
The DIRAC interware

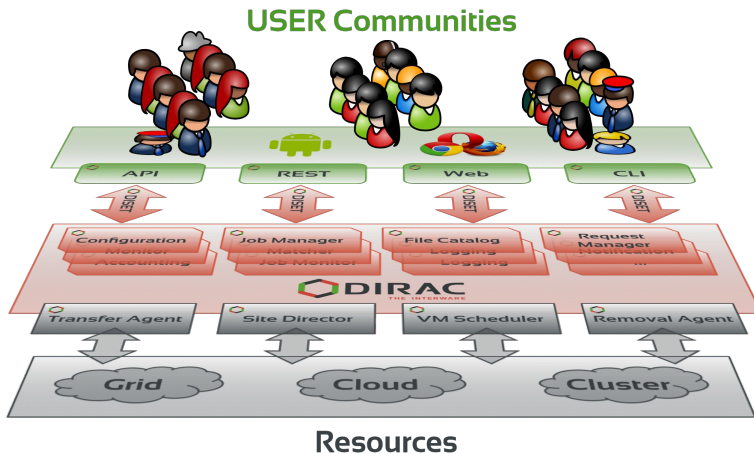
current, upcoming and planned
capabilities and technologies



Federico Stagni
DIRAC technical coordinator
on behalf of the DIRAC consortium

federico.stagni@cern.ch

- 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 in development)
 - No dedicated funding for the development of the “Vanilla” project
 - Publicly [documented](#), active [assistance forum](#), yearly [users workshops](#), open [developers meetings](#) and [hackathons](#)
- The DIRAC consortium as representing body



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

Experiment agnostic

Extensible

Flexible

Distributed data management in Belle2



11:30, R8

Efficient Iterative Calibration on the Grid using iLCDirac

12:00, R3

Dirac-based solutions for JUNO production system

15:00, R3



Jobs and files

WMS and DMS

**Pilots are the
“federators”**

Send them

as “pilot jobs” (via a CE)

Or just **Run them!**

e.g. as part of the contextualization of a (V)M

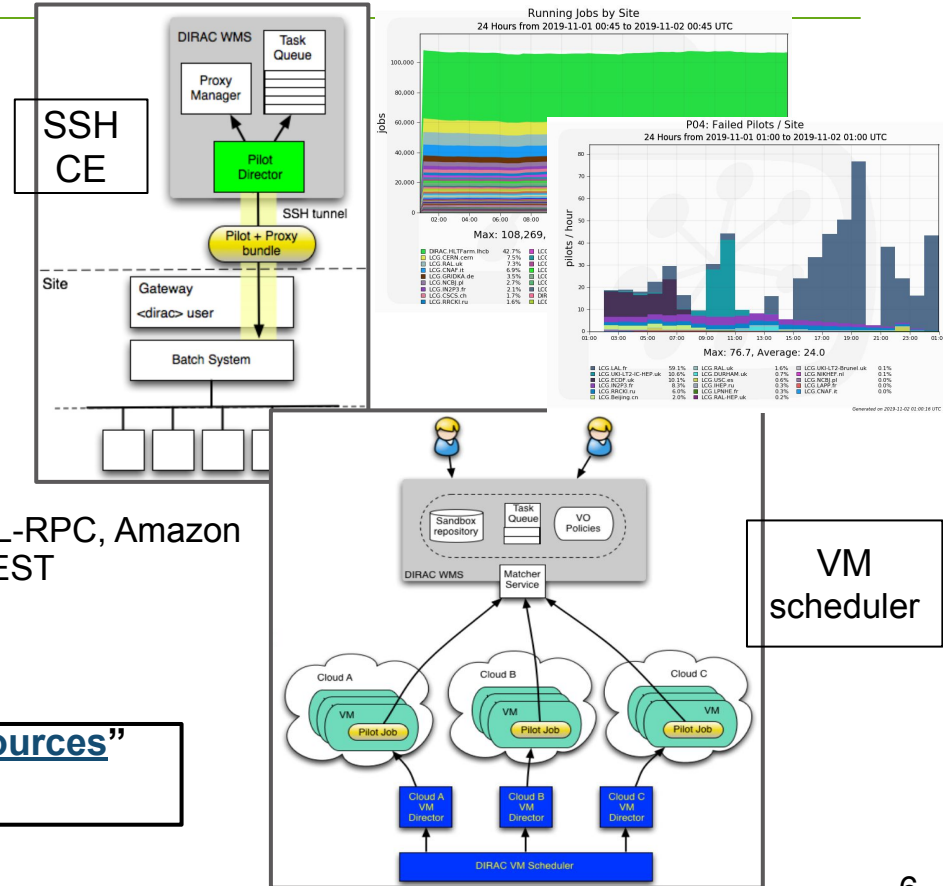
OR

“Make a machine a pilot machine, and you are done”

[WMS] Computing resources

- **Grids** (EGI, OSG, NorduGRID)
 - CREAM, HTCondor, ARC
- **Clusters** behind a BS
 - access through SSH/GSISSH tunnel
 - a really thin layer that we call "SSH CE"
- **Vacuum:**
 - VAC/vcycle resources
 - BOINC Volunteer resources
 - HLT farm (LHCb)
- **VMs scheduler:**
 - Openstack, Keystone v2 & v3, OpenNebula XML-RPC, Amazon EC2 (boto2), Apache libcloud, rocci cli, OCCI REST
 - Contextualization from standard images
 - with, at least, the DIRAC pilot
- **HPC sites**

“Integrating LHCb workflows on HPC resources”
 14:45, R4



Basics of DMS:

- **LFNs**: unique identifier within DIRAC of a file

Logical File Name
(described as paths)

- LFNs are registered in **catalog(s)**.

and there are implementations like the DFC
→ and you can connect as many catalogs as you want
(including the LFC or Rucio catalog)

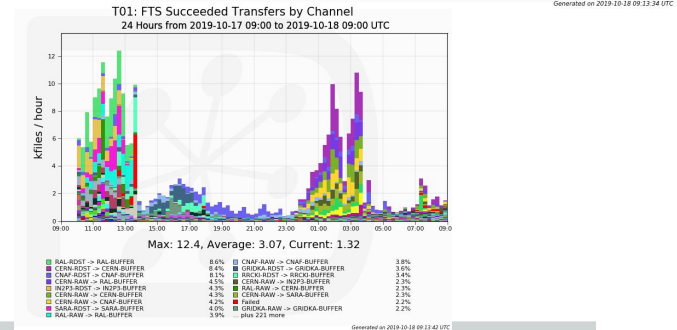
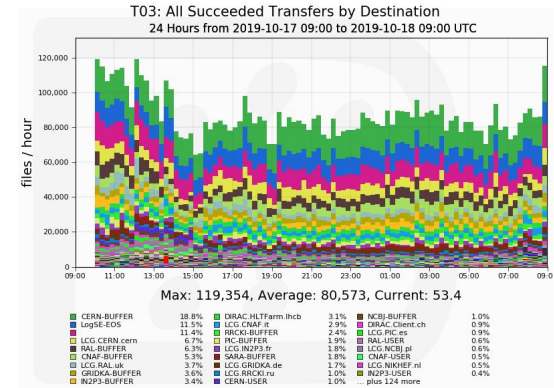
- LFNs *may* have **PFNs**, stored in **SEs**.

Physical File Name on Storage Elements

(and SEs are monitored, within the DIRAC Resource Status System)

- You can access those PFNs with several protocols.

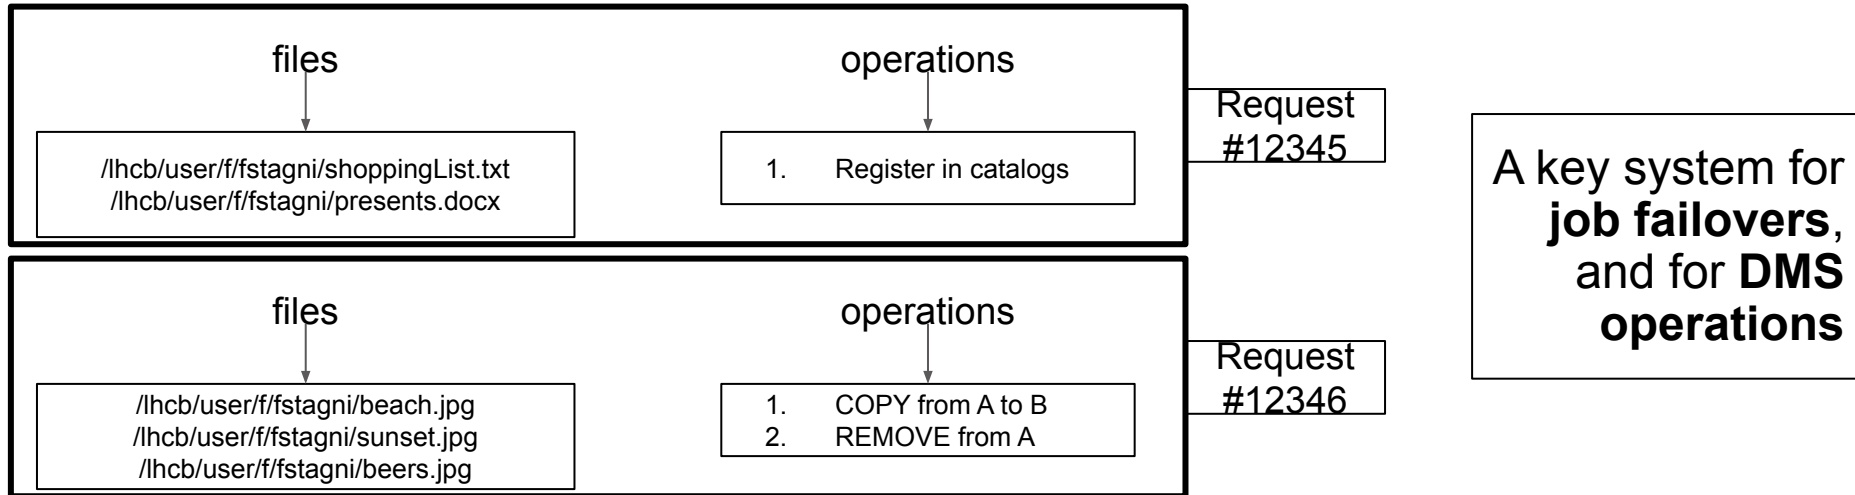
e.g. root, gsiftp, srm, http, file, dip
(and can also be brought online - i.e. staged)



Productions and datasets

Request management system

A generic, flexible system, which can be used for queueing *operations* (on files, but not only)
like a to-do list



Operation types:

- ReplicateAndRegister (e.g. using FTS)
- RemoveFile/RemoveReplica
- Issue an RPC call to a DIRAC service
- ...add your own

A generic system for queueing similar *operation types* on certain *datasets* and forward them to the appropriate *systems*

An *operation type* can be, e.g.:

- a simulation workflow
- a reconstruction workflow
- a replication
- a removal
- ...

A *dataset* is split into groups, based on criterias defined by *plugins*, e.g.:

- split by size
- by destination
- by metadata
- ... [code it]

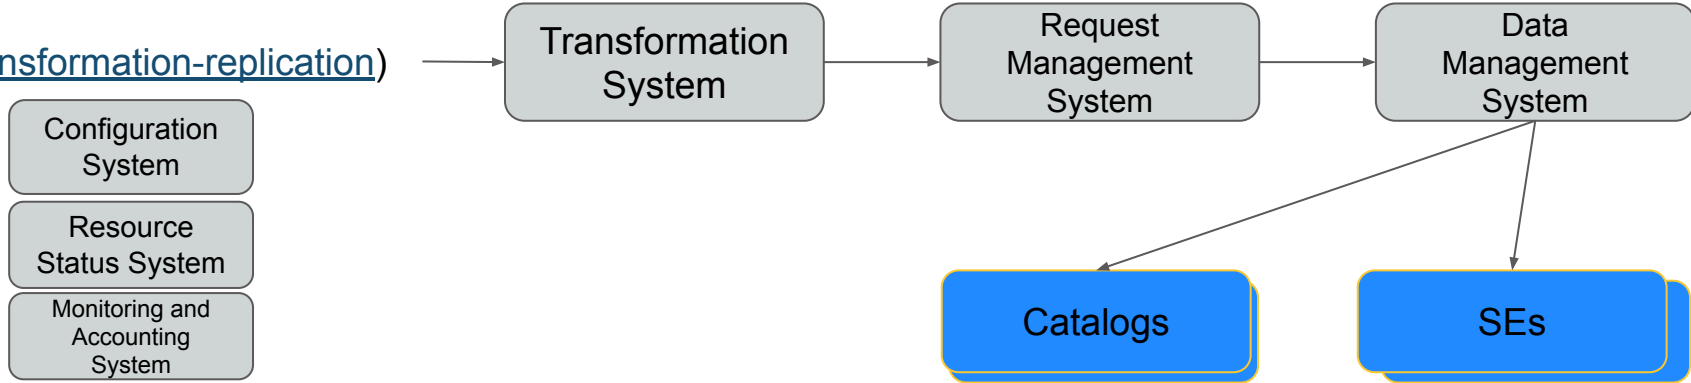
A *system* is either (today) the DIRAC WMS (for productions) or the DIRAC RMS (for dataset management operation types)

[DMS] example (for dataset management): Take all my holidays pictures from 2018 with tag='sunset', make sure that there is one copy on tape and one on disk, distributed on all the sites according to free space, and group the operations by group of at most 100 files.

[WMS] example (for jobs productions): Take all my holidays pictures from 2018 with tag='sunset', make sure to run (only once) the 'red-enhancer' workflow on each one of them, using only Tier2 sites.

[DMS] Dataset management

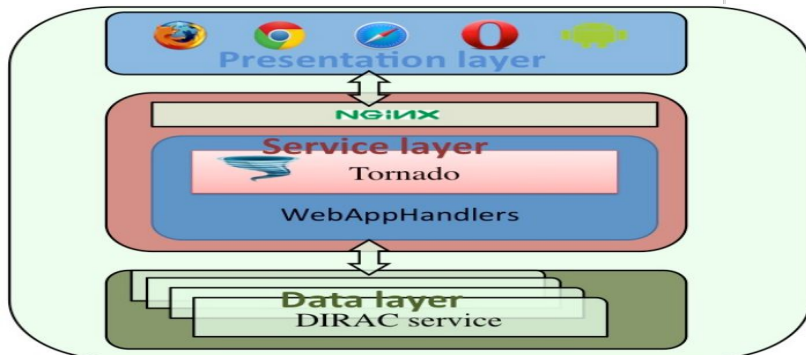
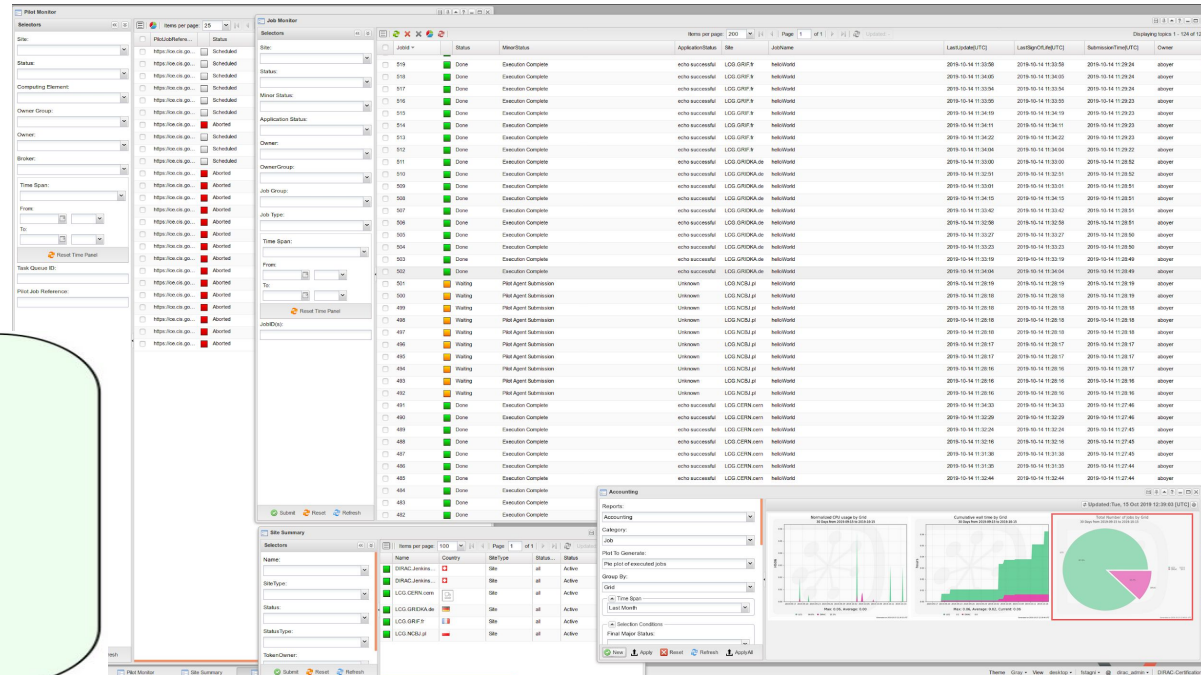
([dirac-transformation-replication](#))



Selectors		Items per page: 100															Page 1 of 3		Showing 1010-1016 of 20200000									
Status	Agent Type	Type	Name	Files	Processed (%)	Created	Total Created	Submitted	Matched	Checking	Waiting	Staging	Rescheduled	Killed	Running	Scheduled	Done	Completed	Failed									
Active	Automatic	Replication	Replication...	186670	57.2	0	186670	0	0	0	0	0	0	0	0	0	1078	0	0									
Active	Automatic	Replication	Replication...	388731	99.1	0	388731	0	0	0	0	0	0	0	0	33	3889	0	2									
Idle	Automatic	Replication	Replication...	29	100.0	0	14	0	0	0	0	0	0	0	0	14	0	0	0									
Idle	Automatic	Replication	Replication...	316	100.0	0	13	0	0	0	0	0	0	0	0	10	0	3	0									
Idle	Automatic	Replication	Replication...	2216	100.0	0	946	0	0	0	0	0	0	0	0	944	0	2	0									
Idle	Automatic	Replication	Replication...	6	100.0	0	6	0	0	0	0	0	0	0	0	6	0	0	0									
Idle	Automatic	Replication	Replication...	6	100.0	0	1	0	0	0	0	0	0	0	0	1	0	0	0									
Idle	Automatic	Replication	Replication...	18	100.0	0	1	0	0	0	0	0	0	0	0	1	0	0	0									
Idle	Automatic	Replication	Replication...	589	100.0	0	10	0	0	0	0	0	0	0	0	10	0	0	0									
Idle	Automatic	Replication	Replication...	914	100.0	0	10	0	0	0	0	0	0	0	0	10	0	0	0									
Idle	Automatic	Replication	Replication...	176	100.0	0	2	0	0	0	0	0	0	0	2	0	0	0	0									
Manual	Manual	Replication	Replication...	322	100.0	0	0	0	0	0	0	0	0	0	0	7	0	0	0									
Com...	Manual	Replication	Replication...	1338	100.0	0	14	0	0	0	0	0	0	0	0	14	0	0	0									
Com...	Manual	Replication	Replication...	16	100.0	0	2	0	0	0	0	0	0	0	0	2	0	0	0									
Idle	Automatic	Replication	Replication...	16	100.0	0	7	0	0	0	0	0	0	0	0	7	0	0	0									
Active	Automatic	Replication	Replication...	83	92.4	0	78	0	0	1	0	0	0	0	6	71	0	0	0									
Active	Automatic	Replication	Replication...	3874	99.9	0	906	1	0	0	0	0	0	0	3	902	0	0	0									
Active	Automatic	Replication	Replication...	11659	96.9	0	141	0	0	0	0	0	0	0	1	140	0	0	0									
Active	Automatic	Replication	Replication...	15	100.0	0	2	0	0	0	0	0	0	0	2	0	0	0	0									
Idle	Automatic	Replication	Replication...	311	100.0	0	7	0	0	0	0	0	0	0	0	7	0	0	0									
Active	Automatic	Replication	Replication...	1302	96.8	0	195	1	0	0	0	0	0	0	1	198	0	0	0									
Idle	Automatic	Replication	Replication...	541	100.0	0	9	0	0	0	0	0	0	0	0	9	0	0	0									
Active	Automatic	Replication	Replication...	1760	97.2	0	943	9	0	0	0	0	0	0	42	492	0	0	0									
Idle	Automatic	Replication	Replication...	349	100.0	0	16	0	0	0	0	0	0	0	0	16	0	0	0									
Idle	Automatic	Replication	Replication...	771	100.0	0	70	0	0	0	0	0	0	0	0	70	0	0	0									
Idle	Automatic	Replication	Replication...	189	100.0	0	4	0	0	0	0	0	0	0	0	4	0	0	0									
Active	Automatic	Replication	Replication...	738	100.0	0	23	0	0	0	0	0	0	0	0	23	0	0	0									
Active	Automatic	Replication	Replication...	23441	99.9	0	866	1	1	0	0	0	0	0	4	861	0	0	0									
Idle	Automatic	Replication	Replication...	40	100.0	0	19	0	0	0	0	0	0	0	0	19	0	0	0									
Idle	Automatic	Replication	Replication...	264	100.0	0	15	0	0	0	0	0	0	0	0	15	0	0	0									
Idle	Automatic	Replication	Replication...	179	100.0	0	22	0	0	0	0	0	0	0	0	22	0	0	0									
Active	Automatic	Replication	Replication...	662	99.6	0	29	0	0	0	0	0	0	0	2	27	0	0	0									
Idle	Automatic	Replication	Replication...	101	100.0	0	11	0	0	0	0	0	0	0	0	11	0	0	0									
Active	Automatic	Replication	Replication...	32	96.8	0	14	0	0	0	0	0	0	0	1	13	0	0	0									
Idle	Automatic	Replication	Replication...	603	100.0	0	36	0	0	0	0	0	0	0	0	36	0	0	0									
Active	Automatic	Replication	Replication...	656	99.8	0	245	0	0	0	0	0	0	0	1	244	0	0	0									
Idle	Automatic	Replication	Replication...	392	100.0	0	288	0	0	0	0	0	0	0	0	288	0	0	0									

- Web users' interface
- Frontend: ExtJS6
- Backend: tornado, NGINX

- Each system has its own Web application

The screenshot displays the DIRAC web application interface, showing a detailed view of job monitoring and a summary view.

Job Monitoring View:

Item	JobID	Status	Message	Application/Task	Site	Page	Last Update(UTC)	Last Sync(UTC)	SubmissionTime(UTC)	Owner
516	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:33:56	2018-10-14 11:33:58	2018-10-14 11:29:24	ajoyee	
518	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:34:05	2018-10-14 11:34:05	2018-10-14 11:29:24	ajoyee	
517	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:33:54	2018-10-14 11:33:54	2018-10-14 11:29:24	ajoyee	
516	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:33:55	2018-10-14 11:33:55	2018-10-14 11:29:23	ajoyee	
515	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:34:19	2018-10-14 11:34:19	2018-10-14 11:29:23	ajoyee	
514	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:34:11	2018-10-14 11:34:11	2018-10-14 11:29:23	ajoyee	
513	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:34:22	2018-10-14 11:34:22	2018-10-14 11:29:23	ajoyee	
512	Done	Execution Complete	echo successful	LOG.GRFP	heidelberg	2018-10-14 11:33:56	2018-10-14 11:33:56	2018-10-14 11:29:23	ajoyee	
511	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:33:00	2018-10-14 11:33:00	2018-10-14 11:29:22	ajoyee	
510	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:32:51	2018-10-14 11:32:51	2018-10-14 11:29:22	ajoyee	
509	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:33:01	2018-10-14 11:33:01	2018-10-14 11:29:21	ajoyee	
508	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:33:42	2018-10-14 11:33:42	2018-10-14 11:29:21	ajoyee	
507	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:33:42	2018-10-14 11:33:42	2018-10-14 11:29:21	ajoyee	
506	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:32:58	2018-10-14 11:32:58	2018-10-14 11:29:21	ajoyee	
505	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:32:27	2018-10-14 11:32:27	2018-10-14 11:29:21	ajoyee	
504	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:32:51	2018-10-14 11:32:51	2018-10-14 11:29:21	ajoyee	
503	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:33:19	2018-10-14 11:33:19	2018-10-14 11:29:49	ajoyee	
502	Done	Execution Complete	echo successful	LOG.GRFXA	heidelberg	2018-10-14 11:34:06	2018-10-14 11:34:06	2018-10-14 11:29:49	ajoyee	
501	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:19	2018-10-14 11:28:19	2018-10-14 11:28:19	ajoyee	
500	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:18	2018-10-14 11:28:18	2018-10-14 11:28:19	ajoyee	
499	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:18	2018-10-14 11:28:18	2018-10-14 11:28:19	ajoyee	
498	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:18	2018-10-14 11:28:18	2018-10-14 11:28:19	ajoyee	
497	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:18	2018-10-14 11:28:18	2018-10-14 11:28:19	ajoyee	
496	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:17	2018-10-14 11:28:17	2018-10-14 11:28:17	ajoyee	
495	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:16	2018-10-14 11:28:16	2018-10-14 11:28:17	ajoyee	
494	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:16	2018-10-14 11:28:16	2018-10-14 11:28:17	ajoyee	
493	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:16	2018-10-14 11:28:16	2018-10-14 11:28:16	ajoyee	
492	Waiting	File Agent Submission	Unknown	LOG.NCBI.jp	heidelberg	2018-10-14 11:28:16	2018-10-14 11:28:16	2018-10-14 11:28:16	ajoyee	
491	Waiting	File Agent Submission	Unknown	LOG.CERN.com	heidelberg	2018-10-14 11:27:59	2018-10-14 11:27:59	2018-10-14 11:27:59	ajoyee	
490	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:32:36	2018-10-14 11:32:36	2018-10-14 11:27:46	ajoyee	
489	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:32:24	2018-10-14 11:32:24	2018-10-14 11:27:45	ajoyee	
488	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:32:36	2018-10-14 11:32:36	2018-10-14 11:27:45	ajoyee	
487	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:31:36	2018-10-14 11:31:36	2018-10-14 11:27:45	ajoyee	
486	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:32:20	2018-10-14 11:32:20	2018-10-14 11:27:44	ajoyee	
485	Done	Execution Complete	echo successful	LOG.CERN.com	heidelberg	2018-10-14 11:32:44	2018-10-14 11:32:44	2018-10-14 11:27:44	ajoyee	

Summary View:

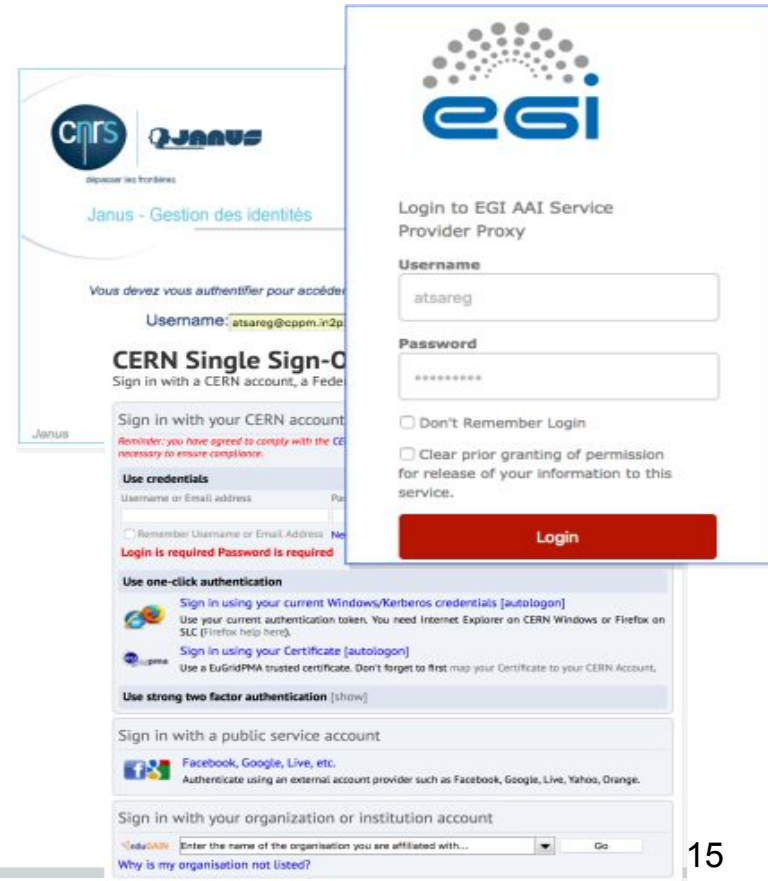
Name	Country	Site	State
DIRAC.Jaipur	IN	site	all
DIRAC.Jinju	KR	site	all
LOG.CERN.com	CH	site	all
LOG.BRE2A.be	BE	site	all
LOG.GRFP	DE	site	all
LOG.NCBI.jp	JP	site	all

The interface also includes a sidebar with navigation options like 'Home', 'Jobs', 'Reports', 'Help', and 'Logout'. A 'Summary' panel on the right shows various charts and reports, including a 'Job Summary' and 'Accounting' section.

Recent and ongoing developments

Until “yesterday”: X509 certificates, DIRAC groups, proxies, VOMS

- DIRAC can delegate AuthN to an external server
 - ensure provisioning of X509 certificate proxies
- Focus: OAuth/OIDC as “industry standards”
 - Use case: [EGI Check-in](#) SSO hub



Janus - Gestion des identités

Vous devez vous authentifier pour accéder à ce service.

Username:

Password:

Don't Remember Login

Clear prior granting of permission for release of your information to this service.

Janus

CERN Single Sign-On
Sign in with a CERN account, a Federated account or your own CERN account.

Sign in with your CERN account

Remember: you have agreed to comply with the CERN security policy, which is necessary to ensure compliance.

Use credentials

Username or Email address:

Remember Username or Email Address:

Login is required Password is required

Use one-click authentication

Sign in using your current Windows/Kerberos credentials [autologon]

Sign in using your current authentication tokens. You need Internet Explorer on a CERN Windows or Firefox on SLC (Firefox help here).

Sign in using your Certificate [autologon]

Use a EuGridPMA trusted certificate. Don't forget to first map your Certificate to your CERN Account.

Use strong two factor authentication [show]

Sign in with a public service account


Facebook, Google, Live, etc.

Authenticate using an external account provider such as Facebook, Google, Live, Yahoo, Orange.

Sign in with your organization or institution account

eduGATE Enter the name of the organisation you are affiliated with...

Why is my organisation not listed?



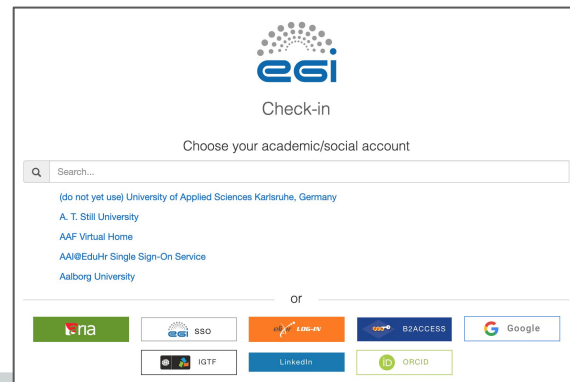
Web portal authentication

CLI authentication

```
[dirac@ce-emi pro]$ python DIRAC/FrameworkSystem/scripts/dirac-proxy-init.py -O CheckIn -g training_user -q
OAuth authentication from CheckIn.
Use link to authentication..
https://ce-emi.bitp.kiev.ua:9943/oauth2/oauth?getlink=M27xN04yMYTz9Vw2wkpBBHm30z8f

Waiting 3.0 minutes when you authenticated... * * * [3-

Proxy generated:
subject      : /DC=org/DC=uprid/O=people/O=BITP/CN=Andrey Litovchenko/CN=3461819742
issuer       : /DC=org/DC=uprid/O=people/O=BITP/CN=Andrey Litovchenko
identity     : /DC=org/DC=uprid/O=people/O=BITP/CN=Andrey Litovchenko
timeleft     : 23:59:59
DIRAC group  : training_user
rfc          : True
path         : /tmp/x509up_u3318
username     : allitov
```



- `dips://` → `https://`
 - `dips`: proprietary protocol for RPC calls
 - `http`: frameworks already exists in python 2&3 for server-side (tornado) and client side (requests)
- Python 3
 - Migration started, first production release next year
 - DIRAC Pilot will move first (also b/c of CentOS8)

- **“DIRACOS: a cross platform solution for Grid tools”**
7/11, 11:00, R3

- **“A gateway between Gitlab CI and DIRAC”**
Poster

- diracgrid.org
- dirac.readthedocs.io
 - including [code documentation](#)
- Ops and general questions: Google [forum](#)
- Dev and DevOps issues: on [github](#)
- Bi-weekly developers meetings (and/or hackathons): [BILD](#)



The 10th
DIRAC Users' Workshop

25th - 29th May 2020

KEK/IPNS
TSUKUBA

KEK  DIRAC
THE INTERWARE

[HTTPS://INDICO.CERN.CH/E/DUW10](https://indico.cern.ch/e/DUW10)

Organizers: FEDERICO STAGNI (CERN)
TAKANORI HARA (KEK/IPNS)
IKUO UEDA (KEK/IPNS)
ANDREI TSAREGORODTSEV (IN2P3)

diracgrid.org DIRACGrid dirac-grid dirac.readthedocs.io

10th DIRAC
Users' Workshop


indico.cern.ch/e/DUW10

25-29 May 2020
KEK, Japan


?

... a few examples of what DIRAC can be used for

- sending jobs to “the Grid”
 - the obvious one...
 - interfacing with different *sites*
 - with different *computing elements*
 - and *batch systems*
 - with different *storage elements*
 - interfacing with different *information systems*
 - interfacing with different *catalogs*
 - interfacing with different *MQs, DBs*
 - authenticate through different *providers*
-
- managing “productions” (e.g. reconstruction, simulation...)
 - managing dataset transfers
 - and removals...
 - providing a failover system
 - your jobs won’t fail because a certain SE is down, nor because of central service are down
 - transfer data from the experiment to a Grid SE
 - monitor your resources with a policy-based system
 - ... and more





- **Computing**
 - **CEs**: ARC, CREAM, HTCondor, “SSH” for standalone BS, ...
 - **Batch**: LSF, BQS, SGE, PBS/Torque, SLURM, Condor,...
 - Clouds, BOINC, HPC, “desktops”
- **Storage**
 - SRM2, GSIFTP, XRoot, http, DIPs, ...
 - EOS, Castor, DPM, dCache, StoRM, ECHO, CTA, ...
- **Catalog**
 - DIRAC FC, LFC, (Rucio), [LHCb Bookkeeping], ...
- **Information services**
 - BDII, GOCDB, CRIC...
- **IdProviders**
- **ProxyProviders**
 - VOMS, OAuth2, PUSP...
- **DBs, MQs, LogBackends**
 - MySQL, Oracle, ElasticSearch
 - stomp → ActiveMQ, RabbitMQ
 - file, MQ, ES
 - and logs centralization is easy to set up




LHCb activities in the last year

- HLTfarm works even during data taking
- LHCb 25%
- CERN
- UK, IT, FR, DE
- RU, PL, CH,



- MCSimulation
- MCFastSimulation
- User
-


Belle II computing performance in a year



- Consumed CPU power is not so different from last year
- Resource usage is getting practical
 - Increasing analysis jobs
 - Increasing raw (beam and cosmic ray) data processing

From last DIRAC users' workshop


Usage




2.3k CPU Years, 11.8M Jobs

DIRAC usage since last year

- MC production and analysis running in parallel
- 144 M HS06 hours
- 1.9 M executed jobs



Business as usual - just more of it (Success!)



Experiment agnostic, and extensibility

