

Challenge(s) of the Migration of the RP-Coflu-Cluster

HT Condor Workshop
23.09.2020

Xavier Ouvrard, Chris Theis



During this talk

- Presentation of the RP-COFLU-Cluster
- Challenge(s) of the migration
- What's next?



During this talk

- **Presentation of the RP-COFLU-Cluster**
- Challenge(s) of the migration
- What's next?

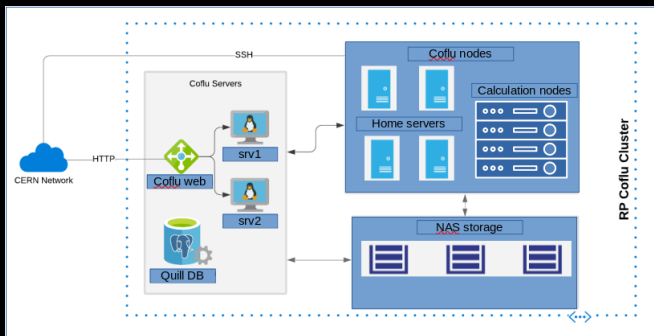


The RP-Coflu-Cluster I

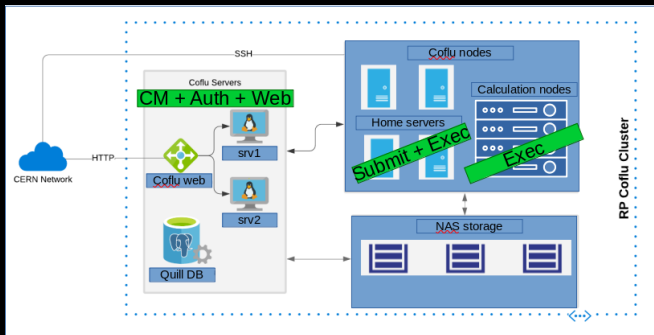
- In **radio-protection**: intensive usage of Fluka, a fully integrated particle physics MonteCarlo simulation package.
More: <https://fluka.cern/home>
- Before 2007: Fluka simulations run individually
- From 2007: Running Fluka simulations in parallel =>
Coflu-Cluster using HTCondor
- Now: Approximately 500 cores; relies on HTCondor as HTC software for the execution of Fluka simulation jobs



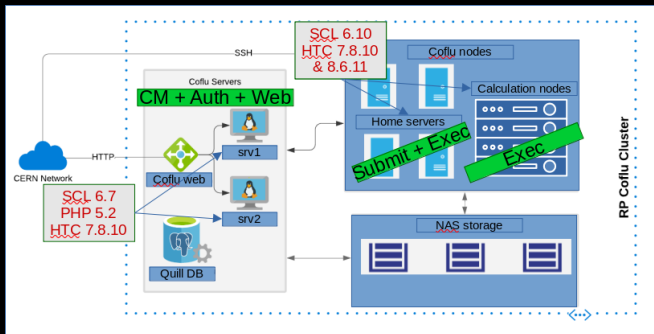
The RP-Coflu-Cluster II



The RP-Coflu-Cluster III



The RP-Coflu-Cluster IV



The RP-Coflu-Cluster V

COFLU v5.0
A Condor cluster front-end for FLUKA simulations.

Home Submissions User Jobs Cluster Status & Running Jobs Job History

Welcome to CoFLU Web

COFLU is a middleware with an interface which aims to facilitate the task of submitting FLUKA jobs into a Condor pool.

This Web-based interface for COFLU besides being more user-friendly provides a set of functionalities not available to the shell version, which can be very valuable.

Live Cluster Diagnostics

Cluster Load: 210 / 314 (185 / 415)

Cluster Health: 4 / 4 (4 / 4)

Jobs Running: 135 / 0 (0 / 0)

Condor
Powered by condor HTC platform.

FLUKA
Fluka - the Monte Carlo simulation software.

IRE
International Reactor Institute
Developed and maintained by IRE RP group.

Home Submissions User Jobs Cluster Status & Running Jobs Job History

Job Submission

Current directory:

Cluster Name:

Job Parameters

FLUKA input file (.inp):

Fluka Qtes [per CPU]:

Number of CPUs:

Test cycles:

Assign to project:

Fluka Dir./FLUPED: Use 4.0.0

Submit debt Do NOT Start Now! (Group only) Hold job on crash

Advanced Parameters

Aux files: Add aux file

FLUKA Executable: Use default

Fluka Script: Use default

CPUs Exps App: Enable

SEED GENERATION

New random seeds

Use seeds from file

Specify next seed

During this talk

- Presentation of the RP-COFLU-Cluster
- **Challenge(s) of the migration**
- What's next?

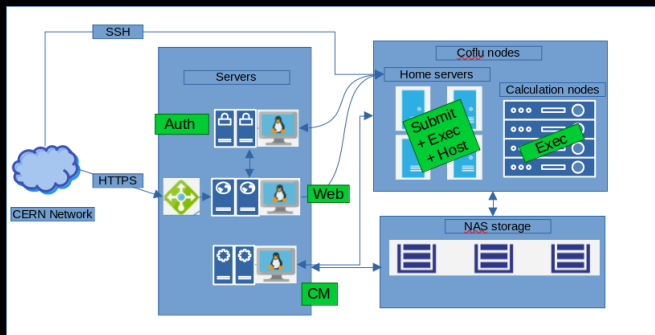


Challenge(s) of the migration I

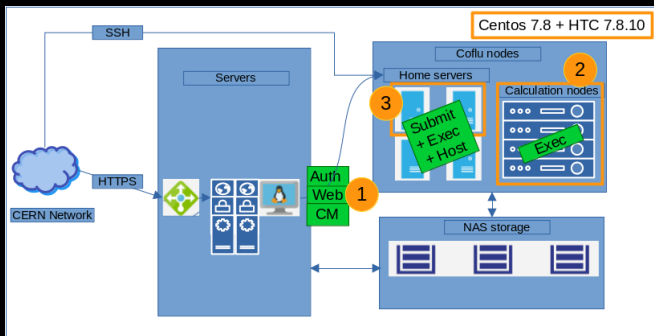
- Fluka was requiring an old fortran (4) until recently => **migration to Centos 7** is then possible
- **Production should have the least possible interruption**
- Coflu-website:
 - **withdraw all Quill Database** links
 - the Quill database was used for job management in the interface and job history
 - use HTCondor commands instead:
condor_q global => available with HTCondor 7.x
condor_history on remote => but we need schedulers to be with HTCondor version 8.x
- Separate as much as possible the roles on server
- Redundancy to ensure fallback in case of failure



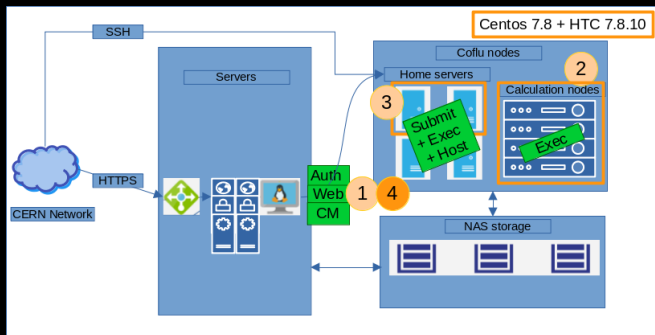
Challenge(s) of the migration II



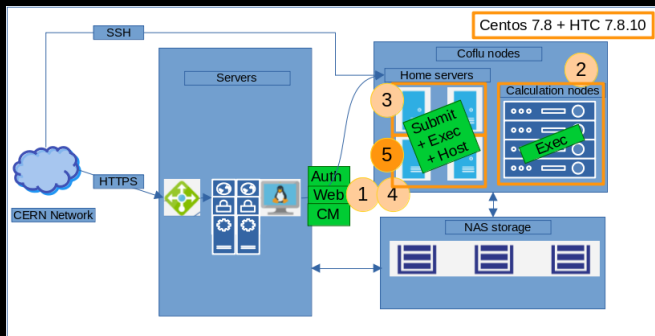
Challenge(s) of the migration III



Challenge(s) of the migration IV



Challenge(s) of the migration V



Challenge(s) of the migration VI

Job Queue Status

Job	Yieldable	Number	Current	Min	Empty	Full	Expired
all-jobs	0	0	0	0	0	0	0
all-jobs[0]	0	0	0	0	0	0	0

Queue details

View details | Stop gracefully jobs | Stop jobs | Kill jobs

Job ID	Queue	Subqueue	Est Time	State	Empty %	Discards	Priority
704	all-jobs	CPUS-01-01	00:00:00	Running	0	0	average[0]
707	all-jobs	CPUS-01-01	00:00:00	Running	0	0	average[0]

Job ID 707

Job ID 707

Subqueue Path /queue/CPUS-01-01

Subqueue Path /queue/CPUS-01-01

Process Group Name /CPUS-01-01-01

Process Name sh

Priority 0

Arguments /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Command /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Command path /bin

Process path /bin/bash

Process arguments /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Estimated remaining time 00:00:00

Notes [Reschedule current job](#) | [Stop gracefully current job](#) | [Kill current job](#) | [Kill current job](#)

Job ID	Queue	Subqueue	Est Time	State	Empty %	Discards	Priority
704	all-jobs	CPUS-01-01	00:00:00	Running	0	0	average[0]
707	all-jobs	CPUS-01-01	00:00:00	Running	0	0	average[0]

Job History Status Per User

User	Total Jobs	Failed	Current
all-jobs	206	04	01
all-jobs[0]	206	04	01

History details

Job ID	Queue	Subqueue	Created	Start Time	State	Exit code	Priority
all-jobs-CPUS-01-01-01-01-01-01	all-jobs	CPUS-01-01	2020-01-01 10:00:00	2020-01-01 10:00:00	Running	0	average[0]
all-jobs-CPUS-01-01-01-01-01-01-01	all-jobs	CPUS-01-01	2020-01-01 10:00:00	2020-01-01 10:00:00	Running	0	average[0]

Job ID 707

Job ID 707

Subqueue Path /queue/CPUS-01-01

Last Execution Node /queue/CPUS-01-01

Start time 2020-01-01 10:00:00

End time 2020-01-01 10:00:00

Process name /bin/bash

Process arguments /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Exit code 0

Priority 0

Arguments /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Command /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Command path /bin

Process path /bin/bash

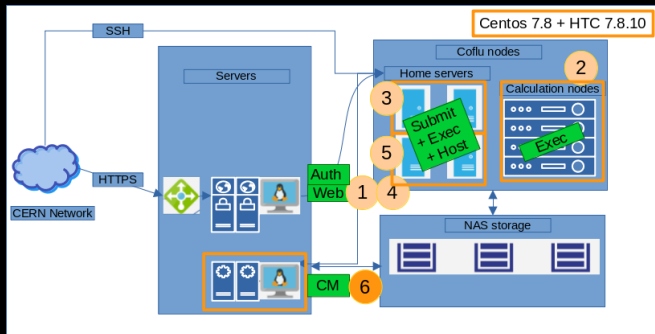
Process arguments /bin/bash -c 'cd /data/CPUS-01-01-01; echo "Job ID: 707" >> /data/CPUS-01-01-01/job.log'

Estimated remaining time 00:00:00

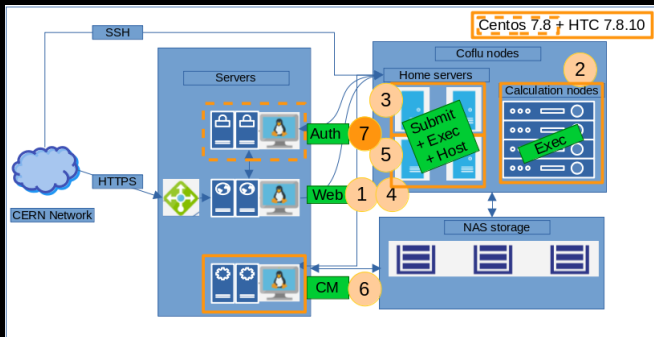
Job ID	Queue	Subqueue	Created	Start Time	State	Exit code	Priority
all-jobs-CPUS-01-01-01-01-01-01	all-jobs	CPUS-01-01	2020-01-01 10:00:00	2020-01-01 10:00:00	Running	0	average[0]
all-jobs-CPUS-01-01-01-01-01-01-01	all-jobs	CPUS-01-01	2020-01-01 10:00:00	2020-01-01 10:00:00	Running	0	average[0]



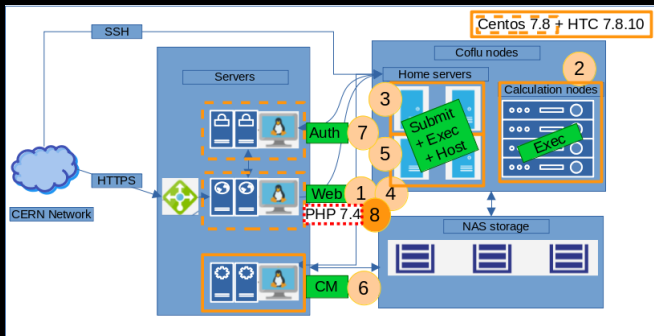
Challenge(s) of the migration VII



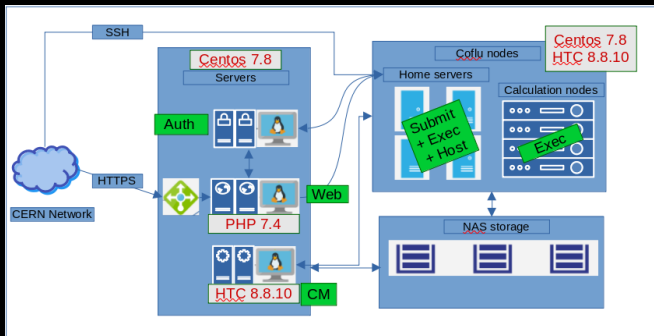
Challenge(s) of the migration VIII



Challenge(s) of the migration IX



Challenge(s) of the migration X



Challenge(s) of the migration XI

- Migration really helped by the **high backward compatibility of HTCondor**
 - **Centrally managed HTCondor configuration files** split in 3 files:
 - 1 file: Common part to all nodes of the cluster
 - 1 file: Configuration of central managers common to all nodes
- => switching of CM is then transparent
- 1 file: Defines daemon list and specific node / CM configuration

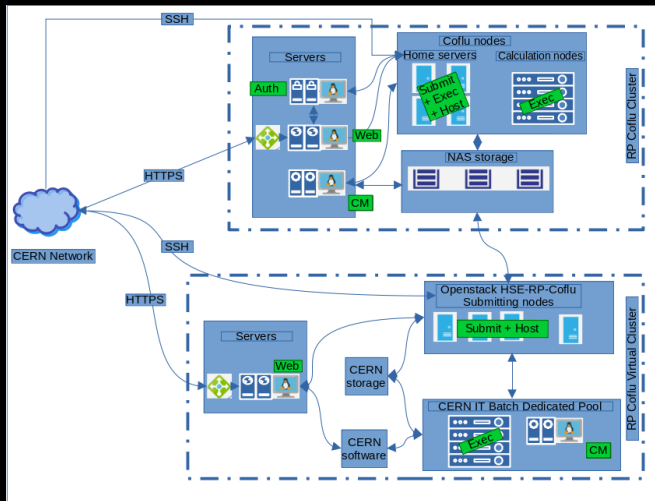


During this talk

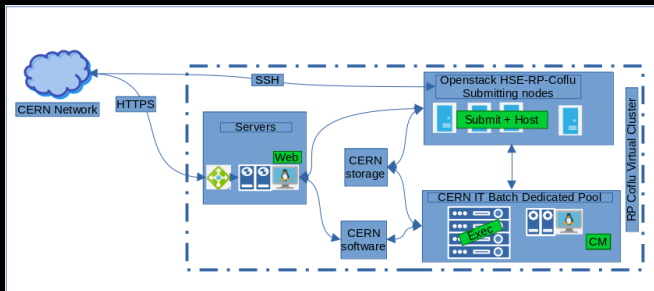
- Presentation of the RP-COFLU-Cluster
- Challenge(s) of the migration
- **What's next?**



What's next? I



What's next? II



Thank you for your attention!



More information:

- <https://hse.cern/services-support/radiation-protection>