





REANA & Compute4PUNCH

Status & Update

Manuel Giffels, 13.05.2024

What is REANA?




[Home](#) [Examples](#) [Get Started](#) [Documentation](#) [News](#) [Roadmap](#) [Contact](#) [Blog](#)



Reproducible research data analysis platform




Flexible

Run many computational workflow engines.





Scalable

Support for remote compute clouds.




Reusable

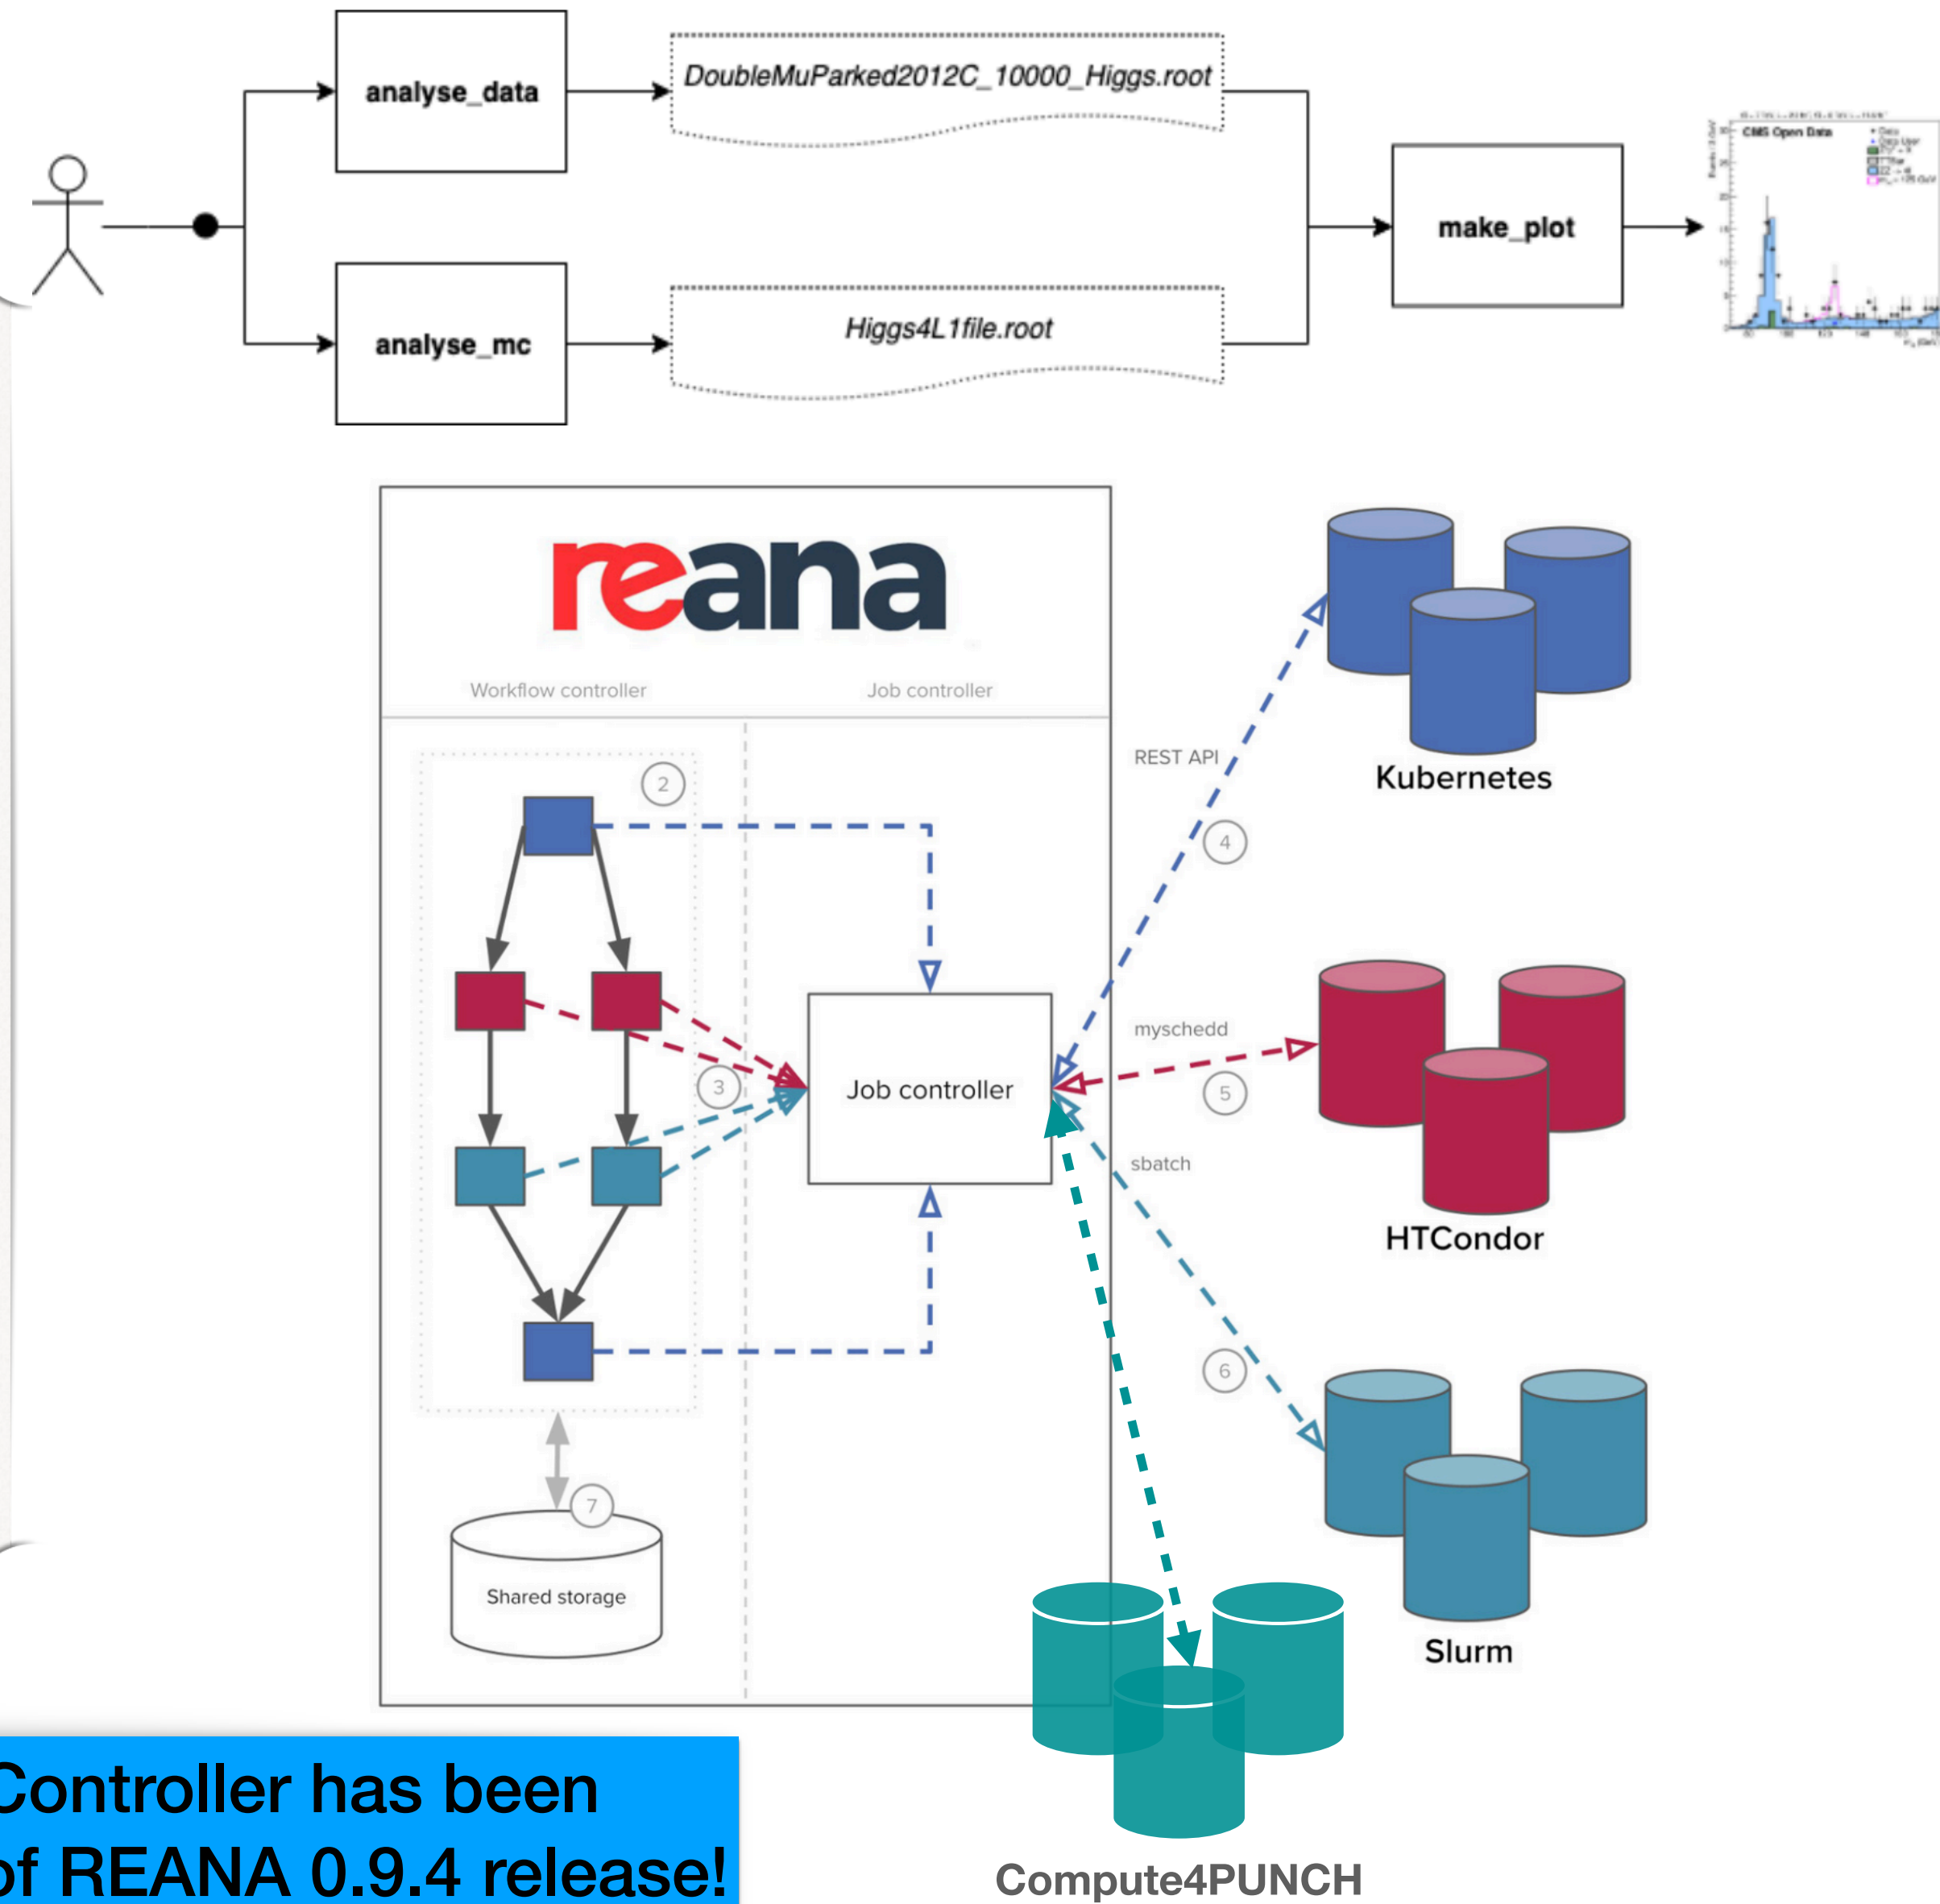
Containerise once, reuse elsewhere. Cloud-native.



Free

Free Software. MIT licence. Made with ❤️ at CERN.



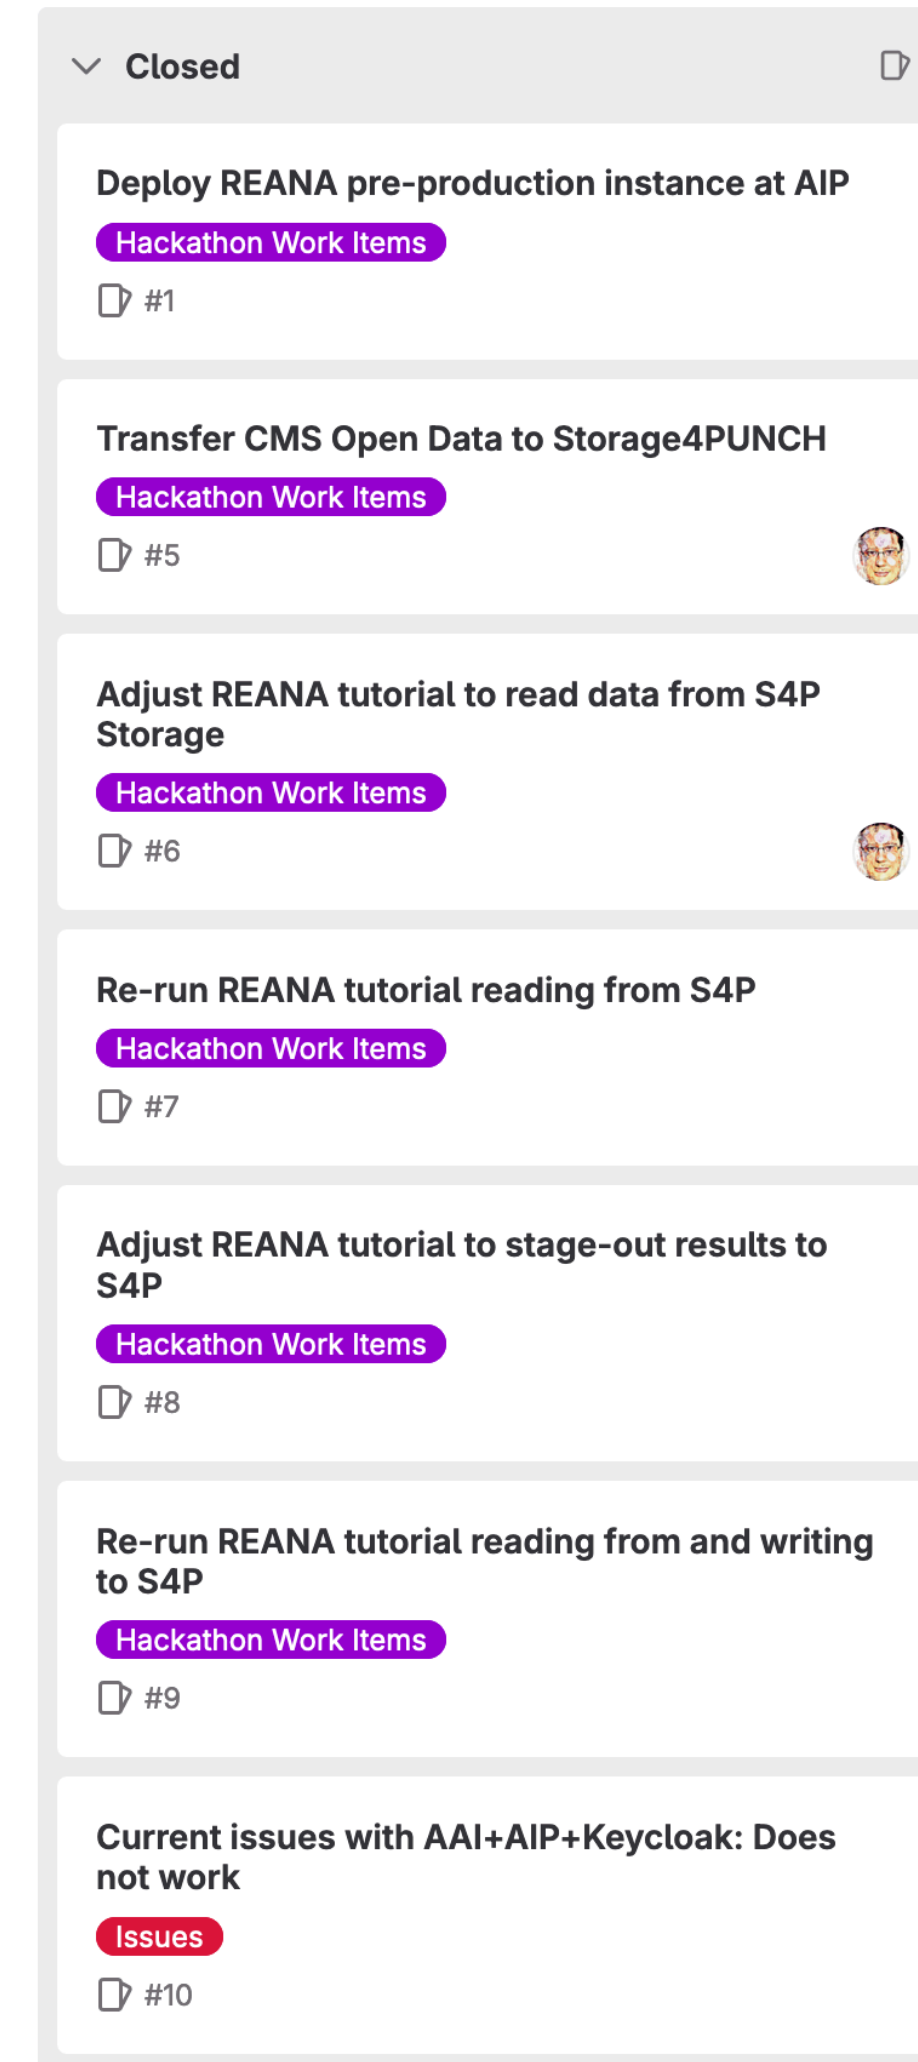


Compute4PUNCH Job Controller has been developed and will be part of REANA 0.9.4 release!

REANA4PUNCH - Hackathon Bonn

Integration of C4P, S4P and REANA

- **Goal:** Put all ingredients together C4P, S4P and REANA
- Deploy REANA development instance at AIP (including C4P Job Controller)
- Transfer CMS open data to S4P
- Run CMS Open Data Higgs tutorial on C4P reading from and writing to S4P
- After two days of collaborative work, we achieved our goal!
- Very successful event. Thanks to the colleagues from Bonn for organizing and hosting.



A screenshot of a GitLab board titled "Closed" with 7 items. The tasks listed are:

- Deploy REANA pre-production instance at AIP (Hackathon Work Items, #1)
- Transfer CMS Open Data to Storage4PUNCH (Hackathon Work Items, #5)
- Adjust REANA tutorial to read data from S4P Storage (Hackathon Work Items, #6)
- Re-run REANA tutorial reading from S4P (Hackathon Work Items, #7)
- Adjust REANA tutorial to stage-out results to S4P (Hackathon Work Items, #8)
- Re-run REANA tutorial reading from and writing to S4P (Hackathon Work Items, #9)
- Current issues with AAI+AIP+Keycloak: Does not work (Issues, #10)

[Link to GitLab Board](#)





Your REANA token

In order to use your token, make sure you have reana-client installed and run:

```
$ export REANA_SERVER_URL=https://reanadev.aip.de
$ export REANA_ACCESS_TOKEN=
```



Your GitLab projects



Connect to GitLab

In order to integrate your GitLab projects with REANA you need to grant permissions.

Connect

Your quota



0s
out of 1h

used 0%



0 Bytes
out of 100 GiB

used 0%

```
Last login: Tue Apr 23 11:00:01 on ttys006
giffler MacBook-Pro2023 ~ ssh -L 30443:127.0.0.1:30443 render.aip.de
Last login: Tue Apr 23 10:56:05 2024 from 129.13.101.141
(base) manuel@arm2armweb:~$
```

```
manuel_giffels@c4p-login:~ (ssh)

Last login: Tue Apr 23 11:01:06 on ttys008
giffler MacBook-Pro2023 ~ venv mccli-env
giffler MacBook-Pro2023 (e) mccli-env ~ all_proxy=socks5://127.0.0.1:5055 mccli ssh c4p-lo
gin.gridka.de
Last login: Tue Apr 23 10:58:51 2024 from 131.220.249.86
This host is managed by Puppet.

Config Management Information
  Satellite Host: https://satellite.scc.kit.edu/hosts/c4p-login.gridka.de
  Hostgroup: PUNCH/login
  Puppet Environment: production
  Satellite Owner: Benoit Roland <benoit.roland@kit.edu>
  i-doit Host: Virtual machines are not registered in i-doit

This machine is a Compute4PUNCH login node. Have fun!

[manuel_giffels@c4p-login ~]$ condor_q

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 11:02:44
OWNER BATCH_NAME      SUBMITTED   DONE    RUN    IDLE   HOLD   TOTAL JOB_IDS

Total for query: 0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
Total for manuel_giffels: 0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
Total for all users: 3 jobs; 0 completed, 0 removed, 0 idle, 0 running, 3 held, 0 suspended
```

reana

Your REANA token

In order to use your token, make sure you have reana-client installed and run:

```
$ export REANA_SERVER_URL=https://reanadev.aip.de
$ export REANA_ACCESS_TOKEN=
```

Your GitLab projects

i

Connect to GitLab

In order to integrate your GitLab projects with REANA you need to grant permissions.

Connect

Your quota

CPU

0s
out of 1h
used 0%

Disk

0 Bytes
out of 100 GiB
used 0%

Docs

Forum

Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

```
9b761fb1-9ab5-4d46-88e5-86195aaca4cc: command not found
(venv-reana) (base) manuel@arm2armweb:~$ export REANA_SERVER_URL=https://reanadev.aip.de
export REANA_ACCESS_TOKEN=
(venv-reana) (base) manuel@arm2armweb:~$ reana-client secrets-add --env HELMHOLTZ_TOP=eyJhbGciOiJFUzUxMiIsInR5cCI6Ikk1UK0pXVCJ9.eyJ2ZXIiOiIwLjYiLCJ0b2t1b190eXB1IjoibXl0b2t1b1IsImVzcyI6Imh0dHBzOi8vbXl0b2t1b15kYXRhLmtpdC5lZHUvIiwic3ViIjoioFVWNC9RN3RIN3RCVFprbnhSNjNRaGVVVGxjSHRqektIYWxjaGRTQUFYZ0iLCJzZXFFbm8iOjEsImF1ZCI6Imh0dHBzOi8vbXl0b2t1b15kYXRhLmtpdC5lZHUvIiwib2lkY19zdWIiOiJmZGM5MGRlYy1kNGEwLTRhNTAtYjI3Ny0zMTRlM2ZmMjY5MzYiLCJvaWRjX2lzcYI6Imh0dHBzOi8vbG9naW4uaGVsbWhvbHR6LmRlL29hdXR0MiIsImNhcGFiaWxpdlGllcyI6WyJBVCIsInRva2VuaW5mbyJdLCJleHAiOjE3MTQzNzgzNzAsIm5iZiI6MTcxMzc3Njg5NiwiWF0IjoxNzEzNzc2ODk2LCJhdXR0X3RpbWUiOjE3MTM3NzY4OTYsImp0aSI6IjVzMzExMDYwLTAwZjYtNDc1OS04MjU1LTZhNWY3ZTA3ZDRjNCIsInJlc3RyaWN0aW9ucyI6W3siaXhwIjoxNzE0Mzc5MzcwV19uANjoJ5V-e4KJh3EpAJlWvcogcLnSfx68KMRsTG-UTpfbAcJf_EH4c4zzxX91C6_p4zAFLYUsIEqkozmTGXhXAncqAestIIP9ij2u_wbotAlmBELKizwy_JPP5qL23-Up-pwYTRNALCj2env24u0-v5IR9CPU4gxzqNn42bsPr2TwkoDX
==> SUCCESS: Secrets HELMHOLTZ_TOP were successfully uploaded.
(venv-reana) (base) manuel@arm2armweb:~$ reana-client secrets-list
NAME          TYPE
HELMHOLTZ_TOP env
(venv-reana) (base) manuel@arm2armweb:~$
```

manuel_giffels@c4p-login:~ (ssh)

```
[manuel_giffels@c4p-login ~]$
```



In order to use your token, make sure you have reana-client installed and run:

In order to integrate your GitLab projects with REANA you need to grant permissions.

 Connect

used 0%



used 0%

reana-cern-open-data-tutorial / X

https://reanadev.aip.de/details/671e9d71-070d-42... Suchen

reana

reana-cern-open-data-tutorial #1

Started 3 minutes ago

running for 2 min 30 sec
step 0/3

Engine logs Job logs Workspace Specification

Step analysis running for 2 min 26 sec compute4punch wlcg-... \$ mkdir -p results && source /cvmfs/sf...

Docs Forum Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

```
==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml
-> SUCCESS: Valid REANA specification file.
==> Verifying REANA specification parameters...
-> SUCCESS: REANA specification parameters appear valid.
==> Verifying workflow parameters and commands...
-> SUCCESS: Workflow parameters and commands appear valid.
==> Verifying dangerous workflow operations...
-> SUCCESS: Workflow operations appear valid.
==> Verifying compute backends in REANA specification file...
-> SUCCESS: Workflow compute backends appear to be valid.
reana-cern-open-data-tutorial.1
==> SUCCESS: File /reana.yaml was successfully uploaded.
==> Uploading files...
==> Detected .gitignore file. Some files might get ignored.
==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.
==> SUCCESS: File /code/stage_out.sh was successfully uploaded.
==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.
==> Starting workflow...
==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued
(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial$
```

manuel_giffels@c4p-login:~ (ssh)

```
-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:03:34
OWNER      BATCH_NAME  SUBMITTED  DONE  RUN  IDLE  TOTAL  JOB_IDS
manuel_giffels ID: 158      4/23 12:03      _    1    _      1 158.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]$ condor_ssh_to_job 158
Welcome to slot1_1@"wwu-os-38ddb54312"@"wwu-os-38ddb54312.local.os.wwu.de!
Your condor job is running with pid(s) 9994.
Singularity> tail -f _condor_stdout
| Welcome to ROOT 6.30/02                                     https://root.cern |
| (c) 1995-2023, The ROOT Team; conception: R. Brun, F. Rademakers |
| Built for linuxx8664gcc on Dec 20 2023, 18:55:48              |
| From heads/master@tags/v6-30-02                             |
| With g++ (GCC) 12.1.0                                         |
| Try '.help'/'?', '.demo', '.license', '.credits', '.quit'/'.' |
|-----|

Processing code/df103_NanoAODHiggsAnalysis.C+...
```

reana-cern-open-data-tutorial / X

https://reanadev.aip.de/details/671e9d71-070d-42... Suchen

reana

reana-cern-open-data-tutorial #1

Started 7 minutes ago

running for 6 min 45 sec
step 1/3

Engine logs

Job logs

Workspace

Specification

Step

plots

running for 30 seconds

compute4punch

wlcg...

source /cvmfs/sft.cern.ch/lcg/views/...

Docs

Forum

Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml
-> SUCCESS: Valid REANA specification file.
==> Verifying REANA specification parameters...
-> SUCCESS: REANA specification parameters appear valid.
==> Verifying workflow parameters and commands...
-> SUCCESS: Workflow parameters and commands appear valid.
==> Verifying dangerous workflow operations...
-> SUCCESS: Workflow operations appear valid.
==> Verifying compute backends in REANA specification file...
-> SUCCESS: Workflow compute backends appear to be valid.
reana-cern-open-data-tutorial.1
==> SUCCESS: File /reana.yaml was successfully uploaded.
==> Uploading files...
==> Detected .gitignore file. Some files might get ignored.
==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.
==> SUCCESS: File /code/stage_out.sh was successfully uploaded.
==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.
==> Starting workflow...
==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued
(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial\$

manuel_giffels@c4p-login:~ (ssh)

Singularity> tail -f _condor_stdout
| Welcome to ROOT 6.30/02 https://root.cern |
| (c) 1995-2023, The ROOT Team; conception: R. Brun, F. Rademakers |
| Built for linuxx86_64gcc on Dec 20 2023, 18:55:48 |
| From heads/master@tags/v6-30-02 |
| With g++ (GCC) 12.1.0 |
Try '.help'/'?', '.demo', '.license', '.credits', '.quit'/'q'

Processing code/df103_NanoAODHiggsAnalysis.C+...
read returned, exiting
Connection to condor-job.wwu-os-38ddb54312.local.os.wwu.de closed.
[manuel_giffels@c4p-login ~]\$ condor_q

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07
OWNER BATCH_NAME SUBMITTED DONE RUN IDLE TOTAL JOB_IDS
manuel_giffels ID: 159 4/23 12:09 _ 1 _ 1 159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]\$

reana-cern-open-data-tutorial / X

https://reanadev.aip.de/details/671e9d71-070d-42... Suchen

reana

🔔 👤

✔ reana-cern-open-data-tutorial #1

Finished a few seconds ago

finished in 10 min 17 sec
step 3/3

⋮

⚙ Engine logs

📄 Job logs

📁 Workspace

📄 Specification

Search...

🔍

Name ↕	Modified ↕	Size ↕
📄 reana.yaml	2024-04-23T10:09:34	1.2 KiB
📄 .root_hist	2024-04-23T10:09:34	38 Bytes
📄 code/df103_NanoAODHiggsAnalysis.C	2024-04-23T10:09:34	22.85 KiB
📄 code/stage_out.sh	2024-04-23T10:09:34	316 Bytes
📄 code/PrintHistos.C	2024-04-23T10:09:34	415 Bytes
📄 code/df103_NanoAODHiggsAnalysis_C.d	2024-04-23T10:09:34	32.77 KiB
📄 code/df103_NanoAODHiggsAnalysis_C_...	2024-04-23T10:09:34	1.14 KiB
📄 code/df103_NanoAODHiggsAnalysis_C.so	2024-04-23T10:09:34	444.23 KiB
📄 logs/158.0.err	2024-04-23T10:09:33	615 Bytes
📄 logs/158.0.out	2024-04-23T10:09:33	617 Bytes
📄 results/higgs_4l.root	2024-04-23T10:09:35	9.59 KiB
📄 results/higgs_4mu.root	2024-04-23T10:09:35	9.42 KiB
📄 results/higgs_2el2mu.root	2024-04-23T10:09:35	9.54 KiB
📄 results/higgs_4el.root	2024-04-23T10:09:35	9.49 KiB
📄 .condor_creds/helmholtz.use	2024-04-23T10:09:35	968 Bytes

Docs

Forum

Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

```
==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml
-> SUCCESS: Valid REANA specification file.
==> Verifying REANA specification parameters...
-> SUCCESS: REANA specification parameters appear valid.
==> Verifying workflow parameters and commands...
-> SUCCESS: Workflow parameters and commands appear valid.
==> Verifying dangerous workflow operations...
-> SUCCESS: Workflow operations appear valid.
==> Verifying compute backends in REANA specification file...
-> SUCCESS: Workflow compute backends appear to be valid.
reana-cern-open-data-tutorial.1
==> SUCCESS: File /reana.yaml was successfully uploaded.
==> Uploading files...
==> Detected .gitignore file. Some files might get ignored.
==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.
==> SUCCESS: File /code/stage_out.sh was successfully uploaded.
==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.
==> Starting workflow...
==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued
(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial$
```

manuel_giffels@c4p-login:~ (ssh)

```
Singularity> tail -f _condor_stdout
| Welcome to ROOT 6.30/02 https://root.cern |
| (c) 1995-2023, The ROOT Team; conception: R. Brun, F. Rademakers |
| Built for linuxx8664gcc on Dec 20 2023, 18:55:48 |
| From heads/master@tags/v6-30-02 |
| With g++ (GCC) 12.1.0 |
| Try '.help'/'?', '.demo', '.license', '.credits', '.quit'/'q' |
|-----|

Processing code/df103_NanoAODHiggsAnalysis.C+...
read returned, exiting
Connection to condor-job.wwu-os-38ddb54312.local.os.wwu.de closed.
[manuel_giffels@c4p-login ~]$ condor_q

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07
OWNER          BATCH_NAME      SUBMITTED   DONE    RUN    IDLE  TOTAL JOB_IDS
manuel_giffels ID: 159        4/23 12:09      -        1      -        1 159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]$
```

reana-ern-open-data-tutorial / X

https://reanadev.aip.de/details/671e9d71-070d-4207-bc5b-7b9aa2e55664

Suchen

reana

reana-ern-open-data-tutorial #1

Finished a few seconds ago

finished in 10 min 17 sec
step 3/3

Engine logs

Job logs

Workspace

Specification

Step

stage_out

finished in 1 min 54 sec

compute4punch

wlcg-wn...

\$ chmod +x code/stage_out.sh && co...

Copying 16177 bytes file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.pdf

event: [1713867114632] BOTH GFAL2:CORE:COPY LIST:ENTER

event: [1713867114632] BOTH GFAL2:CORE:COPY LIST:ITEM file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.pdf

event: [1713867114632] BOTH GFAL2:CORE:COPY LIST:EXIT

event: [1713867114632] BOTH http_plugin PREPARE:ENTER file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.pdf

event: [1713867114754] BOTH http_plugin PREPARE:EXIT file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.pdf

event: [1713867114754] BOTH http_plugin TRANSFER:ENTER file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.pdf

event: [1713867114755] BOTH http_plugin TRANSFER:TYPE streamed

event: [1713867115034] BOTH http_plugin TRANSFER:EXIT file:///srv/results/higgs_2el2mu.pdf => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-ern-open-data-tutorial/

Docs

Forum

Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

==> Verifying REANA specification file... /home/manuel/reana-ern-open-data-tutorial/reana.yaml

-> SUCCESS: Valid REANA specification file.

==> Verifying REANA specification parameters...

-> SUCCESS: REANA specification parameters appear valid.

==> Verifying workflow parameters and commands...

-> SUCCESS: Workflow parameters and commands appear valid.

==> Verifying dangerous workflow operations...

-> SUCCESS: Workflow operations appear valid.

==> Verifying compute backends in REANA specification file...

-> SUCCESS: Workflow compute backends appear to be valid.

reana-ern-open-data-tutorial.1

==> SUCCESS: File /reana.yaml was successfully uploaded.

==> Uploading files...

==> Detected .gitignore file. Some files might get ignored.

==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.

==> SUCCESS: File /code/stage_out.sh was successfully uploaded.

==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.

==> Starting workflow...

==> SUCCESS: reana-ern-open-data-tutorial.1 has been queued

(venv-reana) (base) manuel@arm2armweb:~/reana-ern-open-data-tutorial\$

manuel_giffels@c4p-login:~ (ssh)

Singularity> tail -f _condor_stdout

| Welcome to ROOT 6.30/02 https://root.cern |

| (c) 1995-2023, The ROOT Team; conception: R. Brun, F. Rademakers |

| Built for linuxx8664gcc on Dec 20 2023, 18:55:48 |

| From heads/master@tags/v6-30-02 |

| With g++ (GCC) 12.1.0 |

| Try '.help'/'?', '.demo', '.license', '.credits', '.quit'/'q' |

Processing code/df103_NanoAODHiggsAnalysis.C+...

read returned, exiting

Connection to condor-job.wmu-os-38ddb54312.local.os.wmu.de closed.

[manuel_giffels@c4p-login ~]\$ condor_q

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07

OWNER BATCH_NAME SUBMITTED DONE RUN IDLE TOTAL JOB_IDS

manuel_giffels ID: 159 4/23 12:09 - 1 - 1 159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]\$

reana-cern-open-data-tutorial / X

https://reanadev.aip.de/details/671e9d71-070d-4207-bc5b-7b9aa2e55664

Suchen

reana

reana-cern-open-data-tutorial #1

Finished a few seconds ago

finished in 10 min 17 sec
step 3/3

Engine logs

Job logs

Workspace

Specification

Step

stage_out

finished in 1 min 54 sec

compute4punch

wlwg-wn...

\$ chmod +x code/stage_out.sh && co...

Copying 9764 bytes file:///srv/results/higgs_2el2mu.root => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.root

event: [1713867115578] BOTH GFAL2:CORE:COPY LIST:ENTER

event: [1713867115579] BOTH GFAL2:CORE:COPY LIST:ITEM file:///srv/results/higgs_2el2mu.root => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.root

event: [1713867115579] BOTH GFAL2:CORE:COPY LIST:EXIT

event: [1713867115579] BOTH http_plugin PREPARE:ENTER file:///srv/results/higgs_2el2mu.root => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.root

event: [1713867115647] BOTH http_plugin PREPARE:EXIT file:///srv/results/higgs_2el2mu.root => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.root

event: [1713867115647] BOTH http_plugin TRANSFER:ENTER file:///srv/results/higgs_2el2mu.root => https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664/higgs_2el2mu.root

event: [1713867115647] BOTH http_plugin TRANSFER:TYPE streamed

event: [1713867115864] BOTH http_plugin TRANSFER:EXIT file:///srv/results/higgs_2el2mu.root =>

Docs

Forum

Cluster health

MacBook-Pro2023.local

30443:127.0.0.1:30443 (ssh)

==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml

-> SUCCESS: Valid REANA specification file.

==> Verifying REANA specification parameters...

-> SUCCESS: REANA specification parameters appear valid.

==> Verifying workflow parameters and commands...

-> SUCCESS: Workflow parameters and commands appear valid.

==> Verifying dangerous workflow operations...

-> SUCCESS: Workflow operations appear valid.

==> Verifying compute backends in REANA specification file...

-> SUCCESS: Workflow compute backends appear to be valid.

reana-cern-open-data-tutorial.1

==> SUCCESS: File /reana.yaml was successfully uploaded.

==> Uploading files...

==> Detected .gitignore file. Some files might get ignored.

==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.

==> SUCCESS: File /code/stage_out.sh was successfully uploaded.

==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.

==> Starting workflow...

==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued

(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial\$

~ (-zsh)

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07

OWNER	BATCH_NAME	SUBMITTED	DONE	RUN	IDLE	TOTAL	JOB_IDS
manuel_giffels	ID: 159	4/23 12:09	_	1	_	1	159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]\$ Abgemeldet

Connection to c4p-login.gridka.de closed.

giffler MacBook-Pro2023 (e) mccli-env ~ davix-ls --capath /etc/grid-security/certificates -l -H "Authorization: Bearer \$(oidc-token punch-aai)" https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664

-rwxrwxrwx	0	16177	2024-04-23 12:11:55	higgs_2el2mu.pdf
-rwxrwxrwx	0	9764	2024-04-23 12:11:55	higgs_2el2mu.root
-rwxrwxrwx	0	15531	2024-04-23 12:11:56	higgs_4el.pdf
-rwxrwxrwx	0	9716	2024-04-23 12:11:57	higgs_4el.root
-rwxrwxrwx	0	17292	2024-04-23 12:11:58	higgs_4l.pdf
-rwxrwxrwx	0	9816	2024-04-23 12:11:59	higgs_4l.root
-rwxrwxrwx	0	16109	2024-04-23 12:12:00	higgs_4mu.pdf
-rwxrwxrwx	0	9649	2024-04-23 12:12:00	higgs_4mu.root

giffler MacBook-Pro2023 (e) mccli-env ~

```
-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07
```

OWNER	BATCH_NAME	SUBMITTED	DONE	RUN	IDLE	TOTAL	JOB_IDS
manuel_giffels	TD: 150	4/23 12:00		1		1	150 0

mandel_gtflets	id: 139	4/23 12:09	-	1	-	1 139.0
----------------	---------	------------	---	---	---	---------

```
Total for query: 1 jobs: 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
```

```
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspend
```

Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

```
[manuel_giffels@c4p-login ~]$ Abgemeldet
Connection to c4p-login gridka.de closed
```

reana-cern-open-data-tutorial / ×code/ × +

← → ↺ https://reanadev.aip.de/671e9d71-070d-4207-bc5b-7b Suchen

jupyter

QuitLogout

FilesRunningClusters

Select items to perform actions on them.

UploadNew ↕

0 ▾

/ code

	Name ↓	Last Modified	File size
..	vor ein paar Sekunden		
<input type="checkbox"/> df103_NanoAODHiggsAnalysis.C		vor 6 Minuten	23.4 kB
<input type="checkbox"/> df103_NanoAODHiggsAnalysis_C.d		vor 6 Minuten	33.6 kB
<input type="checkbox"/> df103_NanoAODHiggsAnalysis_C_ACLiC_dict_rdict.pcm		vor 6 Minuten	1.16 kB
<input type="checkbox"/> PrintHistos.C		vor 6 Minuten	415 B
<input type="checkbox"/> stage_out.sh		vor 6 Minuten	316 B

```
MacBook-Pro2023.local
30443:127.0.0.1:30443 (ssh)

==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml
-> SUCCESS: Valid REANA specification file.
==> Verifying REANA specification parameters...
-> SUCCESS: REANA specification parameters appear valid.
==> Verifying workflow parameters and commands...
-> SUCCESS: Workflow parameters and commands appear valid.
==> Verifying dangerous workflow operations...
-> SUCCESS: Workflow operations appear valid.
==> Verifying compute backends in REANA specification file...
-> SUCCESS: Workflow compute backends appear to be valid.
reana-cern-open-data-tutorial.1
==> SUCCESS: File /reana.yaml was successfully uploaded.
==> Uploading files...
==> Detected .gitignore file. Some files might get ignored.
==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.
==> SUCCESS: File /code/stage_out.sh was successfully uploaded.
==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.
==> Starting workflow...
==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued
(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial$
```

```
-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07
OWNER          BATCH_NAME    SUBMITTED   DONE    RUN    IDLE  TOTAL JOB_IDS
manuel_giffels ID: 159      4/23 12:09      _      1      _      1 159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]$ Abgemeldet
Connection to c4p-login.gridka.de closed.
giffler MacBook-Pro2023 (e) mccli-env ~ davix-ls --capath /etc/grid-security/certificates -l -H "Authorization: Bearer $(oidc-token punch-aa)" https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664
-rwxrwxrwx 0      16177      2024-04-23 12:11:55 higgs_2el2mu.pdf
-rwxrwxrwx 0      9764      2024-04-23 12:11:55 higgs_2el2mu.root
-rwxrwxrwx 0     15531      2024-04-23 12:11:56 higgs_4el.pdf
-rwxrwxrwx 0      9716      2024-04-23 12:11:57 higgs_4el.root
-rwxrwxrwx 0     17292      2024-04-23 12:11:58 higgs_4l.pdf
-rwxrwxrwx 0      9816      2024-04-23 12:11:59 higgs_4l.root
-rwxrwxrwx 0     16109      2024-04-23 12:12:00 higgs_4mu.pdf
-rwxrwxrwx 0      9649      2024-04-23 12:12:00 higgs_4mu.root
giffler MacBook-Pro2023 (e) mccli-env ~
```

```
reana-cern-open-data-tutorial / X code/ jovyan@reanadev24-run-sessio X +
https://reanadev.aip.de/671e9d71-070d-4207-bc5b-7b Suchen
jupyter Logout

legend.SetBorderSize(0);
legend.SetTextSize(0.03);
legend.AddEntry(&h_data, "Data", "PE1");
legend.AddEntry(&h_bkg, "ZZ", "f");
legend.AddEntry(&h_cmb, "m_{H} = 125 GeV", "f");
legend.DrawClone();

// Add header
TLatex cms_label;
cms_label.SetTextSize(0.04);
cms_label.DrawLatexNDC(0.16, 0.92, "#bf{CMS Open Data}");
TLatex header;
header.SetTextSize(0.03);
header.DrawLatexNDC(0.63, 0.92, "#sqrt{s} = 8 TeV, L_{int} = 11.6 fb^{-1}");

// Save plot
auto output = std::string{"results/"} + filename;
c->SaveAs(output.c_str());
}

void df103_NanoAODHiggsAnalysis(const bool run_fast = false)
{
    // Enable multi-threading
    ROOT::EnableImplicitMT();

    // In fast mode, take samples from */cms_opendata_2012_nanoaod_skimmed/*, which has
    // the preselections from the selection_* functions already applied.
    //std::string path = "root://eospublic.cern.ch//eos/root-eos/cms_opendata_2012_nanoaod/";
    //if (run_fast) path = "root://eospublic.cern.ch//eos/root-eos/cms_opendata_2012_nanoaod_skimmed/";
    ;
    std::string path = "root://dcachewebdav-kit.gridka.de/pnfs/gridka.de/punch4nfdi/cms_opendata_2012_nanoaod/";
    if (run_fast) path = "root://dcachewebdav-kit.gridka.de/pnfs/gridka.de/punch4nfdi/cms_opendata_2012_nanoaod_skimmed/";

    // Create dataframes for signal, background and data samples

    // Signal: Higgs -> 4 leptons
    ROOT::RDataFrame df_sig_4l("Events", path + "SMHiggsToZZTo4L.root");

    // Background: ZZ -> 4 leptons
    // Note that additional background processes from the original paper with minor contribution were
    // left out for this
    // tutorial.
    ROOT::RDataFrame df_bkg_4mu("Events", path + "ZZTo4mu.root");
    ROOT::RDataFrame df_bkg_4el("Events", path + "ZZTo4e.root");
    ROOT::RDataFrame df_bkg_2el2mu("Events", path + "ZZTo2e2mu.root");

    // CMS data taken in 2012 (11.6 fb^{-1} integrated luminosity)
    ROOT::RDataFrame df_data_doublemu(
        "Events", {path + "Run2012B_DoubleMuParked.root", path + "Run2012C_DoubleMuParked.root"});
    ROOT::RDataFrame df_data_doubleel(
        "Events", {path + "Run2012B_DoubleElectron.root", path + "Run2012C_DoubleElectron.root"});

    // Reconstruct Higgs to 4 muons
    auto df_sig_4mu_reco = reco_higgs_to_4mu(df_sig_4l);
```

```
MacBook-Pro2023.local
30443:127.0.0.1:30443 (ssh)

==> Verifying REANA specification file... /home/manuel/reana-cern-open-data-tutorial/reana.yaml
-> SUCCESS: Valid REANA specification file.
==> Verifying REANA specification parameters...
-> SUCCESS: REANA specification parameters appear valid.
==> Verifying workflow parameters and commands...
-> SUCCESS: Workflow parameters and commands appear valid.
==> Verifying dangerous workflow operations...
-> SUCCESS: Workflow operations appear valid.
==> Verifying compute backends in REANA specification file...
-> SUCCESS: Workflow compute backends appear to be valid.
reana-cern-open-data-tutorial.1
==> SUCCESS: File /reana.yaml was successfully uploaded.
==> Uploading files...
==> Detected .gitignore file. Some files might get ignored.
==> SUCCESS: File /code/df103_NanoAODHiggsAnalysis.C was successfully uploaded.
==> SUCCESS: File /code/stage_out.sh was successfully uploaded.
==> SUCCESS: File /code/PrintHistos.C was successfully uploaded.
==> Starting workflow...
==> SUCCESS: reana-cern-open-data-tutorial.1 has been queued
(venv-reana) (base) manuel@arm2armweb:~/reana-cern-open-data-tutorial$

~ (-zsh)

-- Schedd: c4p-login.gridka.de : <192.108.45.35:9618?... @ 04/23/24 12:10:07
OWNER          BATCH_NAME      SUBMITTED   DONE    RUN    IDLE  TOTAL JOB_IDS
manuel_giffels ID: 159      4/23 12:09      -        1        -        1 159.0

Total for query: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for manuel_giffels: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
Total for all users: 1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended

[manuel_giffels@c4p-login ~]$ Abgemeldet
Connection to c4p-login.gridka.de closed.
giffler MacBook-Pro2023 (e) mccli-env ~ davix-ls --capath /etc/grid-security/certificates -l -H "Authorization: Bearer $(oidc-token punch-aai)" https://dcache-desy-webdav.desy.de:2880//pnfs/desy.de/punch/tutorials/reana-cern-open-data-tutorial/671e9d71-070d-4207-bc5b-7b9aa2e55664
-rwxrwxrwx 0      16177      2024-04-23 12:11:55 higgs_2el2mu.pdf
-rwxrwxrwx 0      9764      2024-04-23 12:11:55 higgs_2el2mu.root
-rwxrwxrwx 0     15531      2024-04-23 12:11:56 higgs_4el.pdf
-rwxrwxrwx 0      9716      2024-04-23 12:11:57 higgs_4el.root
-rwxrwxrwx 0     17292      2024-04-23 12:11:58 higgs_4l.pdf
-rwxrwxrwx 0      9816      2024-04-23 12:11:59 higgs_4l.root
-rwxrwxrwx 0     16109      2024-04-23 12:12:00 higgs_4mu.pdf
-rwxrwxrwx 0      9649      2024-04-23 12:12:00 higgs_4mu.root
giffler MacBook-Pro2023 (e) mccli-env ~
```

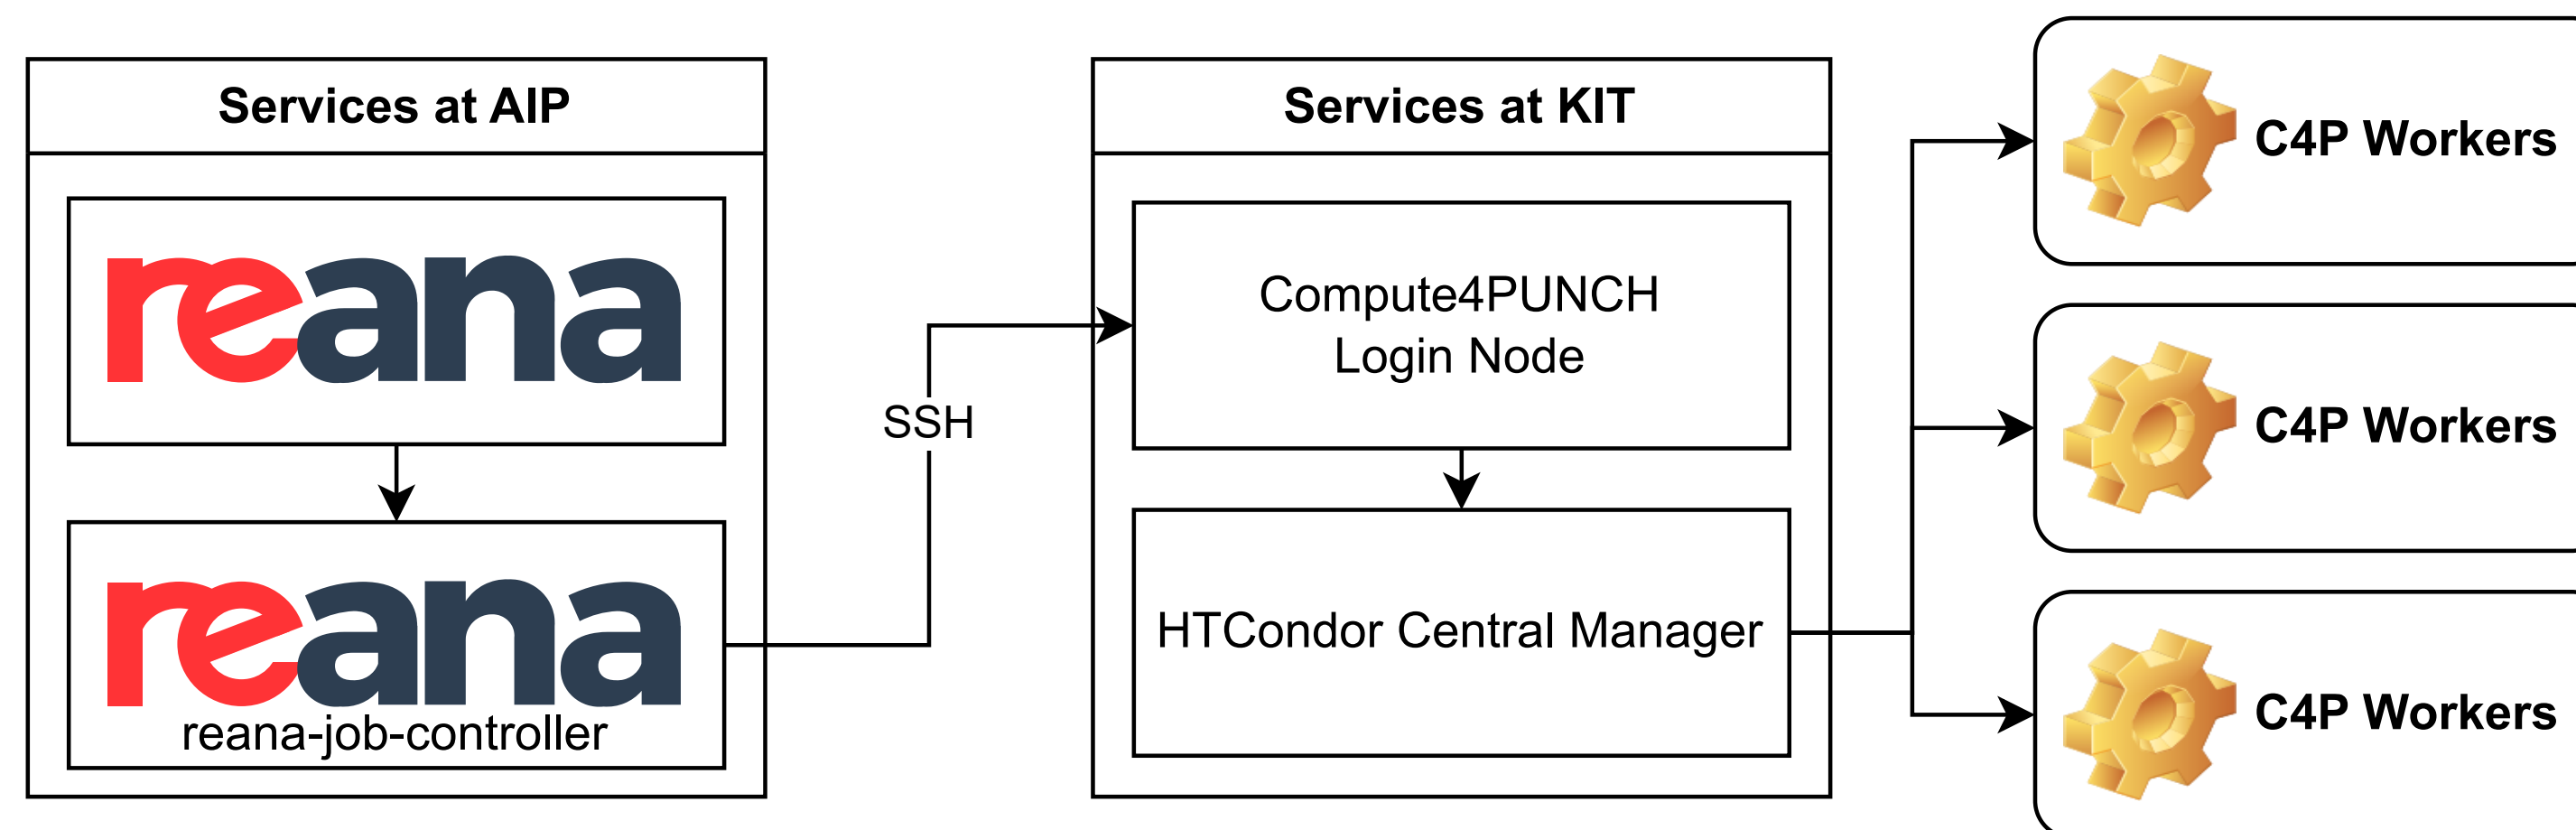
Summary & Outlook

- We have a basic working integration of C4P, S4P and REANA
- REANA C4P Job Controller will be part of REANA 0.9.4 release, so it can be deployed into the production REANA instance at AIP
- So far only, the serial workflow engine is supported and tested
- For the REANA development instance a hacked client is needed (input validation)
- Support for CWL and Snakemake will come with REANA 0.9.4
- Client input validation will be part of REANA 0.9.4 as well

Backup

How to get your jobs on Compute4PUNCH?

- Actual job submission in REANA is handled by a reana-job-controller per workflow (K8S Pod)
- There is now a prototype implementation of the Compute4PUNCH job controller available ([reana-job-controller/pull/430](https://github.com/reana/reana-job-controller/pull/430))
- Prototype utilizes a tokenized (mytoken) ssh connection to the c4p-login node at GridKa to submit, watch and handle jobs on Compute4PUNCH
- Needs a replacement in the future by using e.g. HTCondor's remote submission



How to access Compute & Storage4PUNCH

- An access token is necessary to access Compute & Storage4PUNCH
- Luckily REANA already comprises the management of arbitrary secrets
- Benoit already implement a convenient tool to create and upload mytokens to REANA
- Token can be accessed in the REANA job controller, used for ssh access to c4p-login

Fine for the initial token, but how to deal with renewals?

- C4P's HTCondor will take care of the token renewals

Secret management commands ¶

secrets-add

Add secrets from literal string or from file.

Examples:

```
$ reana-client secrets-add --env RUCIO_USERNAME=ruciouser

$ reana-client secrets-add --file userkey.pem

$ reana-client secrets-add --env VOMSPROXY_FILE=x509up_u1000
                                --file /tmp/x509up_u1000
```

secrets-delete

Delete user secrets by name.

Examples:

```
$ reana-client secrets-delete RUCIO_USERNAME
```

secrets-list

List user secrets.

Examples:

```
$ reana-client secrets-list
```