

CERN ABT's lazy git workflow for equipment-specific designs

Pieter Van Trappen, Léa Strobino

2nd FPGA Developers' Forum – <https://indico.cern.ch/event/1467417>

Presentation content

1. Context
2. Project structure
3. Examples and use cases
4. Future

Motivation: present the current, sub-optimal but relatively lazy efficient, git workflow hoping for ideas for improvement. Or maybe inspire others. All sources available to CERN colleagues, open to share publicly.

Context – CERN SY-ABT

Accelerator Systems, Beam Transfer Group (SY-ABT)

- The ABT group is in charge of the design, development, construction, installation, exploitation and maintenance of injection and extraction related equipment and beam-transfer systems.

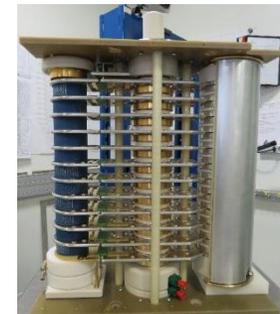
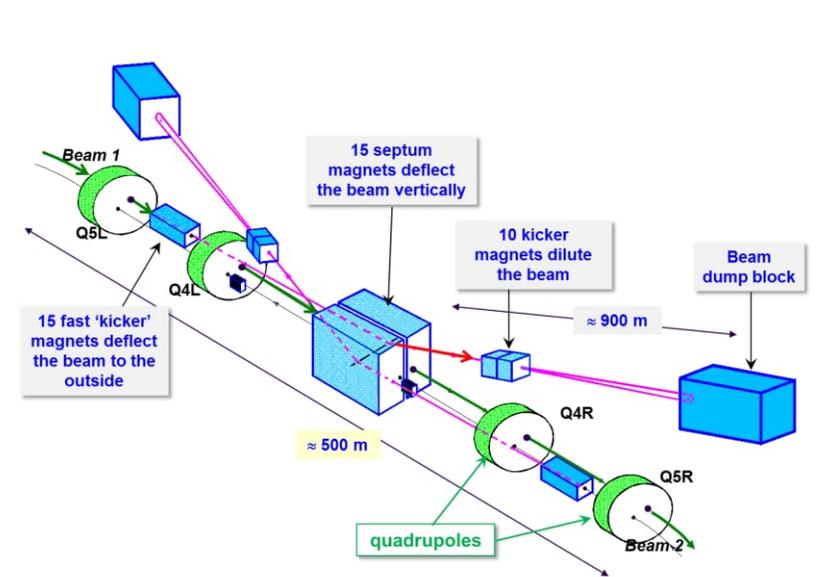
The electronics section designs, among others, cards and gateware for such ‘fast pulsed transfer systems’.

Wide variety of hardware of the ‘70s being consolidated and newer hardware; due to protection latency and timing synchronisation FPGAs in use since 20+ years.

ABT handles about 50 installations (aka equipment) and strives for reusable, modular designs for obvious reasons (limited time, knowledge retention, reduce support vector).

Main focus here will be on two projects, though there’s more with similar structure

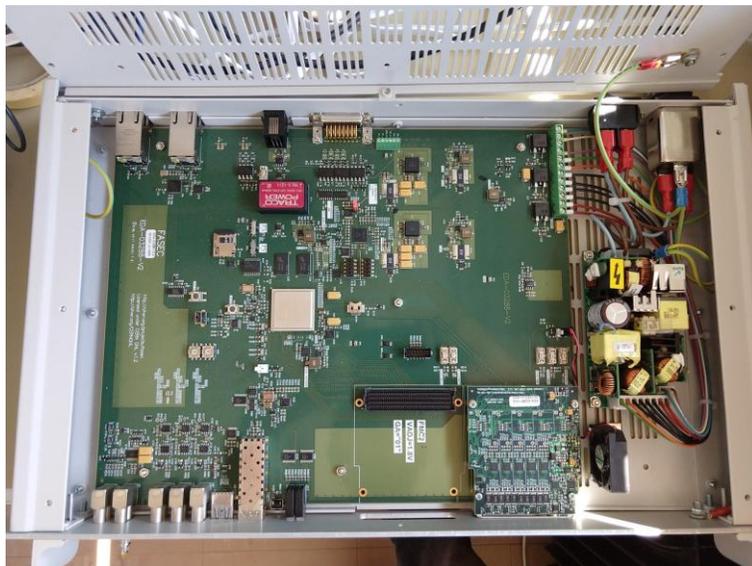
- A Fast Interlock Detection System for high-power switch protection (FIDS) - [paper](#)
- A Generic Timing Software/Gateware for fast pulsed magnet systems (KiTS/R) - [paper](#)



Context – Hardware

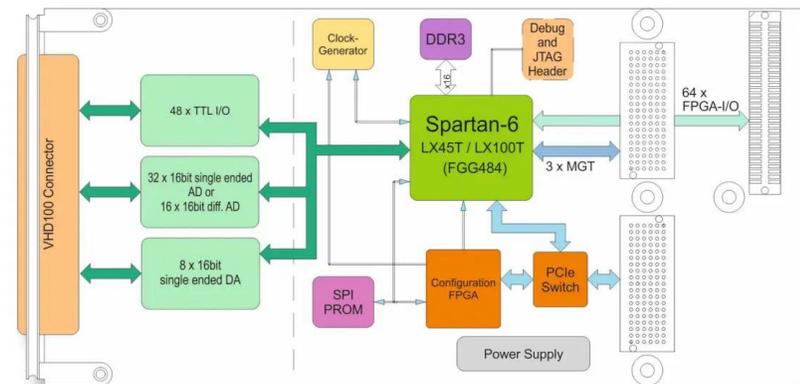
FIDS

- OHWR FASEC (FPGA & ARM SoC FMC Carrier) platform
- AMD Zynq-7000 SoC



KiTR

- Uses TEWS' TXMC635 'mixed GPIO' XMC, hosted on VME or PCIe carriers
- Reconfigurable FPGA with Analog Input (16 bit), Analog Output (16 bit), and TTL I/O
- AMD Spartan-6 FPGA



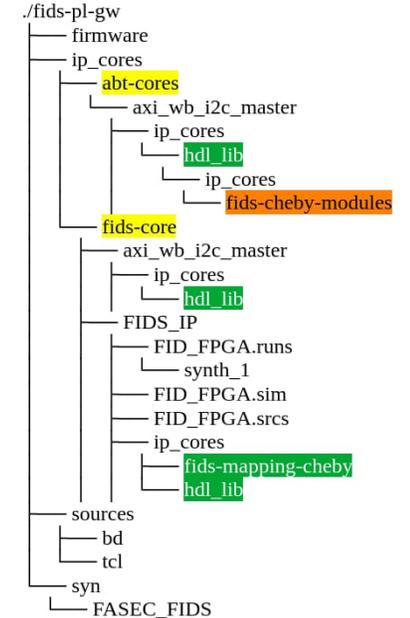
Project structure

- **Our typical gateway project structure is composed as follows**

- one or multiple **common** and **dedicated** IP/HDL cores
- a top-level design, specific for each installation
 - specific: HW IO interconnections, cores instantiation, (cheby) memmap, AXI4 interconnects, etc.
 - common: HW IO interface, abt_card_header, submodules, etc.

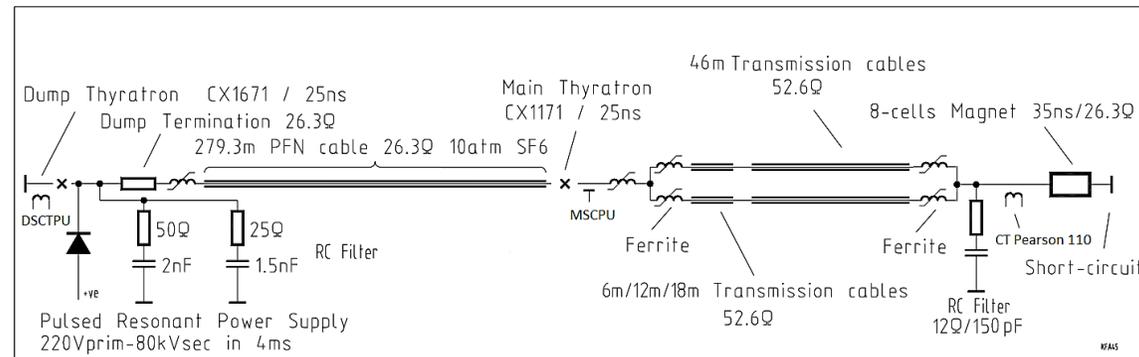
- **Some projects predate wider ABT/CERN git adaptation**

- No use of git or any other version control system; common files were copied around (though often modified).
- Some had an IP-based structure, but such ‘common cores’ were managed at the level of the designer’s HDD with a common, single folder for multiple projects.
 - No use of git submodules thus once such a dependency had advanced for one project, all remaining projects required implementing similar API changes (e.g. modules’ port); zero reproducibility.
- Recently we’ve seen substantial efforts to migrate to the common ‘modern’ approach i.e. **git submodules with a GitLab CI/CD pipeline** to obtain (~reproducible) bitstreams.



Project structure

- **ABT gateway projects are very similar but still different!**
- **The project's top-level design contains both common and project-specific parts (not including the submodules), such as**
 - HW IO interconnections
 - cores instantiation, through (cheby) memmap
 - AXI4 interconnects
 - project peculiarities such as additional gating, muxing, etc.
- **As such, for each project/installation there's a dedicated git branch.**



Project structure

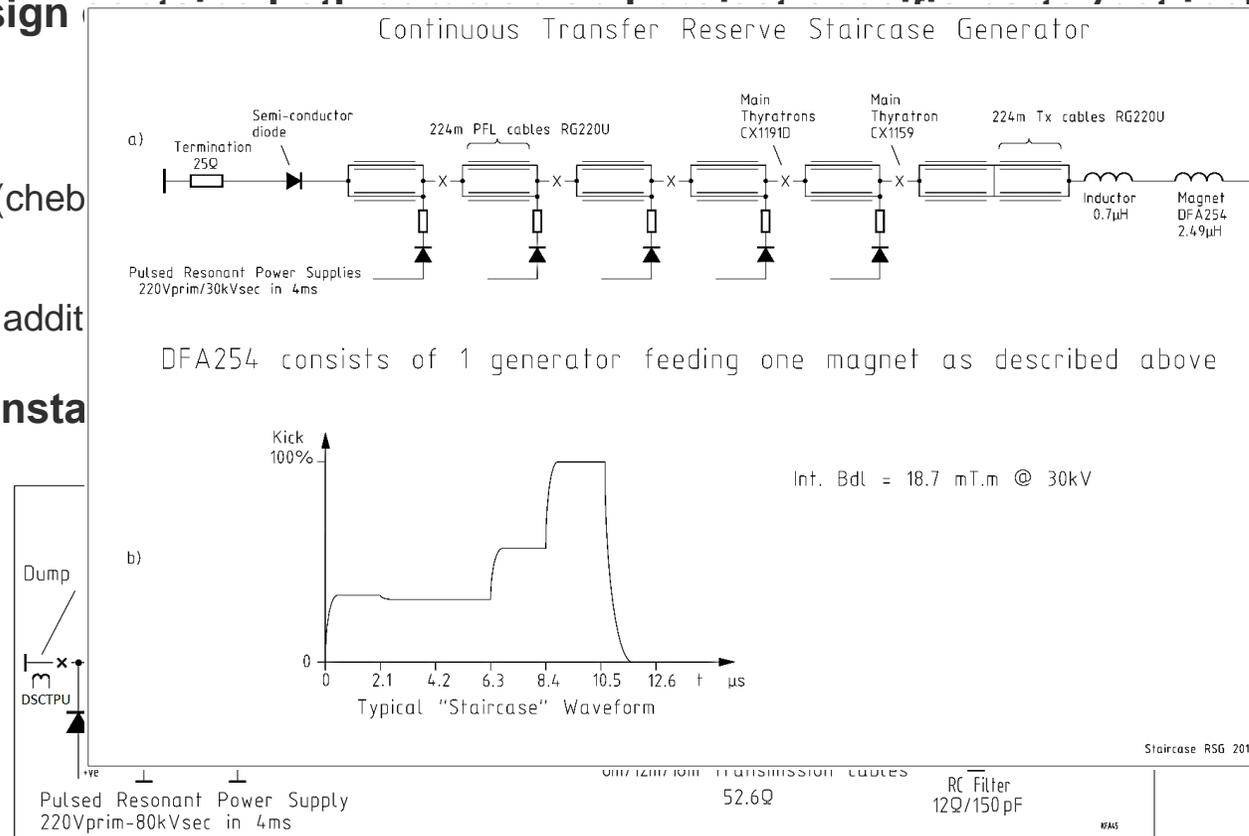
- ABT gateway projects are very similar but still different!

- The project's top-level design such as

- HW IO interconnections
- cores instantiation, through (cheb
- AXI4 interconnects
- project peculiarities such as addit

- As such, for each project/insta

including the submodules),



Project structure

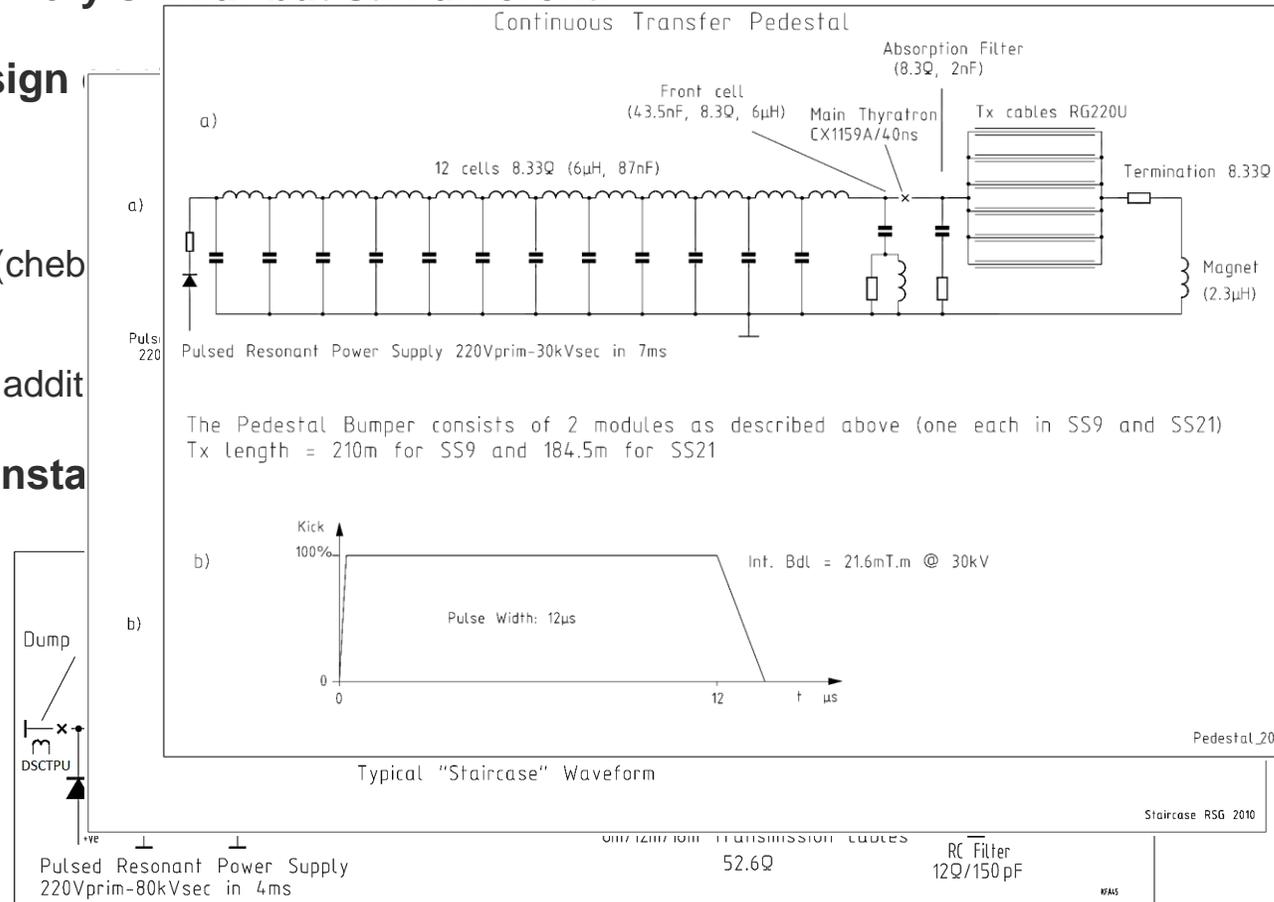
- **ABT gateway projects are very similar but still different!**

- **The project's top-level design such as**

- HW IO interconnections
- cores instantiation, through (cheb
- AXI4 interconnects
- project peculiarities such as addit

- **As such, for each project/insta**

ling the submodules),



Project structure – *branch-per-installation*

Pros

- Apply changes to the common part relatively fast using `git cherry-pick`
- No need to abstract away all peculiarities into a generic module (in a submodule)
- Common GitLab CI/CD repository pipeline configuration

Cons

- Cherry-pick not as fast as rebase and still manual, time consuming e.g. 10+ branches for FIDS
- Not all commits apply cleanly, need to pay attention to git commit best practices
- Difficult to use git tags
 - Tags are unique per repo, thus tags à la `v1.2.3` for such a *branch-per-installation* IP repo cannot be used; workaround being branch-name prefix à la `foo_v1.2.3`.
 - GitLab's **multi-project pipeline** does not work reliably with **tag pipelines** (`if: $CI_COMMIT_TAG`) which one commonly would use to trigger the downstream build → downstream will be triggered but fetch the latest BRANCH artifacts, not necessarily the TAG ones.

Use case – TXMC635 semi-tag pipeline

- Project where gateway commits/tags trigger downstream softcore firmware compilation, to be included in bitstream
- GitLab's multi-project pipeline incompatible with tag pipelines → downstream will be triggered but fetch the latest BRANCH artifacts, not necessarily the latest TAG ones
- Current workaround for such multi-project pipelines
 - Disable tag pipelines
 - Custom job rules
 - Push atomically
 - `git push --atomic origin <branch> <tag>`
- To-do: make a GitLab feature request for tag support for `needs: ref:`, similarly to `$CI_MERGE_REQUEST_REF_PATH` ([docs](#))

Use case – TXMC635 semi-tag pipeline

- Project where gateware commits/tags trigger downstream softcore firmware compilation, to be included in bitstream
- GitLab's multi-project pipeline incompatible with tag pipelines → downstream will be triggered but fetch the latest BRANCH artifacts, not necessarily the latest TAG ones
- Current workaround for such multi-project pipelines
 - Disable tag pipelines
 - Custom job rules
 - Push atomically
 - `git push --atomic origin <branch> <tag>`
- To-do: make a GitLab feature request for tag support for `needs: ref:`, similarly to `$CI_MERGE_REQUEST_REF_PATH` ([docs](#))

```
build:
  rules:
    # Disable tag pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
    - if: $CI_COMMIT_TAG
      when: never
    - when: on_success
  extends:
    - .ise_ublaze
  script:
    # Pass project commit tag, message etc. to downstream project for the release
    # CI_COMMIT_TAG is available only in tag pipeline but we need it in regular ones (if available)
    - TAG=$(git tag -l --contains HEAD)
    # If no tag available, precede commit sha with branch-name for naming consistency
    - [[ -n "$TAG" ]] && VERSION=$TAG-$CI_COMMIT_SHORT_SHA || VERSION=$CI_COMMIT_BRANCH-$CI_COMMIT_SHORT_SHA
    - echo "GW_VERSION=$VERSION" >> build.env
    - echo "GW_COMMIT=$CI_COMMIT_SHORT_SHA" >> build.env
    - echo "GW_TITLE=$CI_COMMIT_TITLE" >> build.env
    - echo "GW_AUTHOR=$CI_COMMIT_AUTHOR" >> build.env
    - !reference [.ise_ublaze, script]
  artifacts:
    reports:
      dotenv: build.env

gen_firmware:
  rules:
    # Disable tag pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
    - if: $CI_COMMIT_TAG
      when: never
    # Don't run on development branches as the downstream job would fail to get the right artifacts
    - if: $CI_COMMIT_BRANCH != $EQUIP_NAME
      when: never
    - when: on_success
  stage: release
  inherit:
    variables: true
  trigger:
    project: abt-projects/gateware/kitr/sdk/txmc63x
    strategy: depend
    branch: kittr_generic
```

Use case ~~TXMC635~~ building pipeline

- Project where (gitlab) softcore firmw

- GitLab's multi-
→ downstream artifacts, not n

- Current workar

- Disable tag pipe
- Custom job rules
- Push atomically
- git push --a

- To-do: make a needs: ref:, S

```
variables:
  GIT_SUBMODULE_STRATEGY: recursive
  WORK_DIR: workspace
  SOCHW: $WORK_DIR/txmc635_soc
  US_GW_PATH: abt-projects/gatewaye/kitr/ise/txmc635-kitr
  #EQUIP_NAME is a global var from upstream

stages:
  - compile
  - convert
  - gen_config
  - release
  - deploy

# Continue passing the upstream project commit tag, message etc.
pass_env_vars:
  rules:
    # only run (once!) when triggered by upstream or web, so EQUIP_NAME is defined
    - if: $CI_PIPELINE_SOURCE == "pipeline"
    - if: $CI_PIPELINE_SOURCE == "web"
    - when: never
after_script:
  - echo "GW_VERSION=$GW_VERSION" >> compile.env
  - echo "GW_COMMIT=$GW_COMMIT" >> compile.env
  - echo "GW_TITLE=$GW_TITLE" >> compile.env
  - echo "GW_AUTHOR=$GW_AUTHOR" >> compile.env
artifacts:
  name: kitr-$EQUIP_NAME-$CI_COMMIT_SHORT_SHA-$CI_JOB_NAME
  reports:
    dotenv: compile.env

compile_kitr:
  stage: compile
  extends:
    - .pass_env_vars
  timeout: 10 min
  image: gitlab-registry.cern.ch/cce/docker_build/vivado:2019.1
  needs:
    - project: abt-projects/gatewaye/kitr/ise/txmc635-kitr
      job: build
      ref: $EQUIP_NAME
      artifacts: true
  script:
    #TODO: build script should take arguments such as $SOCHW
    - echo "Building for gatewaye version ${GW_VERSION}"
    - xsct xsct-build.tcl | tee build.log
    - 'ls build.log && if grep -i error build.log; then echo "build errors found" && exit 1; fi'
  artifacts:
    paths:
      - build.log
      - $WORK_DIR/cpp_boot_ram_launcher/Debug/cpp_boot_ram_launcher.elf
      - $WORK_DIR/cpp_boot_ram_launcher/Debug/cpp_boot_ram_launcher.elf.size
      - $WORK_DIR/AppV0/Debug/AppV0.elf
      - $WORK_DIR/AppV0/Debug/AppV0.elf.size
      # soc system files needed for next stage
      - $SOCHW/

data2mem_boot_ram_launcher:
  stage: convert
  extends:
    - .pass_env_vars
  timeout: 10 min
  image: gitlab-registry.cern.ch/cce/docker_build/ise:14.7
  variables:
    proc: "microblaze_0"
    app: "${WORK_DIR}/cpp_boot_ram_launcher/Debug/cpp_boot_ram_launcher.elf"
  script:
    - echo "adding boot_ram_launcher .elf to bistream BlockRAM"
    - elfcheck -hw $SOCHW/system.xml -mode bootload -mem BRAM -pe $proc $app
    - data2mem -bm $SOCHW/edkBmmFile_bd.bmm -bt $SOCHW/*.bit
    - bd $app tag $proc -o b kitr-$EQUIP_NAME-programmed.bit
  artifacts:
    paths:
      - kitr-$EQUIP_NAME-programmed.bit
```

eam
eam

pipelines
BRANCH

or
(docs)

```
build:
  rules:
    # Disable tag pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
    - if: $CI_COMMIT_TAG
      when: never
    - when: on_success
  extends:
    - .ise_ublaze
  script:
    # Pass project commit tag, message etc. to downstream project for the release
    # CI_COMMIT_TAG is available only in tag pipeline but we need it in regular ones (if available)
    - [TAG=$(git tag -l --contains HEAD)]
    # If no tag available, precede commit sha with branch-name for naming consistency
    - [[ -n "$TAG" ]] && VERSION=$TAG-$CI_COMMIT_SHORT_SHA || VERSION=$CI_COMMIT_BRANCH-$CI_COMMIT_SHORT_SHA
    - echo "GW_VERSION=$VERSION" >> build.env
    - echo "GW_COMMIT=$CI_COMMIT_SHORT_SHA" >> build.env
    - echo "GW_TITLE=$CI_COMMIT_TITLE" >> build.env
    - echo "GW_AUTHOR=$CI_COMMIT_AUTHOR" >> build.env
    - !reference [.ise_ublaze, script]
  artifacts:
    reports:
      dotenv: build.env

gen_firmware:
  rules:
    # Disable tag pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
    - if: $CI_COMMIT_TAG
      when: never
    # Don't run on development branches as the downstream job would fail to get the right artifacts
    - if: $CI_COMMIT_BRANCH != $EQUIP_NAME
      when: never
    - when: on_success
  stage: release
  inherit:
    variables: true
  trigger:
    project: abt-projects/gatewaye/kitr/sdk/txmc63x
    strategy: depend
    branch: kitr_generic
```

Use case

• Project where (softcore firmw

• GitLab's multi- → downstream artifacts, not n

• Current workar

- Disable tag pipe
- Custom job rules
- Push atomically
- git push --a

• To-do: make a needs: ref:, S

```

variables:
  GIT_SUBMODULE_STRAT
  WORK_DIR: workspace
  SOCHW: $WORK_DIR/tx
  US_GW_PATH: abt-pro
  #EQUIP_NAME is a gl

stages:
  - compile
  - convert
  - gen_config
  - release
  - deploy

# Continue passing th
pass_env_vars:
  rules:
    # only run (once!
    - if: $CI_PIPELIN
    - if: $CI_PIPELIN
    - when: never
  after_script:
    - echo "GW_VERSIO
    - echo "GW_COMMIT
    - echo "GW_TITLE=
    - echo "GW_AUTHOR

artifacts:
  name: kitr-$EQUIP
  reports:
    dotenv: compile

compile_kitr:
  stage: compile
  extends:
    - pass_env_vars
  timeout: 10 min
  image: gitlab-regis
  needs:
    - project: abt-pr
    - job: build
    - ref: $EQUIP_NAM
    - artifacts: true
  script:
    #TODO: build scri
    - echo "Building
    - xsct xsct-build
    - 'ls build.log &

artifacts:
  paths:
    - build.log
    - $WORK_DIR/cpp
    - $WORK_DIR/cpp
    - $WORK_DIR/App
    - $WORK_DIR/App
    # soc system fl
    - $$SOCHW/

data2mem_boot_ram_lau
stage: convert
extends:
  - pass_env_vars
  timeout: 10 min
  image: gitlab-regis
  variables:
    proc: "microblaze_u"
    app: "${WORK_DIR}/cpp_boot_ram_launcher/Debug/cpp_boot_ram_launcher.elf"
  script:
    - echo "adding boot_ram_launcher .elf to bistream BlockRAM"
    - elfcheck -hw $$SOCHW/system.xml -mode bootload -mem BRAM -pe $proc $app
    - data2mem -bm $$SOCHW/edkBmmFile_bd.bmm -bt $$SOCHW/*.bit
    -bd $app tag $proc -o b kitr-$EQUIP_NAME-programmed.bit
  artifacts:
    paths:
      - kitr-$EQUIP_NAME-programmed.bit

```

(docs)

ine

```

g pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
MMIT_TAG
r
uccess
e

ct commit tag, message etc. to downstream project for the release
TAG is available only in tag pipeline but we need it in regular ones (if available)
tag -l --contains HEAD)
available, precede commit sha with branch-name for naming consistency
AG ]] && VERSION=$TAG-$CI_COMMIT_SHORT_SHA || VERSION=$CI_COMMIT_BRANCH-$CI_COMMIT_SHORT_SHA'
ERSION=$VERSION" >> build.env
OMMIT=$CI_COMMIT_SHORT_SHA" >> build.env
ITLE=$CI_COMMIT_TITLE" >> build.env
UTHOR=$CI_COMMIT_AUTHOR" >> build.env
[.ise_ublaze, script]

ild.env

g pipelines cause downstream will fetch latest BRANCH artifacts, not necessarily latest TAG ones
MMIT_TAG
r
on development branches as the downstream job would fail to get the right artifacts
MMIT_BRANCH != $EQUIP_NAME
r
uccess

rue

-projects/gateway/kitr/sdk/txmc63x
pend
_generic

```



Use case – fids-ps-sw (old pipeline)

- Downstream pipeline from gateway, one branch per installation
- **80 minutes** to build the software, must run for each branch (10+)
- Triggers one downstream pipeline to package and deploy
- Not reproducible

The screenshot shows a GitHub Actions pipeline for the repository 'te-abt-ec / fids_petalinux_release'. The pipeline is titled 'bootscript - FMC input termination now possible by FESA, activate all input...'. It was created by Pieter Van Trappen for commit 20d09b46, 2 years ago. The pipeline is currently in a 'Passed' state. It is triggered for the 're'lease-mkp_pfn branch, specifically the 'latest' branch. The pipeline consists of two jobs: 'build' and 'deploy'. The 'build' job contains a single step 'build_petalinux'. The 'deploy' job contains a single step 'deploy' with a 'Trigger job' button. Below the main pipeline, there is a 'Downstream' section showing a 'Package Deploy ... #4885148' job, which is a 'Multi-project' job. This job triggers a downstream pipeline with two jobs: 'package' and 'deploy'. The 'package' job contains a single step 'package_petalinux'. The 'deploy' job contains two steps: 'deploy_B867' and 'deploy_operation'.

Use case – fids-ps-sw (new pipeline)

- Single branch, loops over a list of installations
- Configures the Yocto recipes and submodules branches within the recipes
- Saves identification files with git hashes of all dependencies
- 120 minutes to build all 9 “branches”
- Triggers one downstream pipeline per “branch” to package and deploy
- More reproducibility ([docs](#))

The screenshot displays a GitLab CI pipeline for the project 'abt-projects / gateway / fids / fids-ps-sw'. The pipeline is titled '#5233021' and is in a 'Passed' state. It was created by 'Lea Strobino' for the 'master' branch. The pipeline consists of several stages:

- build**: Contains a job 'build:multiple'.
- trigger**: Contains a job 'trigger:multiple'.
- Downstream**: A group of jobs, each representing a 'fids-release' for a specific commit (e.g., #5233955, #5233954, #5233953, #5233952, #5233951, #5233950). Each job is a 'Multi-project' type.
- package**: Contains jobs 'package:develop' and 'package:stable'.
- release**: Contains jobs 'release:develop' and 'release:stable'.
- deploy**: Contains jobs 'deploy:867:develop', 'deploy:867:stable', and 'deploy:operation'.

Use case – fids-ps-sw (new pipeline)

- Additional ‘trigger’ stage
- Downstream pipelines are **created dynamically using GitLab API**, authenticated with the job token
- Also possible to pass variables

```
30 $ for branch in $BRANCHES; do # collapsed multi-line command
31 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233946 for branch bidis-ls2
32 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233947 for branch ct_bfa9p
33 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233948 for branch ct_dfa
34 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233949 for branch kfa45
35 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233950 for branch mki_typeA
36 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233951 for branch mki_typeB
37 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233952 for branch mkp_pfn
38 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233953 for branch sbds_mkdh
39 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233954 for branch sbds_mkdv
40 Created downstream pipeline https://gitlab.cern.ch/abt-projects/gateway/fids/fids-release/-/pipelines/5233955 for branch sbds_ptm
41 Cleaning up project directory and file based variables
42 Job succeeded
```

Use case – fids-release

► package & release jobs

- Main branch containing global CI config
- One branch per installation with installation-specific mappings (gw & sw branches → devices)
- Combines gateway and software artifacts, creates boot images, uploads to GitLab package registry
 - Selects one software artifact from upstream pipeline (all included in the same archive because of the loop)
- Generates release text from identification files of each artifact
- Creates a release using the **job ID as tag**
- Passes this tag to next job in .env

```
20 Downloading artifacts for build:multiple (28889034)...
21 Downloading artifacts from coordinator... ok id=28889034 responseStatus=200 OK token=64_3sCm2
✓ 22 Executing "step_script" stage of the job script 08:08
23 Using docker image sha256:0e068cb435b42abd99a22926bcf76769277ce7f28c9b2ddba042c0132f9ddcc1 for gitlab-registry.cern.ch/cce/
docker_build/vivado:2021.1 with digest gitlab-registry.cern.ch/cce/docker_build/vivado/sha256:0e62fce0ebb7bdcf3362673c4c3b7
bb89f17d03f60caa2ff1b28d68bc24f6c7f ...
24 $ set -e
25 $ echo "Packaging Zynq PL gateway with PS software..." # collapsed multi-line command
26 Packaging Zynq PL gateway with PS software...
27 Gateway from https://gitlab.cern.ch/abt-projects/gateway/fids/fids-pl-gw/-/commit/6fe89c7c (branch: FASEC_FIDS_MKI_TypeB)
28 Software from https://gitlab.cern.ch/abt-projects/gateway/fids/fids-ps-sw/-/commit/62783e20 (branch: develop)
29 Selecting software artifacts for mki_typeB
30 $ bootgen -arch zynq -image boot.bif -split bin && mv *.bit.bin bitstream.bin
31 ***** Xilinx Bootgen v2021.1
32 **** Build date : Jun 10 2021-20:11:31
33 ** Copyright 1986-2021 Xilinx, Inc. All Rights Reserved.
34 [INFO] : Bootimage generated successfully
35 $ bootgen -arch zynq -image boot.bif -w -o boot.bin
36 ***** Xilinx Bootgen v2021.1
37 **** Build date : Jun 10 2021-20:11:31
38 ** Copyright 1986-2021 Xilinx, Inc. All Rights Reserved.
39 [INFO] : Bootimage generated successfully
40 $ bootgen -arch zynq -image boot-no-bitstream.bif -w -o boot-no-bitstream.bin
41 ***** Xilinx Bootgen v2021.1
42 **** Build date : Jun 10 2021-20:11:31
43 ** Copyright 1986-2021 Xilinx, Inc. All Rights Reserved.
44 [INFO] : Bootimage generated successfully
45 $ cp images/linux/image.ub image.ub
✓ 46 Uploading artifacts for successful job 08:08
47 Uploading artifacts...
48 bitstream.bin: found 1 matching files and directories
49 boot.bin: found 1 matching files and directories
50 boot-no-bitstream.bin: found 1 matching files and directories
51 image.ub: found 1 matching files and directories
52 release.txt: found 1 matching files and directories
53 Uploading artifacts as "archive" to coordinator... 201 Created id=28811260 responseStatus=201 Created token=64_3sCm2
✓ 54 Cleaning up project directory and file based variables 08:01
55 Job succeeded
```

Use case – fids-release

► GitLab release

- All releases attached to the 'fids-release' repository, which doesn't contain any code (just config files)
- Direct links to artifacts stored in GitLab package registry
- Table summarizing gw & sw git hashes, incl. dependencies, with links to commits in their respective repository
- Link to pipeline containing deployment jobs

abrt-projects / gateway / fids / fids-release / Releases / kfa45 (v2024.1)

kfa45 (v2024.1) Edit release

Assets 9

- Source code (zip) ↓
- Source code (tar.gz) ↓
- Source code (tar.bz2) ↓
- Source code (tar) ↓

Images

- Zynq PL bitstream ↗
- Zynq PS FIT image ↗

Packages

- U-Boot script file ↗
- U-Boot image ↗
- U-Boot image (without bitstream) ↗

Evidence collection

50197129-evidences-22418.json ↗ ... de9a61ef ↗

Collected 3 months ago

Release notes

Submodule	Commit	Description
fids-pl-gw	1b0c29fb	update fids_ip submodule - minor fids_slowVoltageCheck improvement upon masking; IP version unchanged
fids-ps-sw	b61e55c8	kernel - remove unused CAN and net vendors
fasec-driver	e1034ced	ci - ensure kernel headers can be found for both checks
libsnap7	b6a33e55	Merge branch 'SIL-781_diagnostics_on_rhel9' into 'master'
silecsserv	3c0e8576	v4.2.0 - bump minor due to fidsSlowVoltageCheck new regs

Built by pipeline #10062117 (open for deployment options)

47d1b9c3 50197129 Created 3 months ago by 🧑

Use case – fids-release

► deploy job

- 1-click deployment to operational and test devices, from pre-configured mappings: **no risk to deploy a development version in operation**
- Downloads artifacts from GitLab package registry, using release tag passed from previous job in .env
- Uploads the bitstream and boot image to CERN's TFTP server, for each configured device
 - Example here: same image for devices cfe-365-mkkfa45fidsgen1 to cfe-365-mkkfa45fidsgen4

deploy

⚙️ deploy:867:v2021.1	▶
✅ deploy:867:v2024.1	▶
⚙️ deploy:cfe-865-spare6v2:v2024.1	▶
⚙️ deploy:operation:v2021.1	▶
✅ deploy:operation:v2024.1	▶

```
324 $ echo "$EOS_ACCOUNT_PASSWORD" | kinit -f $EOS_ACCOUNT_USERNAME@CERN.CH
325 Password for btebot@CERN.CH:
326 $ cat <<EOF >> ~/.curlrc # collapsed multi-line command
327 $ set -e
328 $ curl -O $PACKAGE_URL/bitstream.bin
329 $ curl -O $PACKAGE_URL/image.ub
330 $ for device in $DEVICES; do # collapsed multi-line command
331 Deploying image to AIMS server for device cfe-365-mkkfa45fidsgen1 00:04
332 Proceeding with the upload.
333 Image CFE-365-MKKFA45FIDSGEN1 removed from aims.
334 Uploading to server:
335     vmlinuz
336         file: image.ub
337         md5: f62f7f6d3675c47acd207ff70d3c2d69
338         size: 38771564
339     initrd
340         file: bitstream.bin
341         md5: 490d3f1204147f3c63ea0605a04b32e3
342         size: 5979968
343 (this will take 20-30 seconds for typical linux vmlinuz/initrd sizes)
344 Please wait ...
345 Image CFE-365-MKKFA45FIDSGEN1 uploaded to aims.
346 Deploying image to AIMS server for device cfe-365-mkkfa45fidsgen2 00:03
347 Proceeding with the upload.
348 Image CFE-365-MKKFA45FIDSGEN2 removed from aims.
349 Uploading to server:
350     vmlinuz
351         file: image.ub
352         md5: f62f7f6d3675c47acd207ff70d3c2d69
353         size: 38771564
354     initrd
355         file: bitstream.bin
356         md5: 490d3f1204147f3c63ea0605a04b32e3
357         size: 5979968
358 (this will take 20-30 seconds for typical linux vmlinuz/initrd sizes)
359 Please wait ...
360 Image CFE-365-MKKFA45FIDSGEN2 uploaded to aims.
361 Deploying image to AIMS server for device cfe-365-mkkfa45fidsgen3 00:03
362 Proceeding with the upload.
```

Use case – fids-release

► identification

From AIMS boot server:
deployed image with release tag

```
[lstrobin@cwe-513-vm1321 ~]$ aims2client showimage cfe-ua*-*fids*
```

Image Name	Arch	UEFI	Description
CFE-UA23-MKIFIDSA1	armv7	N	FIDS-RELEASE-MKI_TYPEA-44971639
CFE-UA23-MKIFIDSA2	armv7	N	FIDS-RELEASE-MKI_TYPEA-44971639
CFE-UA23-MKIFIDSB1	armv7	N	FIDS-RELEASE-MKI_TYPEB-44971650
CFE-UA23-MKIFIDSB2	armv7	N	FIDS-RELEASE-MKI_TYPEB-44971650
CFE-UA87-MKIFIDSA1	armv7	N	FIDS-RELEASE-MKI_TYPEA-44971639
CFE-UA87-MKIFIDSA2	armv7	N	FIDS-RELEASE-MKI_TYPEA-44971639
CFE-UA87-MKIFIDSB1	armv7	N	FIDS-RELEASE-MKI_TYPEB-44971650
CFE-UA87-MKIFIDSB2	armv7	N	FIDS-RELEASE-MKI_TYPEB-44971650

From FASEC devices (SSH):
gw & sw git hashes (incl. submodules)

```
[lstrobin@cwe-513-vm1321 ~]$ ssh cfe-867-mkfa45fids
```

```
=====
cfe-867-mkfa45fids - 172.18.210.112 - FASEC FIDS v10.6.2 - kfa45_gpiod
=====
FIDS gateway          commit: 1b0c29fb    date:  Fri Jan 31 13:02:25 UTC 2025
PetaLinux              commit: b61e55c8    date:  Mon Oct 21 15:24:00 UTC 2024
FASEC driver          commit: e1034ced    date:  Tue Jun 11 13:11:22 UTC 2024
SILECS server         commit: 3c0e8576    date:  Thu Dec 12 15:18:42 UTC 2024
Snap7 library         commit: b6a33e55    date:  Mon May  6 14:52:47 UTC 2024
=====
Xilinx Zynq(R)-7000 XC7Z030, Kintex(TM)-7 FPGA, 125K logic cells
2x ARM(R) Cortex(R)-A9 MPCore(TM) CPU @ 333 MHz, 1006 MB memory
PetaLinux 2024.1
Linux cfe-867-mkfa45fids.cern.ch 6.6.10-xilinx-v2024.1-g62eec70a4258 armv7l
Up since: Sat Mar 22 07:02:46 UTC 2025

cfe-867-mkfa45fids:~$ sudo peek 0x43c10008

0x1b0c29fb

cfe-867-mkfa45fids:~$ ls -l /etc/revision/

-rw-r--r--  1 root    root          78 Oct 21  2024 fasec-driver
-rw-r--r--  1 root    root          59 Oct 21  2024 libsnap7
-rw-r--r--  1 root    root          53 Oct 21  2024 petalinux
-rw-r--r--  1 root    root          64 Oct 21  2024 silecsserv
```

Future

- **Migration of both projects to DI/OT, little impact on project structure**
 - Distributed Kicker Fast Control (DKFC), using CERN common DI/OT hardware and ABT specific hardware
 - ZynqMP System Board allows for a user script (block design compatible?) – docs
 - *Branch-per-installation* approach still required...
- **Possible improvements to the *branch-per-installation* approach**
 - To be investigated but with considerable design effort, all gateway variations between projects/installations could be embedded in the (bigger) ZynqMP PL/FPGA.
 - Software could configure MUXs and enable/disable modules to cater for installation differences.
 - Software in any case already today requires a hand-crafted (FESA3) Deploy Unit i.e. the instances.
 - Better separation between generic and specific parts so git rebase can be used? Technically not feasible since git uses snapshots and not differences. New version control systems could help:
 - Jujutsu VCS (jj) – ‘experimental’ but stable and well-advanced, most promising
 - can be hosted on git forges
 - simultaneous multi-branch rebasing would allow for more ~~laziness~~ efficiency
 - submodule support lacking but can still be done using git native commands



Thank you for your attention!

Questions?



home.cern