



The Linux & Open Source Company

Fun with artifacts and child-pipelines in GitLab

Include or trigger?

Jan Bundesmann

THE Linux & Open Source Company!



Consulting



Engineering



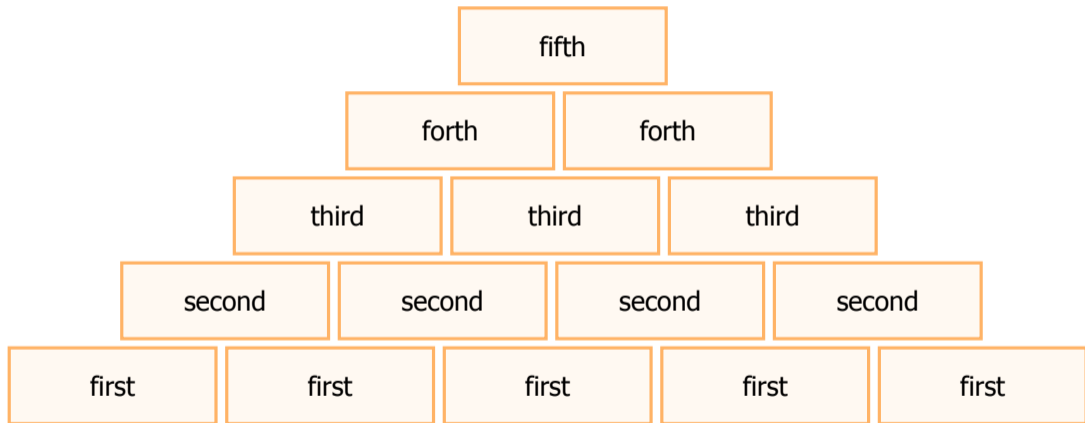
Support



Training



Let's build



Let's build II



Bricks and mortar

Why?

- ▶ Reduce **complexity**
- ▶ Keep system **flexible** and **modular**

How?

- ▶ Write one "glue" pipeline
- ▶ Pass information: **parameters** and **artifacts**
- ▶ Final artifact shall combine all intermediate steps

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Include the first repo

.gitlab-ci.yml:

```
1 stages:
2   - first
3
4 include:
5   - local: variables.yaml
6
7 doit:
8   stage: first
9   script:
10    - 'echo "first: ${information}" >> data.html'
11
12   artifacts:
13     paths:
14       - data.html
```

variables.yaml:

```
1 variables:
2   information: This is the first repo.
```

.gitlab-ci.yml:

```
1 include:
2   - project: gent2024/first
3     file: .gitlab-ci.yml
```

variables.yaml:

```
1 variables:
2   information: This is glue1-include.
```

Quote from docs

From <https://docs.gitlab.com/ee/ci/yaml/#include>

The include files are:

- ▶ Merged with those in the `.gitlab-ci.yml` file.
- ▶ Always evaluated first and then merged with the content of the `.gitlab-ci.yml` file, regardless of the position of the include keyword.

From <https://docs.gitlab.com/ee/ci/yaml/#includeproject>

All nested includes are executed in the scope of the project containing the configuration file with the nested include keyword. You can use local (relative to the project containing the configuration file with the include keyword), project, remote, or template includes.

Take home

```
1 include:
2   - project: gent2024/first
3     file: .gitlab-ci.yml
4   - local: variables.yaml
5
6 variables:
7   information:
8     value: glue1 (.gitlab-ci.yml)
```

- ▶ Included files are merged in the order they are included.
- ▶ Nested includes are evaluated relative to the included file.
- ▶ Properties are unchanged unless explicitly overwritten (e.g. stage)
- ▶ Project `first` is unchanged.

Agenda

- 1 Includes
- 2 Trigger**
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Trigger pipeline from the first repo

.gitlab-ci.yml:

```
1 stages:
2   - first
3
4 include:
5   - local: variables.yaml
6
7 doit:
8   stage: first
9   script:
10    - 'echo "first: ${information}" >> data.
11      html'
12   artifacts:
13     paths:
14       - data.html
```

variables.yaml:

```
1 variables:
2   information: This is the first repo.
```

.gitlab-ci.yml:

```
1 stages:
2   - first
3   - second
4   - third
5
6 first:
7   stage: first
8   trigger:
9     project: gent2024/first
10
11 second:
12   stage: second
13   trigger:
14     include:
15       - project: gent2024/first
16         file: .gitlab-ci.yml
```

Take home

- ▶ Trigger create downstream pipelines
- ▶ Downstream pipelines have their own stages
- ▶ Environments are created parallel to downstream pipeline
- ▶ Artifacts belong to downstream pipeline
- ▶ Variables are not included from downstream pipeline
- ▶ Parameters can be passed to downstream pipelines → also via `parallels:matrix`
- ▶ Two kinds of downstream pipelines:
 - ▶ multi-project ("trigger-project")
 - ▶ child ("trigger-include")

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties**
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Comparison

Triggers:

- ▶ Decoupled from parent pipelines
 - ▶ Pass variables downstream
 - ▶ Child pipelines can use parallel
- Scaling ✓
- artifact handling?

Includes:

- ▶ Merging overwrites prior includes
 - ▶ Setting variables possible but tedious
 - ▶ No steps required for artifacts
- Scaling ?
- artifact handling ✓

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Specs to include pipeline multiple times

in first:

```
1 spec:
2   inputs:
3     name:
4       default: 'first'
5   ---
6   stages:
7     - first
8     - second
9     - third
10
11  include:
12    - local: variables.yaml
13
14  doit${[[ inputs.name ]]}:
15    stage: ${[[ inputs.name ]] }
16    script:
17      - 'echo "${CI_PROJECT_NAME}, stage ${[[
18          inputs.name ]]}: ${information}" >>
19          data.html'
20
21  artifacts:
22    paths:
23      - data.html
```

in glue1-include:

```
1 ---
2 include:
3   - project: gent2024/first
4     file: .gitlab-ci.yml
5     ref: specs
6     inputs:
7       name: first
8   - project: gent2024/first
9     file: .gitlab-ci.yml
10    ref: specs
11    inputs:
12      name: second
13   - project: gent2024/first
14     file: .gitlab-ci.yml
15     ref: specs
16     inputs:
17       name: third
18   - local: variables.yaml
19
20 variables:
21   information:
22     value: glue1 (.gitlab-ci.yml)
```

<https://docs.gitlab.com/ee/ci/yaml/#spec>

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Second brick

```
1 spec:
2   inputs:
3     job-stage:
4       default: 'first'
5   ---
6   stages:
7     - first
8     - second
9     - third
10
11  include:
12    - local: variables.yaml
13
14  doit${[[ inputs.job-stage ]]}:
15    stage: $[[ inputs.job-stage ]]
16    script:
17      - ./doit.sh
18    artifacts:
19      paths:
20        - data.html
```

- ▶ Logic externally (with respect to `.gitlab-ci.yml`)
 - ▶ Ansible roles
 - ▶ Terraform modules
 - ▶ Shell scripts
- ▶ Including: ✗
 - Only CI can be included
- ▶ Trigger: ✗
 - pass and retrieve artifact?

Trigger artifact

.gitlab-ci.yml

```
1 initialize:
2   stage: initialize
3   script:
4     - git https://${CI_SERVER_HOST}/gent2024/
      second.git
5   artifacts:
6     paths:
7       - second
8
9   doit:
10    stage: doit
11    variables:
12      PARENT_PIPELINE_ID: ${CI_PIPELINE_ID}
13    trigger:
14      include:
15        - local: child-pipeline.yaml
16      strategy: depend
17      forward:
18        pipeline_variables: true
19    parallel:
20      matrix:
21        - information:
22          - A
23          - B
```

child-pipeline.yaml

```
1 include:
2   - artifact: second/.gitlab-ci.yml
3     pipeline: ${PARENT_PIPELINE_ID}
4     job: initialize
5
6   doitfirst:
7     needs:
8       - pipeline: ${PARENT_PIPELINE_ID}
9         job: initialize
10    before_script:
11      - cd second
12    after_script:
13      - cp second/data.html .
```

Trigger artifact – why so complicated?

- ▶ child pipeline
 - ▶ artifact not yet there
- ▶ needs
 - ▶ requirement to obtain artifacts
- ▶ before and after script
 - ▶ directory shift
 - ▶ actions happen within `second`
 - ▶ artifacts are stored relative to `second`

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding

Upwards and downwards

- ▶ `needs` is your friend
- ▶ Passing artifacts “downwards” is easy
- ▶ Passing “upwards” from included pipelines is straight forward, as well
- ▶ Open question: How to obtain artifacts from downstream pipelines (triggered pipelines)?

Brickable: retrieve artifacts

```
1 retrieve_artifacts:
2   stage: doit
3   needs: ['doit']
4   script:
5     - >
6       PIPELINE_IDS=$(curl -s --header "Authorization: Bearer ${PROJECT_TOKEN}"
7         "${CI_SERVER_URL}/api/v4/projects/${CI_PROJECT_ID}/pipelines/${CI_PIPELINE_ID}/bridges"
8         | jq -r ".[] | select(.name|test(\"doit\")) | .downstream_pipeline.id")
9     - |
10      OLDIFS=$IFS; IFS=$(echo -en "\n\b")
11      for PREPARE_PIPELINE_ID in ${PIPELINE_IDS}
12      do
13        OUTPUT_JOB_ID=$(curl -s --header "Authorization: Bearer ${PROJECT_TOKEN}" \
14          "${CI_SERVER_URL}/api/v4/projects/${CI_PROJECT_ID}/pipelines/${PREPARE_PIPELINE_ID}/jobs" \
15          | jq ".[] | select(.name == \"doitfirst\") | .id")
16
17        curl -s --header "Authorization: Bearer ${PROJECT_TOKEN}" --output artifacts.zip \
18          "${CI_SERVER_URL}/api/v4/projects/${CI_PROJECT_ID}/jobs/${OUTPUT_JOB_ID}/artifacts" \
19
20        unzip artifacts.zip
21
22        <DO SOMETHING TO ARTIFACTS>
23
24        rm -rf artifacts.zip <ARTIFACTS>
25      done
26      IFS=$OLDIFS
27 artifacts:
28   <...>
```

Agenda

- 1 Includes
- 2 Trigger
- 3 Difficulties
 - Scaling Includes
 - more than just CI
 - Artifacts
- 4 Deciding**

For all bricks

independent of method (includ, trigger, trigger artifact)

Conventions

- ▶ `no before_script`
- ▶ `no after_script`
- ▶ well defined parametrization
 - ▶ `inventory.yaml`
 - ▶ `ssh_config`
- ▶ definition of child pipeline in `second`

Matrix

method	readability	easy to use	scales	artifacts	more than just CI
include	✓	✓	✗✗	local	✗
trigger	✗✗	✓	✓	remote	✓
trigger artifact	✗	✗	✓	local	✓✓

Summary

