Summary of (virtual) DIRAC Users' Workshop 2021

DIRAC THE INTERWARE

Federico Stagni

GDB, June 9th 2021

DIRAC: the interware

- A software framework for distributed computing
- A **complete** solution to one (or more) user community
- Builds a layer between users and resources

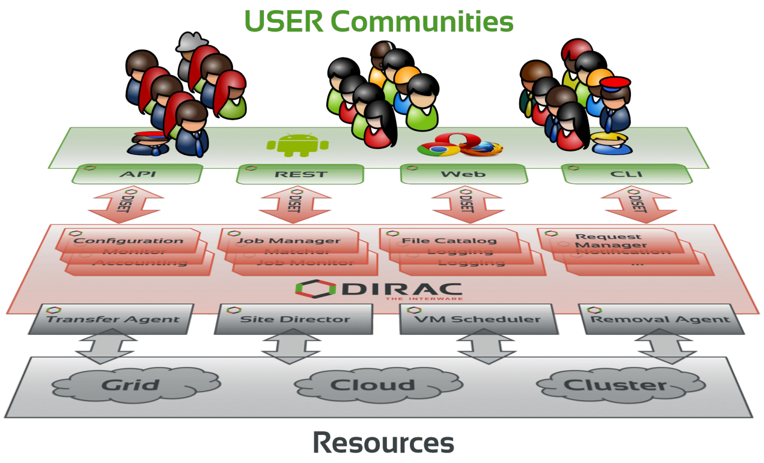

- Started as an LHCb project, experiment-agnostic in 2009
- Developed by communities, for communities
  - Open source (GPL3+), GitHub hosted
  - Python 2.7 → python 3 (see later)
  - No dedicated funding for the development of the “Vanilla” project
  - Publicly documented, active assistance forum, yearly users workshops (when possible…), open developers meetings and hackathons
- The DIRAC consortium as representing body

Users/communities/VOs

LHCb

BESIII

JUNO

CLIC · · · · ilc

cta
cherenkov telescope array

Belle II

GridPP
UK Computing for Particle Physics

egi

FRANCE GRILLES

A *framework* shared by multiple experiments/projects, both inside HEP, astronomy, and life science

<u>Experiment agnostic</u>

<u>Extensible</u>

<u>Flexible</u>

EISCAT 3D

T2K

Fermi
Gamma-ray Space Telescope

NA62

we-nmr

PIERRE AUGER OBSERVATORY

SKA
SQUARE KILOMETRE ARRAY

NICA

FCC
hh ee he

PROJECT 8

CEPC

LSST

DUWs

- 2019: 9th DUW (UK, London, IC)
- ~~2020: 10th DUW (Japan, Tsukuba, KEK)~~
- 2021: mini, virtual DUW
  - hopefully the last virtual one

DUWs are normally 3 to 4 days, ~40% presentations, ~60% discussions, tutorials, hands-on, hackathons

not well suited for online

DIRAC
THE INTERWARE

DUW 2021: setup and participation

- Zoom only meeting, notes on hedgedoc (codiMD), recordings on the agenda
- 80 registered users
  - effectively max ~60 attending at the same time
    - a few more participants on Monday
    - many known names, many new ones
  - participation from the Americas of course problematic for the timezone

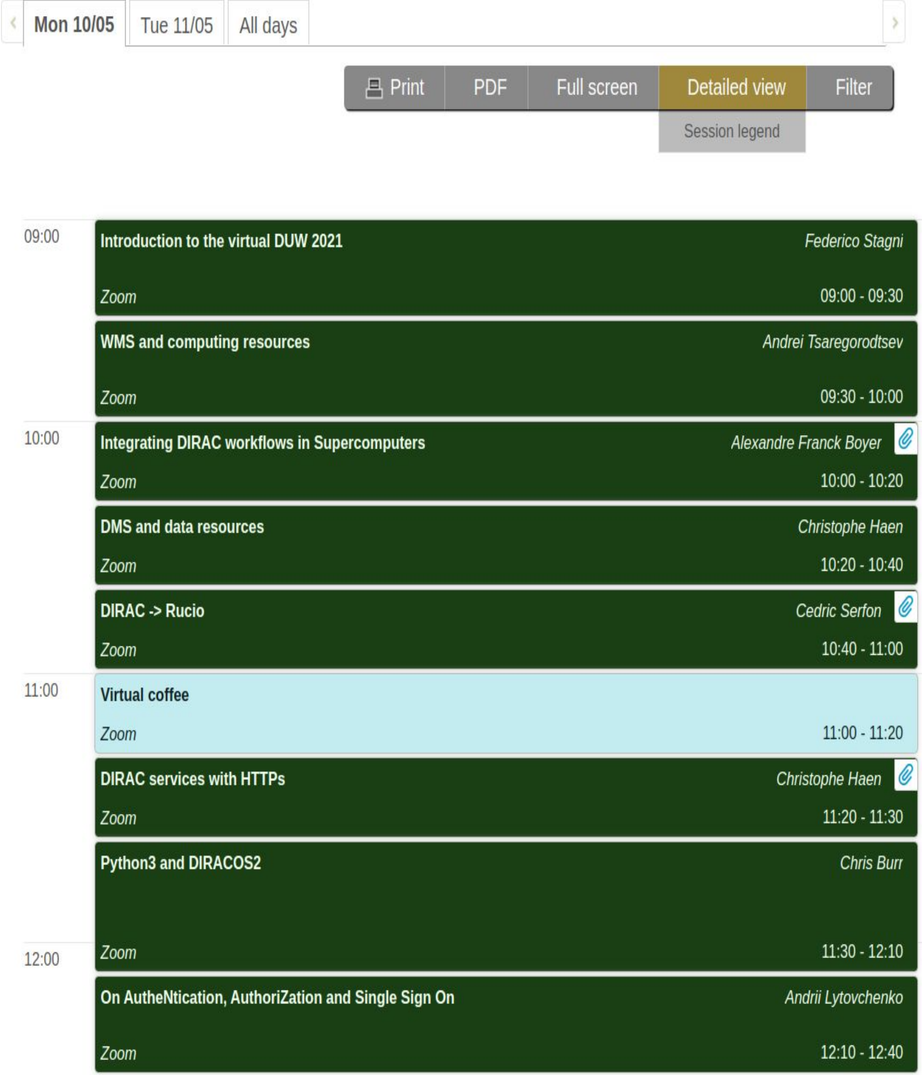

09:00	Introduction to the virtual DUW 2021	*Federico Stagni*	*Zoom*	09:00 - 09:30
	WMS and computing resources	*Andrei Tsaregorodtsev*	*Zoom*	09:30 - 10:00
10:00	Integrating DIRAC workflows in Supercomputers	*Alexandre Franck Boyer*	*Zoom*	10:00 - 10:20
	DMS and data resources	*Christophe Haen*	*Zoom*	10:20 - 10:40
	DIRAC -> Rucio	*Cedric Serfon*	*Zoom*	10:40 - 11:00
11:00	Virtual coffee		*Zoom*	11:00 - 11:20
	DIRAC services with HTTPs	*Christophe Haen*	*Zoom*	11:20 - 11:30
	Python3 and DIRACOS2	*Chris Burr*	*Zoom*	11:30 - 12:10
12:00	On AutheNtication, AuthoriZation and Single Sign On	*Andrii Lytovchenko*	*Zoom*	12:10 - 12:40

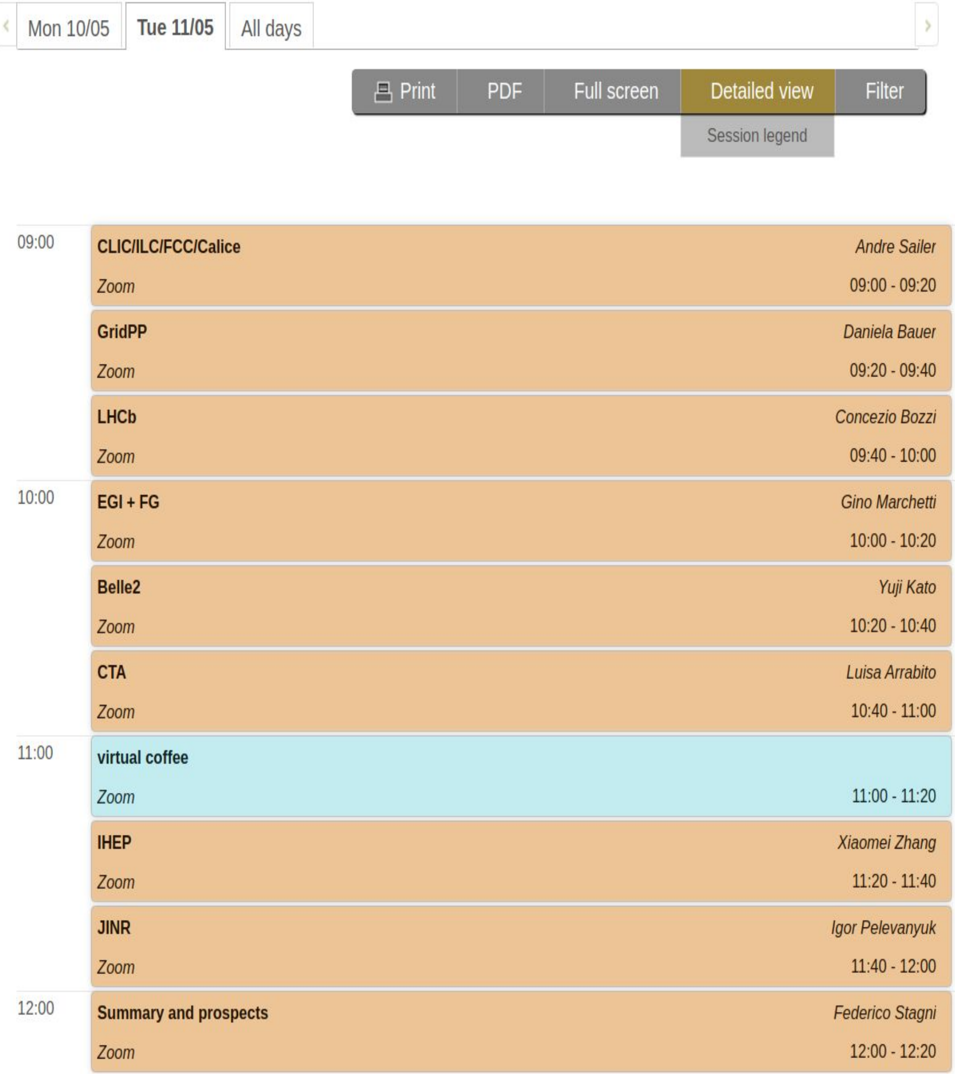

09:00	CLIC/ILC/FCC/Calice	*Andre Sailer*	*Zoom*	09:00 - 09:20
	GridPP	*Daniela Bauer*	*Zoom*	09:20 - 09:40
	LHCb	*Concezio Bozzi*	*Zoom*	09:40 - 10:00
10:00	EGI + FG	*Gino Marchetti*	*Zoom*	10:00 - 10:20
	Belle2	*Yuji Kato*	*Zoom*	10:20 - 10:40
	CTA	*Luisa Arrabito*	*Zoom*	10:40 - 11:00
11:00	virtual coffee		*Zoom*	11:00 - 11:20
	IHEP	*Xiaomei Zhang*	*Zoom*	11:20 - 11:40
	JINR	*Igor Pelevanyuk*	*Zoom*	11:40 - 12:00
12:00	Summary and prospects	*Federico Stagni*	*Zoom*	12:00 - 12:20

DIRAC general principles

- A single server can manage dozens of VOs
- Aim of being a single solution for the whole Grid/computing/“data lake”
- Lots of building blocks, and plugins for “everything”:
    - install/use (VO)DIRAC client → easy
    - install/manage (VO)DIRAC server → NOT easy
        - many functions/options
- Modular, pragmatic: DIRAC is a framework for developing custom grid solutions.

Workload Management

Focus: DIRAC pilots for exploiting all computing resources

- ARC, HTCondorCE, SSH(+batch), clouds, vac, BOINC, ...
- HPCs
  - a variety of solutions for several use cases are possible, others in development, e.g. running on closed WNs is not yet an option

No new “killer feature”: continuously evolving landscape, adaptation is key

Data Management

DMS “machinery” pretty stable ⇒ focus on supporting protocols, SEs, etc.

no multi-hop!

New plugin: Rucio FC

- “yet another catalog plugin”
- developed/used by Belle2
  - PR in DIRAC, reviewed, waiting for developer
    - 3(!) presentations at CHEP

New stuff

- dips:// → https:// (from "DIRAC protocol" to http)
  - based on tornado
- Python3
  - DIRAC Client ready
    - `pip install DIRAC` -- from 7.2 (https://pypi.org/project/DIRAC/)
  - DIRAC server almost there
- DIRACOS2 as environment
  - built using conda constructor and conda-forge (source of packages.
  - support for aarch64 and ppc64le
- Deprecating dirac-install command

New stuff /2

AuthN + AuthZ + SSO, via tokens

- from X509 to tokens
  - EGI Check-in, IaM
  - (but won’t force anyone)
- following and using “standards”
  - design and libraries
- probably python3 only

DIRAC development summary

- Code maintainability is the drive for several DIRAC developments
  - e.g.: m2crypto, py3, DIRACOS2, ...
  - modern code = turn-key for future DIRAC developments
- VOMS → CheckIn/IaM/?
  - drive for re-thinking of DIRAC AuthN/Z

Day 2: Experiments’ reports

Set questions + free contribution

- What do you use DIRAC for, and which DIRAC functionalities you don't use, and why?
- Do you have a DIRAC extension? Why?
  - If yes, do you think some of it could become part of the vanilla projects?
- What is your biggest frustration with DIRAC?
- You can magically add one feature to DIRAC, what is it?
- Any notable operations incident in the last year?
- To support your "Grid", do you have to use other systems than DIRAC?
- How would you rate the communication?
- In the last two years, what has been the DIRAC usage in terms of jobs ran, CPU (or wall time) used, and data transfers?

Experiments’ reports -- answers /1

- Everyone’s using WMS and Accounting
    - VMDIRAC (extension of WMS, for clouds) used by some
        - one day will be fully merged in DIRAC
- Everyone’s using at least partially the DMS
    - everyone (but Belle2) uses the DFC
    - most communities are using FTS (via DIRAC DMS/RMS)
- Several using RMS + TS
    - “Requests Management” and “Transformation” systems
    - for jobs productions + dataset management
    - Vanilla DIRAC ProductionSystem also available, but not widely used yet because it’s too recent

Experiments' reports -- answers /2

- Community-lead DIRAC extensions are common
  - for support of community software, or production requirements
  - for integrating external DBs (e.g. LHCb bookkeeping, Belle2 AMGA)
- Some in-house developments have the potential of becoming part of vanilla DIRAC
  - for component monitoring (i.e. stuck agents) - (Belle2 and ILCDirac)
  - FC ACL plugin (EISCAT)
  - Transformations and DM "bits" (LHCb and CTA)
  - Site Monitoring testing (with RSS) and UserJobs/Tasks grouping (IHEP)
- Some interest in puppet modules

Experiments’ reports -- answers /3

- No “biggest frustration”, but monitoring is never enough
  - Pilots monitoring seems particularly requested (cited in at least 3 reports)
  - Many installations using ElasticSearch -- IMHO this is a great news!
- No documentation requests!
  - and communication is deemed good by all
- LHCb is still the largest experiment when it comes to resources usage.

For Devs

BILD meetings:

“BiWeekly ‘Loyal’ DIRAC Developers meetings”

every 2nd week

Thursday at 10:00 AM CET

LHCb hosted

Clic, Belle2, EGI/FG, BES3/Juno, GridPP, IHEP, Jinr represented

→ *you want to be invited? Just let me know*

Where releases and issues are discussed!

Certification hackathons:

every (other) 2nd week

Thursday at 10:00 AM CET

LHCb hosted

Clic, EGI/FG, GridPP represented

→ *you want to be participate? Just let me know*

lhcbdirac.slack.com + Trello

Development and testing

~6 FTE as core developers, a dozen contributing developers

Tests, certification, integration process is a daily work.

- We use GitHub Actions, GitLab CI/CD (+ Jenkins…)
  - link: instructions on how to run it locally

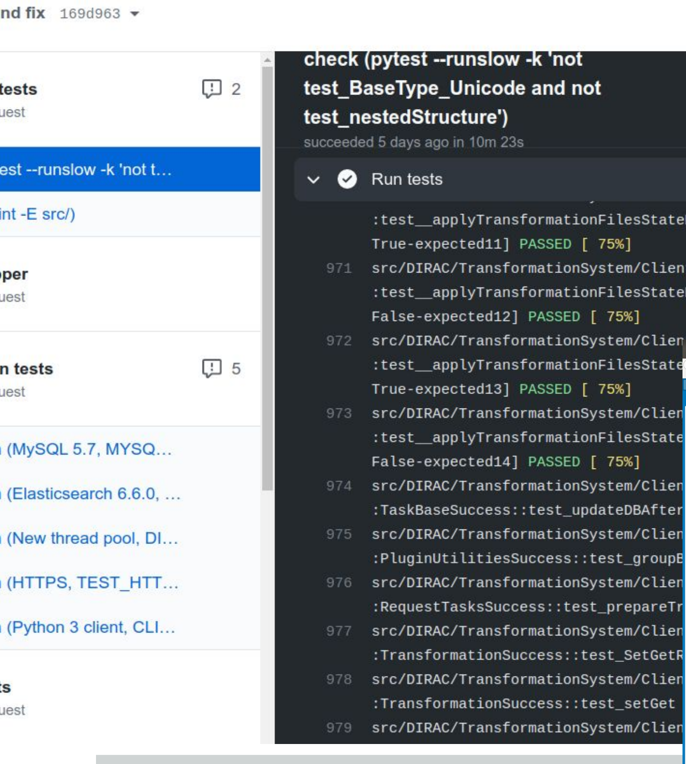

nd fix 169d963

tests 2
uest

est --runslow -k 'not t…

int -E src/)

oper
uest

n tests 5
uest

(MySQL 5.7, MYSQ…

(Elasticsearch 6.6.0, …

(New thread pool, DI…

(HTTPS, TEST_HTT…

(Python 3 client, CLI…

ts
uest

check (pytest --runslow -k 'not test_BaseType_Unicode and not test_nestedStructure')

succeeded 5 days ago in 10m 23s

Run tests

```
     :test__applyTransformationFilesStateMachine[ts
     True-expected11] PASSED [ 75%]
971  src/DIRAC/TransformationSystem/Client/test/Test_Client_TransformationC
     :test__applyTransformationFilesStateMachine[tsFiles12-dictOfNewLFNsStat
     False-expected12] PASSED [ 75%]
972  src/DIRAC/TransformationSystem/Client/test/Test_Client_TransformationC
     :test__applyTransformationFilesState
     True-expected13] PASSED [ 75%]
973  src/DIRAC/TransformationSystem/Clien
     :test__applyTransformationFilesState
     False-expected14] PASSED [ 75%]
974  src/DIRAC/TransformationSystem/Clien
     :TaskBaseSuccess::test_updateDBAfter
975  src/DIRAC/TransformationSystem/Clien
     :PluginUtilitiesSuccess::test_groupB
976  src/DIRAC/TransformationSystem/Clien
     :RequestTasksSuccess::test_prepareTr
977  src/DIRAC/TransformationSystem/Clien
     :TransformationSuccess::test_SetGetR
978  src/DIRAC/TransformationSystem/Clien
     :TransformationSuccess::test_setGet
979  src/DIRAC/TransformationSystem/Clien
```

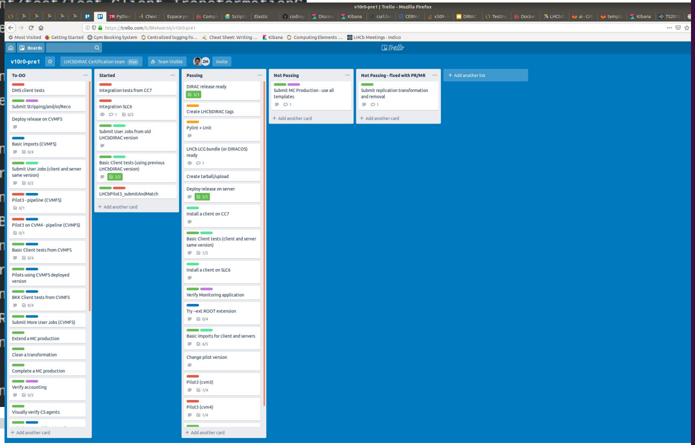

v10r0-pre1 | Trello - Mozilla Firefox

https://trello.com/b/BN4w6r66/v10r0-pre1

v10r0-pre1 · LHCbDIRAC Certification team Free · Team Visible · ZM · Invite

To-DO
- DMS client tests
- Submit Stripping/and/or/Reco
- Deploy release on CVMFS
- Basic imports (CVMFS) 0/4
- Submit User Jobs (client and server same version) 0/2
- Pilot3 - pipeline (CVMFS) 0/1
- Pilot3 on CVM4 - pipeline (CVMFS) 0/1
- Basic Client tests from CVMFS 0/4
- Pilots using CVMFS deployed version
- BKK Client tests from CVMFS 0/4
- Submit More User Jobs (CVMFS)
- Extend a MC production
- Clean a transformation
- Complete a MC production
- Verify accounting 0/5
- Visually verify CS agents
- + Add another card

Started
- Integration tests from CC7
- Integration SLC6 1 0/2
- Submit User Jobs from old LHCbDIRAC version
- Basic Client tests (using previous LHCbDIRAC version) 2/2
- LHCbPilot3_submitAndMatch
- + Add another card

Passing
- DIRAC release ready 3/3
- Create LHCbDIRAC tags
- Pylint + Unit
- LHCb LCG bundle (or DIRACOS) ready 1
- Create tarball/upload
- Deploy release on server 5/5
- Install a client on CC7
- Basic Client tests (client and server same version) 1/2
- Install a client on SLC6
- Verify Monitoring application
- Try --ext ROOT extension 0/4
- Basic imports for client and servers 4/5
- Change pilot version
- Pilot3 (cvm3) 1/4
- Pilot3 (cvm4) 1/4
- + Add another card

Not Passing
- Submit MC Production - use all templates 1
- + Add another card

Not Passing - fixed with PR/MR
- Submit replication transformation and removal 1
- + Add another card

+ Add another list

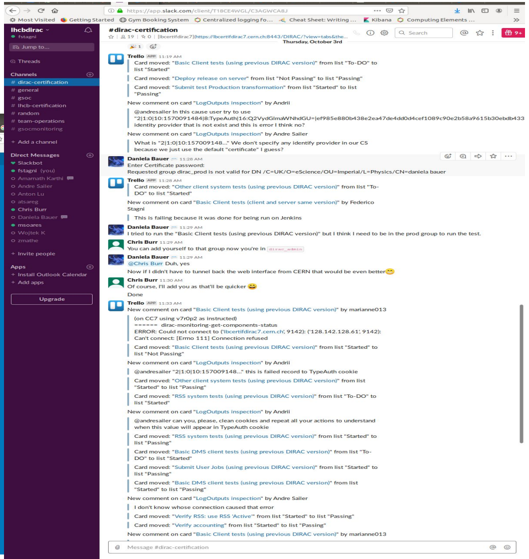

https://app.slack.com/client/T18CE4WGL/C3AGWCA8J

lhcbdirac
fstagni

Jump to…

Threads

Channels
- # dirac-certification
- # general
- # gsoc
- # lhcb-certification
- # random
- # team-operations
- # gsocmonitoring
- + Add a channel

Direct Messages
- Slackbot
- fstagni (you)
- Amarnath Karthi
- Andre Sailer
- Anton Lu
- atsareg
- Chris Burr
- Daniela Bauer
- msoares
- Wojtek K
- zmathe
- + Invite people

Apps
- + Install Outlook Calendar
- + Add apps

Upgrade

#dirac-certification
19 | 0 | [lbcertifdirac7](https://lbcertifdirac7.cern.ch:8443/DIRAC/?view=tabs&the…

Thursday, October 3rd

Trello APP 11:19 AM
Card moved: "Basic Client tests (using previous DIRAC version)" from list "To-DO" to list "Started"
Card moved: "Deploy release on server" from list "Not Passing" to list "Passing"
Card moved: "Submit test Production transformation" from list "Started" to list "Passing"
New comment on card "LogOutputs inspection" by Andrii
@andresailer In this cause user try to use "2|1:0|10:1570091484|8:TypeAuth|16:Q2VydGlmaWNhdGU=|ef985e880b438e2ea47de4dd0d4cef1089c90e2b58a9615b30ebdb433 identity provider that is not exist and this is error I think no?
New comment on card "LogOutputs inspection" by Andre Sailer
What is "2|1:0|10:157009148…" We don't specify any identify provider in our CS because we just use the default "certificate" I guess?

Daniela Bauer 11:28 AM
Enter Certificate password:
Requested group dirac_prod is not valid for DN /C=UK/O=eScience/OU=Imperial/L=Physics/CN=daniela bauer

Trello APP 11:28 AM
Card moved: "Other client system tests (using previous DIRAC version)" from list "To-DO" to list "Started"
New comment on card "Basic Client tests (client and server same version)" by Federico Stagni
This is failing because it was done for being run on Jenkins

Daniela Bauer 11:29 AM
I tried to run the "Basic Client tests (using previous DIRAC version)" but I think I need to be in the prod group to run the test.

Chris Burr 11:29 AM
You can add yourself to that group now you're in `dirac_admin`

Daniela Bauer 11:29 AM
@Chris Burr Duh, yes
Now if I didn't have to tunnel back the web interface from CERN that would be even better😁

Chris Burr 11:30 AM
Of course, I'll add you as that'll be quicker 😄
Done

Trello APP 11:33 AM
New comment on card "Basic Client tests (using previous DIRAC version)" by marianne013
(on CC7 using v7r0p2 as instructed)
====== dirac-monitoring-get-components-status
ERROR: Could not connect to ('lbcertifdirac7.cern.ch', 9142): ('128.142.128.61', 9142): Can't connect: [Errno 111] Connection refused
Card moved: "Basic Client tests (using previous DIRAC version)" from list "Started" to list "Not Passing"
New comment on card "LogOutputs inspection" by Andrii
@andresailer "2|1:0|10:157009148…" this is failed record to TypeAuth cookie
Card moved: "Other client system tests (using previous DIRAC version)" from list "Started" to list "Passing"
Card moved: "RSS system tests (using previous DIRAC version)" from list "To-DO" to list "Started"
New comment on card "LogOutputs inspection" by Andrii
@andresailer can you, please, clean cookies and repeat all your actions to understand when this value will appear in TypeAuth cookie
Card moved: "RSS system tests (using previous DIRAC version)" from list "Started" to list "Passing"
Card moved: "Basic DMS client tests (using previous DIRAC version)" from list "To-DO" to list "Started"
Card moved: "Submit User Jobs (using previous DIRAC version)" from list "Started" to list "Passing"
Card moved: "Basic DMS client tests (using previous DIRAC version)" from list "Started" to list "Passing"
New comment on card "LogOutputs inspection" by Andre Sailer
I don't know whose connection caused that error
Card moved: "Verify RSS: use RSS 'Active'" from list "Started" to list "Passing"
Card moved: "Verify accounting" from list "Started" to list "Passing"
New comment on card "Basic Client tests (using previous DIRAC version)" by marianne013

Message #dirac-certification

DIRAC THE INTERWARE

DIRAC Consortium news

Board of representatives on May 10th

- Consortium Agreement expires in Feb 2022
  - Agreed to prepare an Amendment extending the Consortium Agreement for 5 more years
- Election of the coordinators for the next 2 years:
  - Andrei Tsaregorodtsev confirmed as DIRAC coordinator
  - Federico Stagni confirmed as DIRAC technical coordinator

In conclusion

Despite the lack of face to face discussions and hands-on, it was a useful workshop.

Hope to **see** DIRAC users and developers at the next DUW.

next time, for real!

Questions/comments

?