my-team/my-package.git --revision me/some-feature --s
```

spynl dev.install can also install non-Python dependencies for you or do any other things pre- or post-installation. See *setup.sh.template* in the main Spynl repo. (TODO: point to actual documentation of *setup.sh*).

Short operations tutorial

We show what steps are necessary to get your Spynl-based app built, deployed and smoke-tested.

3.1 Building & Deploying Tutorial

(work in progress)

At some point, the *my-package* app should be built, so it can run somewhere else than on localhost:6543. In Spynl, this means starting a job on Jenkins, where the following stages happen:

- All tests are run
- A Docker image is built and deployed to your dev container registry
- Smoke tests are run against a freshly-started container based on that new image

There is only one *spynl* command necessary here: *spynl dev.start_build*. However, some services need to be set up and configured, namely Jenkins and one or two container registries.

3.1.1 Creating a Jenkins job

Jenkins is an open-source build server. We'll assume in this tutorial that you installed one locally and have it Configuring your Spynl plugin for Deployment running at <http://localhost:8080>. Jenkins needs Docker engine installed on it, plus *aws-utils* (?) and any libraries you need for tests to be run.

Spynl uses the Pipeline feature of Jenkins. Here is how Jenkins displays your build history:

		Unit Tests	Deploy	Smoke Tests
Average stage times:		5min 43s	6min 30s	2min 28s
#394	Feb 03 15:43 No Changes	6min 5s	7min 24s	2min 43s
#393	Feb 03 15:08 No Changes	6min 23s	8min 2s	3min 8s

If any error happens, your build has failed - the stage is coloured red and the pipeline aborts.

After you have set up a Jenkins server, you need to create a *Pipeline* job called “Spynl”. In the Jenkins web interface, navigate to *Jenkins* -> *New Item* -> *Pipeline*.

Configure the Pipeline “Definition” to be a *Pipeline script from SCM*, the “SCM” to be *Git* and the “Repository URL” to be `https://github.com/SoftwearDevelopment/spynl.git`. under “Branches to build”, add `refs/heads/$spynlbranch`.

Finally you need to add a few String parameters to the job:

- `scm_urls`
- `revision`
- `fallbackrevision` (default: `$revision`)
- `task`
- `spynlbranch`

That’s it! Save the Jenkins job.

3.1.2 Configuring your Spynl plugin for Jenkins

Add the following ini-setting to `development.ini`:

```
[app:main]
spynl.ops.jenkins_url: http://localhost:8080
```

Now we can start a build:

```
$ spynl ops.start_build
```

This will build the *master* branch on *spynl* as well as *my-package*. See `spynl -help ops.start_build` for more options.

3.1.3 Configuring your Spynl plugin for Deployment

(TODO: more verbose, this is an outline)

At the moment, Jenkins will be able to build a Docker image, but will fail to push it anywhere.

Set at least these two ini-settings:

- `spynl.ops.ecr.dev_tasks`
- `spynl.ops.ecr.dev_url`

You also need to set the `task` parameter to `spynl ops.start_build`. It needs to be one (or more) of `spynl.ops.ecr.dev_tasks` and a task that exists in your development container service (e.g. [in AWS](#)):

```
spynl ops.start_build --task dev
```

Now Jenkins can deploy that image to your development container registration and for that task and restart that task so it will serve your new container.

3.1.4 Configuring your Spynl plugin for Smoke Testing

The third stage in the pipeline is the smoke test. Out of the box, Spynl checks if a container actually exists at the location you want your container service to serve them. Add the following ini-setting:

```
spynl.ops.dev_url
```

Spynl also checks if this image has been built within the last 15 minutes.

Your app can specify it’s own smoke tests (TODO: write one in *my-package*)

3.1.5 Deploying to your production environment

(TODO: more verbose)

Set this ini-setting:

```
spynl.ops.ecr.prod_url:
```

Run the *dev.start_build* task

```
spynl ops.start_build --task production
```

TODO: At the moment, Splyn only pushes the new image to that registry. We could make it an ini-setting if Splyn should try to restart a task there.

In-depth documentation

Shining more light on a few important topics (work in progress):

4.1 ini-settings

4.1.1 /about/ini

4.1.2 Spynl checks required settings on startup

4.1.3 Documenting your own settings

```
from spynl.main.docs.settings import ini_doc
my_ini_doc = ...
ini_doc.extend(my_ini_doc)
```

4.2 Your production environment(s)

You'll probably have a dev environment and at least one production(-like) environment. Spynl helps you to:

- keep the code consistent between deploys to each of them
- Make sure test do not affect real users and unfinished things are turned off in production
- allow a third party to control your production pipeline

4.2.1 Docker

We use Docker to ship Spynl. You can be sure there that you look at the same code in dev and production. (*/about/versions* can help you to look up precisely which code is in there).

/about/build helps you to see when the Image was made (built on Jenkins) and when it was started.

The image is Ubuntu-based.

Custom pre-install and post-install hooks are possible in *setup.sh* (also works for *dev.install*)

prepare-docker-run.sh can influence *production.ini* or other relevant things in the Docker container right before it is run.

4.2.2 Possibilities to turn off endpoints and whole resources on production

(TODO: add issue about a more generic approach)

4.2.3 Do not send emails to real addresses from non-production environments

TODO: link to the email section

4.2.4 Using different container registries for dev and for production

Useful if you keeo them separated or a thrid party manages your production pipeline (e.g. when using a DTAP approach).

4.3 Parameter Handling

Assumption: both GET and POST work (reason for this or scrap it)

all parameters get collected in request["args"]

4.4 Serialisation

JSON is default choice

methods of selecting a different type: Content-Type Header, file extension, ...

supported types: XML, CSV, HTML, YAML

Incoming HTTP data is decoded and outgoing data encoded. Special data type (de)serialisations are easy to add, useful e.g. for DateTime objects.

4.5 Translations

Maintaining a translation (i18n) infrastructure is basically a solved problem. But setting it up is tedious and making sure all strings get translated, no matter where they live, can present pitfalls.

4.5.1 Translating a string in Spynl

4.5.2 dev.translate to simplify the workflow

4.5.3 A word about when Spynl translates

Spynl translates everything at the end (so you capture all strings, e.g. also constants)

4.6 Routing

Spynl takes control over many aspects of routing.

4.6.1 URLDispatch

(no traversal) default routes `/ {resource} / {method} / {method}`

4.6.2 Endpoint registration

`config.add_endpoint`

It's custom (do not use Pyramid's `config.add_view`) - why? (at least to have a grip on documentation of endpoints, TODO: look for other reasons)

4.6.3 Custom resource registration

`config.add_resource`

A resource class at least needs a `paths` attribute, can also have `available_in_production`.

Multiple paths (aliases) is possible in Spynl.

4.6.4 Custom routes

Basically adding meta data and a function to `spynl.resource_routes_info`. TODO: show tutorial, argue why it is better to do it this way than simple using Pyramid's `config.add_route` directly. Hint: It has to do with applying routes to resources unknown in the current plugin. It has/had a use case but maybe everyone is better off now without it. Research.

4.7 HTML emails

4.7.1 Custom base template

(show this in my-package?)

4.7.2 Write a content template

(show this in my-package?)

4.7.3 Non-production email behaviour

4.8 Custom `get_user_info`

A lot of code needs information about the current (authenticated) user, e.g. in endpoints or for logging.

4.8.1 The in-built `get_user_info` function

4.8.2 Writing your own

```
def my_user_info(request, purpose=None)
    pass # TODO

config.add_settings(user_info_function=my_user_info)
```

4.9 Validation per JSON Schema

4.10 Aggregation of Logs

4.10.1 NewRelic

spynl.newrelic.key

4.10.2 Sentry

- *spynl.sentry.project*
- *spynl.sentry.key*