
developer.skatelescope.org
Documentation
Release 0.1.0-beta

Marco Bartolini

Jan 19, 2022

1	Structure	3
2	Local Development	5
2.1	General Workflow	5
2.2	Linting and code-style	6
2.3	Building and Running	6
2.4	Adding a new cicd automation service	7
2.5	Testing With FastAPI and Pytest-BDD	8
2.6	Publishing/Releasing	9
3	SKA CI/CD Automation Services MR Checks	11
3.1	Checks	11
3.2	How to Add a New Check	12
4	SKA Slack Integration	15
4.1	Environment	15
4.2	Slack Bolt API	15
4.3	Plugin Features	16
4.4	Plugin implementation	16
5	Merge Request Checks	17

This project is the template projects that is intended for use by SKA CICD Services API Applications (*mainly bots*). The project and generated projects from this template should follow the [SKA Developer Portal](#) closely. Please check the latest guidelines, standards and practices before moving forward!

STRUCTURE

Project structure is following a basic python file structure for [FastAPI](#) as below:

```
.
├── Dockerfile
├── LICENSE
├── Makefile
├── README.md
├── app
│   └── ...
├── build
│   └── ...
├── charts
│   └── ...
├── conftest.py
├── docs
├── pyproject.toml
├── tests
│   └── ...
```

Basically, this project uses the following technologies:

- Docker: To build a docker image that exposes port *80* for API endpoints.
- Kubernetes and Helm: Project also includes a helm chart to deploy the above image in a loadbalanced kubernetes cluster.
- Python Package: Project also includes a python package so that it can be downloaded as such. (**The usability of this capability highly depends on the actual implementation!**)

LOCAL DEVELOPMENT

2.1 General Workflow

Install **Poetry** for Python package and environment management. This project is structured to use this tool for dependency management and publishing, either install it from the website or using below command:

```
curl -sSL https://raw.githubusercontent.com/python-poetry/poetry/master/get-poetry.py |  
↵python -
```

However if you want to use other tools (**Pipenv** or any other tool that can work with `requirements.file`) run the following script to generate a requirements file.

```
pip install Poetry  
make exportlock
```

This will generate `requirements.txt` and `requirements-dev.txt` files for runtime and development environment.

2.1.1 *What is Poetry and Why are we using it*

Poetry is a python dependency management and packaging tool. It is similar to **pipenv** for dependency management but it is more actively developed. It organizes dependencies in separate sections of the same file using `pyproject.toml`, described in [PEP 518](#) so it can specify publishing information and configure installed tools (like `black`, `isort`, `tox` etc.) which makes it easy to both configure and manage dependencies and publishing.

Then, you can install all the dependencies with:

```
make requirements
```

Next, you need to define the variables that are used in each plugins (see plugins respected README files for the all variables) in your `PrivateRules.mak` (needed for makefile targets) and, `.env` file (needed tp tests from vscode and interactive docker development). i.e. for GitLab MR and JIRA Support services:

```
PRIVATE_TOKEN=...  
REQUESTER=...  
GITLAB_TOKEN=...  
GITLAB_HEADER=...  
JIRA_URL=...  
JIRA_USERNAME=...  
JIRA_PASSWORD=...  
SLACK_BOT_TOKEN=...  
UNLEASH_API_URL=...
```

(continues on next page)

(continued from previous page)

```
UNLEASH_INSTANCE_ID=...
UNLEASH_ENVIRONMENT=...
RTD_TOKEN=...
GOOGLE_API_KEY=...
GOOGLE_SPREADSHEET_ID=...
NEXUS_HMAC_SIGNATURE_SECRET=...
OPENID_SECRET=...
GITLAB_CLIENT_ID=...
GITLAB_CLIENT_SECRET=...
```

Now, the project is ready for local development.

Note: depending on the IDE (*vscode is suggested*), `PYTHONPATH` may need to be adjusted so that IDE picks up imports and tests correctly. Please refrain from changing the main structure (*top level folders*) as it may break the CI/CD pipeline, make targets and the very fabric of the universe may be at stake.

2.2 Linting and code-style

The project follows PEP8 standard closely.

The linting step uses `black`, `flake8`, `isort` and `pylint` tools to check to code. It maybe useful to adjust your local environment as such as well.

Run `make lint` to check your code to see any linting errors. `make apply-formatting` could also be used to auto adjust the code style with `black` and `isort`. Note this also includes tests as well.

2.3 Building and Running

k8s note: if you are deploying or testing locally using Minikube, you should first run `eval $(minikube docker-env)` before you create your images, so that Minikube will pull the images directly from your environment.

To build the project for local development (and releasing in later stages), run `make build`. This will build a docker image (if a tag is present it will also tag it accordingly, or a *dev* tag will be used) and will build the python package as well.

To run/deploy the project, you can use Docker and Kubernetes as described below.

2.3.1 With Docker

Testing with Docker only is also possible: `make development` starts the latest built docker image (`make docker-build`) with `app` folder mounted into it and the server is set to `--reload` flag. This enables local development by reflecting any change in your `app/` folder to loaded into the api server as well.

2.3.2 With Kubernetes

An example minikube installation with loadbalancer enabled could be found [here](#) - this is the suggested starting point for testing locally with Minikube.

You want to install charts using the docker image created with `make docker-build`. If you ran `eval $(minikube docker-env)` before building, the image will be pulled from your local cache.

Next, you want to deploy your charts. `make install-chart` deploys the helm chart into a k8s environment using the [default configuration](#) with the following ingress controllers:

- NGINX
- Traefik

By default, it uses `nginx` for local development and testing. You can override this by providing `INGRESS` variable like `make install-chart INGRESS=traefik`. In deployment correct ingress is automatically selected.

Using `make template-chart` it is possible to inspect the actual deployment that will happen with `make install-chart`.


2.3.3 With VSCode

To run the app directly from VSCode for debugging purposes. Create a `launch.json` under your workspace configuration folder (`.vscode` by default):

```
{
  // Use IntelliSense to learn about possible attributes.
  // Hover to view descriptions of existing attributes.
  // For more information, visit: https://go.microsoft.com/fwlink/?linkid=830387
  "version": "0.2.0",
  "configurations": [
    {
      "name": "Python: FastAPI",
      "type": "python",
      "request": "launch",
      "module": "uvicorn",
      "args": ["app.main:app"],
      "jinja": true,
      "envFile": "${workspaceFolder}/.env"
    }
  ]
}
```

2.4 Adding a new cicd automation service

In order to add a new plugin to the main host, it is required to develop it according to the following structure:



(continues on next page)

Because FastAPI is being used, tests are done by using the FastAPI TestClient - you can read more about it [here](#).

To run all tests:

```
$ make unit_test
.
.
.
platform linux -- Python 3.6.9, pytest-6.1.2, py-1.9.0, pluggy-0.13.1
rootdir: /home/clean/ska-cicd-automation/testing, configfile: pytest.ini
plugins: bdd-4.0.1
collected 3 items

test_server.py::test_check_server PASSED
↪ [ 33%]
test_server.py::test_add_member_with_json PASSED
↪ [ 66%]
test_server.py::test_get_member_id PASSED
↪ [100%]
```

To run tests for an individual plugin, pass the PLUGIN name:

```
$ make unit_test PLUGIN=gitlab_mr
```

2.6 Publishing/Releasing

All the publishing should happen from the pipeline.

TL;DR: run `make release` to learn what you have to do!

When you are ready to publish a new version, you need to run `make update-x-release` where `x` is either `patch`, `minor` or `major`. So if you want to update the patch version, you just run `make update-patch-release`.

This will update the necessary version labels in `.release` (for docker image) , `pyproject.toml` (for python package) files and will make a commit and tag it accordingly. At this stage, you can use `make push` to manually push the docker image to your configured registry although it is not encouraged.

Finally, run `make release`. Once the CI job has completed in the [pipeline](#), make sure you trigger the manual step on the tag ref for publishing either for docker/python or deploying the helm chart.

SKA CI/CD AUTOMATION SERVICES MR CHECKS

This plugin is used to check MR quality and provide feedback on the MR window by making comments and updating them.

It uses FastAPI to create webhook that can be set to listen for the GitLab Projects. The following environment variables must be present, the token should have API access to the project:

```
PRIVATE_TOKEN=...
REQUESTER=...
JIRA_URL=...
JIRA_USERNAME=...
JIRA_PASSWORD=...
GITLAB_TOKEN=...
GITLAB_HEADER=...
UNLEASH_API_URL=...
UNLEASH_INSTANCE_ID=...
RTD_TOKEN=...
```

3.1 Checks

Each check should have:

- Feature Toggle Name: name of the check for runtime configuration
- Result Type: If the check is not successful, whether it should be marked as FAILURE, WARNING, or INFO
- Description: Brief description about what the check is about
- Mitigation Strategy: How to take corrective action to fix the broken check

Currently the plugin checks the MR are:

3.1.1 Automatic Fixing of Wrong Merge Request Settings

Marvin will attempt to automatically check the delete source branch and uncheck the squash commits on merge settings. Next to each of the other Wrong Merger Request Settings messages is a [Fix](#) link, which will trigger Marvin to attempt to fix that setting after the user is authenticated. Only users that are assigned to the merge request can trigger this automatic setting fix feature.

3.1.2 Marvin MR Approval

Marvin after creating the table will verify if there is any checks under the failure category failed, if so Marvin does not approve the MR, and in the case that that MR was already approved before by him he unapproves it. If none of the checks under the failure category failed Marvin will approve the MR.

3.1.3 Runtime Configuration

This service is using feature toggles to determine which checks to enable/disable at the runtime. It uses [Unleash](#) integration provided by GitLab to achieve this.

For the project level configuration, a project could be disabled using `Project Tags/Topics`. The service uses a blacklist to determine whether it should run the checks as well.

Precedence of configuration

- Project Level Configuration with Tags/Topics
- Feature Toggle Strategies

3.2 How to Add a New Check

Each new check must use the abstract base class, `Check`, to ensure to define its type, description, mitigation strategy and check action, which performs the actual checking on the MR and returns a boolean indicating the result.

Base Class:

```
class Check(ABC):

    feature_toggle: str = NotImplemented

    @abstractmethod
    async def check(self, mr_event, proj_id, mr_id) -> bool:
        pass

    @abstractmethod
    async def type(self) -> MessageType:
        pass

    @abstractmethod
    async def description(self) -> str:
        pass

    @abstractmethod
    async def mitigation_strategy(self) -> str:
        pass
```

Example Check:


```
class CheckAssigneesComment(Check):
    feature_toggle = "check-assignees-comment"

    def __init__(self, api: GitLabApi, logger_name):
        self.api = api
        self.logger = logging.getLogger(logger_name)

    async def check(self, mr_event, proj_id, mr_id):
        mr = await self.api.get_merge_request_info(proj_id, mr_id)
        self.logger.debug("Retrieved MR: %s", mr)
        return len(mr["assignees"]) > 0

    async def type(self) -> MessageType:
        return MessageType.FAILURE

    async def description(self) -> str:
        return "Missing Assignee"

    async def mitigation_strategy(self) -> str:
        return "Please assign at least one person for the MR"
```

Then the necessary tests for the added checks should be added in `tests` folder. These tests should get picked up by the main frameworks testing.

Finally, each check should be initialised and called in the `mrevents` file to be included into the list of checks that are performed for the MRs.

SKA SLACK INTEGRATION

4.1 Environment

To develop a Slack app, it is recommended to create your own Slack workspace and test against it.

4.1.1 Marvin Use Case

If the Slack app you are working on is Marvin itself, then you can use a pre-made Slack app named *Marvin-Test* to conduct your code changes, without the need for setting up your own Slack workspace.

In that case, you should use the following `.env` variables from the *Marvin-Test* app.

```
SLACK_BOT_TOKEN=...  
SLACK_SIGNING_SECRET=...
```

Following the steps in this repository README, start your local testing with `make development` and expose it outside your local development machine. You can do so using `ngrok` for instance: `ngrok http 3000`

Change the various *Marvin-Test* endpoints to point at your local deployment (i.e.: `https://da0dd48a7db7.ngrok.io/jira/support/slack/events`)

You should now have the *Marvin-Test* Slack app connected to your local development.

4.2 Slack Bolt API

This plugin was developed using a FastAPI implementation of the Slack Bolt framework. The documentation is available on a [thorough tutorial](#).

Note that our plugin uses asynchronous methods, which are available as part of the Bolt SDK.

Example apps are available on the [Bolt Github repository](#).

4.3 Plugin Features

This plugin uses two integration points for Slack: the Slack Events API, and the Message Shortcuts integration.

4.3.1 Slack Events

The Bolt API listens on the `/slack/events/` endpoint, and redirects all traffic through this endpoint. On Slack, we can however configure certain Slack Events to trigger specific functionality. For more information, visit the [Slack documentation on Events](#).

For this example app we listen for the `@mention` of the bot and respond with a simple message.

4.3.2 Message Shortcuts

Slack provides two different shortcuts: a Global shortcut menu, and a Message shortcut menu. For the Jira Support Issue plugin, the idea was to use the message contents to generate some of the contents of the Jira issue automatically, so that the user can have a pre-populated form and quickly submit the data to Jira. It therefore made more sense to use the Message Shortcut.

A modal is opened using the list of users that is used for populating the drop-down menu, for authorization checking first (the user opening the modal should be on the list of users on the sheet). This list of users is created using a Google Sheet as [external data source](#) (see below section). The list of projects in which the is also populated from this spreadsheet.

4.3.3 Google Sheets

The Google Sheets API provides a rudimentary data source for management of the list of users and projects that can be assigned and populated with Jira tickets, respectively. The API Key and Sheet ID are both stored as environment variables and called in the [handler](#).

For the *Marvin* Slack app you can request permission to access the spreadsheet [here](#).

4.4 Plugin implementation

The plugin basically follows the architecture of the SKA CICD framework, but since the Bolt API is also a framework in itself, it is easy to get confused. The plugin has the normal `routers` and `models` directories. All requests are handled by the endpoints declared under `routers/jira_support_ticket`. The asynchronous call to the `SlackAppHandler` is awaited, and the internal authentication with Slack is handled by the Bolt API.

Important to note is Slack's requirement to get a response from the Web service within three seconds. This is accomplished by the `ack()` calls found in all the methods in `/models/slack_handler.py`.

The APIs created to integration with Jira etc are imported and used as with all other plugins.

MERGE REQUEST CHECKS

These are all the packages, functions and scripts that form part of the project.

- *SKA CI/CD Automation Services MR Checks*
- *SKA Slack Integration*