

# Container Federation Use Cases

Ricardo Rocha <[ricardo.rocha@cern.ch](mailto:ricardo.rocha@cern.ch)>

CERN IT-CM-RPS

# Motivation for Federation

Periodic Load Spikes

*International Conferences, Reconstruction Campaigns*

Simplification

*Monitoring, Lifecycle, Alarms*

Deployment

*Uniform API, Replication, Load Balancing*



*AcctGroup = "ATLAS"*  
*JobPrio = 0*  
*RequestCpus = 2*  
*RequestMemory = 4260*  
 ...

*CERNEnvironment = "production"*  
*Datacenter = "meyrin"*  
*HasMPI = true*  
*TotalCpus = 8*  
*TotalMemory = 22500*  
 ...



*Matchmaking with ClassAds*

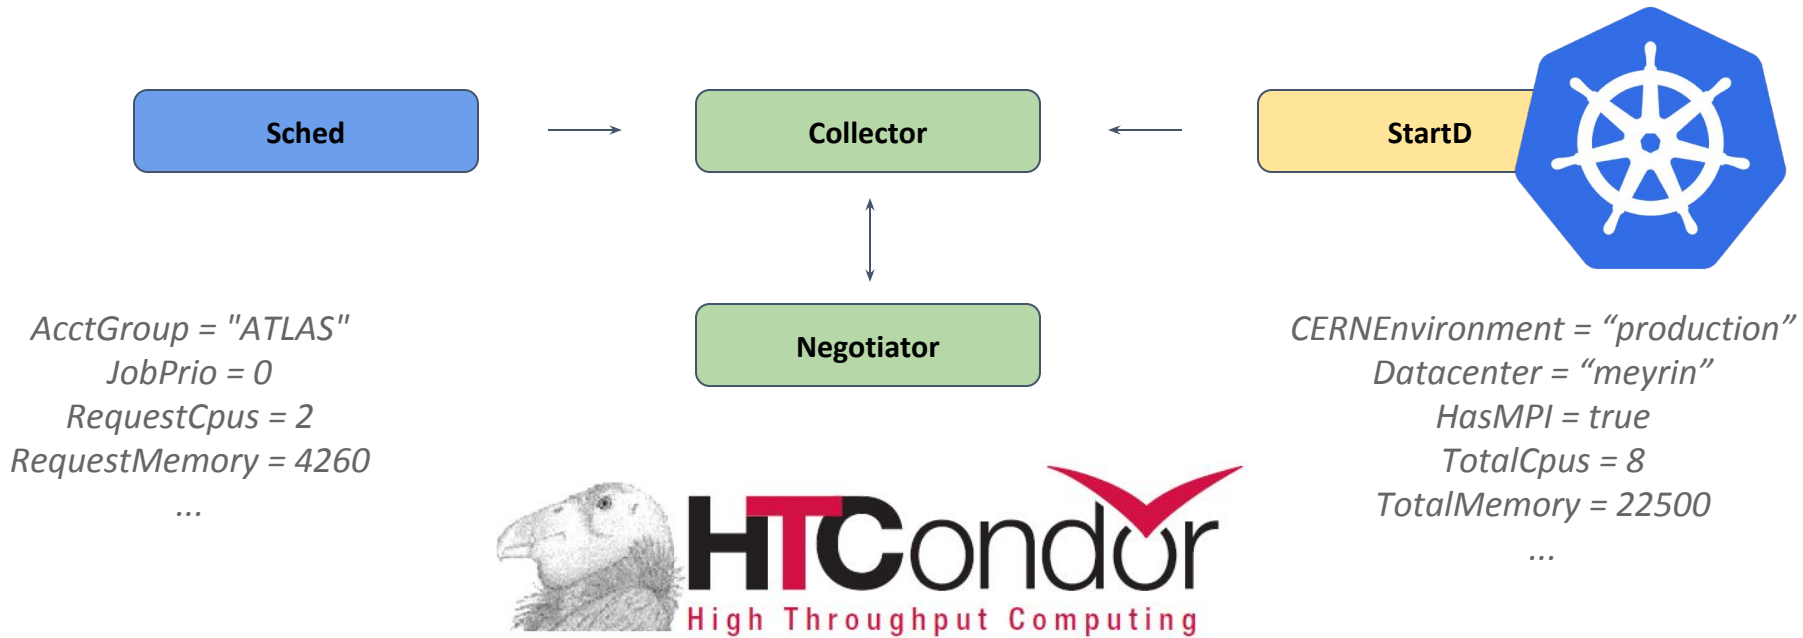
*Extensive Experience in HEP*

*Fair Share*

*Running Virtualized*

*Preemption*

*External Storage and Networking*



*Matchmaking with ClassAds*

*Fair Share*

*Preemption*

*Extensive Experience in HEP*

*Running Virtualized*

*External Storage and Networking*

**Host**



Sched

Collector




Negotiator

*kubefed init fed --host-cluster-context=condor-host ...*


StartD

...



StartD

...



*kind: DaemonSet*

...

*hostNetwork: true*

*containers:*

- *name: condor-startd*

*image: .../cloud/condor-startd*

*command: ["/usr/sbin/condor\_startd", "-f"]*

*securityContext:*

*privileged: true*

*livenessProbe:*

*exec:*

*command:*

- *condor\_who*

StartD

...



**T Systems**

**Host**



Sched

Collector



Negotiator

*kubefed init fed --host-cluster-context=condor-host ...*

*kubefed join --context fed tsystems \  
--host-cluster-context condor-host --cluster-context tsystems*

# REANA / RECAST

*Reusable Analysis Platform*

Workflow Engine (Yadage)

Each step a Kubernetes Job

Integrated Monitoring & Logging

Centralized Log Collection

Workflow Monitor for b907a91b-7f02-4b3e-afa3-e9298b888e2d

Status: **ACTIVE**

Workflow Visualization

[redraw](#) [relayout](#) [collapse](#) [expand](#)



Last seen: 2018-04-25 15:57:13

Log

Messages from the request processor will appear below.

```
2018-04-25T15:53:30. INFO:wflowbackend.process:acquiring context from wflowserver for wflowid b907a91b-7f02-4b3e-afa3-e9298b888e2d
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:running analysis on worker: wflow-nonint-b907a91b-7f02-4b3e-afa3-e9298b888e2d-76skn
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:acquiring wflow context from http://wflow-server.default.svc.cluster.local
2018-04-25T15:53:30. INFO:wflowbackend.process:processing ['setupFromURL', 'generic_onsuccess', 'cleanup', '--config-from-server', 'b907a91b-7f02-4-
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:running analysis on worker: wflow-nonint-b907a91b-7f02-4b3e-afa3-e9298b888e2d-76skn
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:setting up for context {'u'shipout_spec': {'u'suser': 'u'root', 'u'host': 'u'wflow-shiptarget.default-
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:prepared workdir workdirs/b9/07/a9/1b-7f02-4b3e-afa3-e9298b888e2d
2018-04-25T15:53:30. INFO:wflowbackend.backendtasks:preparing workdir workdirs/b9/07/a9/1b-7f02-4b3e-afa3-e9298b888e2d
2018-04-25T15:53:30. WARNING:wflowbackend.backendtasks:No input archive specified, skipping download
2018-04-25T15:53:30. INFO:wflowbackend.process:setting up entry point wflowyadageworker.backendtasks:run_workflow
2018-04-25T15:53:30. INFO:wflowyadageworker.backendtasks:running yadage workflow for context: {'u'shipout_spec': {'u'suser': 'u'root', 'u'host': 'u'wflow-
2018-04-25T15:53:30. INFO:wflowyadageworker.backendtasks:connecting backend: py:sharedstatekube:backend
2018-04-25T15:53:30. INFO:yadage.utils:setting up backend py:sharedstatekube:backend with opts {'optsyaml': '/yadageconfig/backendopts'}
2018-04-25T15:53:40. INFO:wflowyadageworker.backendtasks:yadage base settings {'updateinterval': 30.0, 'dataarg': 'workdirs/b9/07/a9/1b-7f02-4b3e-a-
2018-04-25T15:53:40. INFO:wflowyadageworker.simple_workflow:running workflow on context {'u'shipout_spec': {'u'suser': 'u'root', 'u'host': 'u'wflow-shipt-
2018-04-25T15:53:40. INFO:wflowyadageworker.simple_workflow:no workflow parameters in init dir
2018-04-25T15:53:40. INFO:wflowyadageworker.backendtasks:load workflow from source
```

# Summary

- Federation support in Kubernetes is ready
  - Ongoing development for the v2 API, with significant changes
  - Work with the community to have our use case taken into account
- Initial use cases
  - Started with a legacy application, limited integration
  - Expanded to a *cloud native* implementation, with great results
- Ongoing work to expand its usage
  - Use built-in “Kubernetes As a Service” capabilities
  - Investigations on how to best handle data distribution

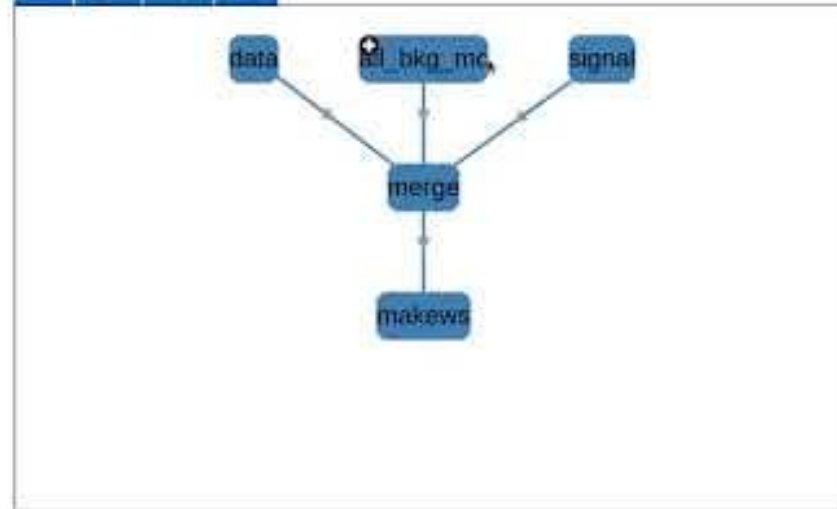


## Workflow Monitor for bf95b4cf-1baa-4320-b4c2-2a2bd5bc6cac

Status: **ACTIVE**

### Workflow Visualization

[refresh](#) [refresh](#) [collapse](#) [expand](#)



Let's see: 2018-05-25 22:37:05

<https://www.youtube.com/watch?v=jNyd97LiTXk>