DIRAC and CLouds

A.Tsaregorodtsev,

CPPM, Marseille

Outline

- Pilot based WMS
- Application to Cloud resources
- VMDIRAC components
- Status and outlook
- Conclusions

DIRAC WMS

- Jobs are submitted to the DIRAC Central Task Queue with credentials of their owner (VOMS proxy)
- Pilot Jobs are submitted by specific Directors to a Grid WMS with credentials of a user with a special Pilot role
- The Pilot Job fetches the user job and the job owner's proxy
- The User Job is executed with its owner's proxy used to access SE, catalogs, etc

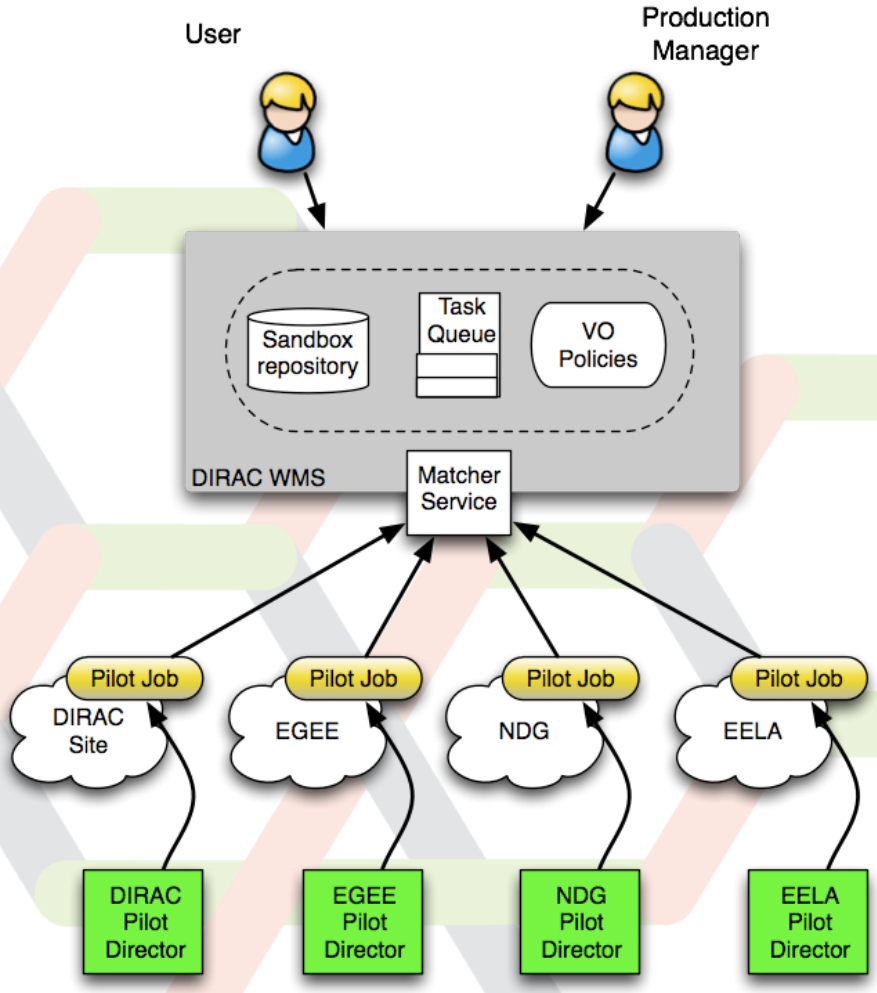

WMS: using heterogeneous resources

- Including resources in different grids and standalone clusters is simple with Pilot Jobs
  - Needs a specialized Pilot Director per resource type
  - Users just see new sites appearing in the job monitoring

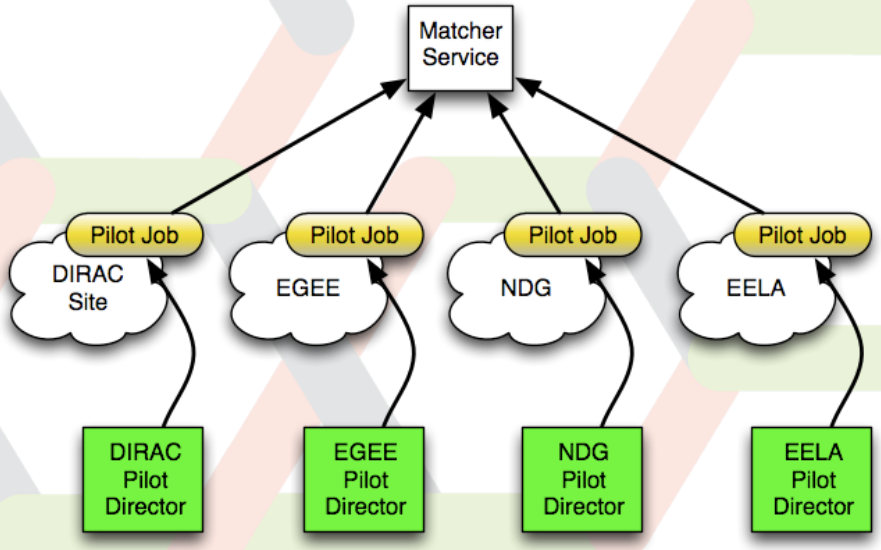

Clouds

- VM scheduler initially developed for the Belle MC production system
  - Dynamic VM spawning taking Amazon EC2 spot prices and Task Queue state into account
  - Discarding VMs automatically when no more needed
- The VM at boot time starts the Pilot Agent
  - This makes the instantiated VMs behave as any other WN with respect to the DIRAC WMS

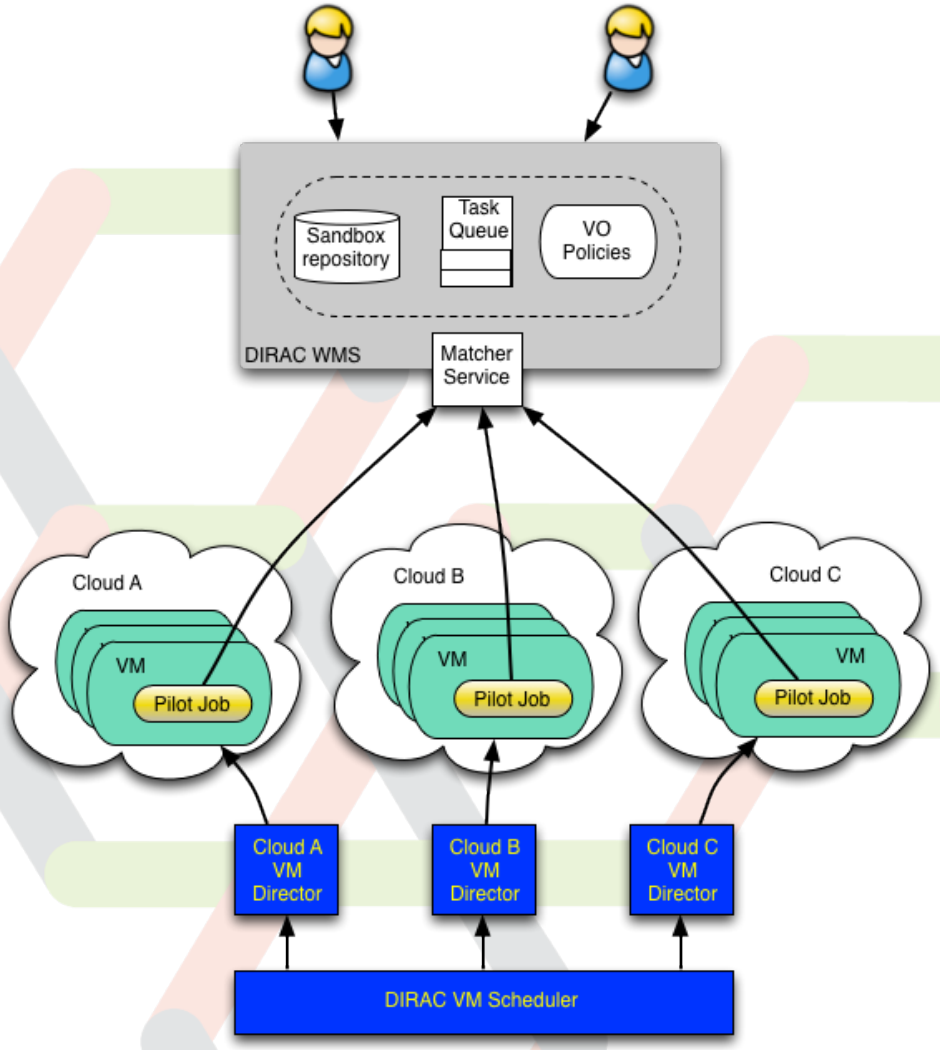

Contextualization

- Standard VM image, e.g. CERNVM, is complemented by a context image
  - Containing the Pilot software
  - Security token, e.g. pilot certificate (proxy)
- The contextualization mechanism allows to configure the VM to start the Pilot script at boot time
  - Provides also the necessary configuration parameters to the pilot
- The contextualization mechanism is rather complicated
  - Different for different cloud managers

Standalone computing clusters

- Dedicated Pilot Director per group of sites
- Off-site Director
  - Site delegates control to the central service
  - Site must only define a dedicated local user account
  - The payload submission through the SSH tunnel
- The site can be a single computer or a cluster with a batch system
  - LSF, BQS, SGE, PBS/Torque, Condor
  - More to come:
    - OAR, SLURM, LoadLeveler. etc
- The user payload is executed with the owner credentials
  - No security compromises with respect to external services

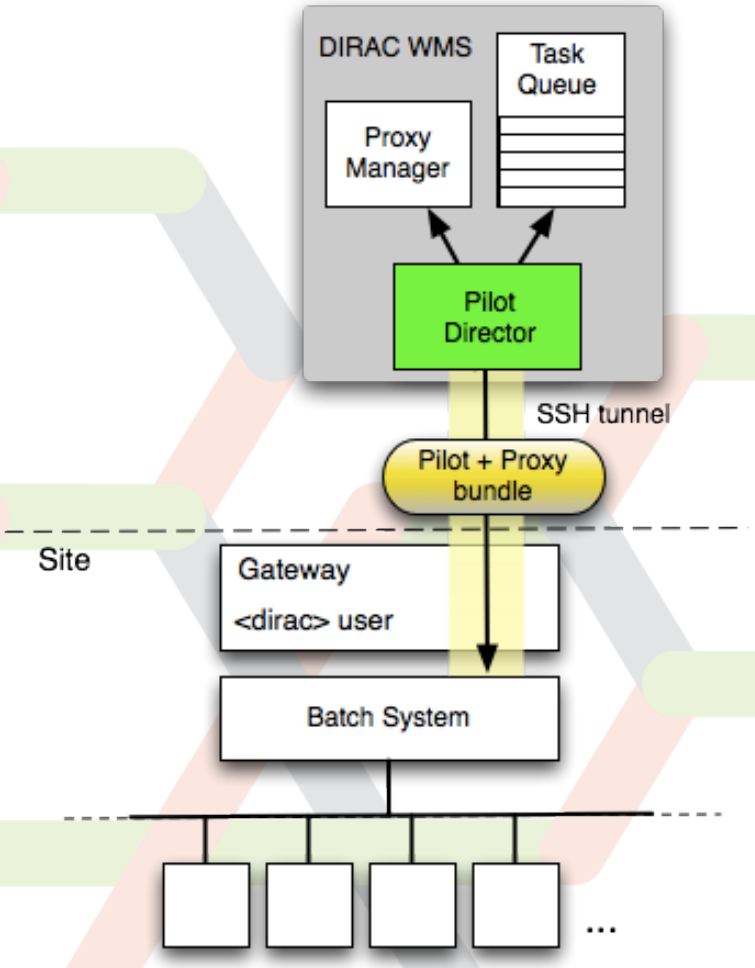

SSH sites

- SSHComputingElement (*L.Dimitriu*)
  - Multiple IP addresses can be specified
  - Number of slots per IP can be specified
  - Share the same access credentials
    - Public key
- Allows to create a computing cluster even without a batch system
- Can be useful together with Cloud resources:
  - VMs are often accessible with public ssl keys

Simple contextualization

- VMs are instantiated with the simplest contextualization
  - Just user account with the public key for ssh login
- The IPs of the VMs are collected to form a DIRAC Site of type SSH
  - Entered dynamically into the DIRAC CS
- Standard SSH Director starts to send pilots

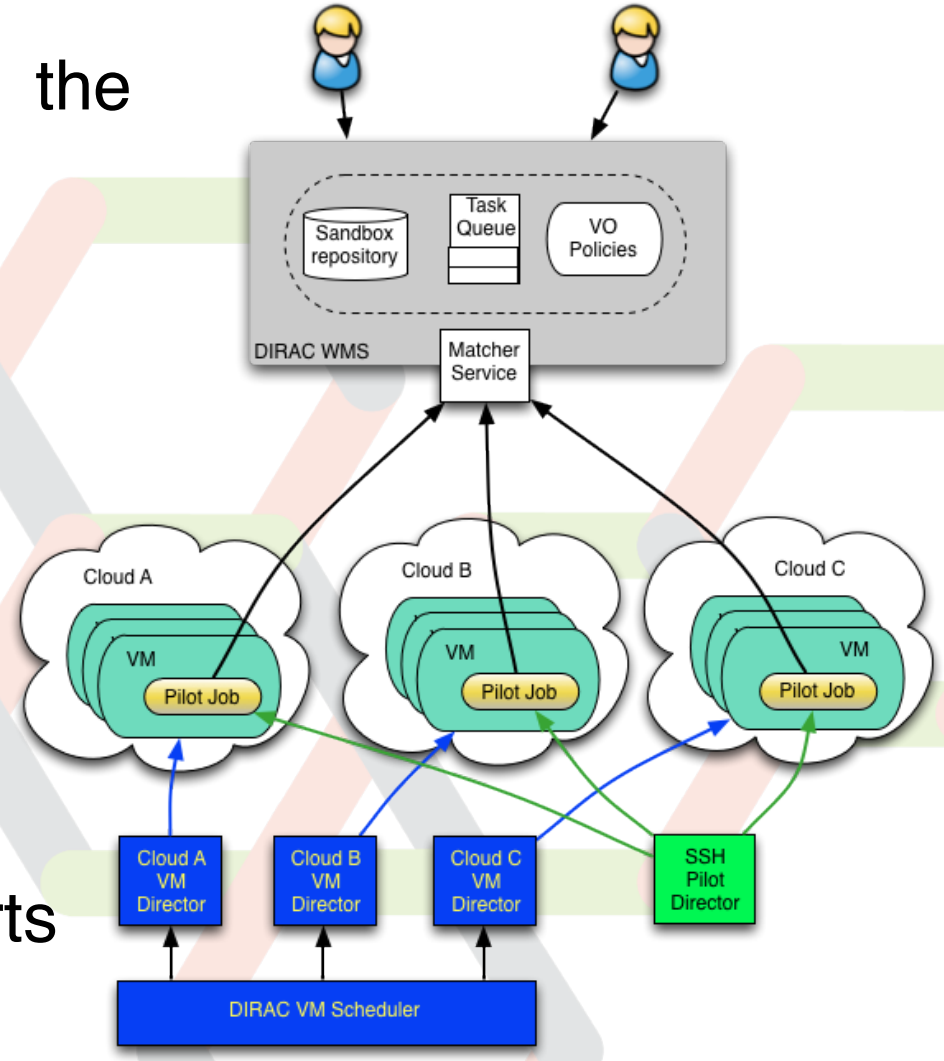

Simple contextualization

- Advantages of simple contextualization
  - No complicated contextualization ;-)
  - No special requirements for the VM images
  - No time limit for the security token
- Disadvantages
  - VMs must have public IPs
    - At least in the same network as the corresponding Pilot Director

Current status *(V.Mendes, V.Fernandez)*

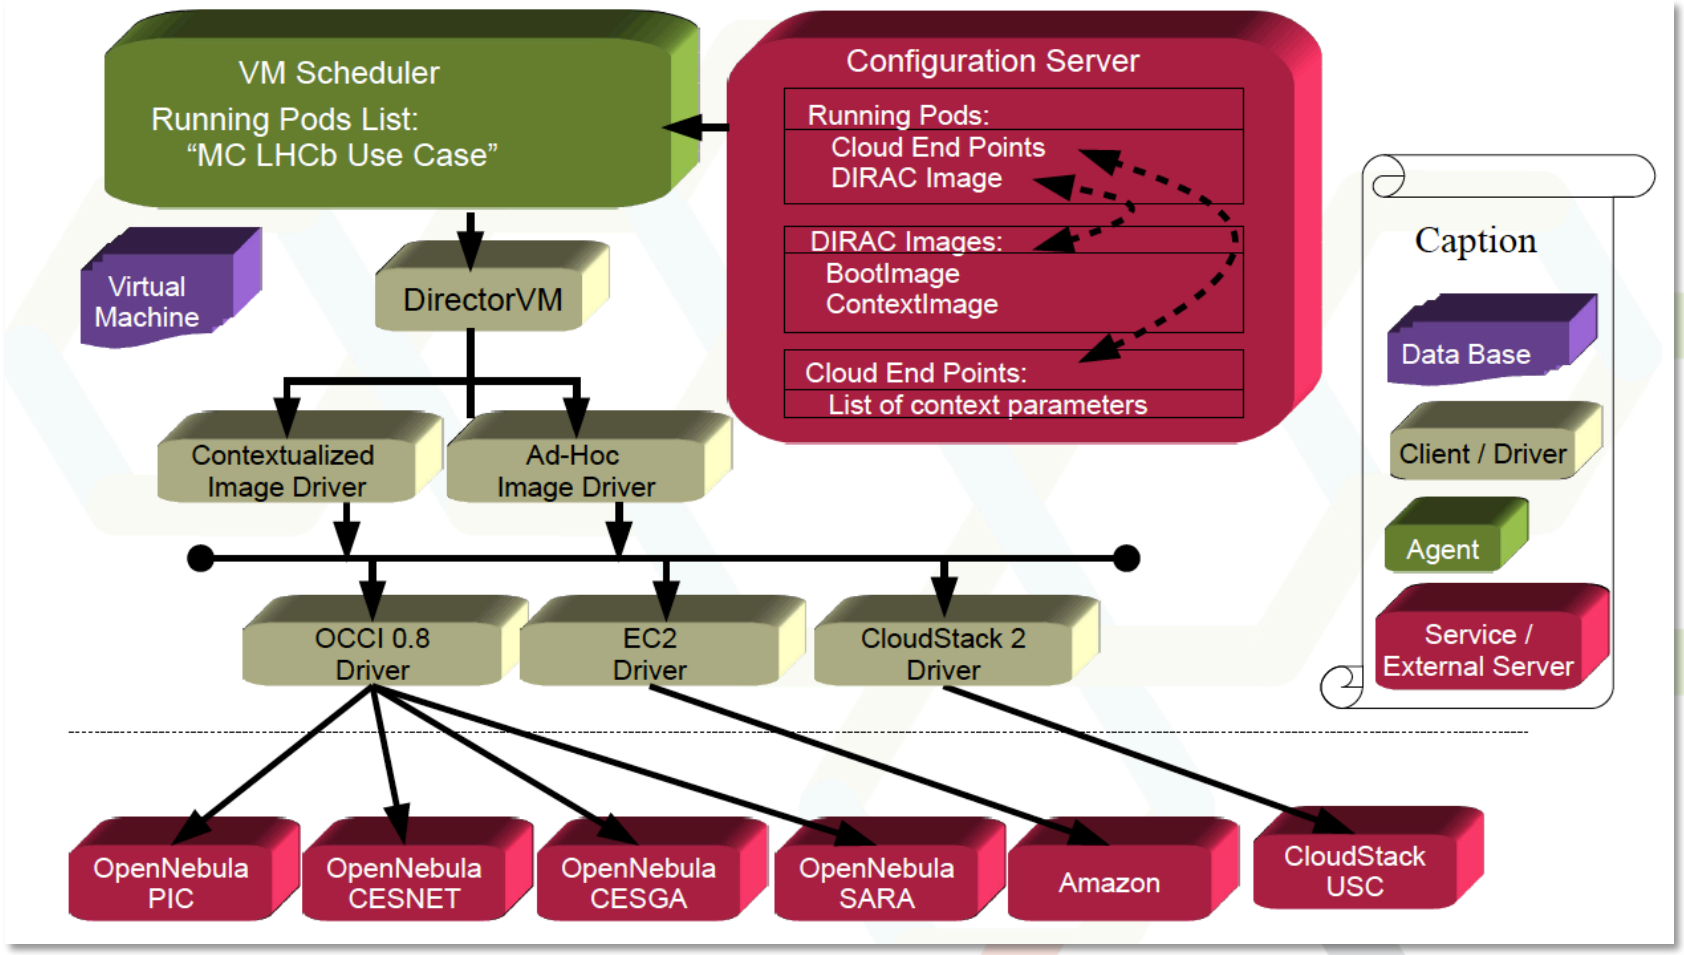

VMDIRAC components

- VMDIRAC v0r4 (released)
  - VMManager service – frontend to the VM Database
    - keeps track of VM instances
  - VM Scheduler triggers new VM instantiations according to the state of the WMS Task Queue, cloud resources availability, etc
  - VM Monitor – follows the VM status, stops idle VMs
  - VM Director instantiates VMs using specific cloud plugins:
    - OCCI 0.8 Driver (OpenNebula)
    - EC2 Driver (Amazon)
    - CloudStack 2 Driver
- VMDIRAC v0r5 (next)
  - OCCI 1.1 Driver (rOCCI)
  - EC2 contextualization

Tests

- Belle MC Production using Amazon EC2 resources (2010)
- LHCb MC Production using multiple OpenNebula (CESNET, CESGA, SARA) and CloudStack (USC) cloud end-points
  - Regular LHCb Gauss jobs 100 events each
  - High success rate – 96%
  - Large latencies in VM submissions (hours)

Work in progress

- OCCI 1.1 Driver (OpenNebula, OpenStack)
- CloudStack 3 Driver
- Stratuslab Marketplace incorporation as a common VM image repository
- Participation in the Federated Cloud EGI Task Force
  - With the LHCb MC Production use case

Conclusions

- WMS with Pilot Jobs offers a simple and efficient mechanism to provide heterogeneous resources transparently to users
- Cloud resources can be incorporated with different mechanisms for running Pilots
  - R&D activity is ongoing
- Work ongoing to exploit new cloud interfaces, Federated Cloud services (Marketplace, Info systems, etc)