

How to improve the developer experience in Heptapod/GitLab

Automate the dull works to focus on development



Noé Gaumont : ngaumont@logilab.fr

Logilab

Summary

- Logilab presentation and our code management
- What problems are we facing within Logilab ?
- Existing solutions
- Tools designed and used internally:
 - `assign-bot` to have review
 - `gitlab-ci-template` to have fine-tuned job definition
 - `release-new` to have proper release (to pypi)
 - `cube-doctor` to update dependencies

- small IT service and consulting company (~25 persons) focused on:
 - Web Semantic,
 - Scientific computing,
 - Python training.
- Work almost only with open source softwares
- Contribute either with code or money.







Purpose:

- a semantic web application framework,
- developed in Python (LGPL),
- explicit data model,
- reusable components (called cubes),
- a front agnostic way to serve your data.

Our go-to tech stack is composed of :



Code organisation within Logilab

- **mercurial**, a vcs written in python
- **Heptapod**, a friendly fork of GitLab with mercurial support
- **multi-repo** approach:
 - CubicWeb
 - public cubes (~**200 repo**)
 - private cubes (mainly client related)
 - opensource project (~ 40 repos)
 - internal projects

Problems caused by the multi-repo approach

Keeping one repo tidy is ~~simple~~, keeping several hundred repos is ~~hard~~.

Both are hard, problems occurs only **faster** within multi-repo architecture :

Problems caused by the multi-repo approach

Keeping one repo tidy is ~~simple~~, keeping several hundred repos is ~~hard~~.

Both are hard, problems occurs only **faster** within multi-repo architecture :

- Make sure all the test/lint are **green**
- Assign a reviewer for each Merge request
- Coherent code everywhere
- Maintain good practices in CI configuration
- Properly release everything to Pypi, npm, ...
- Up-to-date docker images with correct tag

Existing solutions

Within Github, there is the probot solution such as:

- Auto-Assign
- Dependabot
- Release-drafter
- PR-triage

Existing solutions

Within GitLab, custom CI or **danger-bot** which is good for :

- Coding style
- Database review
- Documentation review
- Merge request metrics
- Reviewer roulette
- Single codebase effort

A tale of a CubicWeb MergeRequest

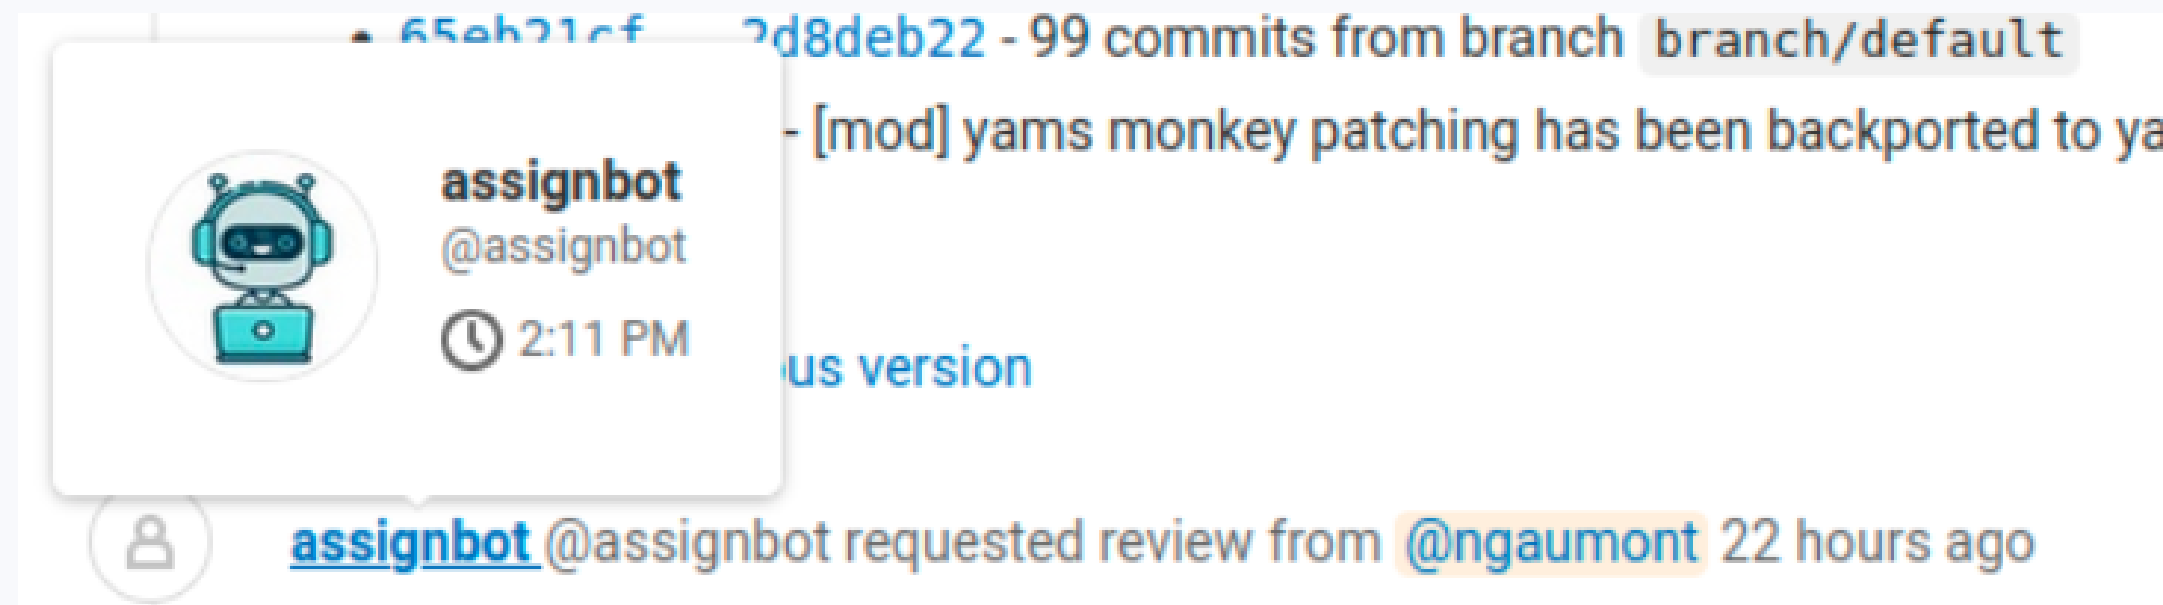
Classic use case: A new feature introduced in CubicWeb induced a deprecation warning.

Goal: The MR is reviewed, the tests pass and the repo depending on CubicWeb are updated, all version are published.

MergeRequest creation

A scheduled job of **assign-bot** will assign a reviewer base on user preferences:

```
naumont:  
  max_auto_review_per_week: 15  
  max_auto_review_per_day: 3  
username_2:  
  max_auto_review_per_week: ZZZ  
  max_auto_review_per_day: WWW
```



It will also comment on the MR inactive for one week.

Running the correct set of tests

A coherent set of up-to-date jobs are run, thanks to a shared definition, [gitlab-ci-template](#):

It takes care of:

- reducing the CI load by using custom image
- defining rules to prevent CI on some branches
- define some stages

```
1      include:
2        - project: "open-source/gitlab-ci-templates"
3          ref: "branch/default"
4          file:
5            - "templates/no-duplicated-ci-pipelines.yml" # use workflow to avoid duplicated pipelines
6            - "templates/lint/black.yml" # will do the equivalent of 'tox -e black'
7            - "templates/tests/py3.yml" # will do the equivalent of 'tox -e py3'
8            - "templates/create-release-on-heptapod.yml" # this will create a release on heptapod
9            - "templates/build-docker-image.yml"
10         # uncomment and uses to customize/extend the configuration here if needed
11         # (it needs to be at the same level than "- project")
12         - ".gitlab-ci-extended.yml"
```

It is always possible to customise the jobs within `.gitlab-ci-extended.yml`

Running the correct set of tests

A coherent set of up-to-date jobs are run, thanks to a shared definition, [gitlab-ci-template](#):

It takes care of:

- reducing the CI load by using custom image
- defining rules to prevent CI on some branches
- define some stages

```
1      include:
2        - project: "open-source/gitlab-ci-templates"
3          ref: "branch/default"
4          file:
5            - "templates/no-duplicated-ci-pipelines.yml" # use workflow to avoid duplicated pipelines
6            - "templates/lint/black.yml" # will do the equivalent of 'tox -e black'
7            - "templates/tests/py3.yml" # will do the equivalent of 'tox -e py3'
8            - "templates/create-release-on-heptapod.yml" # this will create a release on heptapod
9            - "templates/build-docker-image.yml"
10         # uncomment and uses to customize/extend the configuration here if needed
11         # (it needs to be at the same level than "- project")
12         - ".gitlab-ci-extended.yml"
```

It is always possible to customise the jobs within `.gitlab-ci-extended.yml`

Example of a python test jobs

```
py3:
  interruptible: true
  image: ${CI_REGISTRY}/cubicweb/dockerfiles/buster-slim-pg11
  stage: tests
  script: tox -e py3 -- --junitxml=report.xml
  tags:
    - py3
  needs: []
  artifacts:
    when: always
    reports:
      junit: report.xml
  paths:
    - py3-deprecated-warnings.json
```

It uses:

- a custom docker image with all the dependencies already installed
- store some artifacts generated during the tests
- make the jobs interruptible

Pushing the image to the registry

Once the MR is accepted, another job defined in [gitlab-ci-template](#) is used to build and publish a Docker image:

```
1  image_build_latest:
2    [...]
3    image:
4      name: gcr.io/kaniko-project/executor:debug
5      entrypoint: [""]
6    script:
7      - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
8      - /kaniko/executor
9      --context $CI_PROJECT_DIR
10     --dockerfile $CI_PROJECT_DIR/Dockerfile
11     --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_HG_SHORT_SHA
12     --destination $CI_REGISTRY_IMAGE:latest
13    rules:
14      - if: '$CI_COMMIT_REF_NAME == "branch/default"'
15
16  image_build_tag:
17    [...]
18    image:
19      name: gcr.io/kaniko-project/executor:debug
20      entrypoint: [""]
21    script:
22      - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
```

The image tag is **correctly tag** inside the job.

Pushing the image to the registry

Once the MR is accepted, another job defined in `gitlab-ci-template` is used to build and publish a Docker image:

```
1  image_build_latest:
2    [...]
3    image:
4      name: gcr.io/kaniko-project/executor:debug
5      entrypoint: [""]
6    script:
7      - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
8      - /kaniko/executor
9      --context $CI_PROJECT_DIR
10     --dockerfile $CI_PROJECT_DIR/Dockerfile
11     --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_HG_SHORT_SHA
12     --destination $CI_REGISTRY_IMAGE:latest
13    rules:
14      - if: '$CI_COMMIT_REF_NAME == "branch/default"'
15
16  image_build_tag:
17    [...]
18    image:
19      name: gcr.io/kaniko-project/executor:debug
20      entrypoint: [""]
21    script:
22      - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
```

The image tag is **correctly tag** inside the job.

Pushing the image to the registry

Once the MR is accepted, another job defined in `gitlab-ci-template` is used to build and publish a Docker image:

```
3      image:
4        name: gcr.io/kaniko-project/executor:debug
5        entrypoint: [""]
6      script:
7        - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
8        - /kaniko/executor
9          --context $CI_PROJECT_DIR
10         --dockerfile $CI_PROJECT_DIR/Dockerfile
11         --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_HG_SHORT_SHA
12         --destination $CI_REGISTRY_IMAGE:latest
13      rules:
14        - if: '$CI_COMMIT_REF_NAME == "branch/default"'
15
16      image_build_tag:
17        [...]
18      image:
19        name: gcr.io/kaniko-project/executor:debug
20        entrypoint: [""]
21      script:
22        - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS
23        - /kaniko/executor
24        --context $CI_PROJECT_DIR
```

The image tag is **correctly tag** inside the job.

Pushing the image to the registry

Once the MR is accepted, another job defined in [gitlab-ci-template](#) is used to build and publish a Docker image:

```
11      --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_HG_SHORT_SHA
12      --destination $CI_REGISTRY_IMAGE:latest
13      rules:
14        - if: '$CI_COMMIT_REF_NAME == "branch/default"'
15
16      image_build_tag:
17        [...]
18        image:
19          name: gcr.io/kaniko-project/executor:debug
20          entrypoint: [""]
21        script:
22          - echo "{\"auths\":{\"$CI_REGISTRY\":{\"username\":\"$CI_REGISTRY_USER\",\"password\":\"$CI_REGISTRY_PASS\"}}}"
23          - /kaniko/executor
24          --context $CI_PROJECT_DIR
25          --dockerfile $CI_PROJECT_DIR/Dockerfile
26          --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_HG_SHORT_SHA
27          --destination $CI_REGISTRY_IMAGE:$CI_COMMIT_TAG
28      rules:
29        - if: "$CI_MERGE_REQUEST_ID"
30          when: never
31
32        - if: "$CI_COMMIT_TAG"
33          when: on_success
```

The image tag is **correctly tag** inside the job.

A new version is released

The MR is accepted and a new version has to be released. `release-new` is used for that :

```
$ release-new
Automatic release guesser decided to release the version '3.35.0' (minor)
Are you ok with that? [Y/n]: Y
🎉 Congratulation, we have made a new minor release 3.35.0 \o/ 🌟

✨ 🍰 ✨

Now you need to hg push the new commits
```

Under the hood, it :

- detects the version number using semantic versioning
- updates the version number in the source
- generates a changelog from commit message and let you modify it
- update packaging MANIFEST.in if needed
- create a commit and then tag it

A new version is released

Once pushed, a job on the tag releases to pypi, with [gitlab-ci-template](#).

```
pypi-publish:
  [...]
  - if: '$CI_MERGE_REQUEST_ID'
    when: never
  - if: '$CI_COMMIT_TAG'
    when: on_success
script:
  - tox -e pypi-publish
```

Repositories are updated

cube-doctor will update repositories:

- Rebase the MR in same project
- In other project, create MR updating the package version
- Create MR for lint configurations

Manually, it can run automatic refactoring (handle deprecation warning).

The screenshot shows a GitLab merge request (MR) titled "feat(pkginfo): upgrade cubicweb[pyramid] to >=3.33.0,<3.34.0" created 5 months ago by "cube-doctor". The MR has 14 unresolved threads and 1 change. The "Changes" tab is selected, showing a diff for the file "cubicweb_logilabfr/__pkginfo__.py". The diff highlights a change in the "cubicweb" dependency from ">=3.31.0,<3.32.0" to ">=3.33.0,<3.34.0". The interface includes navigation links for Overview (14), Commits (1), Pipelines (6), and Changes (1). It also features a "Compare" section with dropdowns for "branch/default" and "latest version". At the bottom, there are links to "Show all unchanged lines" and "Show 20 lines".

Open Created 5 months ago by cube-doctor Options ▾

feat(pkginfo): upgrade cubicweb[pyramid] to >=3.33.0,<3.34.0

14 unresolved threads 🔍 🗨 ⤴

Overview 14 Commits 1 Pipelines 6 **Changes 1**

🔍 Compare branch/default ▾ and latest version ▾

▾ cubicweb_logilabfr/__pkginfo__.py 🔍 +1 -1 ☐ Viewed ⋮

↑ Show all unchanged lines ↑ Show 20 lines

29	29]
30	30	
31	31	__depends__ = {
32	-	"cubicweb": ">=3.31.0,<3.32.0",
	32	+ "cubicweb": ">=3.33.0,<3.34.0",
33	33	"cubicweb-file": ">=3.0.0,<3.1.0",
34	34	"cubicweb-blog": ">=1.14.0,<1.15.0",
35	35	"cubicweb-link": ">=1.9.0,<1.10.0",

↓ Show 20 lines ↑ Show all unchanged lines

Conclusion

- GitLab/Heptapod are a great tools battery-included:
 - scheduled job
 - registry
 - shared job definition.
- A Good CI jobs helps **a lot** to keep repositories clean and up-to-date.
- However, manual effort are still needed. In our case:
 - forge.extranet.logilab.fr/open-source/assign-bot
 - forge.extranet.logilab.fr/open-source/gitlab-ci-template
 - forge.extranet.logilab.fr/open-source/release-new
 - forge.extranet.logilab.fr/cubicweb/cube-doctor



Noé Gaumont : ngaumont@logilab.fr

Logilab