

Déploiement d'un HTCONDOR-CE au CC

V. HAMAR

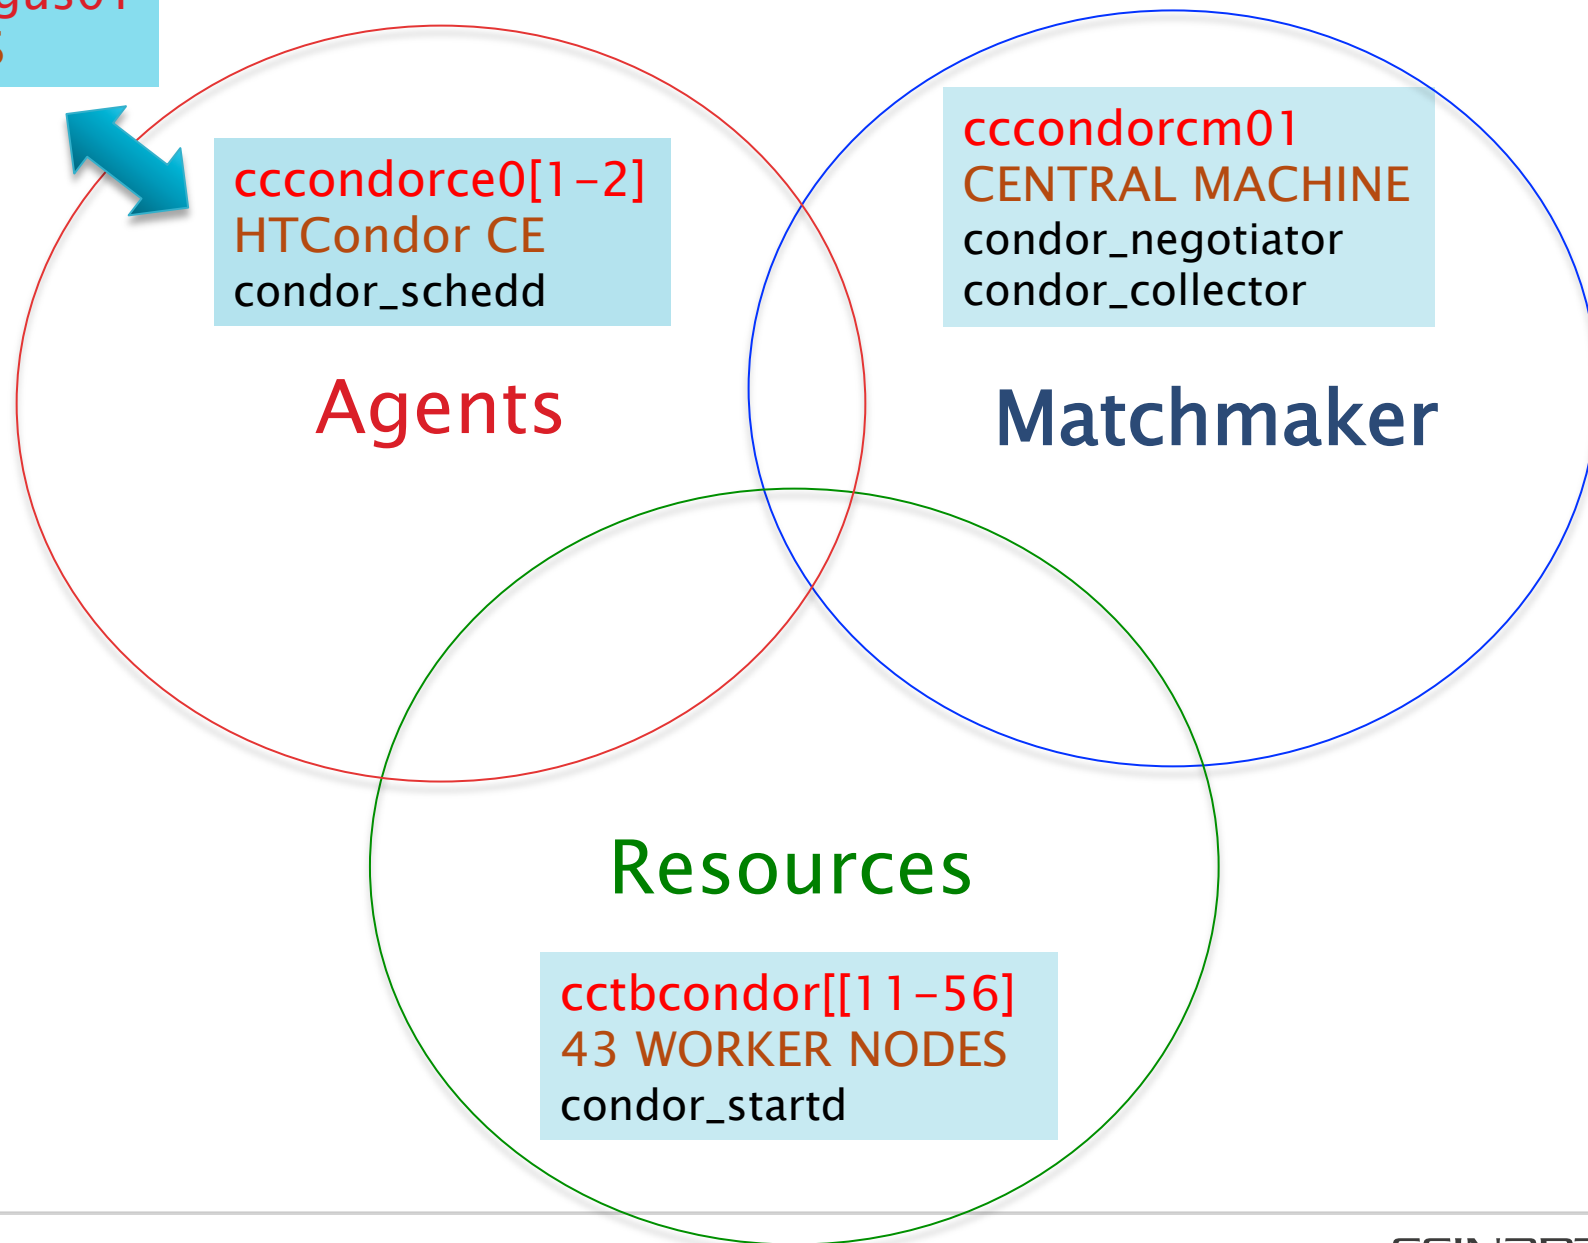
- ▶ Starting point
- ▶ Architecture
- ▶ Installation / Configuration
- ▶ Tests
- ▶ Future work

« POC - Comment se positionne HTCondor, à minima pour couvrir le périmètre de l'activité calcul (HTC, HPC, GPGPU) ? »

- ▶ Trying to:
 - Find an efficient and simple configuration
 - Be open to new ideas

- ▶ Before we know that CREAM-CE will be discontinued, the logic was:
- ▶ If we are going to evaluate condor ... why not to evaluate also HTCondor-CE ??

cctbargus01
ARGUS





UNIVERSITY OF OXFORD

CCIN2P3

University of BRISTOL

HTCondor configuration with puppet

Oliver Freyermuth, Luke Kreczko, Kashif Mohammad, Frederic Schaar, Peter Wienemann

cea

UNIVERSITÄT BONN

▶ Module htcondor

- <https://github.com/HEP-Puppet/htcondor>
- Latest commit [7062f1f](#) 22/05/2019

HTCondor Puppet module

- <https://github.com/HEP-Puppet/htcondor>
- Almost 10,000 downloads ([Puppet forge](#))
- Covers the **most common use cases**
 - Setup managers (w/o high-availability), schedulers, and workers
 - Kerberos authentication
 - Singularity configuration
 - Fully configurable via [hiera](#) (YAML files)
- Since 2.1.0: no (condor) knob left behind
 - Big thanks to @ccnifo

Table of Contents

1. Overview - What is the htcondor module?
2. Module Description - What does the module do?
3. Setup - The basics of getting started with htcondor
4. Singularity container support
5. Kerberos authentication support
6. Additional logging parameters
7. Additional custom parameters
8. Limitations - OS compatibility, etc.
9. Development - Guide for contributing to the module
 - Contributing to the htcondor module
 - Running tests - A quick guide

▶ Module htcondor-ce

- https://github.com/cernops/puppet-htcondor_ce
- Latest commit 8a982f2 on Dec 2, 2016 IainSteers Merge pull request #7 from IainSteers/master

▶ HTCondor-CE

- Version:
 - htcondor-ce-condor: 3.1.2-3
 - condor-8.6.12-0.445603
- Work in progress:
 - Automatic creation of grid users
 - Accounting
 - Apel
- Problem:
 - dependency of htcondor-ce-condor and condor versions to upgrade.

Added Grid Universe support for Azure, SLURM, Cobalt

- ▶ Speak to **Microsoft Azure**
- ▶ Speak **native SLURM protocol**
- ▶ Speak to **Cobalt Scheduler**
 - Argonne Leadership Computing Facilities



Also HTCondor-CE "native" package

- › HTCondor-CE started as an OSG package
- › IN2P3 wanted HTCondor-CE without all the OSG dependencies....
- › Now HTCondor-CE available stand-alone in HTCondor repositories

▶ Parameterising the module with Hiera

```
# HTCONDOR CE
htcondor_ce::pool_collectors:
  - 'cccondorc01.in2p3.fr'
htcondor_ce::condor_view_hosts: []
htcondor_ce::ce_version: '3.1.2-3.el7'
htcondor_ce::lrms_version: '8.6.12-0.445603.el7'
htcondor_ce::uid_domain: 'in2p3.fr'
htcondor_ce::gsi_regex: '^VO\=GRID-FR\C\=FRVO\=CNRSVOU\=CC-IN2P3\CN\=([A-Za-z0-9.\-]*)$'
htcondor_ce::gsi_backend: 'argus'
htcondor_ce::argus_server: 'cctbargus01.in2p3.fr'
htcondor_ce::argus_port: 8154
htcondor_ce::argus_resourceid: 'http://cc.in2p3.fr/ce'
htcondor_ce::use_static_shadow: false

# BDII
htcondor_ce::install_bdii: true
htcondor_ce::supported_vos:
  - atlas
  - cms
  - dteam
  - ops
htcondor_ce::goc_site_name: 'IN2P3-CC'
htcondor_ce::benchmark_result: '10.26-HEP-SPEC06'
htcondor_ce::execution_env_cores: 32
```


- ▶ 01-ce-auth.conf
- ▶ 01-ce-router.conf <- Hold, remove jobs
- ▶ 01-common-auth.conf <- Common authorisation
- ▶ 02-ce-condor.conf <- Job router to vanilla
- ▶ 03-managed-fork.conf <- Limit number of jobs
- ▶ 06-ce-bdii.conf <- VOs
- ▶ 59-site-security.conf <- Autorisation methods
- ▶ 60-configured-attributes.conf <- Statistics by DN
- ▶ 61-job-routes.conf <- set environment variables
- ▶ 99-disablelegsicache.conf <- Remove argus cache

▶ Grid Universe

- Authentication GSI ✓
- Users mapped
 - by hand ✓
 - using ARGUS
 - Dynamic pool accounts ✓
 - Fix mapping (atlas100) ✓
- Interfaces used:
 - Atlas pilot factory: atlas ✓
 - CMS ✓
 - DIRAC: vo.france-grilles.fr, vo.formation.idgrilles.fr ✓
- cgroups
 - Memory ✓
 - CPU ✓
- Job types:
 - Sequential ✓
 - Multicores ✓

▶ Problems:

- Mapping: errors remains after CE reconfiguration

▶ Solution:

- GSS_ASSIST_GRIDMAP_CACHE_EXPIRATION=0 added to CE configuration

▶ BDII

- CERN RPM (htcondor-ce-bdii) : OK

```
ldapsearch -LLL -x -h ccosvm1009.in2p3.fr:2170 -b "o=glue"  
dn: o=glue  
objectClass: organization  
o: glue
```

```
dn: GLUE2GroupID=grid,o=glue  
objectClass: GLUE2Group  
GLUE2GroupID: grid
```

```
dn: GLUE2GroupID=resource,o=glue  
objectClass: GLUE2Group  
GLUE2GroupID: resource
```

```
dn: GLUE2ServiceID=cctbcondor10.in2p3.fr,GLUE2GroupID=resource,o=glue  
GLUE2ServiceID: cctbcondor10.in2p3.fr  
objectClass: GLUE2Entity  
objectClass: GLUE2Service  
objectClass: GLUE2ComputingService  
GLUE2ServiceType: org.opensciencegrid.htcondorce  
GLUE2ServiceQualityLevel: production  
GLUE2ServiceComplexity: endpointType=2, share=3, resource=1  
GLUE2ServiceCapability: executionmanagement.jobexecution  
GLUE2EntityName: Computing Service cctbcondor10.in2p3.fr  
GLUE2ServiceAdminDomainForeignKey: IN2P3-CC
```

```
dn: GLUE2ManagerID=cctbcondor10.in2p3.fr_Manager,GLUE2ServiceID=cctbcondor10.i  
n2p3.fr,GLUE2GroupID=resource,o=glue  
GLUE2ComputingManagerComputingServiceForeignKey: cctbcondor10.in2p3.fr  
GLUE2ManagerProductName: HTCondor  
objectClass: GLUE2Entity  
objectClass: GLUE2Manager  
objectClass: GLUE2ComputingManager  
GLUE2ComputingManagerTotalPhysicalCPUs: 800  
GLUE2ManagerProductVersion: 8.7.8  
GLUE2ComputingManagerTotalLogicalCPUs: 800  
GLUE2ManagerServiceForeignKey: cctbcondor10.in2p3.fr  
GLUE2ManagerID: cctbcondor10.in2p3.fr_Manager
```

```
dn: GLUE2ResourceID=cctbcondor10.in2p3.fr_standard,GLUE2ServiceID=cctbcondor10.i  
.in2p3.fr,GLUE2GroupID=resource,o=glue  
GLUE2ExecutionEnvironmentPlatform: amd64  
GLUE2ExecutionEnvironmentTotalInstances: 0  
GLUE2ExecutionEnvironmentLogicalCPUs: 8  
GLUE2ResourceManagerForeignKey: cctbcondor10.in2p3.fr_Manager  
objectClass: GLUE2Entity  
objectClass: GLUE2Resource  
objectClass: GLUE2ExecutionEnvironment  
GLUE2ExecutionEnvironmentCPUMultiplicity: singlecpu-multicore  
GLUE2ExecutionEnvironmentOSFamily: linux  
GLUE2ExecutionEnvironmentConnectivityOut: TRUE  
GLUE2ExecutionEnvironmentMainMemorySize: 2000  
GLUE2ExecutionEnvironmentVirtualMemorySize: 2048  
GLUE2ResourceID: cctbcondor10.in2p3.fr_standard  
GLUE2ExecutionEnvironmentConnectivityIn: TRUE  
GLUE2ExecutionEnvironmentPhysicalCPUs: 8  
GLUE2ExecutionEnvironmentComputingManagerForeignKey: cctbcondor10.in2p3.fr_Man  
ager
```


Algorithm for calculating the Site and Service Availability

cccondorce01.in2p3.fr [Link to data](#)

No data returned by DB query. Test history ccosvms0201.in2p3.fr using ATLAS_CRITICAL


Test history cccondorce02.in2p3.fr using ATLAS_CRITICAL

720 hours from 2019-04-23 17:00 to 2019-05-23 17:25



Test history cccondorce01.in2p3.fr using ATLAS_CRITICAL

720 hours from 2019-04-23 17:00 to 2019-05-23 17:25




Algorithm for calculating the Site and Service Availability

cccondorce01.in2p3.fr [Link to data](#)


Test history cccondorce02.in2p3.fr using CMS_CRITICAL

720 hours from 2019-04-23 17:00 to 2019-05-23 17:27



Test history cccondorce01.in2p3.fr using CMS_CRITICAL

720 hours from 2019-04-23 17:00 to 2019-05-23 17:27



- ▶ Accounting – APEL
- ▶ BDII publication verification
- ▶ Monitoring (Nagios, Kibana)
- ▶ Configuration improvements
- ▶ Publish in production the HTCondor-CE
- ▶ Add Alice and LHCb VOs



- ▶ <https://bbockelm.github.io/docs/compute-element/htcondor-ce-overview/>
- ▶ <https://wlcg-ops.web.cern.ch/htcondor-ce>
- ▶ <https://twiki.cern.ch/twiki/bin/viewauth/LCG/HtCondorCeAccounting>