

Cloud as resource for DIRAC

A. Tsaregorodtsev

CPPM-IN2P3-CNRS, France,

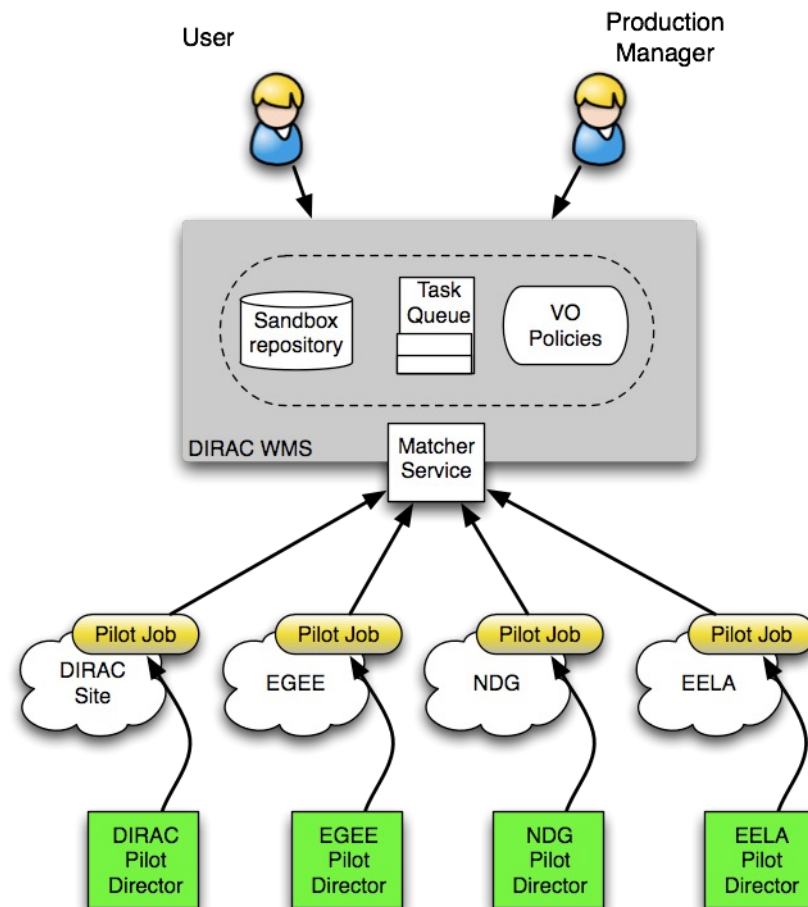
20th Feb 2024,

Jennifer2 Workshop



- ▶ DIRAC WMS
- ▶ CloudComputingElement
- ▶ Pilots in VMs
- ▶ VM life cycle
- ▶ Conclusion

- ▶ Pilot jobs are submitted to computing resources by specialized Pilot Directors
- ▶ Pilots retrieve user jobs from the central Task Queue and steer their execution on the worker nodes including final data uploading
- ▶ Pilot based WMS advantages:
 - ▶ increases efficiency of the user job execution
 - ▶ allows to apply efficiently community policies at the Task Queue level
 - ▶ allows to integrate heterogeneous computing resources

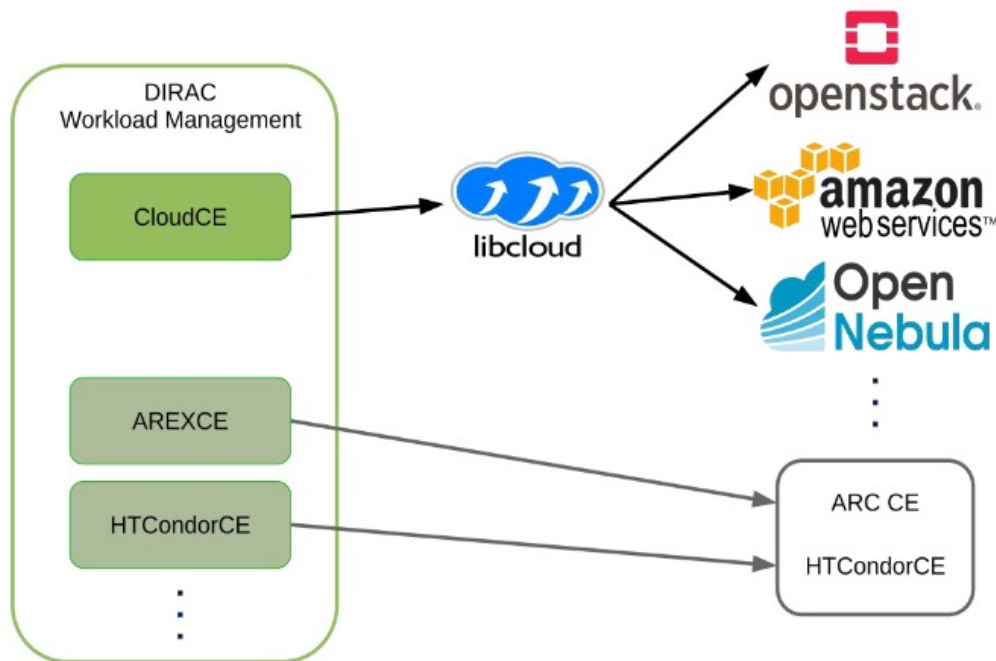


- ▶ DIRAC uses ComputingElement plugins to access different types of computing resources
 - ▶ Encapsulate access protocol (HTCondor, ARC, SSH)
 - ▶ Manages the pilot life cycle

- ▶ Recently CloudComputingElement was developed by the GridPP team (*Daniela, Simon*)
 - ▶ Access clouds in a similar way as other computing resources
 - ▶ Documentation:
<https://dirac.readthedocs.io/en/latest/CodeDocumentation/Resources/Computing/CloudComputingElement.html>

- ▶ CloudComputingElement is based on the Apache libcloud
 - ▶ <https://libcloud.apache.org>
- ▶ Apache libcloud is an open source collection of python based cloud interfaces, maintained by the Apache foundation
 - ▶ General cloud access functionality (computing and storage)
 - ▶ Driver plugins for (almost) all public and private clouds

- ▶ Inherits the DIRAC ComputingElement interface
 - ▶ Works with the standard pilot factory (SiteDirector)
- ▶ The pilot payload script and data are added as instance metadata in **cloud-init** format
- ▶ Any VM image containing **cloud-init** to decode and start the DIRAC pilot bootstrap scripts.



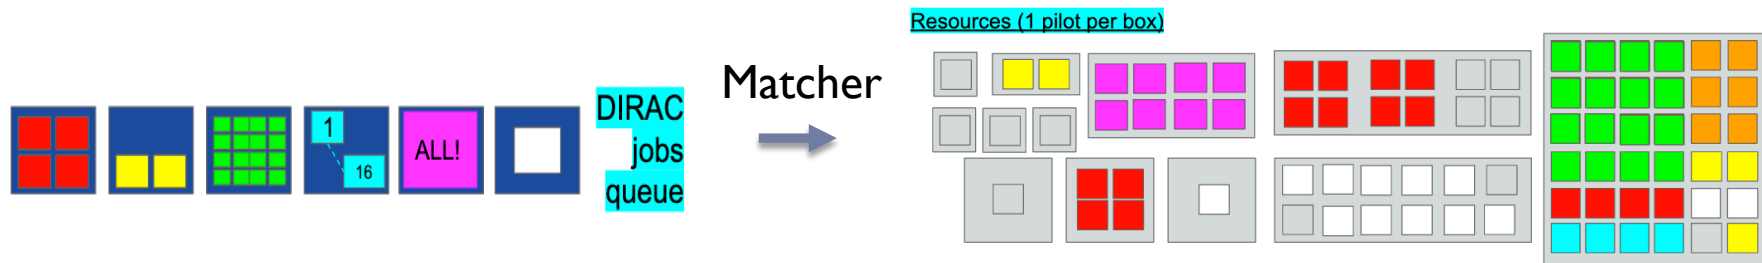
```
└─ federation.cstcloud.cn
  ├── CType = Cloud
  ├── LocalCType = Pool/InProcess
  ├── CloudType = OPENSTACK
  ├── NumberOfProcessors = 4
  ├── Driver_ex_force_auth_url = https://federation.cstcloud.cn:5000
  ├── Driver_ex_force_auth_version = 3.x_appcred
  ├── Driver_ex_domain_name =
  ├── Driver_ex_force_service_region = RegionOne
  ├── Instance_Image = name:CentOS7-image
  ├── Instance_Flavor = name:egi-integration-system-flavor
  ├── Instance_SSHKey = DIRAC-Cloud
  ├── Instance_Networks = name:public_network
  └─ Queues
     ├── eiscat-queue
     └─ enmr-queue
        ├── NumberOfProcessors = 4
        ├── VO = enmr.eu
        ├── CPUTime = 9999999
        ├── MaxTotalJobs = 10
        ├── CloudAuth = /vo/dirac/etc/cloud/enmr.eu.auth
        ├── Driver_ex_tenant_name = enmr_project
        └── Context_ExtPackages = csh, time
```

- ▶ Access point
- ▶ Auth method
 - ▶ Password, Application Credentials
- ▶ Project/Tenant
- ▶ VM parameters
 - ▶ Image, flavor
- ▶ Max number of VMs to create
- ▶ Extra software to install
 - ▶ CVMFS is always installed

- ▶ **Openstack authentication**
 - ▶ Login/password
 - ▶ Certificates
 - ▶ Application Credentials

- ▶ **Accessing clouds with tokens recipe**
 - ▶ Login to the cloud dashboard with a token
 - ▶ Choose the proper project
 - ▶ Set up Application Credentials
 - ▶ Store the AppCred ID and secret in the pilot factory configuration

- ▶ Pilots exploit multi-core VMs using **PoolICE** “inner” Computing Element
 - ▶ On-WN batch system
 - ▶ Flexible strategy with prioritized job requests to the Matcher, e.g.:
 - ▶ First, ask for jobs requiring WholeNode tag
 - ▶ If none, ask for jobs requesting as many cores as available
 - ▶ If none, ask for jobs with MultiProcessor requirement
 - ▶ If none, ask for single-core jobs
 - ▶ The goal is to fill the nodes with payloads fully exploiting there multi-core capacity



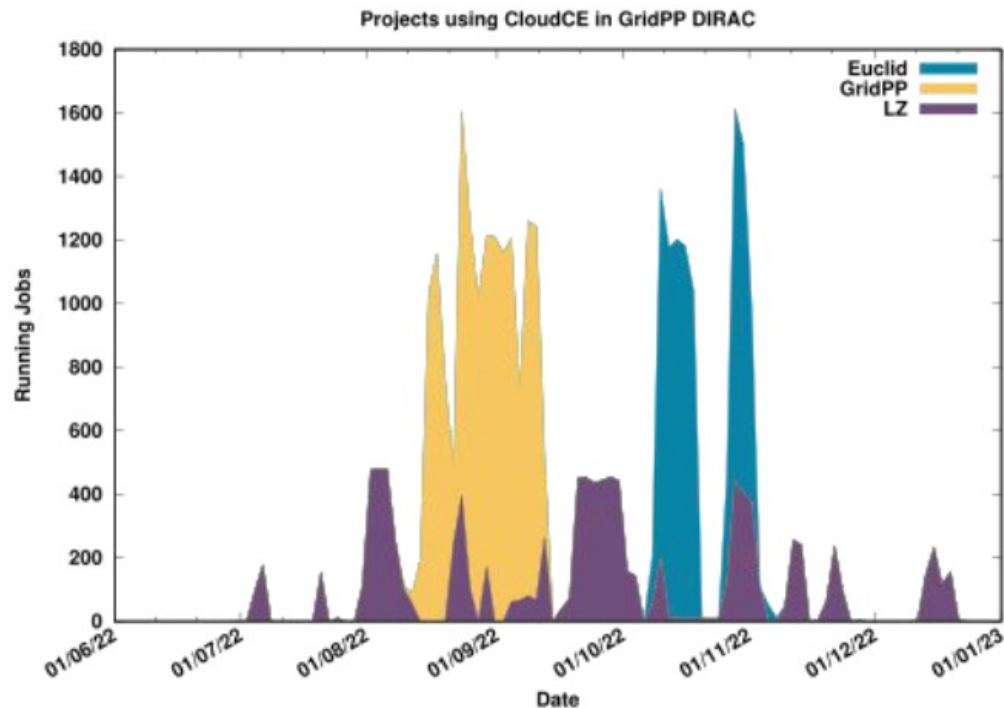
- ▶ User payloads execution options (configurable)
 - ▶ Execution by a pilot user process
 - ▶ **InProcessCE** « inner » Computing Element
 - ▶ No barriers between parallel user jobs
 - ▶ No barriers between the user job and the pilot
 - ▶ Execution in separate **Singularity** containers
 - ▶ Good payloads separation

- ▶ The pilot life cycle finishes when
 - ▶ The limit of the number of user jobs is reached
 - ▶ The limit of fruitless attempts to get user jobs is reached
 - ▶ Typically 10 attempts within 20 minutes
 - ▶ The hard time limit for the VM is reached
 - ▶ 2 weeks by default
 - ▶ Stuck VM will be cleaned when reaching this limit

- ▶ Pilot logs are not easily available right now
 - ▶ No inbound connectivity to VMs from the DIRAC services
 - ▶ One has to log in to the running VM for debugging
 - ▶ Work in progress on the system to push pilot log messages to a central service

- ▶ Pilots use credentials to connect to the DIRAC central service
 - ▶ X.509 proxy certificates
- ▶ No mechanism for the pilot proxy renewal
 - ▶ Pilots are instrumented with proxies as long as the VM time limit, e.g. 2 weeks
- ▶ The user payload proxies are renewed by pilots as needed

- ▶ Used quite intensively in GridPP
- ▶ Few cloud sites in EGI
 - ▶ 15 configured, ~6 actually used
 - ▶ Not so many resources available through clouds
 - ▶ Available sites are quite stable



See Daniela's presentation at the
DIRAC&Rucio Workshop, KEK, 2023
<https://indico.cern.ch/event/1252369/>

- ▶ CloudComputingElement works well for Openstack clouds
 - ▶ With Application Credentials set up for each cloud
- ▶ OpenNebula cloud access is being tested
 - ▶ Other clouds can be tried out as soon as this will be requested by users
- ▶ Centralized pilot logging is in the certification phase
 - ▶ Will help all the pilots, not only those in the clouds