

DIRAC for IHEP distributed computing

Xiaomei ZHANG
Institute of High Energy Physics

Virtual DIRAC User Workshop

May 11, 2021

What do you use DIRAC for?

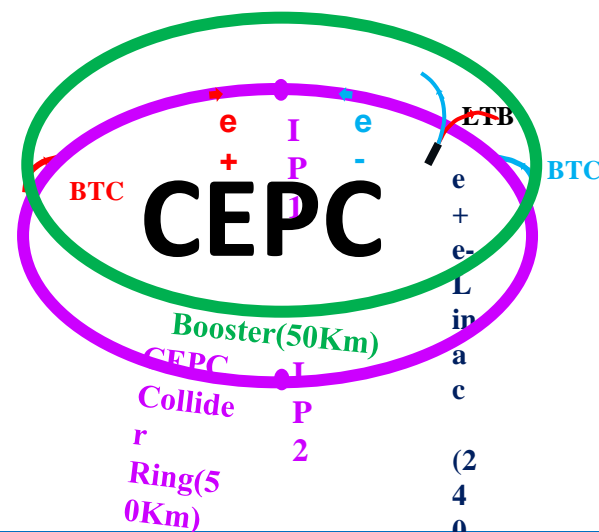
- **The IHEP Distributed Computing built on DIRAC is used for three experiments now**
 - BESIII : Beijing Spectrometer III at BEPCII
 - JUNO : Jiangmen Underground Neutrino Observatory
 - CEPC : Circular Electron Positron Collider
- **Mainly support Monte Carlo production and data transfer, also for reconstruction and user analysis**

The logo for BESIII, featuring the letters 'B', 'E', 'S', and 'III' in a stylized font. 'B' is blue, 'E' is red, 'S' is green, and 'III' is black.

BESIII (Beijing Spectrometer III at BEPCII)

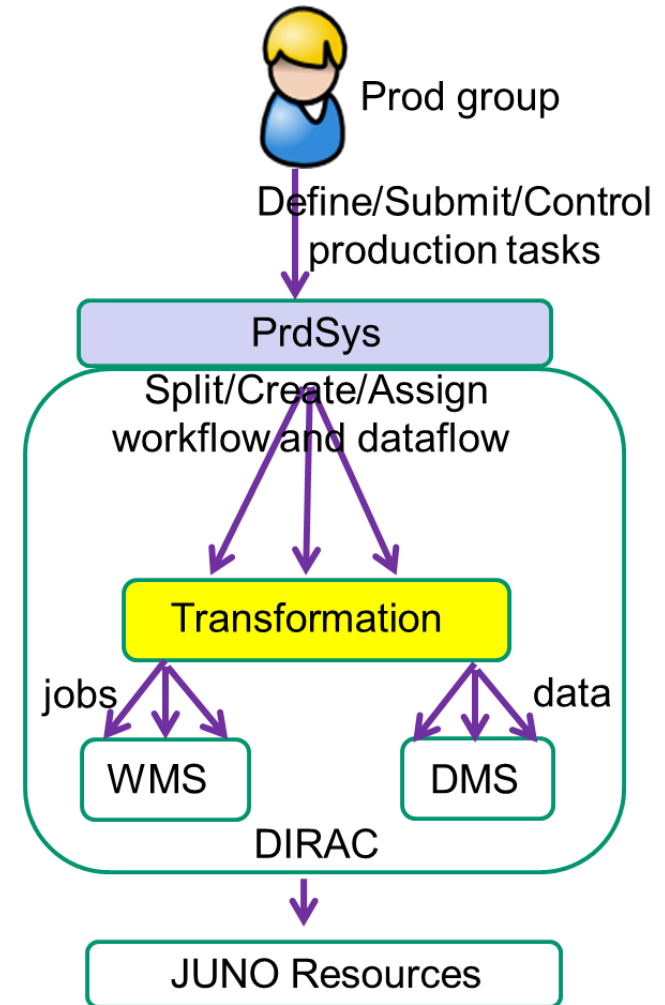


JUNO (Jiangmen Underground Neutrino Observatory)



The DIRAC functionalities used (1)

- Most of DIRAC functionalities are used by IHEP
- Workload management system for integrating resources
 - Cluster: HTCondor, SLURM, PBS
 - Grid: ARCCE, HTCondorCE
 - Cloud: OpenNebula, OpenStack
- DIRAC File Catalog (DFC) for metadata and replica catalog
- DFC + RMS + Transformation + WMS
 - Production system for MC simulation and raw reconstruction
 - Bulk data replication/removal



The DIRAC functionalities used (2)

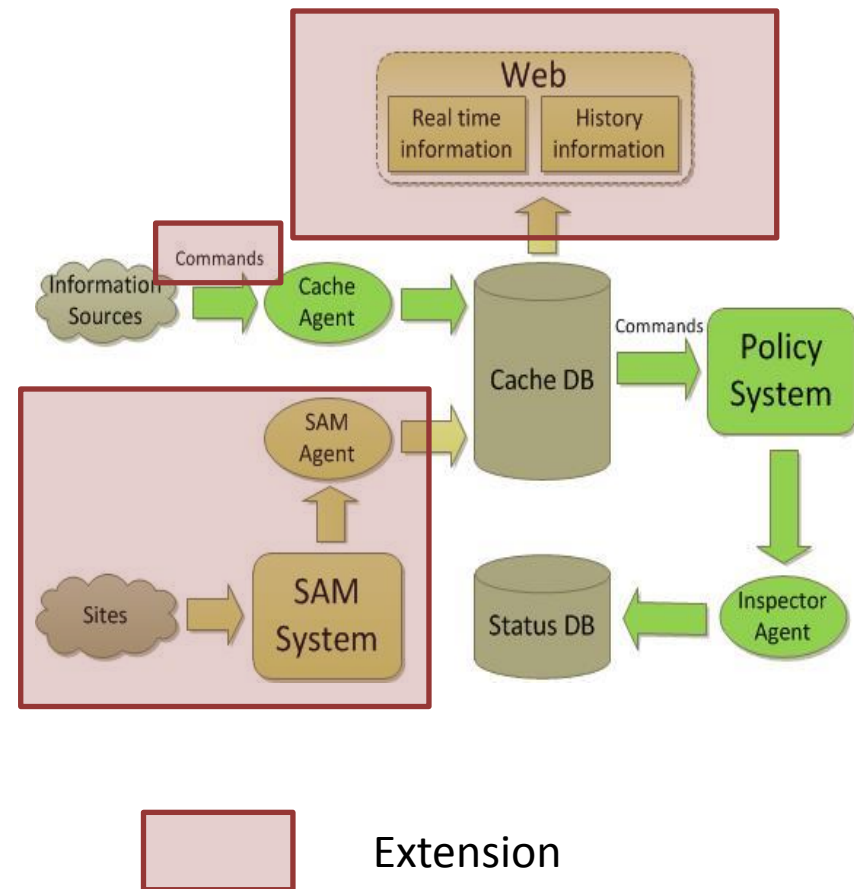
- Monitoring and Accounting for tracking problems and providing operation reports
- Resource Status System with extension to monitor site status and usage
- Multi-VO feature to support three experiments in one instance
- SingularityCE feature to give a unique environment for all the sites
- PoolCE and tag feature to support GPU resource and multi-core jobs (Not in production yet)

DIRAC functionalities not used, and why?

- Production System
 - Already has own production interface
 - After upgrading to v7, will give a try
- Workflow
 - Currently a simple implementation will do
 - Will improve code with it when manpower available
- Storage Management System
 - Depend on site choices
 - Currently use dCache, EOS, StORM, not DIRAC SE

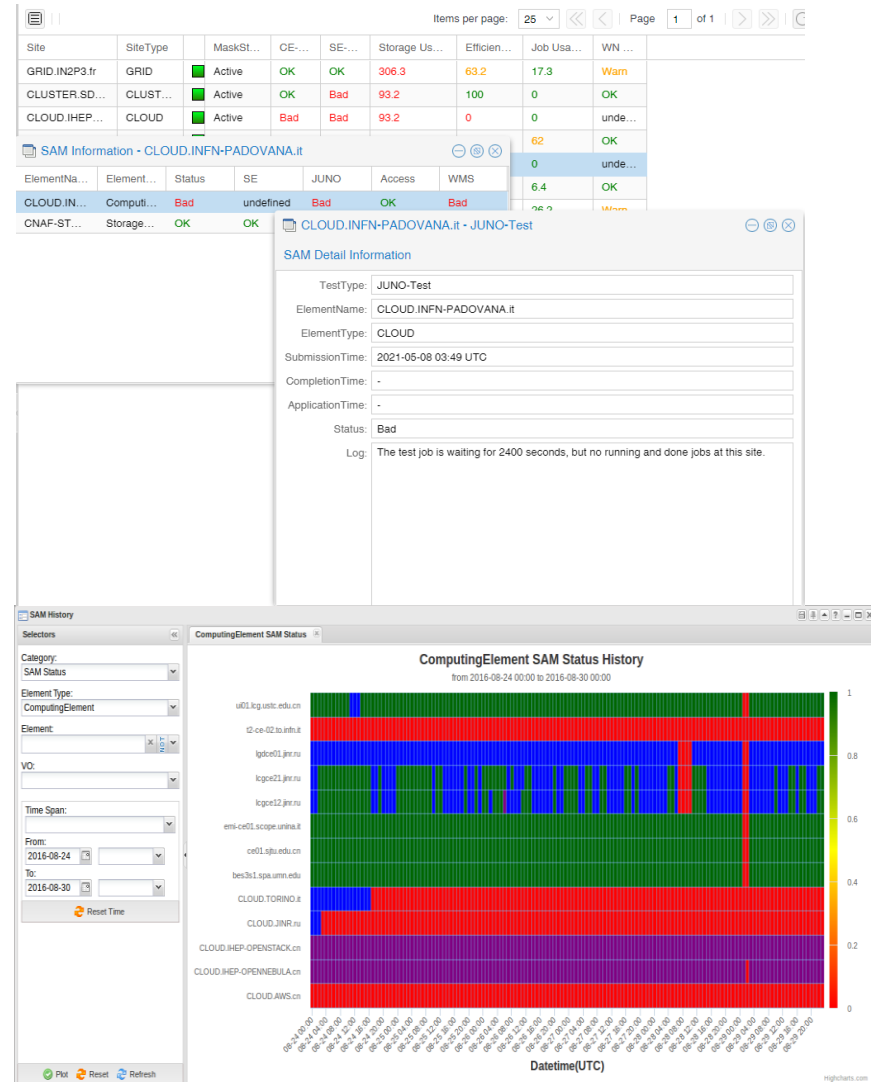
Do you have a DIRAC extension? Why?

- IHEP DIRAC extension – IHEPDIRAC
 - Site status monitoring - extension to RSS
 - Purpose:
 - Monitor status of site and resource
 - Track site problems
 - Give summary reports
 - Reuse architecture of RSS
 - Add SAM tests
 - Routinely send tests using experiment data processing examples, eg MC, Rec, Ana...
 - Return detailed logs to help track problems of tests



Site status monitoring

- Add specific information collecting commands
 - Depending on experiment resource needs, eg. some sites are not standard grid sites
- Add web pages
 - Show detailed site information from tests
- Add historic views
 - Give regular and annual reports of site status
 - Help track site problems
 - Think about using Elasticsearch in future



User task monitoring

- Monitor, summarize and operate individual user tasks
 - Collect statistic info of use tasks in one page
 - Provide reschedule and delete in one command
 - Used together with user submission tools

The screenshot displays a web-based interface for monitoring user tasks. On the left, a 'Selectors' sidebar allows filtering by status (Ready, Processing, Finished, Expired, Init), owner, and time span. The main area is a table of tasks with columns for TaskId, TaskName, Status, Jobs, Progress, CreationTime, UpdateTime, and Site. A context menu is open over task 41, offering actions like Progress, Information, History, Show Jobs, Jobs Information, Activate, Rename, Reschedule Failed Jobs, Reschedule All Jobs, and Delete. Two pop-up windows are visible: 'Information for task 41' showing metadata like BossVersion (6.6.4.p03) and Dataset (Prod_jpsi_664p03_rhopi_round02_9947_9970_stream063_root), and 'Progress for task 41' showing job status counts: Total (43), Done (0), Failed (0), Running (17), Waiting (26), and Deleted (0).

TaskId	TaskName	Status	Jobs	Progress (D F R W O)	CreationTime[UTC]	UpdateTime[UTC]	Site
41	aws_test_c3	Processing	43/43	0 0 17 26 0	2015-06-01 02:09:59	2015-06-01 03:24:45	CLOUD.AWS.cn
40	aws_test_m3	Finished	43/43	34 9 0 0 0	2015-05-31 06:34:33	2015-06-01 03:24:45	CLOUD.AWS.cn
39	aws_test_c3	Finished	23/23	20 3 0 0 0	2015-05-29 03:42:35	2015-06-01 03:24:45	CLOUD.AWS.cn
38	aws_test	Finished	1/1	1 0 0 0 0	2015-05-26 04:43:41	2015-06-01 03:24:45	CLOUD.AWS.cn
37	aws_test	Expired	0/4	0 0 0 0 4	2015-05-26 04:42:23	2015-05-26 04:44:29	CLOUD.AWS.cn
34	aws_test_c3	Finished					
33	aws_test_c3	Finished					
31	aws_test_m3	Expired					
30	aws_test_m3	Expired					
29	aws_test_m3	Expired					
28	sra_rhopi	Expired					
27	aws_test_t2	Expired					
26	aws_test_t2	Expired					
25	sra_rhopi	Expired					
24	sra_rhopi	Expired					
23	sra_rhopi	Expired	0/13	0 0 0 0 13			
22	sra_rhopi	Expired	0/13	0 0 0 0 13			
21	sra_rhopi	Expired	0/0	0 0 0 0 0			
20	sra_rhopi	Expired	0/6	0 0 0 0 6			
19	sra_rhopi	Expired	0/6	0 0 0 0 6			
18	sra_rhopi	Expired	0/6	0 0 0 0 6	2015-05-07 05:13:43	2015-05-21 05:30:48	GRID.INFN-ReCas.it,GRID.JINR.ru
17	sra_rhopi	Expired	0/6	0 0 0 0 6	2015-05-07 04:53:47	2015-05-21 05:48:48	CLUSTER.UCAS.cn

Do you think some of it could become part of the vanilla projects?

- Not much in extension
- Maybe part of site status monitor can be used for others
 - Regularly send SAM tests and give history reports
 - In future should combine with RSS

What is your biggest frustration with DIRAC?

- We are happy with DIRAC when everything goes smoothly
- When problems, dealing with FTS3Agent and RequestExecutingAgent could be a pain
 - Some logs are not clear to understand problems, confused with status of transfers when there are problems

add one feature to DIRAC, what is it?

- No magic, Just some needs
- Computing resource information system to supervise CPU resources pledged by sites in MoU
 - Store pledges of sites
 - Collect and count CPU resource in HEPSpec06 with DIRAC standard benchmark?
 - Check site service level with pilots?
- Provide a unique user-customized software environment
 - With extensions to SingularityCE?

Any notable operations incident in the last year?

- The CERN web server problem causes all the jobs failed
 - Connect problems to diracproject.web.cern.ch

To support your "Grid", do you have to use other systems than DIRAC?

- For us, now we just use DIRAC for anything except
 - Basic grid facilities, eg. CA, VOMS, FTS3, SE, CVMFS, ES
 - Frontier/Squid system for accessing condition data
 - Job submission tools for user analysis

How would you rate the communication?


- We are satisfied with current communication
- When I met problems
 - Search through diracgrid-forum@googlegroups.com to see if the problem has been solved before
 - If not clear enough, I try to submit the problem to the forum
 - Also I directly write to DIRAC experts
 - Also bi-weekly meeting is a good place if you understand and explain your problem well
- User workshop is good for communications between experiments
 - Not only learning from DIRAC experts, but learning from experiments is also useful

DIRAC usage in last two years (1)

- DIRAC is used mainly in the following activities in the last two years
 - Official MC production
 - Data transfers among sites
 - Private MC production and analysis
- There are ~3M jobs done in DIRAC and Normalized CPU usage is 1.12MHS06 days for the last two years

Cumulative Jobs by Site

105 Weeks from Week 17 of 2019 to Week 18 of 2021



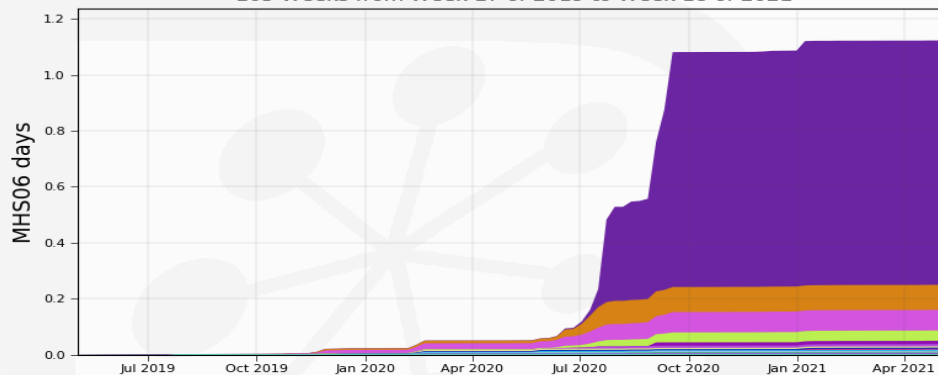
Max: 3.13, Min: 0.01, Average: 1.41, Current: 3.13

■ CLOUD.JINRONE.ru	0.5	■ CLOUD.INFN-PADOVANA.it0.2	■ GRID.INFN-ReCas.it	0.1	■ DIRAC.Client.it	0.0	
■ GRID.INFN-CNAF.it	0.3	■ GRID.JINR-CONDOR.ru	0.2	■ CLOUD.JINR.ru	0.1	■ DIRAC.Client.cn	0.0
■ CLUSTER.USTC.cn	0.3	■ GRID.MANCHESTER.uk	0.1	■ CLUSTER.UMN.us	0.0	■ ANY	0.0
■ CLOUD.IHEPCLOUD.cn	0.3	■ GRID.JINR.ru	0.1	■ CLUSTER.SJTU.cn	0.0	■ Multiple	0.0
■ GRID.IN2P3.fr	0.3	■ GRID.QMUL.uk	0.1	■ CLUSTER.SDU.cn	0.0		
■ CLUSTER.NEUI.tr	0.2	■ CLUSTER.IPAS.tw	0.1	■ GRID.LANCASTER.uk	0.0		
■ GRID.IHEP.cn	0.2	■ CLUSTER.IHEP-CONDOR.cn0.1	■ GRID.JINR-ARC.ru	0.0			

Generated on 2021-05-07 07:45:58 UTC

Normalized CPU used by Site

105 Weeks from Week 17 of 2019 to Week 18 of 2021



Max: 1.12, Min: 0.00, Average: 0.42, Current: 1.12


■ CLOUD.JINRONE.ru	0.9	■ CLOUD.JINR.ru	0.0	■ CLUSTER.SJTU.cn	0.0
■ GRID.IN2P3.fr	0.1	■ CLUSTER.USTC.cn	0.0	■ GRID.LANCASTER.uk	0.0
■ CLOUD.IHEPCLOUD.cn	0.1	■ GRID.INFN-ReCas.it	0.0	■ CLUSTER.IHEP-CONDOR.cn	0.0
■ CLOUD.INFN-PADOVANA.it	0.0	■ CLUSTER.SDU.cn	0.0	■ Multiple	0.0
■ GRID.IHEP.cn	0.0	■ CLUSTER.NEUI.tr	0.0	■ GRID.MANCHESTER.uk	0.0
■ GRID.INFN-CNAF.it	0.0	■ GRID.JINR-CONDOR.ru	0.0	■ DIRAC.Client.it	0.0
■ CLUSTER.UMN.us	0.0	■ GRID.QMUL.uk	0.0	■ DIRAC.Client.cn	0.0
■ GRID.JINR.ru	0.0	■ ANY	0.0		
■ GRID.JINR-ARC.ru	0.0	■ CLUSTER.IPAS.tw	0.0		

Generated on 2021-05-07 08:05:32 UTC

DIRAC usage in last two years (2)

- Transfers through FTS3 with DIRAC DMS are 655TB
- The maximum speed can reach beyond 6Gb/s

Transferred data by Channel
105 Weeks from Week 17 of 2019 to Week 18 of 2021

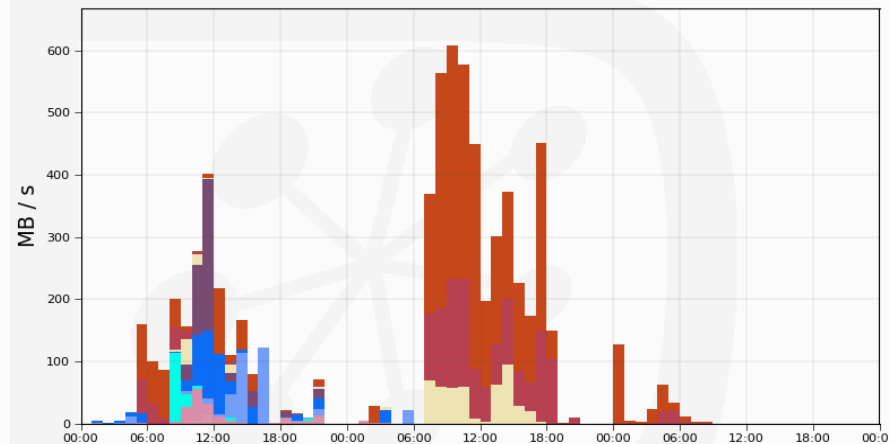


Max: 655, Min: 0.28, Average: 261, Current: 655

JJNR-EOS -> IHEP-STORM	183.9	CNAF-STORM -> IN2P3-DCACHE	12.4	JJNR-JUNO -> IN2P3-DCACHE	4.6
CNAF-STORM -> IHEP-STORM	137.8	IHEP-STORM -> JJNR-JUNO	11.5	JJNR-JUNO -> IHEP-STORM	0.5
IHEP-STORM -> JJNR-EOS	126.9	IHEP-STORM -> IN2P3-DCACHE	10.6	IN2P3-DCACHE -> JJNR-JUNO	0.0
IN2P3-DCACHE -> IHEP-STORM	63.3	JJNR-JUNO -> CNAF-STORM	7.2	CNAF-STORM -> JJNR-JUNO	0.0
IHEP-STORM -> CNAF-STORM	38.3	JJNR-EOS -> CNAF-STORM	7.1	JJNR-EOS -> JJNR-JUNO	0.0
IN2P3-DCACHE -> CNAF-STORM	24.6	CNAF-STORM -> JJNR-EOS	6.0	IHEP-STORM -> IHEP-EOS	0.0
JJNR-EOS -> IN2P3-DCACHE	15.3	IN2P3-DCACHE -> JJNR-EOS	5.2		

Generated on 2021-05-07 07:50:05 UTC

Throughput by Channel
72 Hours from 2021-01-14 00:00 to 2021-01-17 00:00 UTC



Max: 608, Average: 97.4

IN2P3-DCACHE -> IHEP-STORM	49.6%	IN2P3-DCACHE -> JJNR-EOS	4.8%
CNAF-STORM -> IHEP-STORM	20.8%	CNAF-STORM -> IN2P3-DCACHE	2.1%
JJNR-EOS -> IHEP-STORM	7.7%	CNAF-STORM -> JJNR-EOS	1.6%
IHEP-STORM -> IN2P3-DCACHE	6.4%	JJNR-EOS -> IN2P3-DCACHE	0.8%
IHEP-STORM -> JJNR-EOS	6.2%		

Generated on 2021-01-19 06:41:44 UTC

Summary

- DIRAC is quite convenient for experiments without enough manpower but want to use “Grid”
 - Provide almost everything for us
- With it, we can have more time to focus on our own experiment specific needs and problems
- Thank help from DIRAC society, especially Federico, André, Christophe, Luisa, Ueda, Andre.....
- We are happy to be in the DIRAC family!