

Cloud access with DIRAC

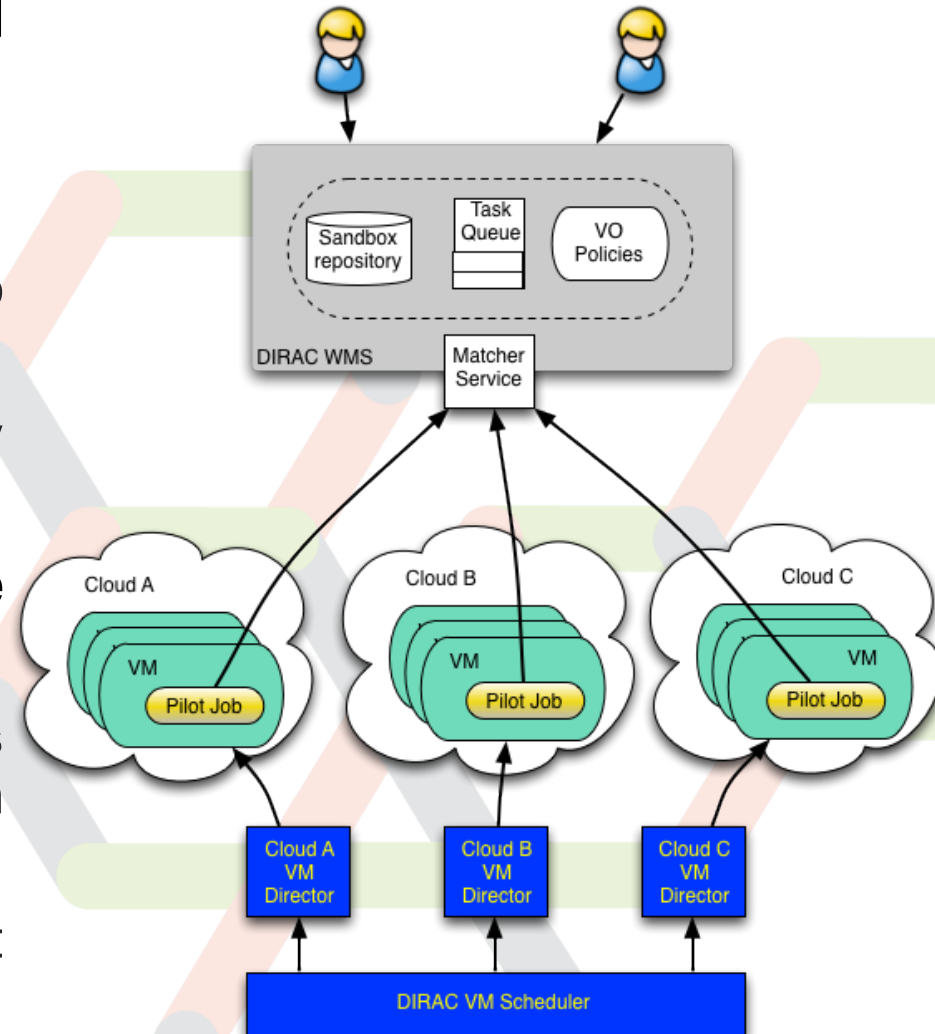
A.Tsaregorodtsev,
CPPM, Marseille
LCG-France, 28 March, LLR



- ▶ VMDIRAC
- ▶ Federated cloud tests
- ▶ Prospects for the french cloud federation
- ▶ Conclusions

- ▶ This is the work mostly done by Victor Mendez (PIC), Victor Fernandez (USC), Mathieu Puel (CC/IN2P3)

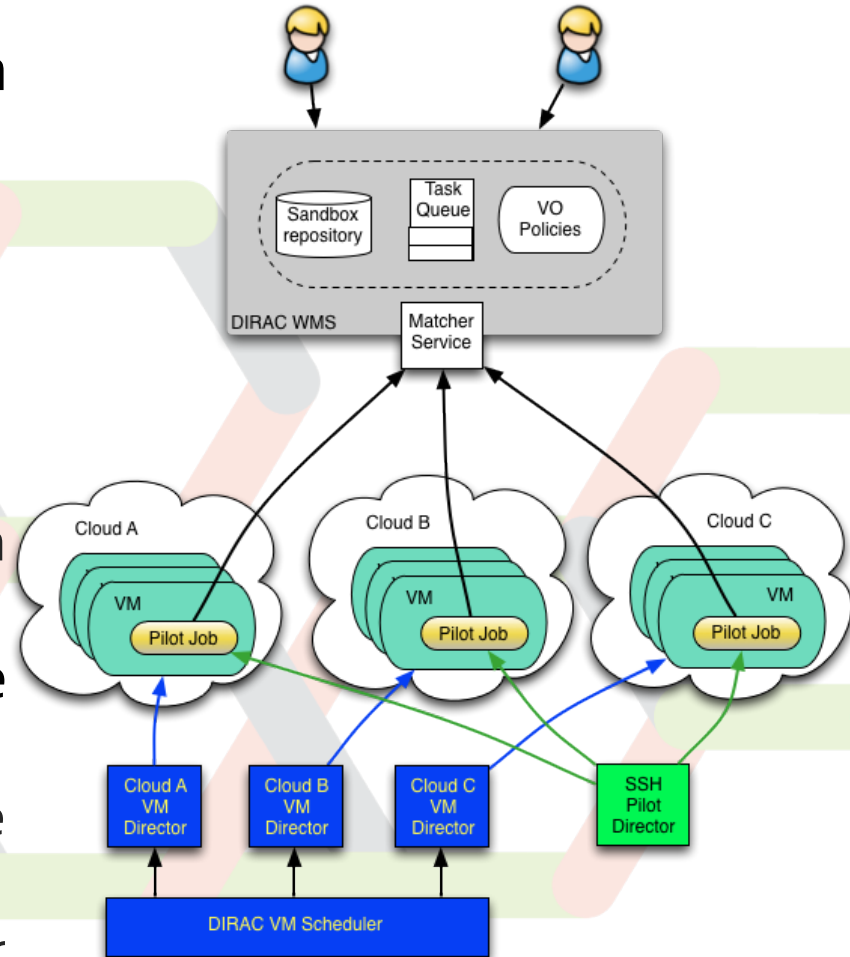
- ▶ VM scheduler initially developed for the Belle MC production system (VMDIRAC)
- ▶ Dynamic VM spawning taking the Task Queue state into account
- ▶ Discarding VMs automatically when no more needed
- ▶ The VM at boot time starts the “Pilot Job”
- ▶ This makes the instantiated VMs behave as any other WN with respect to the DIRAC WMS
- ▶ JobAgent and VMMonitorAgent are started inside the “Pilot Job”

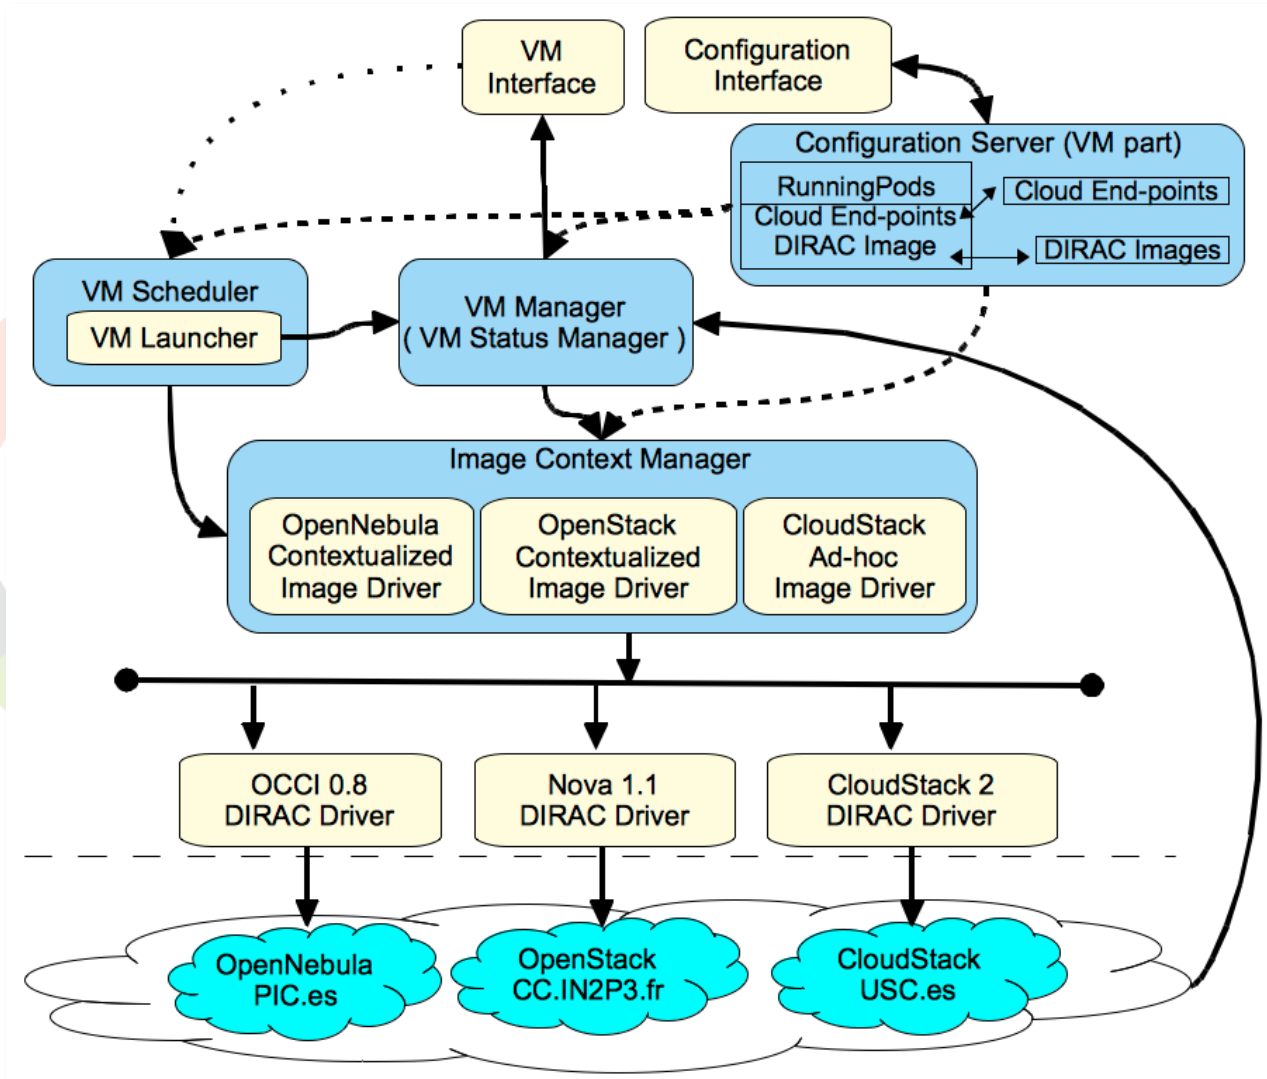


- ▶ Standard (golden) VM image, e.g. CERNVM, is complemented by a context image
 - ▶ Containing the DIRAC software
 - ▶ Security token, e.g. VM host certificate
- ▶ The contextualization mechanism allows to configure the VM to start the Pilot script at boot time
 - ▶ Provides also the necessary configuration parameters to the pilot
- ▶ The contextualization mechanism is rather complicated
 - ▶ Different for different cloud managers
 - ▶ cernvm
 - ▶ amiconfig
- ▶ Ad hoc images
 - ▶ No contextualization, everything is put in the image for a particular endpoint (software, certificates)

Simple SSH contextualization

- ▶ VMs are instantiated with golden images
 - ▶ Plus user account with a public key for ssh login
- ▶ The IPs of the VMs are should be available to the VMDIRAC scheduler
 - ▶ Polling the VMs to establish connection
- ▶ Use SSH access to contextualize the image
 - ▶ Copy the software + certificate bundle via scp to the VM
 - ▶ Start the JobAgent and VM Monitor Agent





- ▶ DIRAC + VMDIRAC installation at PIC (dirac.pic.es)
 - ▶ Configuration
 - ▶ VM metadata from MarketPlace (copy by hand)
 - ▶ VM images uploaded to end-points manually
 - ▶ Cloud manager configuration static in the DIRAC CS
 - ▶ VM scheduler:
 - ▶ Decisions to schedule VMs based on the status of the DIRAC Task Queue, VMs already launched, prices, etc.
 - ▶ VM Launcher instantiates VMs using specific cloud plugins
 - ▶ VM Manager follows the VM status, stops idle VMs

- ▶ Total 250 VM slots at 3 sites
 - ▶ USC, PIC, CC/IN2P3 with 3 different cloud managers
 - ▶ 1 virtual CPU, 2GB RAM
- ▶ LHCb standard MC production application
 - ▶ Software distribution by CVMFS
 - ▶ 2000 jobs each producing 50 events

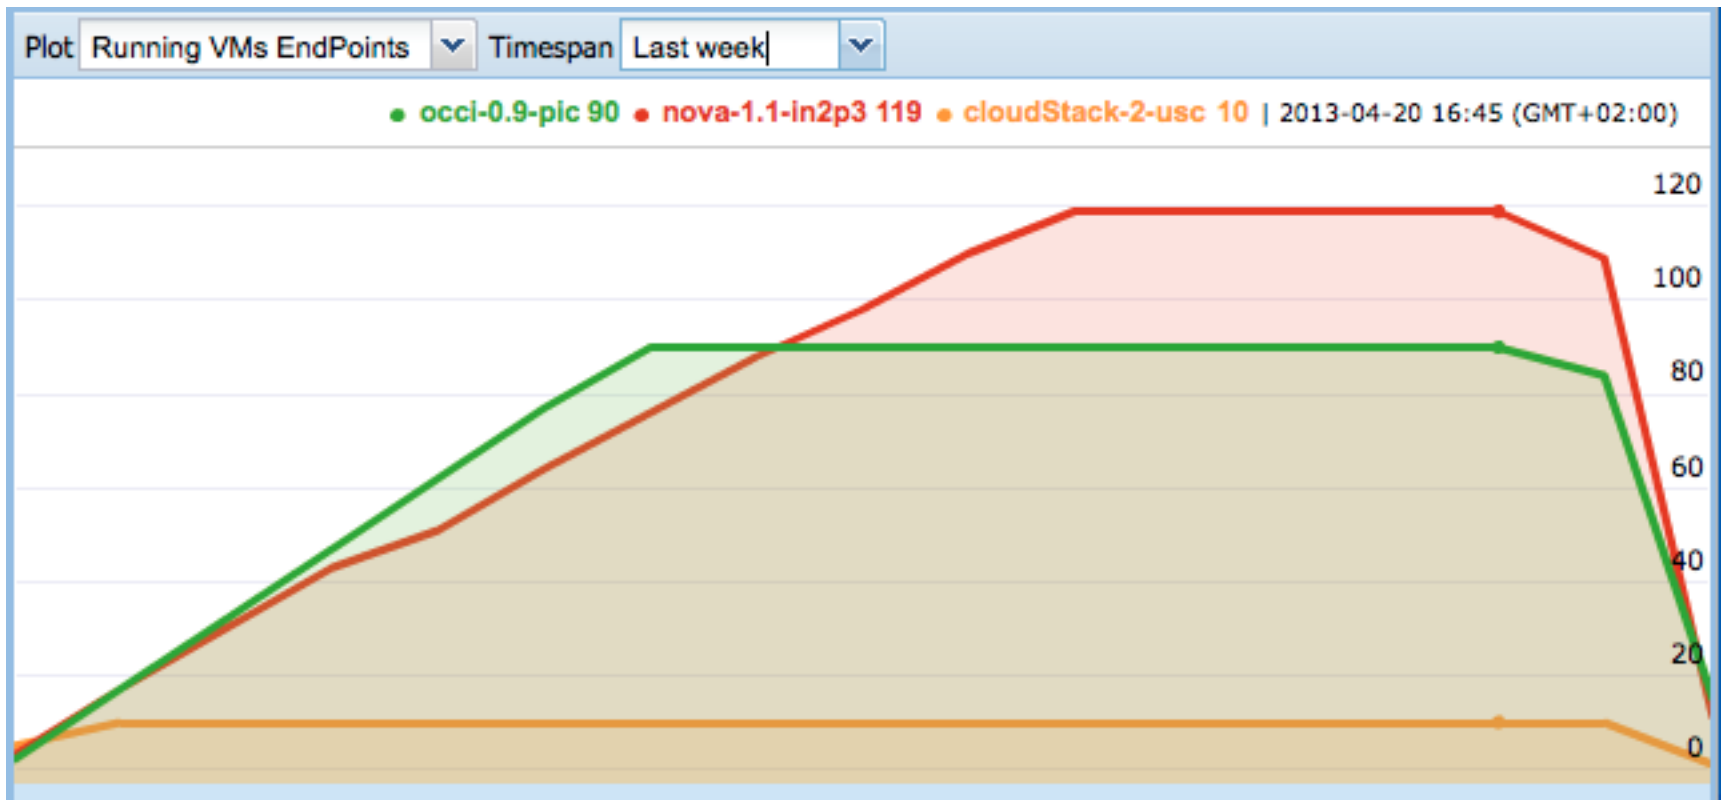
▶ PIC

- ▶ 90 VM slots
- ▶ OpenNebula, OCCl-0.8 driver
- ▶ VM image: cernvm-batch-node 2.6
- ▶ Contextualization
 - ▶ OpenNebula .ISO context (software and tools)
 - ▶ context section on the OpenNebula VM creation (endpoint)

- ▶ USC
 - ▶ 10 VM slots
 - ▶ CloudStack, CloudStack 2 driver, libcloud
 - ▶ ad hoc VM image: CentOS 5.5 + preinstalled DIRAC + CVMFS
 - ▶ Contextualization: none

- ▶ CC/IN2P3
 - ▶ 150 VM slots
 - ▶ OpenStack, libcloud + nova 1.1 python client
 - ▶ VM image: cernvm-batch-node 2.6
 - ▶ Contextualization:
 - ▶ Amiconfig: *userdata* (software and tools), *metadata* (endpoint))
 - ▶ Golden image with DIRAC public key + ssh contextualization

- ▶ VM scheduling compromise
 - ▶ Each VM launching has a price
 - ▶ The more VMs scheduled the faster user job turnaround but less efficiency
- ▶ VM scheduler:
 - ▶ 1 VM instantiated per minute, per cloud
 - ▶ *CPUPerInstance* parameter to limit the number of VMs to instantiate
 - ▶ VM stopped after 5 minutes of “no job” waiting time
- ▶ With more experience more sophisticated VM scheduling policies will be elaborated



- ▶ 219 VM at plateau
 - ▶ No failed jobs
 - ▶ Several failed VMs in OpenStack – timeout while setting up network interfaces (floating IP's)

- ▶ **French Federated cloud setup**
 - ▶ VMDIRAC is installed as part of FG-DIRAC
 - ▶ Configuration is to be done following the PIC example
 - ▶ Update to vOr7 (includes fixes from the tests)
 - ▶ **StratusLab resources incorporation**
 - ▶ Based on libcloud driver
 - ▶ Straightforward to implement
 - ▶ **OpenNebula sites**
 - ▶ OCCI 0.8 (1.1) driver
 - ▶ Contextualization ?
 - ▶ **Tests with other applications**
 - ▶ Biomed ?
 - ▶ **Manpower is a problem**

- ▶ **MarketPlace incorporation**
 - ▶ VM image metadata catalog dynamic look up
 - ▶ VM image provisioning service
- ▶ **Cloud access authorization**
 - ▶ X509, VOMS in addition to login/password
- ▶ **Cloud status information monitoring**
 - ▶ EGI Federated Cloud TF standard (?), Helix Nebula (?)
- ▶ **Cloud usage accounting**
 - ▶ Incorporate into the DIRAC Accounting service

- ▶ VMDIRAC tests accessing 3 different clouds are successful including OpenStack/CC
- ▶ This is a good basis to incorporate other clouds in France: StratusLab, OpenNebula clouds
- ▶ Demonstrations with other applications than the ones of LHCb are to be done
- ▶ Manpower is an issue