



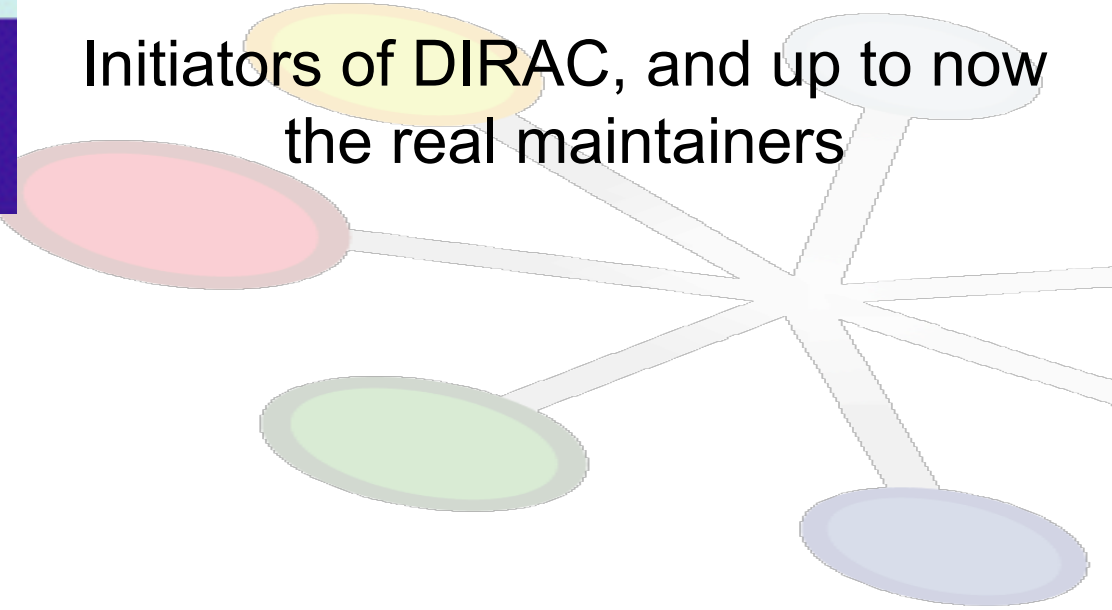
LHCb

Federico Stagni
(on behalf of the LHCb collaboration)



The “Beauty” experiment at CERN

Initiators of DIRAC, and up to now
the real maintainers



- Accounting

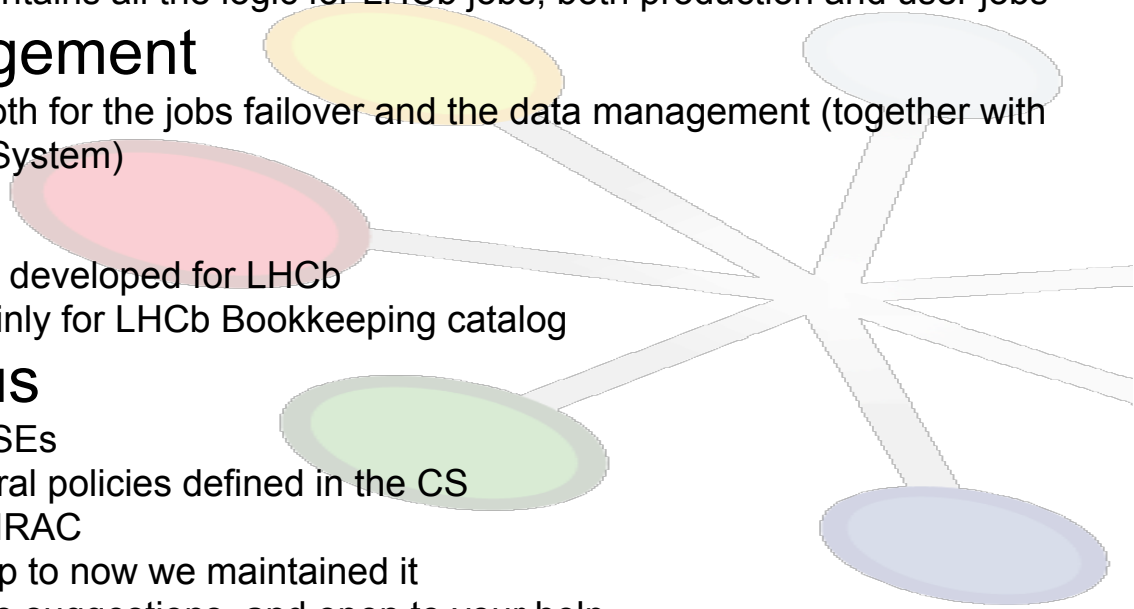
- All the accounting types of DIRAC are used
- Extended in LHCbDIRAC with few new types:
 - Both for storage usage and jobs
- WMSHistory is a “special” accounting
 - can become way too big if you have a lot of jobs in the JobDB
 - We decided to split WMSHistory from the rest of the accounting
 - requires new DB, and will require DIRAC v6r12
 - started working on a generic “monitoring” solution
 - requirements: <https://github.com/DIRACGrid/DIRAC/wiki/Monitoring-service>
 - everybody free to participate

- Data Management

- We use almost everything in it right now, but, NOT the DIRAC File Catalog
 - Several technical reasons make it not practical for us right now
 - Also worries about its scalability
 - We have proposals for it: see DMS talk later



- **Interfaces (API)**
 - We extended most of the APIs
 - the LHCbJob API contains all the logic for LHCb jobs, both production and user jobs
- **Request Management**
 - Used extensively, both for the jobs failover and the data management (together with the Transformation System)
- **Resources**
 - Almost all resources developed for LHCb
 - Few extensions, mainly for LHCb Bookkeeping catalog
- **Resource Status**
 - Used for managing SEs
 - We have several policies defined in the CS
 - Extended in LHCbDIRAC
 - We created it, and up to now we maintained it
 - We are open to suggestions, and open to your help
 - There are some tasks in GitHub and more might come



- **Storage Management**

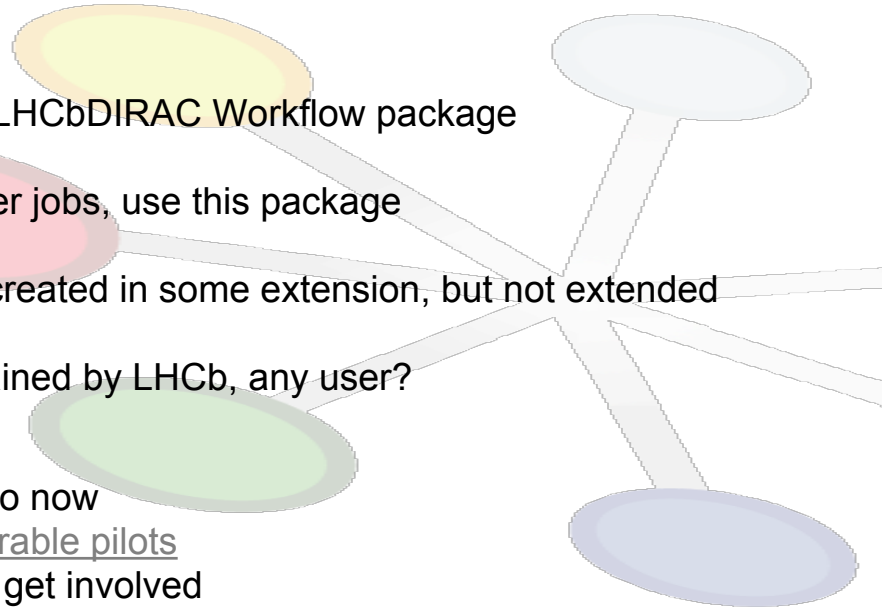
- Staging of files used by the jobs
- This system is not extended

- **Workflow**

- What is in DIRAC is a porting from LHCbDIRAC Workflow package
- Well extended in LHCbDIRAC
- All our jobs, either production or user jobs, use this package
- The Job API use this package
- Apparently this package has been created in some extension, but not extended directly yet
 - Up to now, it has been maintained by LHCb, any user?

- **Workload Management**

- Very minor extensions in LHCb up to now
- Started working on generic, configurable pilots
 - VERY important RFC, please get involved
 - First step: refactoring the pilot script



- Transformation

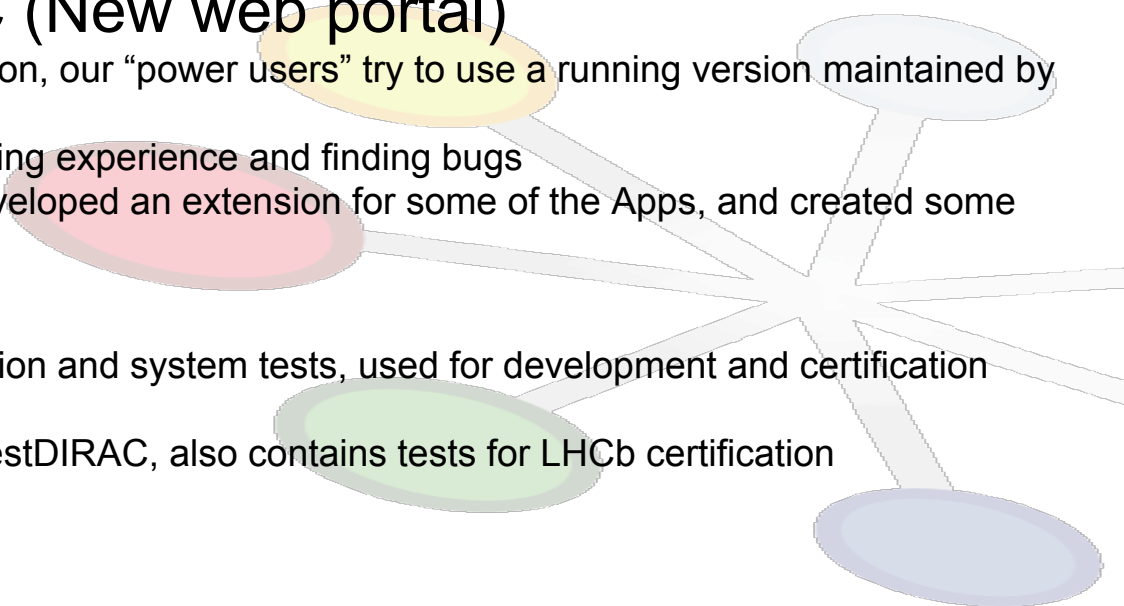
- Extensively used and extended
- All our production activities based on it
 - Sending jobs to the WMS
 - Jobs are created using the Job API
 - So, the Workflow package
- All our data management activities based on it
 - Sending requests to the RMS
- Productions created using a LHCb extension
- For more info: LHCb presentation at the [3rd DUW](#)
- Scaling for the system is the real issue that we had to face
 - Soon we will need to migrate to use parametric job submission
 - In any case the real problem is the lack of real data driven: many agents pulling for files/tasks instead of pushing system
 - No big RFC on it, only [RFC 15](#)
 - So, up to now it has moved on mostly because of LHCb
 - and contributions from ILC (Stephane)
- Surprised to see that up to now it has not been used by several VOs





DIRAC extensions in use

- **Web (current web portal)**
 - Also extended in LHCbWebDIRAC
- **WebAppDIRAC (New web portal)**
 - Not really in production, our “power users” try to use a running version maintained by the developer
 - Mostly for gaining experience and finding bugs
 - We have already developed an extension for some of the Apps, and created some LHCb-only
- **TestDIRAC**
 - Collection of integration and system tests, used for development and certification process
 - Extended in LHCbTestDIRAC, also contains tests for LHCb certification

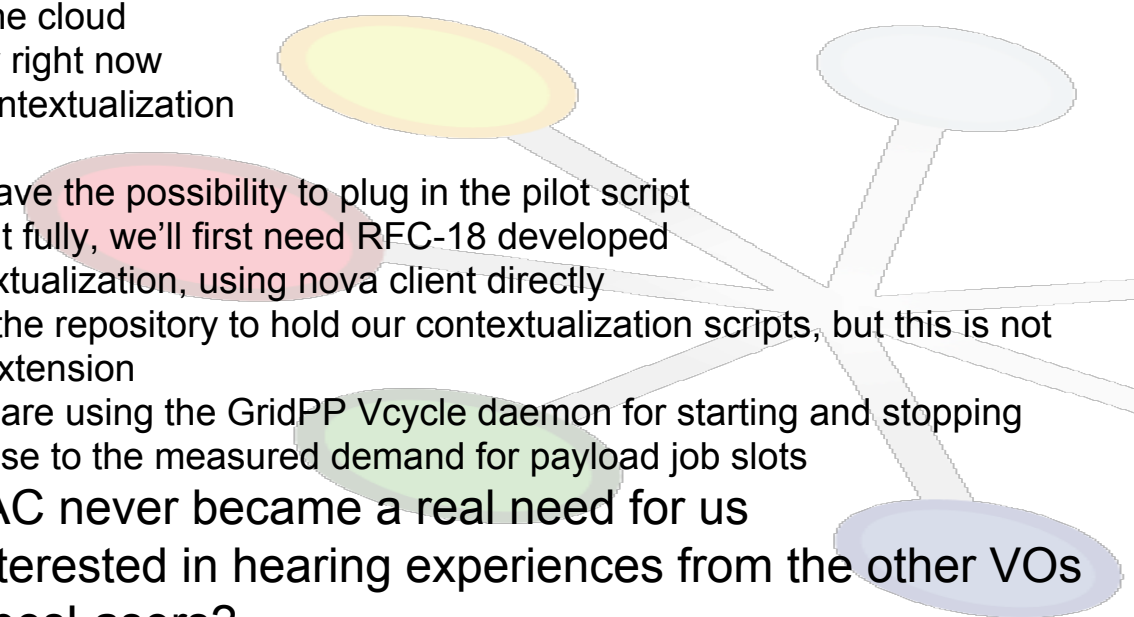


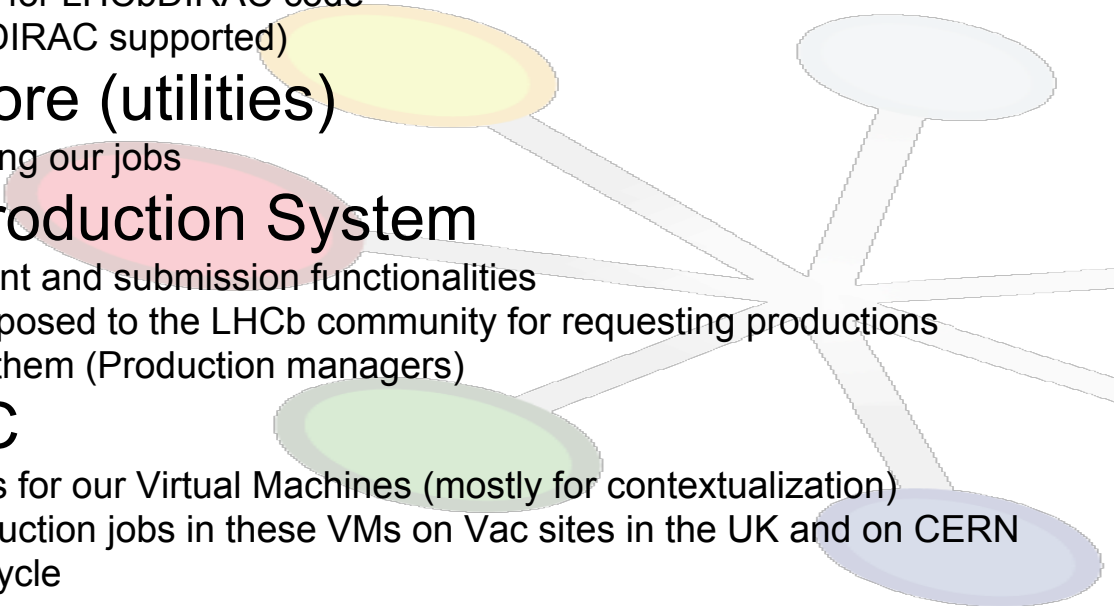


DIRAC systems not in production

We have looked into using VMDIRAC

- Not so much recently, we have to say
- We run already jobs on the cloud
 - On OpenStack only right now
 - We use our own contextualization
 - Shell scripts
 - We already have the possibility to plug in the pilot script
 - To use it fully, we'll first need RFC-18 developed
 - No ssh contextualization, using nova client directly
 - LHCbVMDIRAC is the repository to hold our contextualization scripts, but this is not (yet) a VMDIRAC extension
 - On OpenStack, we are using the GridPP Vcycle daemon for starting and stopping machines in response to the measured demand for payload job slots
- Up to now, VMDIRAC never became a real need for us
 - We are **very** interested in hearing experiences from the other VOs
 - Belle2? AlpesLasers?



- **LHCbDIRAC Bookkeeping system**
 - A file metadata and provenance system
 - Often, the real driver for LHCbDIRAC code
 - Backend is Oracle (DIRAC supported)
 - **LHCbDIRAC Core (utilities)**
 - Just utilities for running our jobs
 - **LHCbDIRAC Production System**
 - Provides management and submission functionalities
 - A web frontend is exposed to the LHCb community for requesting productions
 - And to launch them (Production managers)
 - **LHCbVMDIRAC**
 - A collection of scripts for our Virtual Machines (mostly for contextualization)
 - We are running production jobs in these VMs on Vac sites in the UK and on CERN OpenStack using Vcycle
- 

BeautyDIRAC Clients


- End users (physicists) can submit jobs:
 - using directly the LHCbDIRAC APIs
 - using Ganga to prepare and send jobs
 - this is what most of them use
 - Ganga then submits jobs to DIRAC using the LHCbDIRAC APIs
- Productions requesters
 - every user, in theory
 - but productions go through an approval mechanism
- Production managers, GEOCs
- Production shifters
- Data managers
- Resources (sites) managers

In practice, all our distributed computing activities go through BeautyDIRAC



Running jobs by JobType

25 Weeks from Week 46 of 2013 to Week 20 of 2014



MCSimulation user	85.3%	DataReprocessing	1.2%	MCReconstruction	0.1%	test	0.0%
DataStripping	2.2%	DataSwimming	0.3%	WGProduction	0.1%	sam	0.0%
		Merge	0.3%	MCStripping	0.0%	unknown	0.0%


Generated on 2014-05-22 14:38:17 UTC

The system in action /1

- About 30K jobs constantly running
- Peak of ~50K jobs running

CPU days used by Grid

25 Weeks from Week 46 of 2013 to Week 20 of 2014



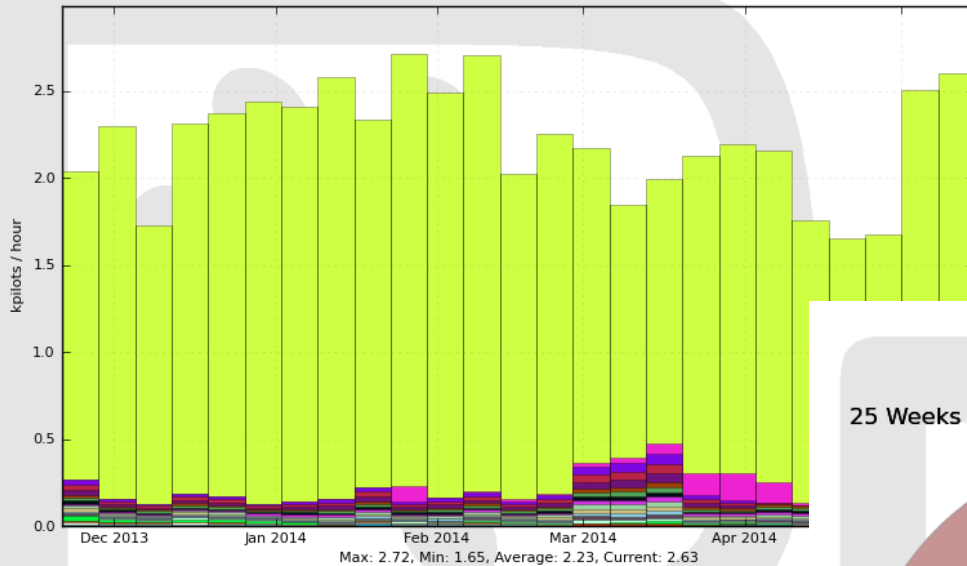
Generated on 2014-05-22 14:44:12 UTC

Using several types of computing resources

Still dominated by “LCG” (CREAM), others growing fast

Pilots by GridResourceBroker

25 Weeks from Week 46 of 2013 to Week 19 of 2014

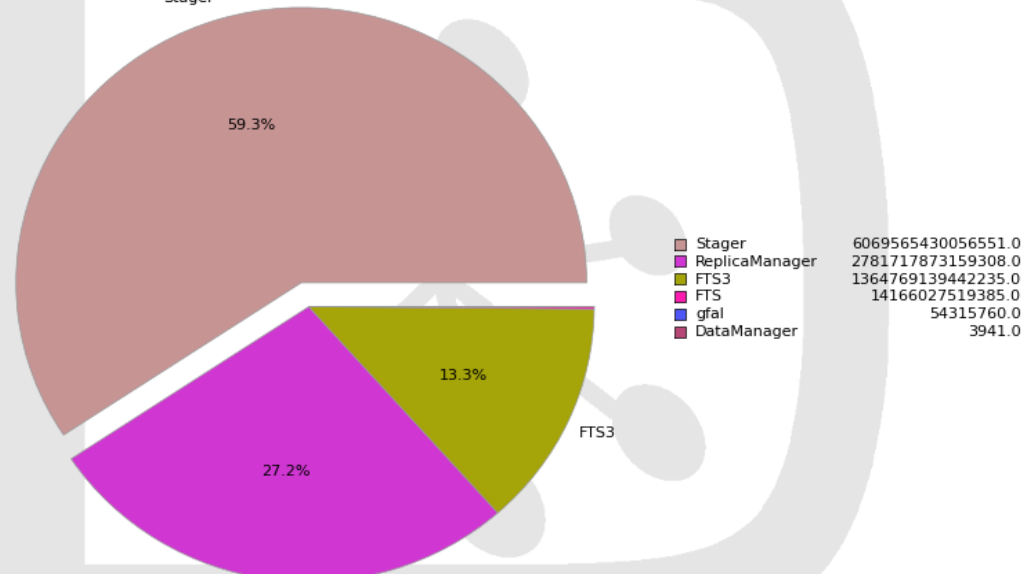


The system in action /2

Almost all the pilots sent through SiteDirectors

Total data transferred by Protocol

25 Weeks from Week 46 of 2013 to Week 20 of 2014



volhcb20.cern.ch	90.0%	wms1.grid.sara.nl	0.6%	wms-1-kit.gridka.de	0.5%	wms-6-kit
volhcb03.cern.ch	2.1%	wms01.pic.es	0.6%	wms014.cnaf.infn.it	0.4%	wms-3-kit
wms303.cern.ch	0.8%	wms02.pic.es	0.6%	wms004.cnaf.infn.it	0.4%	wms-4-kit
wms302.cern.ch	0.7%	graszode.nikhef.nl	0.6%	wms-2-kit.gridka.de	0.4%	lbvobox01
wms301.cern.ch	0.7%	graspol.nikhef.nl	0.6%	volhcb30.cern.ch	0.3%	

GenerateX

FTS3 ramping up...



- 3 setups (Production, Certification, Development)
- ~15 machines for services, agents and MySQL DBs
 - some real, some virtual, moving all to VMs
 - ask for details if interested
- 1 “dbod” (MySQL -> MariaDB)
 - we moved some DBs over there ~6 months ago, and we’ll move more
 - recent version of MySQL
- Oracle for the Bookkeeping



Questions, comments

Beauty DIRAC

