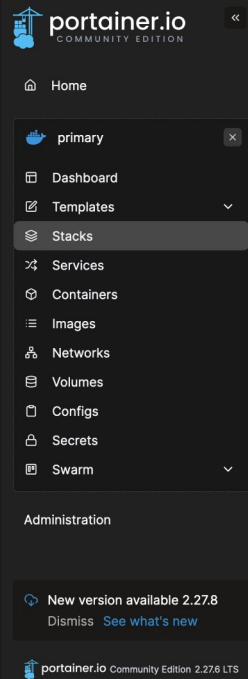


A Distributed Agata Emulator



replay_local_initializergitlab-registry.in2p3.fr/ip2gamma/docker_images:prodreplicated1 / 1Scale-2025-07-01 10:44:30

StatusFilterTaskActionsSlotNodeLast Update

runningtux501nqeallrbftt08mizljn1agata-analysis-62025-07-01 10:44:30

replay_local_processorgitlab-registry.in2p3.fr/ip2gamma/docker_images:prodreplicated11 / 11Scale-2025-07-01 10:44:40

StatusFilterTaskActionsSlotNodeLast Update

running1w0mlb4etmynn2zv6stwycvu6agata-analysis-62025-07-01 10:44:40

running2apgy1qxulpkdoktfyuybyfgo9agata-analysis-82025-07-01 10:44:40

running8oncw3mya4jvblxolgtkezy2b2agata-analysis-82025-07-01 10:44:40

runningafmcv3i7ihz35flqphm97k68210agata-analysis-82025-07-01 10:44:40

runningcn8f4dsqpm6cna2ktmg9lgkrf5agata-analysis-72025-07-01 10:44:40

runningj7jzdpd2fz14ui2amd7acefq1agata-analysis-62025-07-01 10:44:40

runningtkokogplj5gglgw1qm2p4535z4agata-analysis-72025-07-01 10:44:40

runningw4rx96do7xnu1fe3hgnxdlccp11agata-analysis-72025-07-01 10:44:40

runningxdxmvj081e3yw6w73nj6bc3wc8agata-analysis-62025-07-01 10:44:40

runningxqjjwoveqwoi4f1ajvlt71vn3agata-analysis-82025-07-01 10:44:40



Cluster and Scalability
Local level Processing
Global Level Processing

G. Baulieu

New analysis servers are being set up in Legnaro (P. Le Jeannic)

- 4 new machines
 - 2 x Xeon Silver 4310 (48 threads) @ 2.1 GHz
 - 128 GB of RAM
 - 10 Gbs access to data (anodeds5)

New analysis servers are being set up in Legnaro (P. Le Jeannic)

- 4 new machines
 - 2 x Xeon Silver 4310 (48 threads) @ 2.1 GHz
 - 128 GB of RAM
 - 10 Gbs access to data (anodeds5)

Processing will become more demanding

- From ~33 crystals to 135 (up to 180?) : at least a factor 4

New analysis servers are being set up in Legnaro (P. Le Jeannic)

- 4 new machines
 - 2 x Xeon Silver 4310 (48 threads) @ 2.1 GHz
 - 128 GB of RAM
 - 10 Gbs access to data (anodeds5)

Processing will become more demanding

- From ~33 crystals to 135 (up to 180?) : at least a factor 4

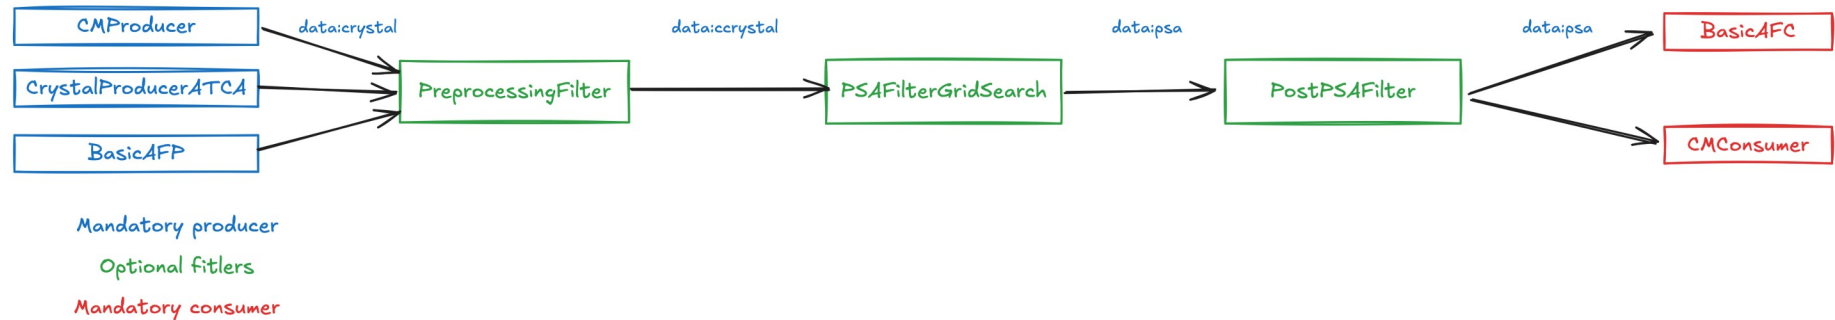
We need to be able to use computing power where it is : dispatched on different servers

Distributed Agata Emulator (DAEmule)

- A new AGAPRO emulator
- Designed to be scalable and able to run on clusters
- Using the same actors and configuration folders (from genconf.py)
- Different instances can communicate through a Central Memory system (REDIS implemented)
- Still under development, first tests performed on the Legnaro new Analysis cluster.

Local level processing

- Subset of actors available



- Each actor runs in its own thread
- Producer : choice between Central Memory, .cdat files or ADF files
- Filters : On/Off on Preprocessing, PSA and PostPSA
- Consumer : choice between ADF file and Central Memory
- 3 run modes :
 - Simple
 - Parallel
 - Batch

Local level processing

- Simple mode :

A single emulator on a single crystal folder, on a single machine.

- Parallel mode :

As many emulators as crystal folders, on a single machine. Define the number of emulators run in parallel. (~FEMUL behaviour)

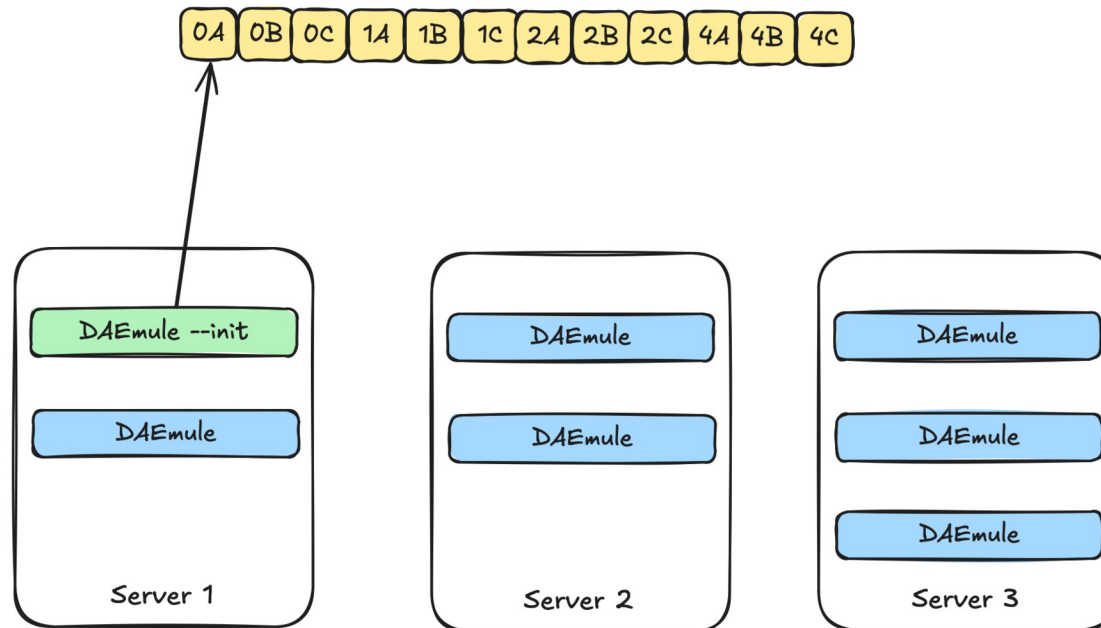
- Batch mode :

A single emulator per DAEmule instance but many instances on different servers.

Each DAEmule instance run on one crystal folder and then ask for a new one.

Local level processing

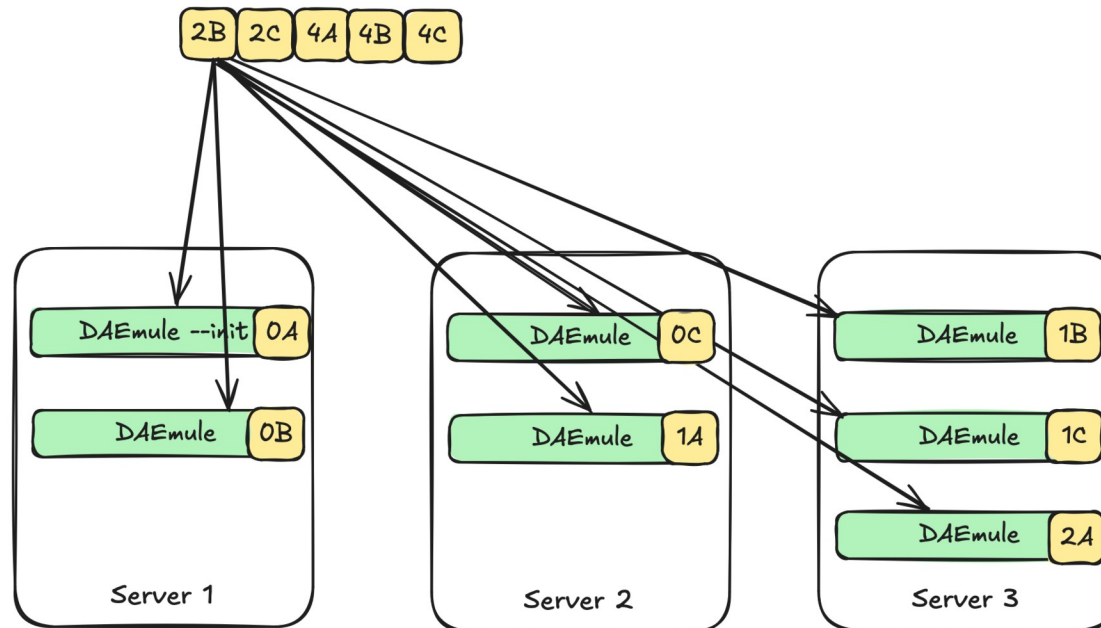
- Batch mode overview :



→ from configuration file to Central Memory list

Local level processing

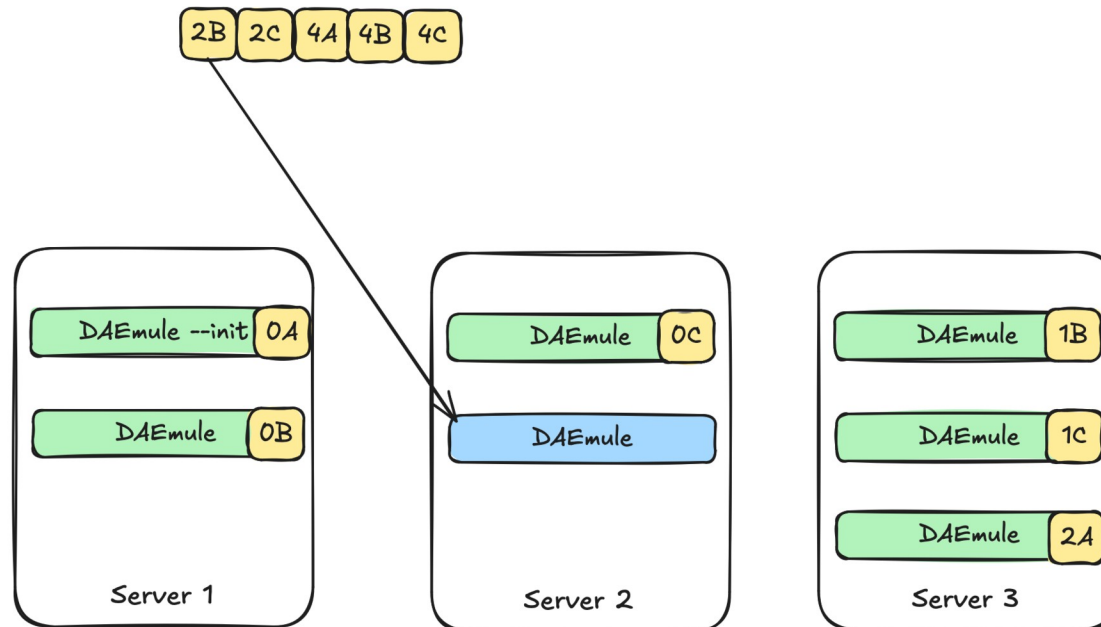
- Batch mode overview :



→ each instance fetches a workload from central memory (atomic)

Local level processing

- Batch mode overview :

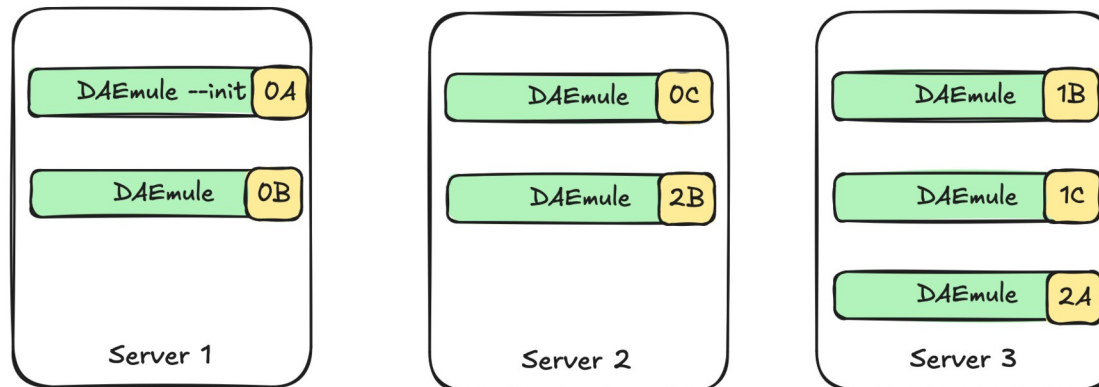


→ when an instance is done with its workload, it asks for a new one

Local level processing

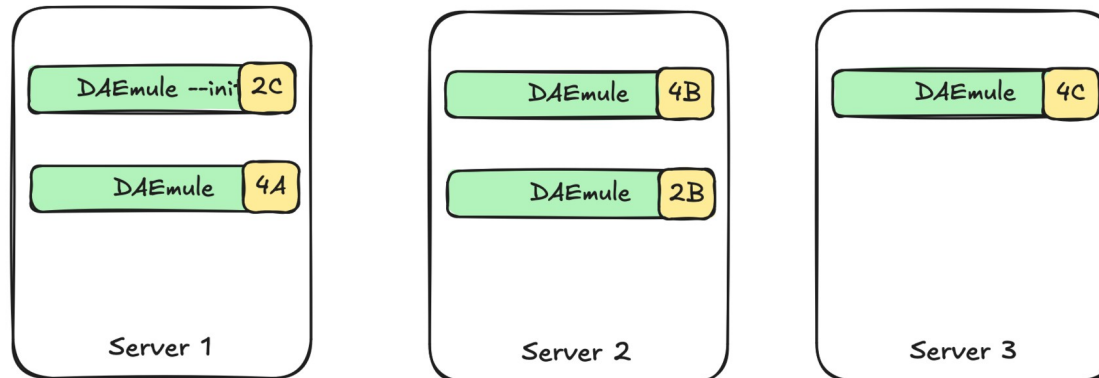
- Batch mode overview :

2C 4A 4B 4C



Local level processing

- Batch mode overview :



→ if the workload list is empty, the instance stops

Handling the cluster

- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**

Handling the cluster

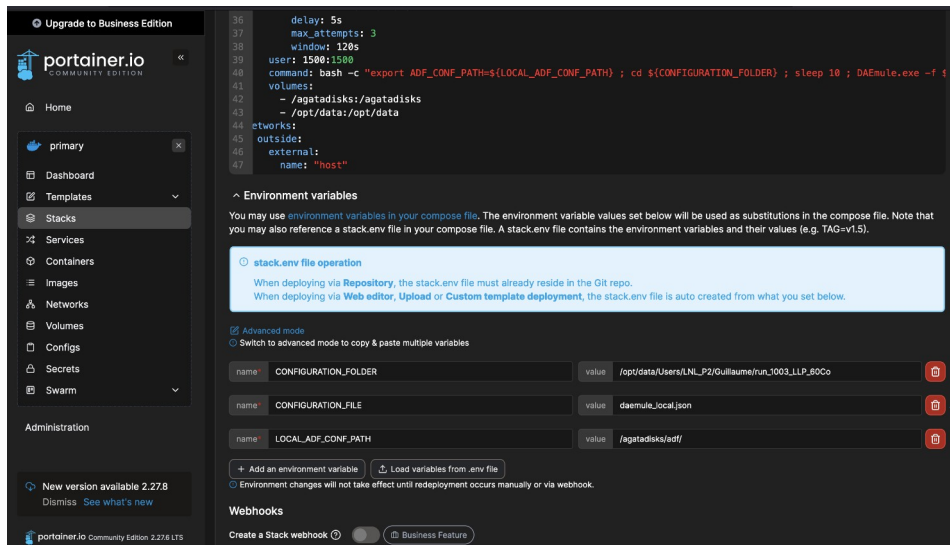
- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**
- Docker Swarm will handle the processes launches on different servers using Docker containers

Handling the cluster

- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**
- Docker Swarm will handle the processes launches on different servers using Docker containers
- Portainer is a web interface to Docker Swarm

Handling the cluster

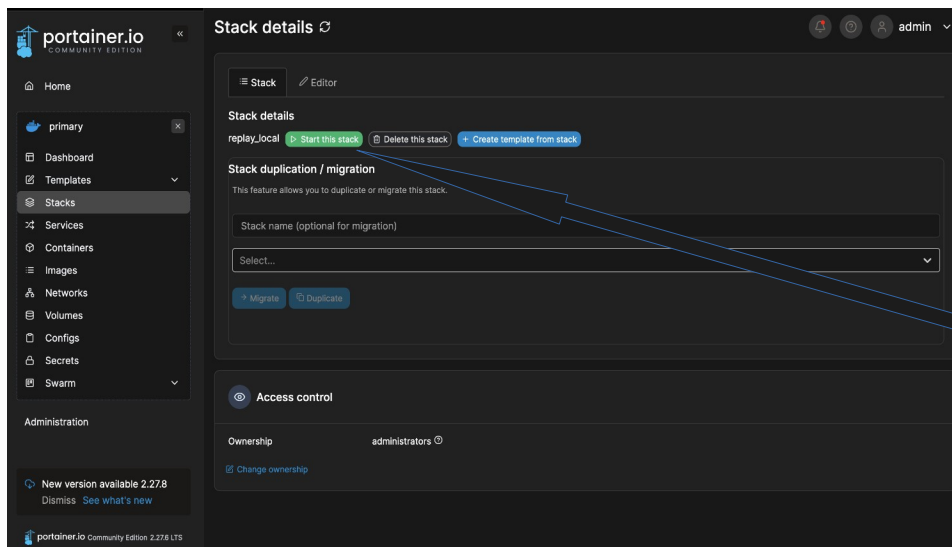
- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**
- Docker Swarm will handle the processes launches on different servers using Docker containers
- Portainer is a web interface to Docker Swarm



Define the
configuration file
location

Handling the cluster

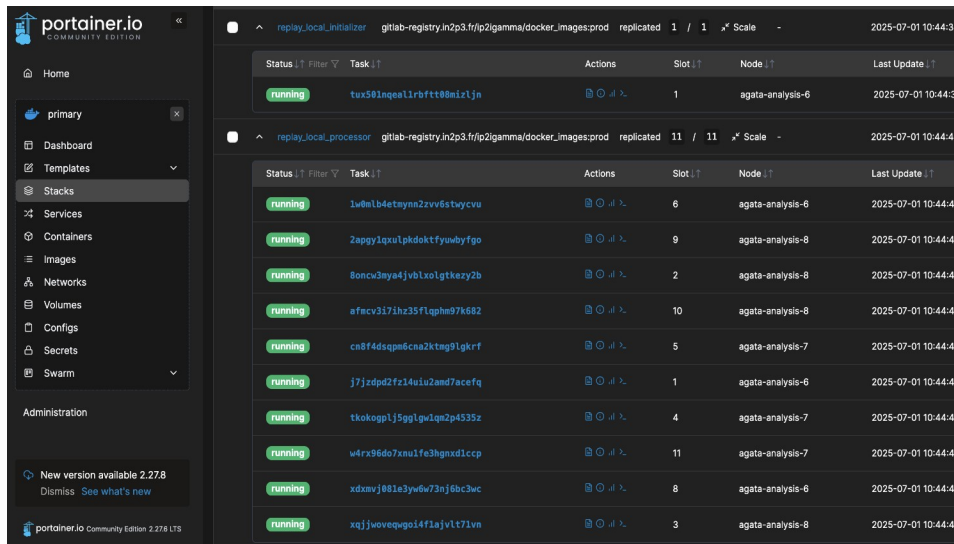
- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**
- Docker Swarm will handle the processes launches on different servers using Docker containers
- Portainer is a web interface to Docker Swarm



Click on the Start button

Handling the cluster

- Launching many instances on different servers is a pain
 - we can automatize this process using **Docker Swarm** and **Portainer**
- Docker Swarm will handle the processes launches on different servers using Docker containers
- Portainer is a web interface to Docker Swarm

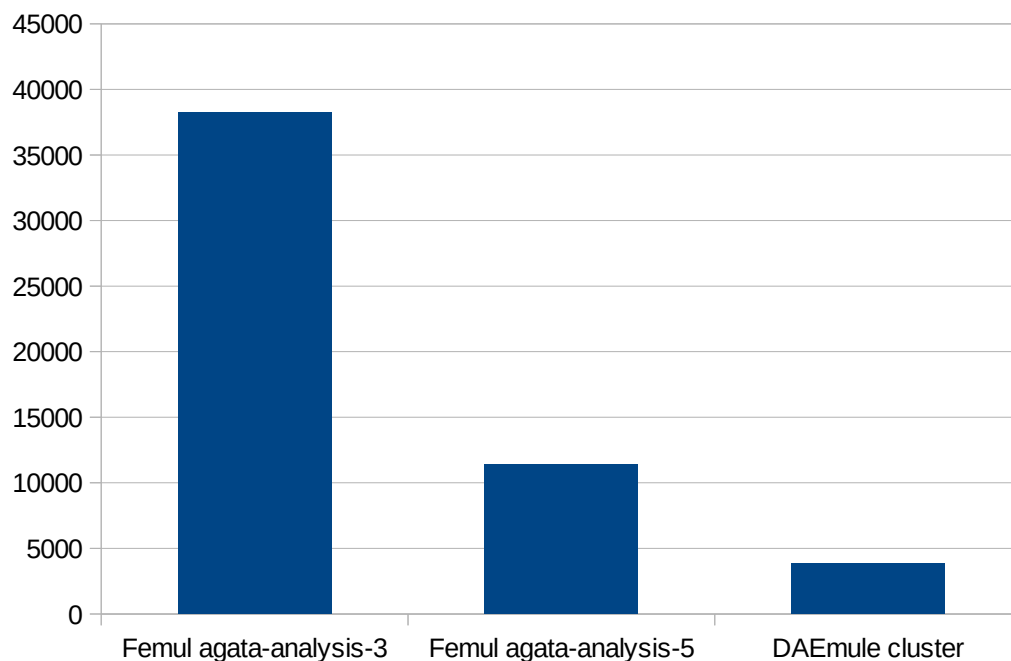


The screenshot displays the Portainer.io web interface. On the left is a sidebar with navigation options: Home, primary (selected), Dashboard, Templates, Stacks, Services, Containers, Images, Networks, Volumes, Configs, Secrets, and Swarm. The main area shows two stacks. The first stack, 'replay_local_initializer', has one task 'tux501nqeal1rbftt0mlz1jn' in a 'running' state. The second stack, 'replay_local_processor', is expanded and shows 11 tasks in a 'running' state. These tasks are distributed across nodes named 'agata-analysis-6' through 'agata-analysis-8'. Each task entry includes a status indicator (green 'running'), a task ID, a 'Slot' number, the node name, and the 'Last Update' timestamp.

It's running !

First tests on analysis cluster @ Legnaro

- Data from run_1003_LLP_60Co (33 crystals)

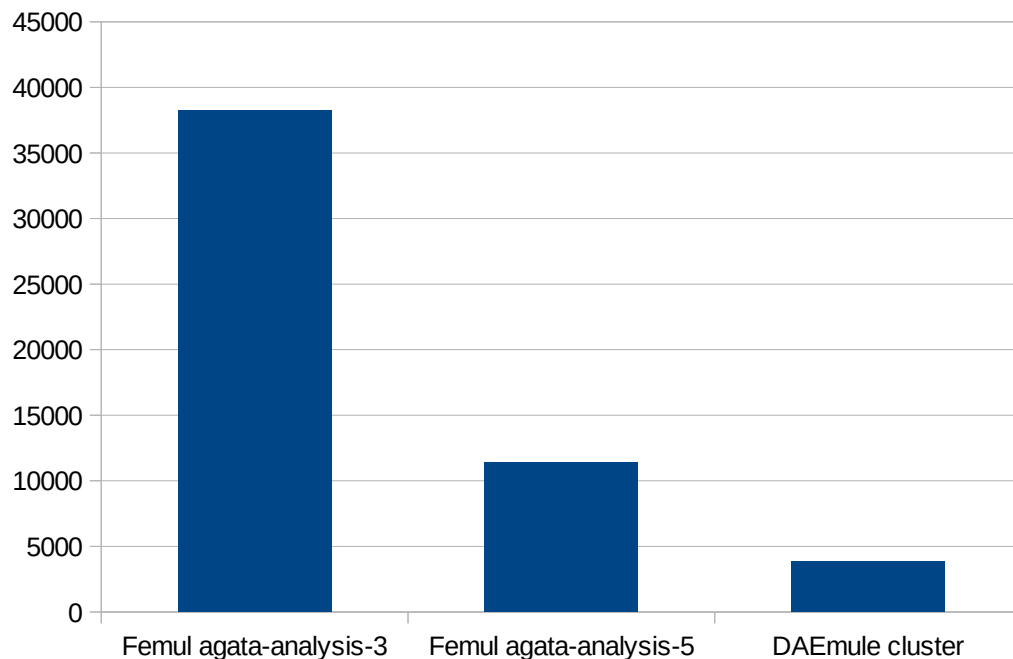


■ Time (s)

Hardware	Time (s)
Femul agata-analysis-3	38209
Femul agata-analysis-5	11405
DAEmule cluster	3828

First tests on analysis cluster @ Legnaro

- Data from run_1003_LLP_60Co (33 crystals)



■ Time (s)

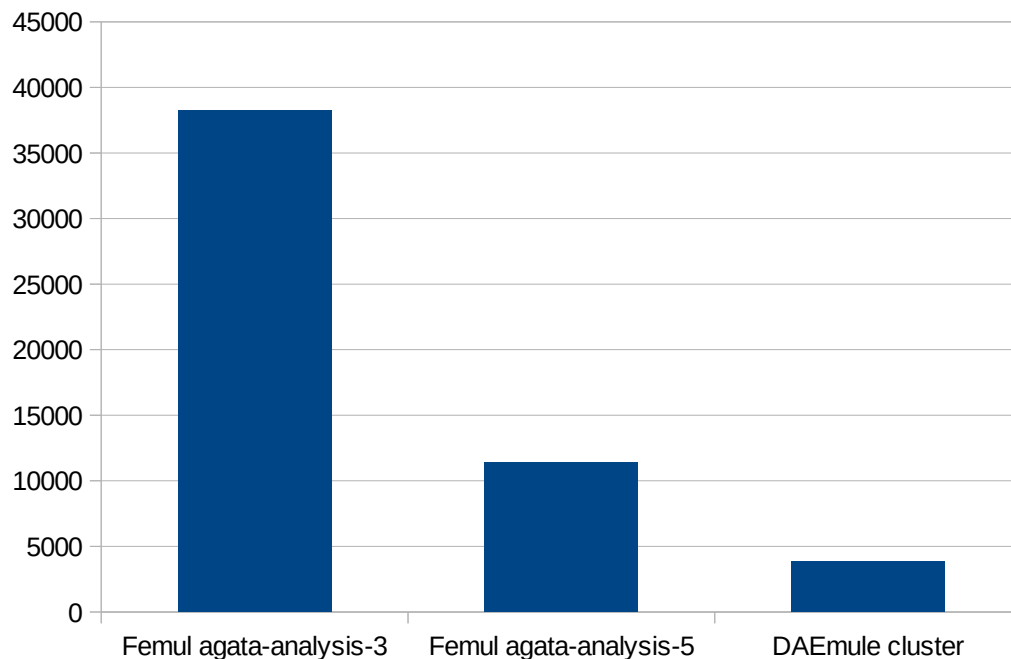
Hardware	Time (s)
Femul agata-analysis-3	38209
Femul agata-analysis-5	11405
DAEmule cluster	3828

Limited by computing power on a single server



First tests on analysis cluster @ Legnaro

- Data from run_1003_LLP_60Co (33 crystals)



■ Time (s)

Hardware	Time (s)
Femul agata-analysis-3	38209
Femul agata-analysis-5	11405
DAEmule cluster	3828

Limited by data access bandwidth

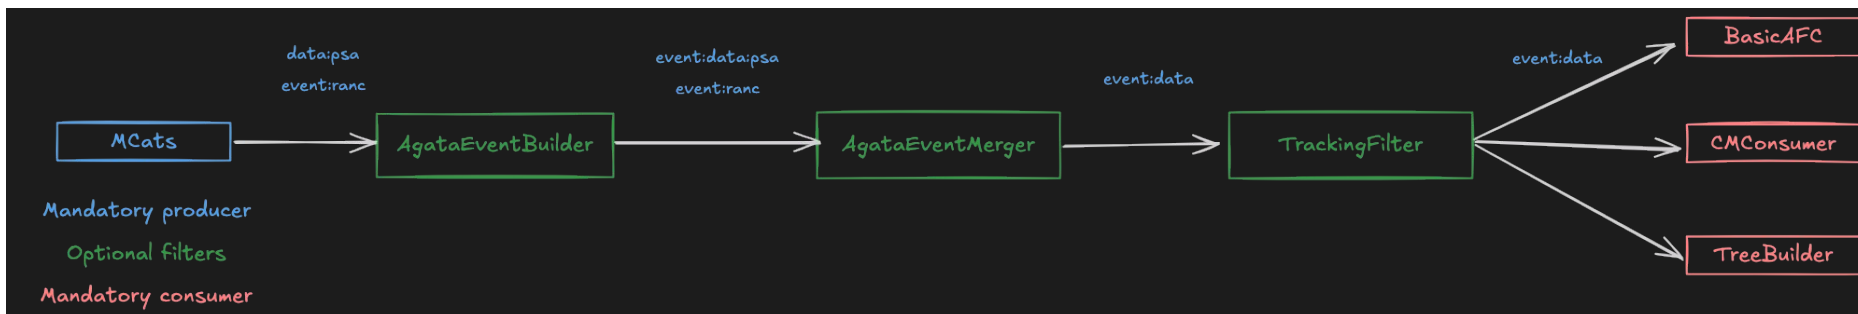


Global level processing

- New actors added to Agapro to manage global level :
 - *Multi-Channels Agata Time Sorter* (MCats) : N entries from Central Memory
→ One sorted output
 - *AgataEventBuilder* : from data:psa to event:data:psa
 - *AgataEventMerger* : from data:psa + event:ranc to event:data

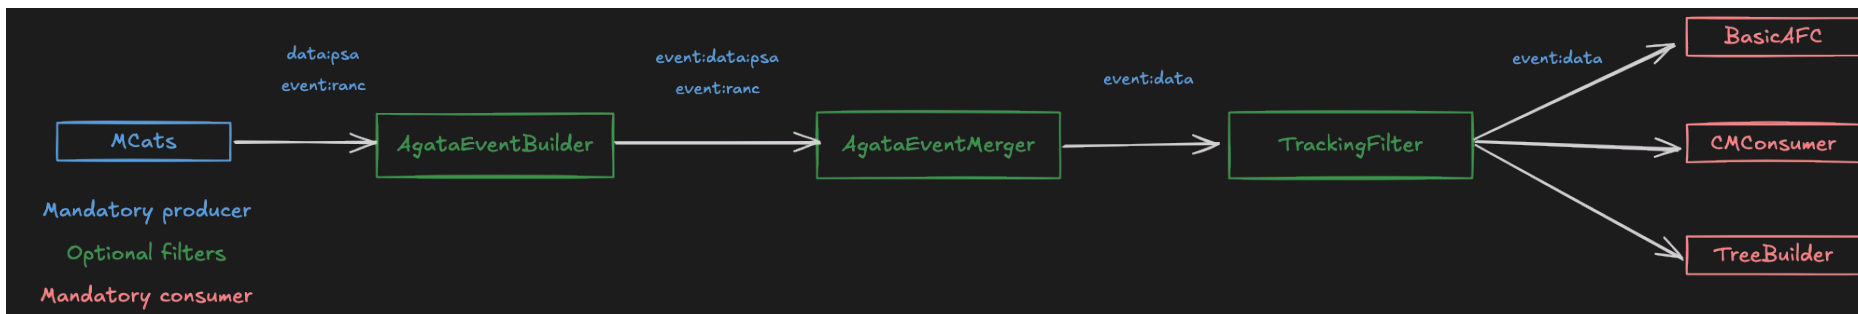
Global level processing

- New actors added to Agapro to manage global level :
 - *Multi-Channels Agata Time Sorter* (MCats) : N entries from Central Memory
→ One sorted output
 - *AgataEventBuilder* : from data:psa to event:data:psa
 - *AgataEventMerger* : from data:psa + event:ranc to event:data
- Different topology



Global level processing

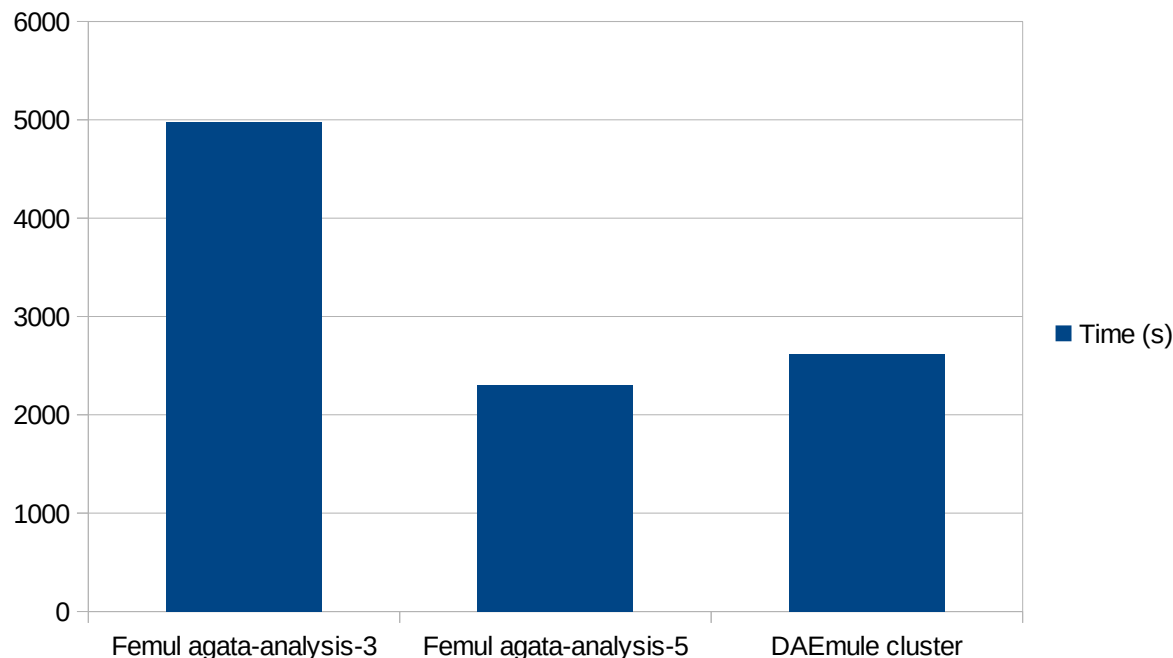
- New actors added to Agapro to manage global level :
 - *Multi-Channels Agata Time Sorter* (MCats) : N entries from Central Memory
→ One sorted output
 - *AgataEventBuilder* : from data:psa to event:data:psa
 - *AgataEventMerger* : from data:psa + event:ranc to event:data
- Different topology



...should be able to run online

First tests on analysis cluster @Legnaro

- Data from ancillary_exp_046_run_50 (31 crystals + ancillaries)



Hardware	Time (s)
Femul agata-analysis-3	4982
Femul agata-analysis-5	2301
DAEmule cluster	2621

DAEmule slower mainly because of timeouts handling
Still under test, coherent results :

```
#####
Number of events sent to tracking      42833573
Number of hits sent to tracking       175884340
Number of gammas out of tracking      61902124
Number of events with tracked gammas  38867883
#####
```

```
#####
Number of events sent to tracking      42833572
Number of hits sent to tracking       175884337
Number of gammas out of tracking      61902218
Number of events with tracked gammas  38868166
#####
```

Conclusion

- First version of an Agapro emulator able to run on clusters under test at Legnaro
- Local level processing running fine
- Global level processing under test – first results ok
- Need to improve integration with Docker Swarm and Portainer.io to ease usage
- Some documentation at <https://gbaulieu.pages.in2p3.fr/handbook-dev/binaries/DAEmule/>