

# ILCDirac Status and Plans

André Sailer (CERN)  
on Behalf of the CLICdp Collaboration


International Workshop on Future Linear Colliders (LCWS)  
Belgrade, Serbia  
October 6–10, 2014

# Content



- 1 ILCDirac
- 2 Recent Grid Use
- 3 Moving the Services
- 4 Issues with (ILC)Dirac
- 5 Future Plans
- 6 Support Channels
- 7 Thanks
- 8 Summary

# What is ILCDirac



ILCDirac is based on the DIRAC interware originally developed for LHCb

- Dirac (Distributed Infrastructure with Remote Agent Control): High level interface between users and distributed resources
- ILCDirac: Additional functionality to provide simple interface for the users to the LC Software (Whizard, Marlin, Mokka, org.lcsim, SLIC, ROOT)

```
from DIRAC.Core.Base import Script
Script.parseCommandLine()
import UserJob
import Marlin
import DiracILC
d = DiracILC()
j = UserJob()
j.setOutputSandbox("recEvents.slcio")
m = Marlin()
m.setVersion("0116")
m.setSteeringFile("Steering.xml")
m.setInputFile("SimEvents.slcio")
j.append(m)
j.submit(d)
```

- ```

23 random.shuffle(self.failoverSEs)
24 self.log.info("Attempting to store file %s to the following SE(s):\n%s" % (tarFileName,
24                                     ', '.join(self.failoverSEs )))
24 result = failoverTransfer.transferAndRegisterFile(tarFileName, '%s/%s' % (tarFileDir, tarFileName),
24                                                 self.logLFNPath,
25                                                 self.failoverSEs, fileGUID=None,
25                                                 fileCatalog = ['FileCatalog', 'LcgFileCatalog'])
25 if not result['OK']:
25     self.log.error('Failed to upload logs to all destinations')
25     self.setApplicationStatus('Failed To Upload Logs')
25     return S_OK() #because if the logs are lost, it's not the end of the world.
25
25 #Now after all operations, retrieve potentially modified request object
25 result = failoverTransfer.getRequestObject()
25 if not result['OK']:
26     self.log.error(result)
26
27 ---- UploadLogFile.py 56% (258,48) Git: (Python +3 Rope AC Flymake:11/7/0)
27 error (E1101, no-member, UploadLogFile.finalize) Instance of 'FailoverTransfer' has no 'getRequestObject' member

```

# Recent Grid Use



Despite some interventions requiring downtime of the system made good use of the system.


## ■ Mostly used for CLICdp productions:

- ▶ 1.2 million jobs since April, 1000 CPU years
- ▶ 86 million events full simulation and reconstruction with overlay for the Higgs studies

## ■ Some users from SiD, ILD, CLICdp and Calice using ILCDirac as well

CPU used by Site


26 Weeks from Week 12 of 2014 to Week 38 of 2014



Generated on 2014-10-05 12:28:54 UTC

Cumulative Jobs by User

26 Weeks from Week 12 of 2014 to Week 38 of 2014



Generated on 2014-10-05 12:46:25 UTC

Problems on Grid sites starting in June, when the VOMS server moved to more secure SHA2 certificates

- Contacted Grid sites in July and by September all sites running again
- Mostly they just had to restart a certain program after an update

There are some sites supporting ILC VO not used via ILCDirac

- Sites banned long time ago that disappeared from the site summary in the web interface
- Need to make sure all available resources can be used
  - ▶ This is not only a matter of ILCDirac configuration, in some cases sites have incompatible configuration,
  - ▶ Will need some time

# Moving the Services



## Reasons for moving the services

- ILCDirac hardware had mostly run out of warranty
  - Any failure of hardware would have been the end
- Configuration management software (quattor) used to configure the machines was coming to end-of-life
  - Cern wide effort to replace all quattor managed machines by end of October
  - New configuration management system based on “Puppet” and almost exclusively virtual machines

## Plan:

- 1 Learn puppet
- 2 Create VMs
- 3 Move services

## Example puppet listing

```
file { '/opt/dirac/etc/grid-security/certificates':  
  ensure => link,  
  target => '/etc/grid-security/certificates',  
  require => Mount['/opt/dirac'],  
}
```

# ...when a Plan Comes Together



Originally wanted to start moving things in August, but postponed several times

- Familiarized myself with puppet, creating VMs, etc.
- Mostly skipped step 1, because Joel Closier (CERN, LHCb) provided the “puppet manifests” from their Dirac setup
- Started seriously working on step 2 and 3 middle of September

Created all VMs by September 12, finished moving services by September 22.

- Three machines for the main services: 8 cores, 16GB RAM
- Three machines: 4 cores, 8GB RAM, 1TB extra storage volume for: Sandbox SE\*, Production Logfile SE and, Dirac SE (CERN-DIP-4)
  - Storage volumes independent of VM, can be moved to different VM if necessary
- One machine: 2 cores for the web front-end
  - All by itself, reachable via http(s) from the outside
- 2 more machine for testing and development
- All databases (except one) moved to Cern DB-on-Demand service
  - Includes daily backup of all the databases

---

\*Storage Element



- No clear guideline how to distribute services across machines
  - ▶ Used to have one big workhorse with 24 cores running everything
  - ▶ Grouped them by system (DataManagement, Transformation, Request, Accounting, WorkloadManagement, ...)
  - ▶ Put resource heavy systems on different machines
- Initially problematic to obtain the large VMs
- Could not attach the volumes to VMs until manual intervention by IT
- A few small tweaks to the puppet manifests
- After machines had been taken out of production and officially given back to IT to retire: Noticed important files not reachable any longer

- No clear guideline how to distribute services across machines
  - ▶ Used to have one big workhorse with 24 cores running everything
  - ▶ Grouped them by system (DataManagement, Transformation, Request, Accounting, WorkloadManagement, ...)
  - ▶ Put resource heavy systems on different machines
- Initially problematic to obtain the large VMs
- Could not attach the volumes to VMs until manual intervention by IT
- A few small tweaks to the puppet manifests
- After machines had been taken out of production and officially given back to IT to retire: Noticed important files not reachable any longer
  - 1 Production steering files only on CERN-DIP-2 (which totally was not on my radar, as I thought there was only CERN-DIP-3 used)

- No clear guideline how to distribute services across machines
  - ▶ Used to have one big workhorse with 24 cores running everything
  - ▶ Grouped them by system (DataManagement, Transformation, Request, Accounting, WorkloadManagement, ...)
  - ▶ Put resource heavy systems on different machines
- Initially problematic to obtain the large VMs
- Could not attach the volumes to VMs until manual intervention by IT
- A few small tweaks to the puppet manifests
- After machines had been taken out of production and officially given back to IT to retire: Noticed important files not reachable any longer
  - 1 Production steering files only on CERN-DIP-2 (which totally was not on my radar, as I thought there was only CERN-DIP-3 used)
  - 2 Oh ?@\*%, oh ?@\*%, oh ?@\*%

- No clear guideline how to distribute services across machines
  - ▶ Used to have one big workhorse with 24 cores running everything
  - ▶ Grouped them by system (DataManagement, Transformation, Request, Accounting, WorkloadManagement, ...)
  - ▶ Put resource heavy systems on different machines
- Initially problematic to obtain the large VMs
- Could not attach the volumes to VMs until manual intervention by IT
- A few small tweaks to the puppet manifests
- After machines had been taken out of production and officially given back to IT to retire: Noticed important files not reachable any longer
  - 1 Production steering files only on CERN-DIP-2 (which totally was not on my radar, as I thought there was only CERN-DIP-3 used)
  - 2 Oh ?@\*%, oh ?@\*%, oh ?@\*%
  - 3 Machine was still up and I could retrieve the files

# Problems Along the Way and Lessons Learned



- No clear guideline how to distribute services across machines
  - ▶ Used to have one big workhorse with 24 cores running everything
  - ▶ Grouped them by system (DataManagement, Transformation, Request, Accounting, WorkloadManagement, ...)
  - ▶ Put resource heavy systems on different machines
- Initially problematic to obtain the large VMs
- Could not attach the volumes to VMs until manual intervention by IT
- A few small tweaks to the puppet manifests
- After machines had been taken out of production and officially given back to IT to retire: Noticed important files not reachable any longer
  - 1 Production steering files only on CERN-DIP-2 (which totally was not on my radar, as I thought there was only CERN-DIP-3 used)
  - 2 Oh ?@\*%, oh ?@\*%, oh ?@\*%
  - 3 Machine was still up and I could retrieve the files

## Lessons learned:

- 1 Do things well in advance to ensure overlap of old and new services
- 2 Always keep a backup of production software

# Virtual Machine Overview



## Summary



Instances  
Used 9 of 50



VCPUs  
Used 50 of 50



RAM  
Used 98.5 GB of  
100.0 GB



Available Volumes  
Used 3 of 10



Available Volume  
Storage  
Used 2.8TB of 2.9TB

Active Instances: 9 Active RAM: 98GB This Period's VCPU-Hours: 112.82 This Period's GB-Hours: 20433.73

## Usage Summary

Download CSV Summary

| Instance Name                    | VCPUs | Disk | RAM  | Uptime            |
|----------------------------------|-------|------|------|-------------------|
| <a href="#">voicdirac01</a>      | 8     | 230  | 15GB | 3 weeks, 2 days   |
| <a href="#">voicdirac02</a>      | 8     | 230  | 15GB | 3 weeks, 2 days   |
| <a href="#">voicdirac03</a>      | 8     | 230  | 15GB | 3 weeks, 2 days   |
| <a href="#">voicdiracdipse01</a> | 4     | 150  | 8GB  | 3 weeks, 5 days   |
| <a href="#">voicdiraclogse01</a> | 4     | 150  | 8GB  | 3 weeks, 5 days   |
| <a href="#">voicdiracsbse03</a>  | 4     | 150  | 8GB  | 3 weeks, 5 days   |
| <a href="#">voicdiractest04</a>  | 8     | 230  | 15GB | 4 months, 3 weeks |
| <a href="#">voicdiractest05</a>  | 4     | 150  | 8GB  | 2 weeks, 4 days   |
| <a href="#">voicdiracweb</a>     | 2     | 110  | 4GB  | 1 month, 2 weeks  |
| Displaying 9 items               |       |      |      |                   |

- Using 50 cores
- Have to see how system load develops
- Maybe rearrange

- Old RequestManagement System (for asynchronous operations) caused some serious degradation of the system
  - ▶ 10 seconds to submit jobs
  - ▶ Manually cleaned up database clogged with broken requests
  - ▶ Will be replaced by new system soon
- Certification Authority certificates used to be taken from AFS
  - ▶ First replaced by nightly cron job copying from AFS
  - ▶ Now deployed via package manager
- StorageManagementSystem not checking for faulty user input
  - ▶ Controls staging files from tape
  - ▶ Cron job to remove bogus entries from DB
  - ▶ Note: new version of Dirac changes order of arguments
- Job scheduling not configured correctly
  - ▶ Noticed problems with reconstruction jobs using way too much memory, had to limit where these jobs are running
  - ▶ Now properly controlling what kind of and how many (production) jobs run where

- Move ILCDirac to latest Dirac release
  - ▶ Currently jobs running on Dirac v6r8, servers on Dirac v6r9
  - ▶ Will probably move straight to Dirac v6r11, because the new RequestManagement System does not work otherwise
- Need new features from up-to-date Dirac
  - ▶ Interface to new grid middleware that sites are moving towards (ARC CEs)
  - ▶ Fixes in FileCatalog interface (permissions)
- Use the software installations on cvmfs
  - ▶ Tarballs extracted on to cvmfs are already used, if available
- Interface for Whizard 2
- ILD productions with ILCDirac (with C. Calancha, S. Lu, F. Gaede)
- ...



- New member of the ILCDirac team, Thibault Frisson(CERN), to answer support requests
- If there are problems (ordered list):
  - 1 FAQ: <https://twiki.cern.ch/twiki/bin/view/CLIC/DiracForUsers>
  - 2 Submit a ticket to the issue tracker  
<https://its.cern.ch/jira/browse/ILCDIRAC> or top middle of [ilcdirc.cern.ch](https://ilcdirc.cern.ch), also has search function. (see the TWIKI for access)
  - 3 [forum.linearcollider.org](https://forum.linearcollider.org)
  - 4 Email: [ilcdirc-support@cern.ch](mailto:ilcdirc-support@cern.ch)

Registration: [ilcdirc-register@cern.ch](mailto:ilcdirc-register@cern.ch)

# Thanks



- Thanks to sites for fixing issues and providing the resources
- Thanks to Joel for the puppet manifests
- Thanks to the Dirac (ILCDirac, LHCbDirac) developers (Stephane, Andrei, Federico, Adri, et al.)
- Thanks to CERN IT, Agile Infrastructure team

- CLICdp making good use of Grid resources through ILCDirac
- Easy to use, at least I rarely hear from users beyond the registration
  - But I do see them use the system
- Moved to new server infrastructure
- Need to upgrade the software, test, and put into production