

DIRAC Data Production System

Andrew C. Smith

CERN

October 21, 2010

- 1 Introduction to DIRAC
- 2 DIRAC Data Production System
- 3 Using the Data Production System



- 1 Introduction to DIRAC
- 2 DIRAC Data Production System
- 3 Using the Data Production System

Community Grid Solution

The DIRAC project is a complete Grid solution.

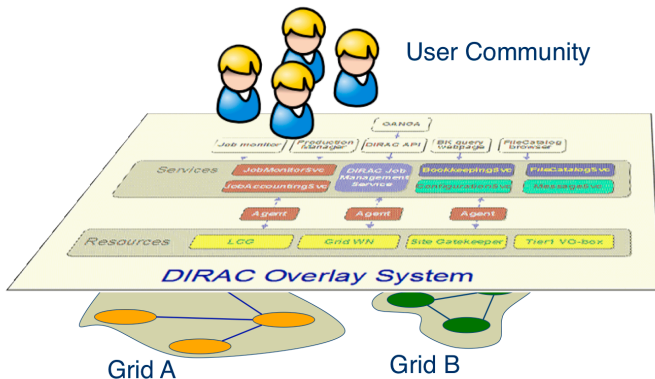
Designed to be used by a community of users.

Services and agents of DIRAC overlay resources.

Transparent use of different grids e.g. gLite, NorduGrid etc.

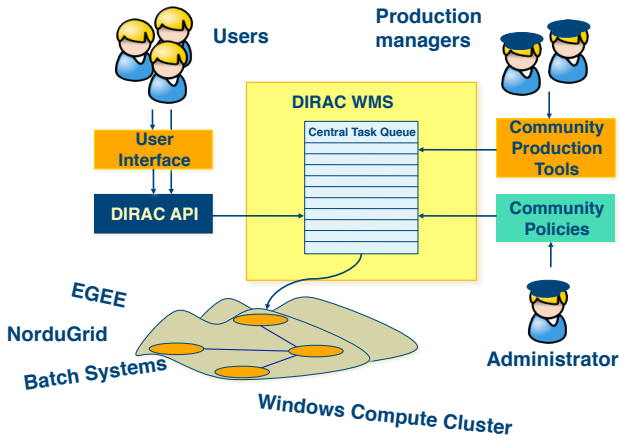
Integration of non-grid resources e.g. local, clouds, batch systems etc.

Grid compliant security framework (OpenSSL with X509 certificates).



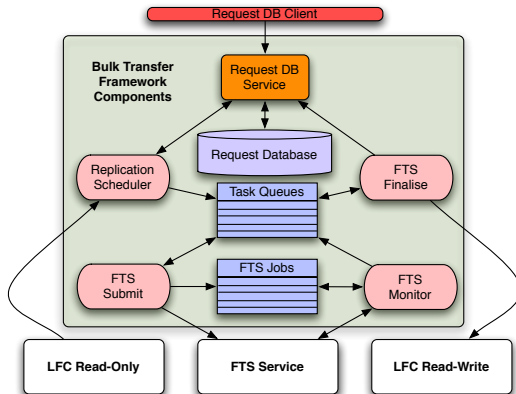
Workload Management

The DIRAC WMS pioneered the use of Pilot Agents.
The Workload is optimised in a central Task Queue.
Allows community priorities and shares to be applied.



DIRAC Data Management System

DIRAC DMS provides tools for interacting with catalogs and storage. Upload, download and replication of data to/from/between storages. Includes bulk transfer framework based on FTS.



Interfaces

- DIRAC API is a Python wrapper used by external projects (for example Ganga).
- DIRAC CLI is based on DIRAC API in Python.
- DIRAC GUI is secure web based interface build in JS.
- Mimics desktop application for short learning curve.

The screenshot displays the DIRAC GUI Job Monitoring interface in a Safari browser. The main window shows a table of jobs with columns for JobID, Status, MinorStatus, and ApplicationStatus. A context menu is open over job 12101554, listing actions like 'JDL', 'Attributes', 'Parameters', and 'Logging info'. A secondary window titled 'JDL for JobID: 12101554' is open, showing the Job Definition Language (JDL) configuration for that job. The JDL includes details such as the executable path, software modules, job name, requirements, owner information, site, and CPU time.

JobID	Status	MinorStatus	ApplicationStatus
12101571	Waiting	Pilot Agent Sub...	Unknown
12101570	Waiting	Pilot Agent Sub...	Unknown
12101569	Waiting	Pilot Agent Sub...	Unknown
12101568	Waiting	Pilot Agent Sub...	Unknown
12101568	Running	Application	Executing Run...
12101565	Waiting	Pilot Agent Sub...	Unknown
12101564	Waiting	Pilot Agent Sub...	Unknown
12101561	Waiting	Pilot Agent Sub...	Unknown
12101559	Running	Downloading In...	Unknown
12101556	Waiting	Pilot Agent Sub...	Unknown
12101554	Waiting	JDL	
12101552	Waiting	Attributes	
12101546	Waiting	Parameters	
12101528	Waiting	Logging info	
12101527	Waiting	Peak StandardOutput	
12101526	Waiting	Get Logfile	
12101525	Waiting	Get PendingRequest	
12101524	Waiting	Get StageReport	
12101523	Waiting	Get Sandbox file	
12101522	Waiting	Actions	
12101521	Waiting	Pilot	

```
JDL for JobID: 12101554
{
  Executable = "/DIRAC007/scripts/dirac-jobexec";
  SoftwareDiatModule = "LHCbDIRAC.Core.Utilities.CombinedSoftwareInstallation";
  JobName = "myPwInsi_[Ganga_PwInsi_(92-107)_v01c1]";
  Priority = 1;
  JobRequirements =
  {
    OwnerDN = "/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=serranoj/CN=67";
    Setup = "LHCb-Production";
    LHCbPlatform =
    {
      "x86_64-slc5-gcc43-opt"
    };
    Site =
    {
      "LOG_CNAP.it",
      "LOG_PAL.us",
      "LOG_ORISWA.de",
      "LOG_CERN.ch"
    };
    CPUtime = 172800;
    OwnerGroup = "lhcb_user";
    UserPriority = 1;
    JobType =
    {
      "User"
    };
  };
}
```

- 1 Introduction to DIRAC
- 2 DIRAC Data Production System**
- 3 Using the Data Production System

Each user community has specific workflows of varying complexity:

- Single application jobs
- Multi-application jobs
- Jobs that require input from previous jobs
- Replication of output data
- Removal of transient data
- etc. etc. etc.

The LHCb computing model outlines several activities:

- Real-time RAW data reconstruction and stripping.
- Regular re-reconstruction and re-stripping exercises.
- Massive Monte Carlo campaigns.
- Replication of data for redundancy and availability.

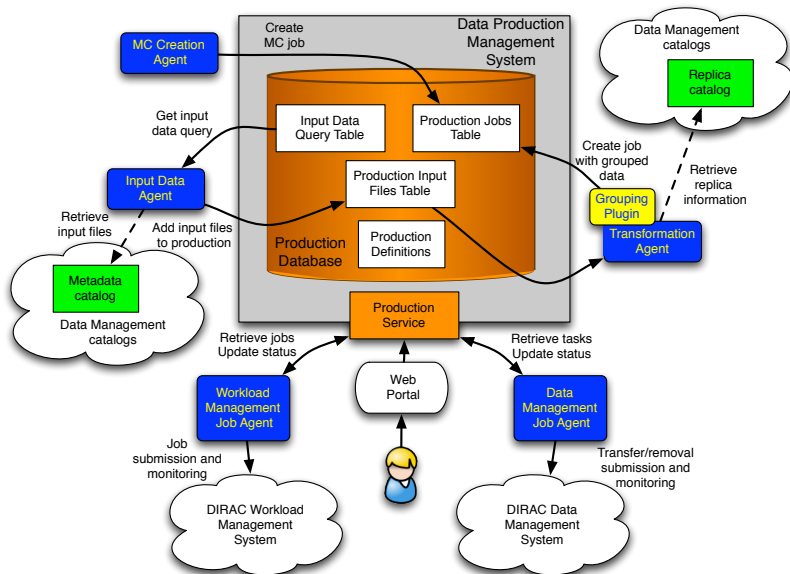
To support this the Data Production System provides:

- Job description library for chaining applications in a single job.
- Job creation based on input data availability.
- Chaining jobs based on output/input data.
- Submission and monitoring of workload and data management jobs.
- Automatic creation of jobs to reach requested amount.
- Post workflow data integrity checking suite.
- Web monitoring and control interface.

What does this system look like

The screenshot displays the DIRAC Production Monitoring system interface. The main window shows a table of production jobs with columns: ID, Status, Agent, Type, Name, Files, Processed (%), Cnstrnd, Submitted, Waiting, Done, and Failed. The table is filtered by 'Request: 1891'. A context menu is open over the first row (ID 8149), showing options like 'Show Jobs', 'Show Request', 'Logging Info', 'Run Status', 'File Status', 'File Error Count', 'Input Data Query', 'Additional Params', 'Show Details', and 'Actions'. The 'Actions' menu is further expanded, showing options like 'Start', 'Stop', 'Extend', 'Flush', 'Complete', and 'Clean'. Three pop-up windows are overlaid on the table, showing detailed statistics for 'Production: 8149'. The first pop-up shows a table of Retries (0, 1, 2) with Count and Percentage. The second pop-up shows a table of Name, Value, and other configuration details. The third pop-up shows a table of Status (Assigned, Processed, Total) with Count and Percentage. The interface includes a navigation bar at the top with 'Safari', 'Archivo', 'Edición', 'Visualización', 'Historial', 'Favoritos', 'Desarrollo', 'Ventana', and 'Ayuda'. The browser address bar shows the URL 'https://lhcbweb.pic.es/DIRAC/LHcb-Production/lhcb_prod/jobs/ProductionMonitor/display'. The status bar at the bottom shows '6 items' and 'Page 1 of 3'.

What does this system look like (really)



Key components - Production database

The central component is the **Production Database**; a simple relational database that persists a full description of a production, its input data and jobs created. It is polled and updated (through the service) by the system agents and clients (including the web GUI).

The key tables are:

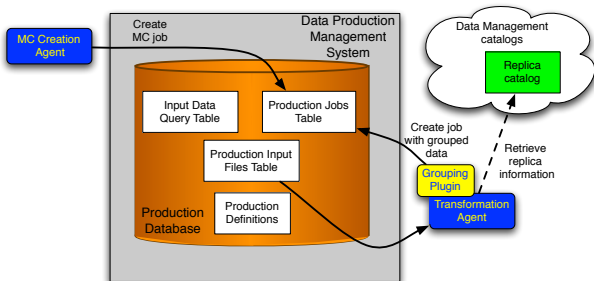
- **Production** - production parameters and job description template.
- **Production Input Query** - metadata query describing the input data.
- **Production Input** - input files and their statuses.
- **Production Jobs** - created jobs, their parameters and status.
- **Production Job Inputs** - mapping production inputs to jobs.

Further auxiliary tables also exist to store any additional parameters and log the production history.

Key components - Job creation agents

The **MC Creation Agent** creates jobs for productions without input data.

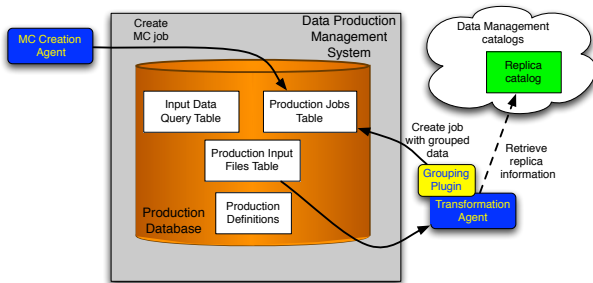
- Retrieves required number of successful jobs from definition.
- Determines the number of jobs to create based on current status.
- Maximum number of jobs created per iteration to avoid flooding.



Key components - Job creation agents

The **Transformation Agent** creates jobs for productions with input data.

- Retrieves unused input data for productions.
- Queries replica catalog for input data location.
- Applies *transformation plugin* to group input data.
- e.g. by size, by run, by file type, according to shares.
- Jobs are created for each of the resulting groups.

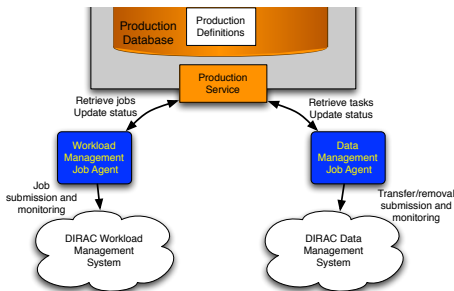


Key components - Job submission and monitoring agents

The **Workload Management Job Agent** and **Data Management Job Agent** submit and monitor jobs destined for the WMS and DMS respectively.

- Created jobs are retrieved for their respective productions.
- Job specific parameters are applied to the job description template.
- Submitted job identifiers are inserted into the database.
- Identifiers used for subsequent updates of statuses.

The code required for interacting with the WMS/DMS is contained in specific plugins. Allows easy extension for use with other systems.



The mutual consistency of data management catalogs is imperative.

To reduce data integrity problems a full consistency check is performed:

- The output directories for the production are retrieved
- For each directory the catalog contents are verified against the storage.
- The storage contents are verified against the catalog.
- Problematic files are stored and invisible to user queries.
- Resolution methods are applied to problematic files.

Checks can be triggered at any time by changing production status.

- 1 Introduction to DIRAC
- 2 DIRAC Data Production System
- 3 Using the Data Production System

Current users of DIRAC

DIRAC developed by LHCb to support their computing model.

Evaluated by other projects including ILC/LCD, Belle, EELA/GISELA.

New communities bring new use-cases:

- Amazon EC2 resource integration for Belle
- MPI job support for EELA

Adoption by new communities simple because of modular design.

Become part of the production infrastructure in some grid segments:

- Latin American - GISELA
- NGI France
- NGI Spain

Monte-Carlo generation

Requires three productions (bullets 1-3).


- 1 Three application-step jobs to generate, digitise and reconstruct events.
- 2 Single application-step jobs submitted to merge files.
- 3 Replication of merged files to two Tier1s.
- 4 Data integrity check of all merged files.
- 5 Removal of all un-merged files.

Productions 2 and 3 take their input data from the previous production.

LHCb workflows - Monte-Carlo generation

Last year MC productions

52 Weeks from Week 41 of 2009 to Week 41 of 2010



00005606	3.6%	00006994	1.4%	00006381	1.0%	00006548	1.0%
00005608	3.1%	00006991	1.3%	00006384	1.0%	00006545	1.0%
00005456	3.0%	00006988	1.3%	00006378	1.0%	00006327	1.0%
00007264	2.5%	00006985	1.3%	00007304	1.0%	00007098	1.0%
00005502	2.2%	00005499	1.1%	00007307	1.0%	00007101	1.0%
00006366	2.0%	00006375	1.1%	00005951	1.0%	00005895	0.9%
00006372	2.0%	00006341	1.1%	00005984	1.0%	00006363	0.9%
00007267	1.6%	00006344	1.1%	00005987	1.0%	00005474	0.8%
00007270	1.5%	00006387	1.0%	00007095	1.0%	...	plus 590 more

Generated on 2010-10-20 11:59:33 UTC

Real-time RAW data taking

Requires four productions (bullets 1,3,5,6).


- 1 Replication of RAW data from CERN to Tier1s.
- 2 Target sites chosen by shares (controlled by dedicated plugin).
- 3 Two application-step jobs submitted on RAW to reconstruct and strip.
- 4 Jobs sent to sites according pledged CPU.
- 5 Two application-step jobs submitted on stripped files to merge and tag.
- 6 Replication of merged and tag files to all LHCb Tier1s.
- 7 Data integrity check of all merged and tag files.
- 8 Removal of all un-merged files.

Productions 3, 5 and 6 take their input data from the previous production.

LHCb workflows - RAW data processing

Last six months processing productions

26 Weeks from Week 16 of 2010 to Week 42 of 2010



00007590	14.5%	00006566	1.5%	00006565	1.0%	00008121	0.5%
00007541	12.1%	00006394	1.4%	00006980	1.0%	00008019	0.4%
00007951	9.8%	00008117	1.4%	00006981	0.9%	00008003	0.4%
00007554	8.8%	00006567	1.4%	00008087	0.9%	00007943	0.4%
00007200	7.7%	00006510	1.3%	00006563	0.8%	00008016	0.4%
00007140	7.3%	00006982	1.2%	00006359	0.8%	00007085	0.4%
00007567	6.3%	00006508	1.2%	00008037	0.6%	00006357	0.4%
00006392	5.0%	00006979	1.2%	00006312	0.6%	00006473	0.4%
00007421	1.5%	00006335	1.1%	00006234	0.6%	...	plus 92 more

Generated on 2010-10-20 12:02:29 UTC

The ILC/LCD group at CERN required massive production of MC events.

Adoption of DIRAC for grid production and analysis.

- Configuration of workflow was done in a couple of weeks.
- Initially, small scale productions submitted as “user” jobs.
- Allowed rapid development and testing of job descriptions.
- After six months production workflows moved to Data Production System.

Flexibility of the system allows to add new functionality very quickly.

- Users from outside CERN now using it.
- Prospects to make it the general grid solution for the ILC community.

For more information see here:

- [DIRAC project page](#)
- [DIRAC Overview at CHEP09](#)
- [DIRAC Pilot Framework](#)
- [DISET Secure Framework](#)
- [Secure Web Interface](#)
- [DIRAC MPI in EELA2](#)
- [Usage of DIRAC by other communities](#)
- [Belle experience with Amazon EC2 - PS50-2-483/PO-WED-054](#)

Email DIRAC project leader, Andrei Tsaregorodtsev, atsareg@in2p3.fr