

Belle II

Ueda I. and Miyake H.
Montpellier, 23 Mai 2016

BelleDIRAC

Working since quite a while

Scaling up with increasing resources from sites

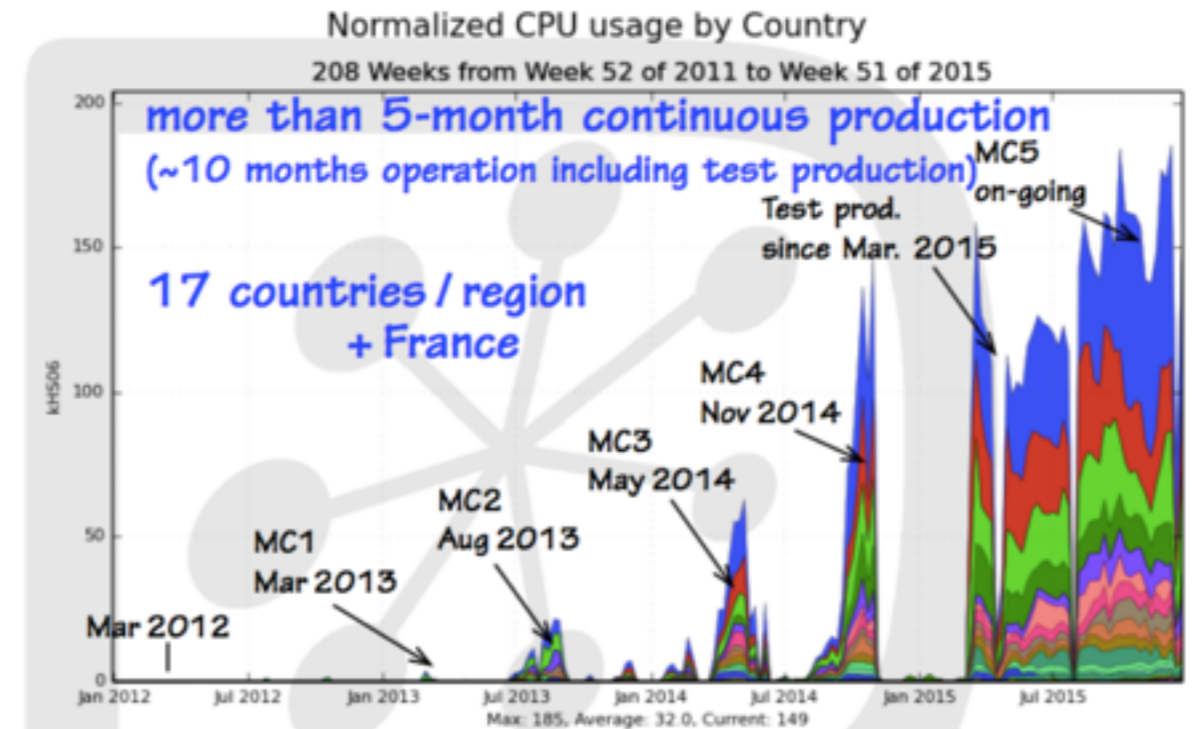
- ~25k concurrent jobs, >150 “HS06” (*)

Evolving towards the start of the experiment

- Automation in production
- Automation in data distribution
- Automation in issue detecting (and operation)

Just upgraded to v6r14

No intention to repeat what we already presented last year



Upgrade to v6r14

Done in Apr 2016

- Avoided major system changes during the “MC5” campaign (Aug 2015 - Apr 2016)

Followed the instructions, and then, ...

- An issue around IPv6 needed a hack — ([DIRACGrid] “*Can we disable IPv6 support?*”)
- Monitoring service works on SL6.7 but not on SL6.5 — ([DIRACGrid] “*ps -q option on SL5/6*”)
 - A list of “certified OS”, or information of the certification environment would be useful
- Short modification to AccountingSystem — manually disabled “rebin” feature
 - no good idea on what to suggest, yet
- DIRAC went to “unicode” (or, is moving to, for it is not done everywhere, eg. SE?)
 - Adapting the DIRAC codes for LFC is on Belle II, in consultation with DIRAC dev. — (PR2891)
- Issues with web portal
 - Needed modification in `dirac.cfg` ([DIRACGrid] “*How to load WebAppDIRAC module?*”)
 - Some bugfix, typo suggested ([DIRACGrid] “*WebAppDIRAC configuration mistake*”)
- The latest `lcgBundles` don’t work (both for SL5 and for SL6)
 - Currently using the most recent “working version” (SL5: 2014-03-20, SL6: 2015-06-10)
- BelleDIRAC adaptations, of course

Now we have seen the v6r15 is out. Should we go for v6r15 soon?

Belle II Production System

In Belle II, we are building a “production system” on top of DIRAC

- Some workflows similar to LHCb, transformations look alike
- Others different, Belle II specific components implemented
 - No Oracle+Bookkeeping, non-T1 SEs distributed over different continents
- Making good use of DIRAC components (TransformationSystem, RMS, ...)

Composition

- ProductionManagement to define/manage what to produce (MC, reprocessing, ...)
 - under development, no report today
- Fabrication to define/manage jobs
 - prototype in production, utilising TransformationSystem
- DistributedDataManagement to distribute/relocate data
 - initial implementation being integrated
- Monitoring and more...

BelleDIRAC components named avoiding those that already exist in DIRAC

- Not always easy... (eg. DDM vs DMS)

Data Management Block

Datasets

- Belle II produces various types of MC data
 - Organised as “datasets” (defined as a part of LFN path)
- “Runs” can also be considered as “datasets”

Data Management Blocks

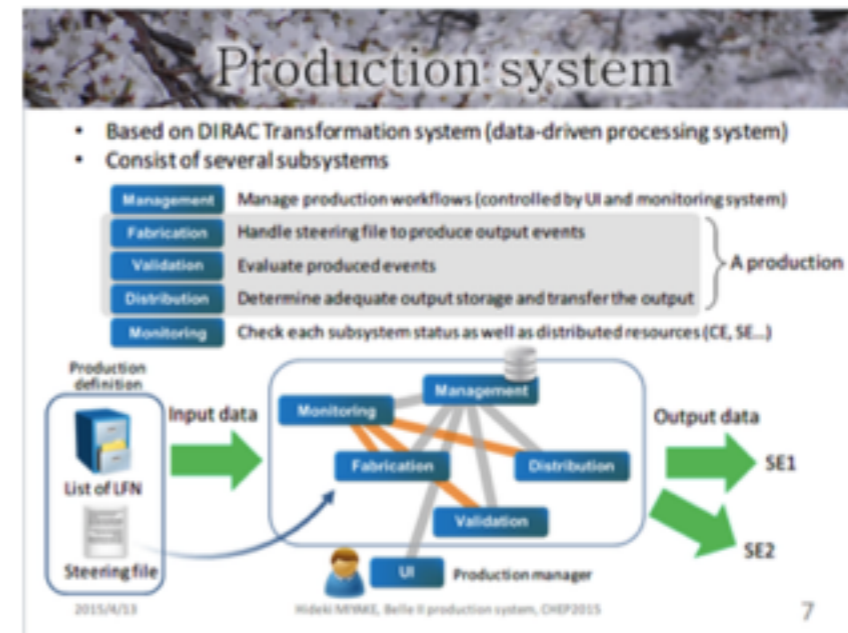
- Clustering files of the same MC type onto the same SE, to some extent, would ease some workflows
 - multiple input of the same type — merge, analysis, ...
 - possible data management at directory-level
- A dataset can contains millions of files — too many as a unit of data management
- “Data block” as a unit of data management
 - max 1000 files as initial implementation, so far so good. May tune with experiences
- “Dataset” is the unit of production, but files are organised in “data blocks”
 - Some system impelementation based on “data blocks”
- A subdirectory under “dataset” path: `/belle/...dataset...name.../subNN/files`

ProductionSystem: Fabrication

Belle II production system (1)

- We are switching to new production scheme using production system based on Transformation System (TS)
- Automatized production operation for both workload and data management

- Production job (gbasf2 job) is registered with TS and gradually submitted through MCExtensionAgent (customized for Belle II)
- Our Agent watches job status and failure jobs are automatically rescheduled
- Concept was proven at our previous MC mass production campaign
- Plan to change behavior not to reschedule but submit new job to keep log of failure jobs



Belle II production system concept

| Total Created | Submitted | Waiting | Running | Done | Completed | Failed |
|-----------------|-----------|-------------|-------------|-----------------|-----------|--------|
| 374000 (+10...) | 0 | 2958 (+82) | 4138 (+873) | 368812 (+16...) | 16 (+7) | 2 |
| 10000 | 0 | 0 | 0 | 10000 | 0 | 0 |
| 374000 (+10...) | 0 | 4381 (+842) | 788 (+80) | 368767 (+12...) | 8 (+1) | 0 (+2) |
| 10000 | 0 | 0 | 0 | 10000 | 0 | 0 |
| 374000 (+10...) | 0 | 4011 (+8) | 1454 (+842) | 368480 (+15...) | 9 (+2) | 1 (+8) |
| 10000 | 0 | 0 | 0 | 10000 | 0 | 0 |
| 374000 (+10...) | 0 (+11) | 4363 (+193) | 803 (+82) | 368720 (+13...) | 13 | 0 (+2) |
| 10000 | 0 | 0 | 0 | 10000 | 0 | 0 |

Automatic failure recovery in Transformation

2015/05/27

Belle II experiment report

10

ProductionSystem: Fabrication

Job definition/submission system based on the TransformationSystem

- A transformation per production definition
 - created manually today, by the other system component in future
- A dataset per production definition, data blocks defined according to the number of files to produce
- Job definition/submission with TransformationSystem to fill data blocks
 - worked fine usually, with occasional “over-creation”
- Job creation throttled, not to overload the system, and to feed enough “waiting” jobs to the system
- For “failed jobs”, new jobs created rather than rescheduling the same jobs (as presented last year), file name uniqueness assured by “taskIDs” (as suggested by FS last year)
 - Works perfectly. Thanks!
- Some issues reported/consulted. Thanks for the help.
- Hoping even smoother production with v6r14

Fabrication system

- A part of the Belle II production system
- manage event processing and initial validation
- Based on Transformation System (TS)
- Workflow organised per “data management block”

Event processing

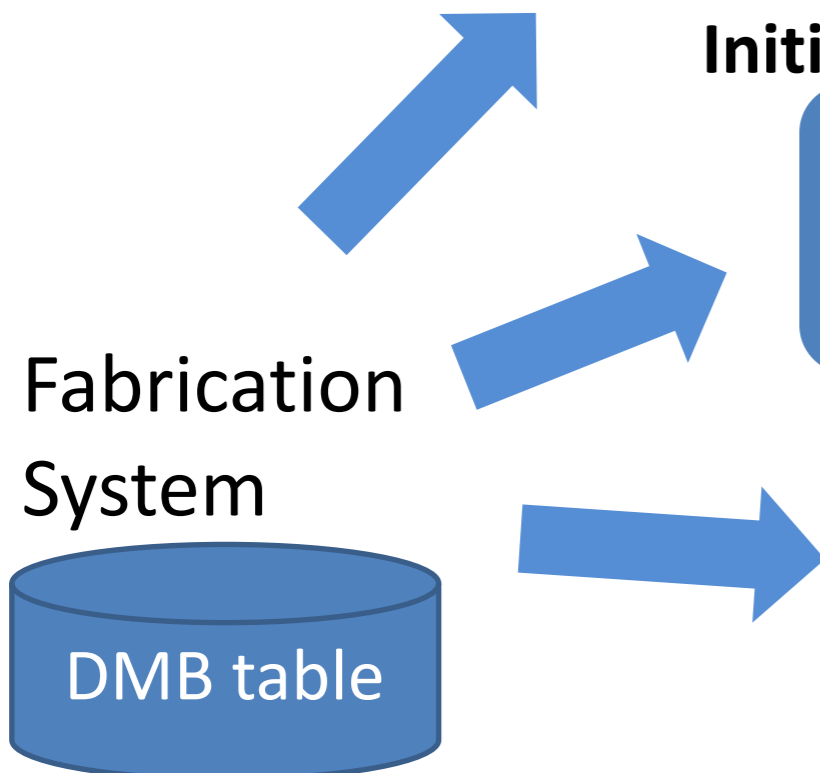
- Associate TS Task to DMB
- Handle failure jobs (if not managed by TS Task)

Initial validation

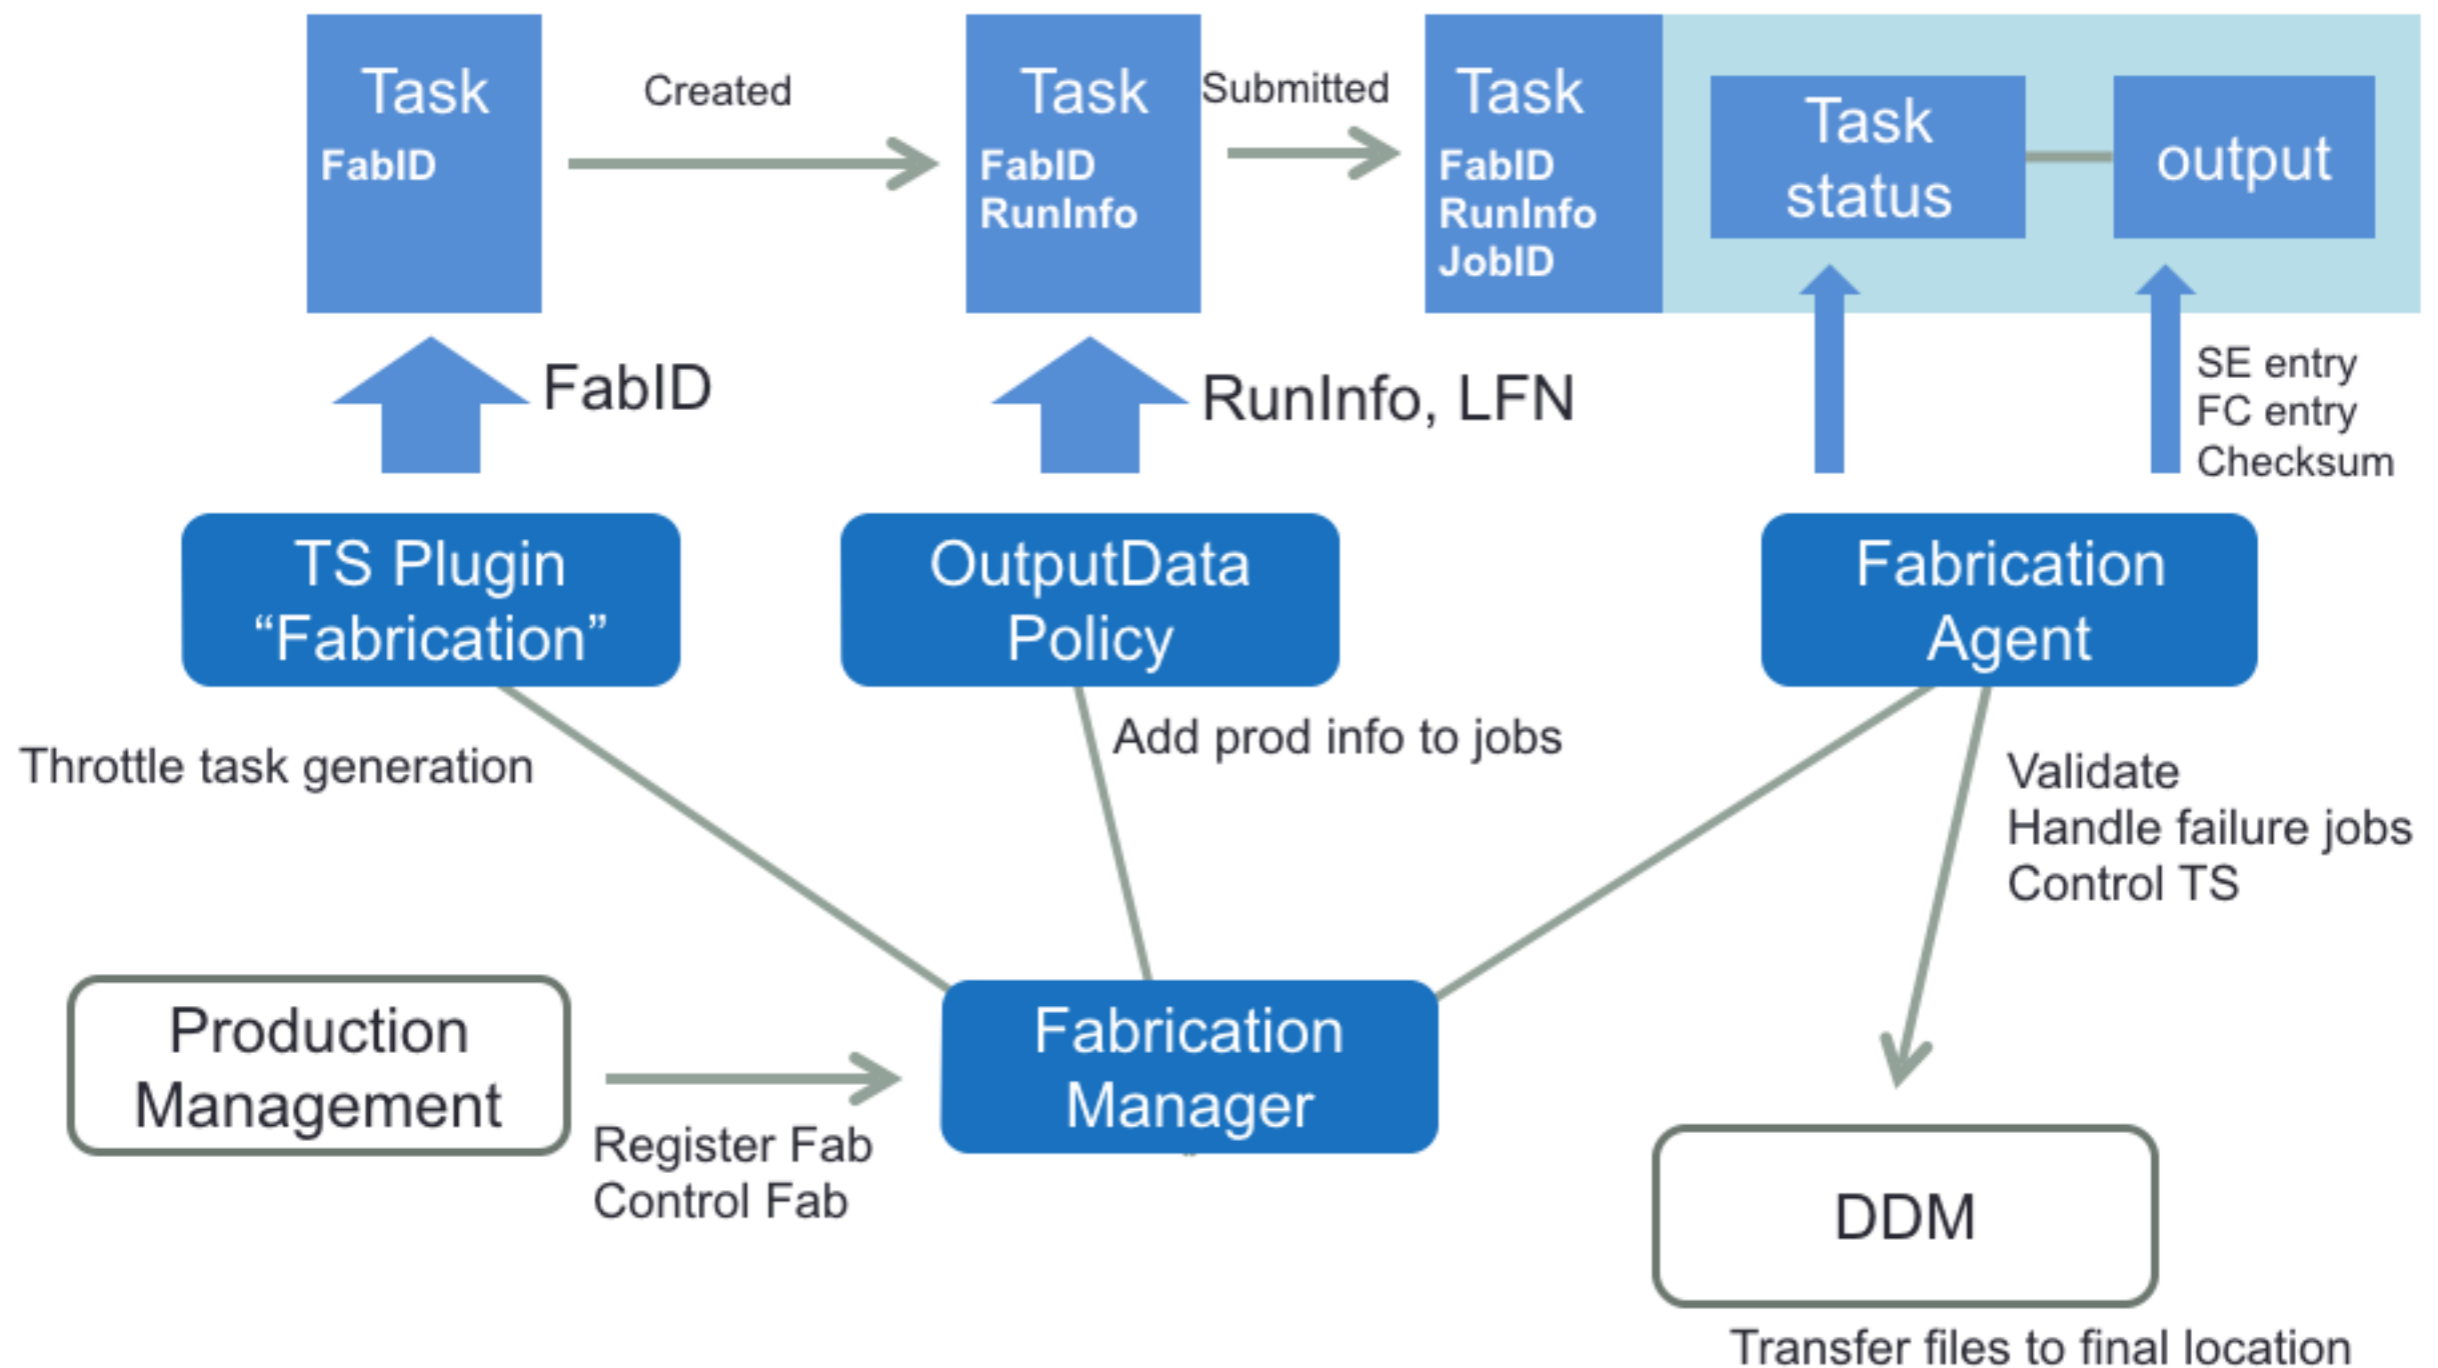
- Any mismatch among SE, FC, metadata
 - Entry, size, checksum...
- as well as production info (e.g. # event)

Transformation control

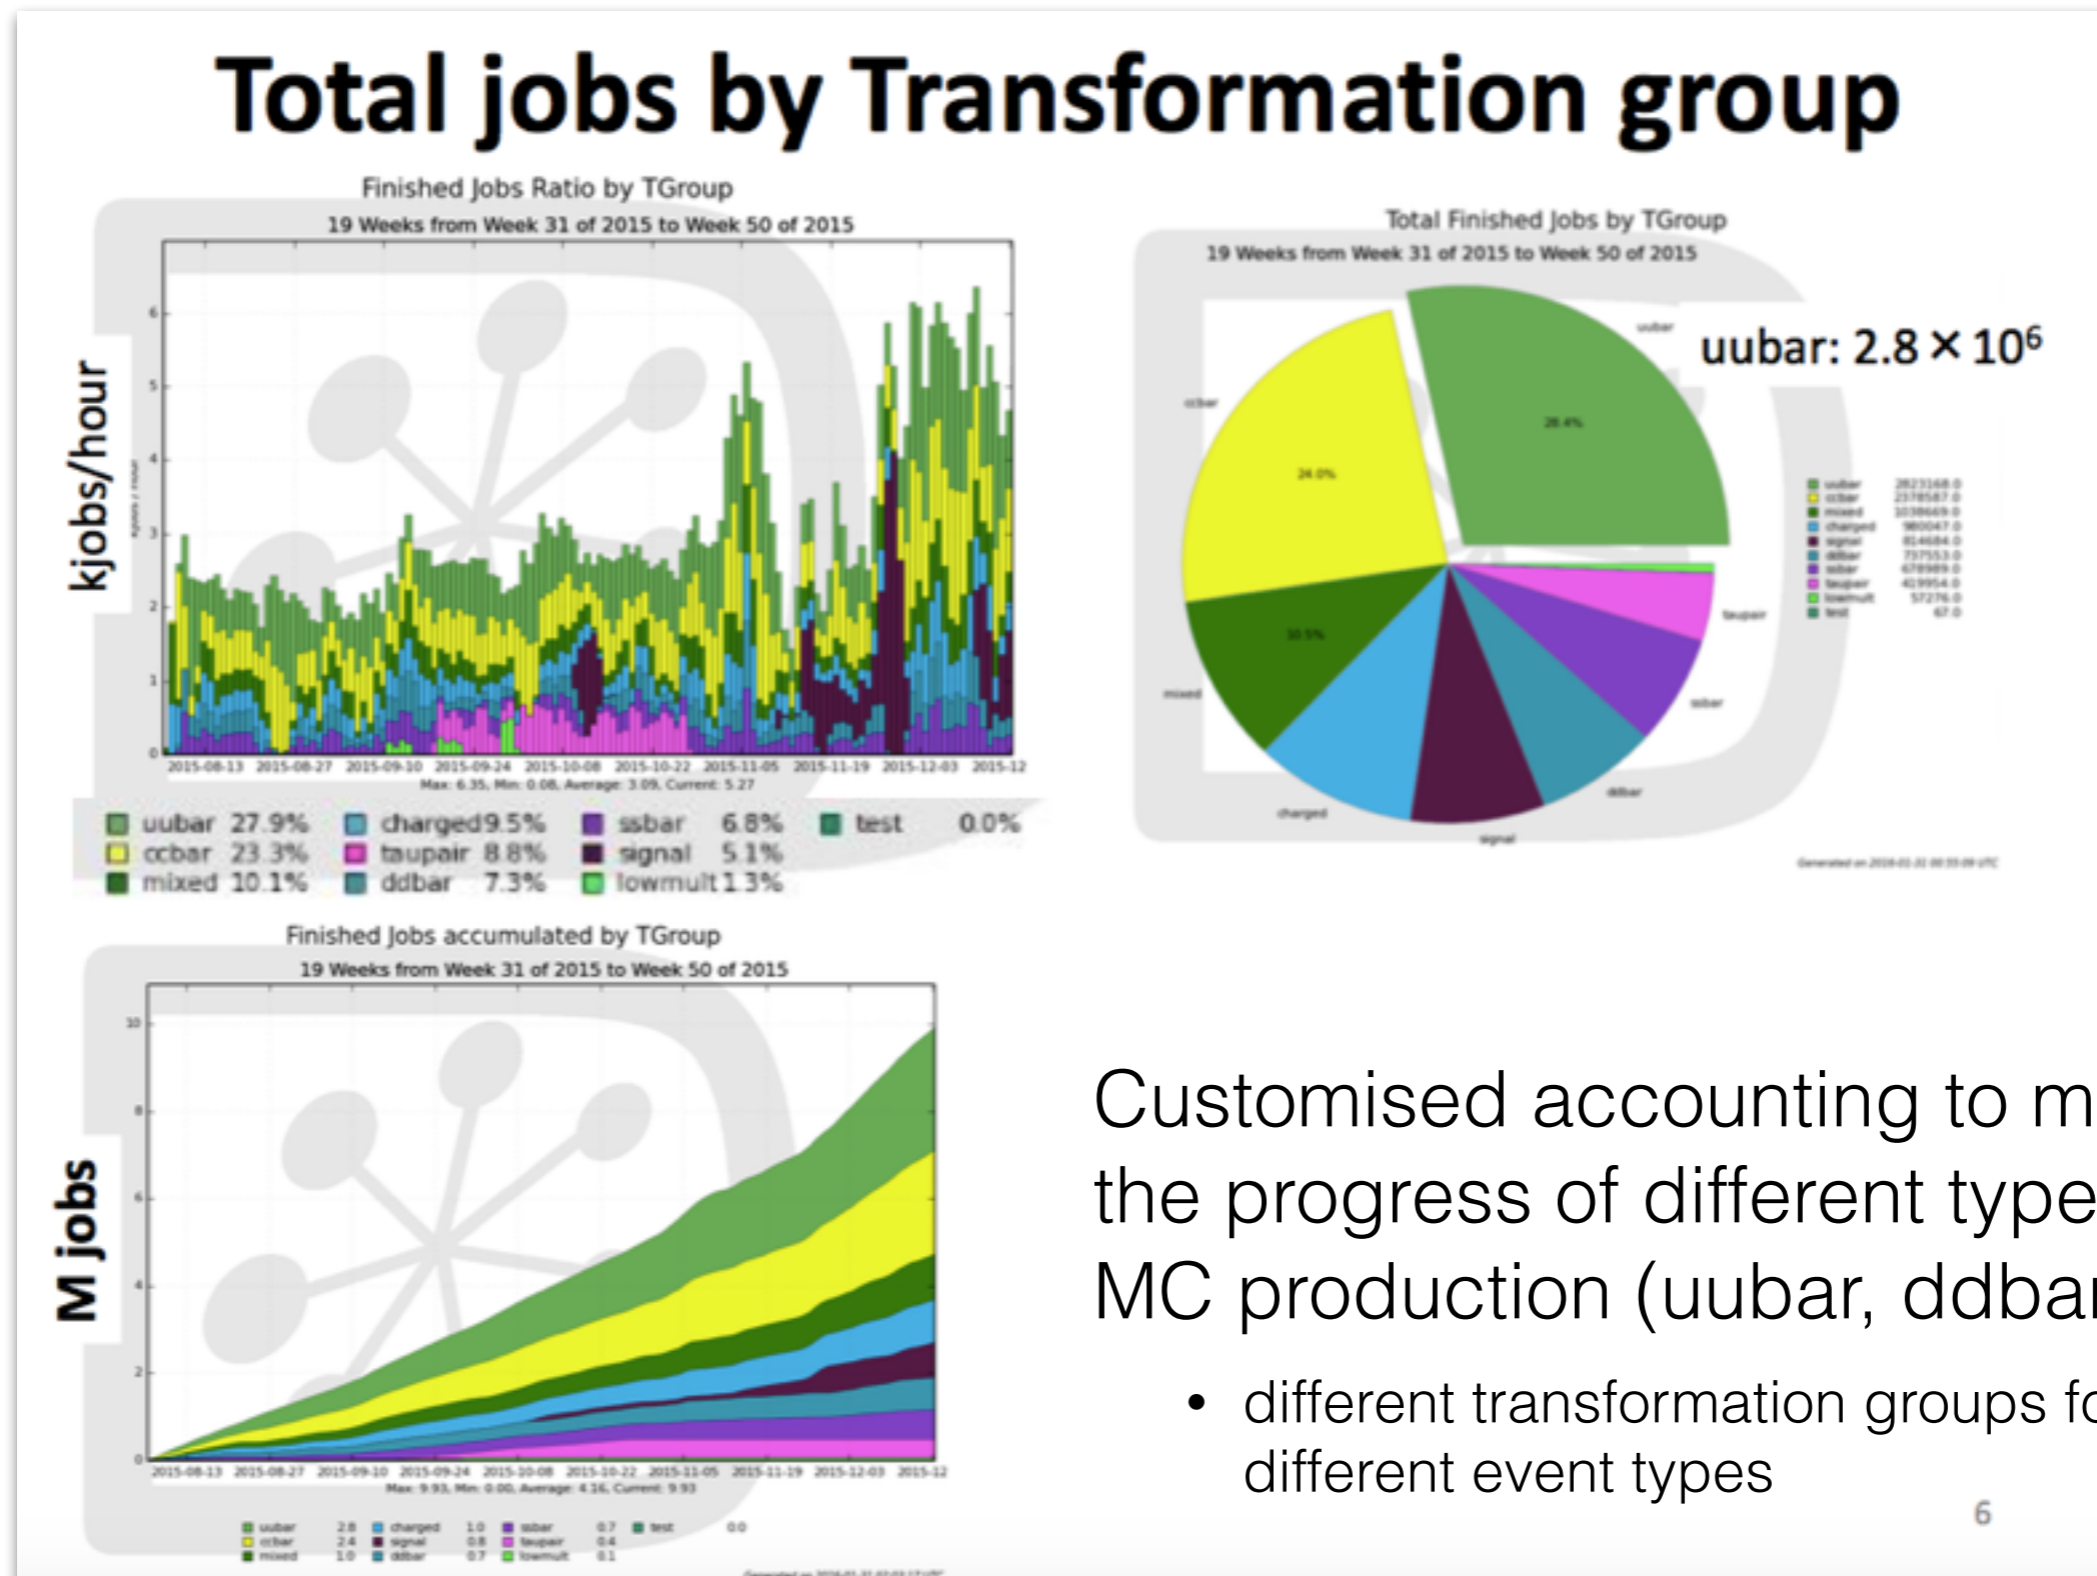
- De-activation and re-activation of Transformation by DMB status



Fabrication workflow



Monitoring Production Progress



Customised accounting to monitor the progress of different types of MC production (uubar, ddbar, ...)


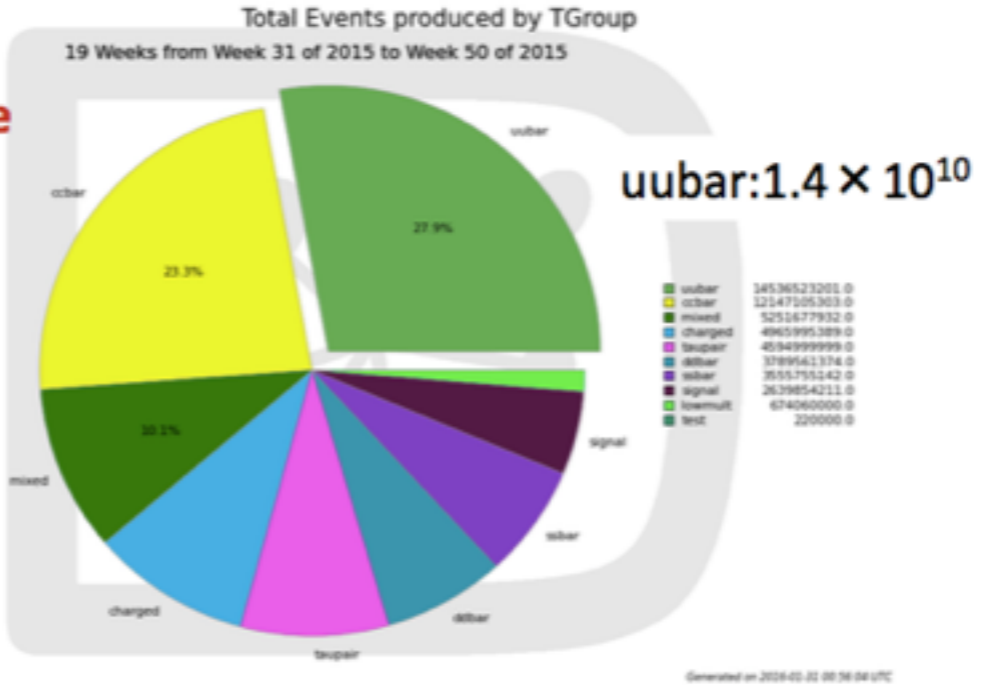

- different transformation groups for different event types

Monitoring Production Progress

Total events by Transformation group

Differential

Pie



and with number of events stored in the “production system”

ProductionSystem: DDM

Belle II production system (2)

- Once job completed, the output is stored on local SE
- Production system collects the outputs and transfers to final destination through FTS
- The basic concept was proven by simple Transformation (Broadcast)



Belle II production system;
Data distribution logic

- In near future, this part will be replaced by smart data distribution system (Belle II DDM)
- Belle II DDM collects environmental information as well as its basic information like data type
 - Environmental information: e.g. SE free space (Spacetoken), down time (RSS), network condition (perfSONAR)
 - We can then dynamically avoid inadequate SE and choose best one

Development is ongoing

Latency Mesh (perfSONAR)



ProductionSystem: DDM

Belle II Distributed Data Management System

- manages data with “data blocks” as the unit of operations
- utilises RMS/DMS to treat files

Transfer component

- throttles “requests” to RMS, not to overload the system, and to keep possibility to prioritise the actions (feed requests to RMS according to the priority/weight)
- initial implementation with TransformationSystem (as was suggested last year)
 - worked fine, but only for “one-time” distribution
 - difficult (or impossible) to “relocate” data after initial replication
- current implementation directly uses RMS, not via TransformationSystem

Deletion component

- Deletion is as challenging as transfers
- A special component under development, not relying on RMS, for we definitely need different throttling for each SE and for FC

Data block replica catalog

- To ease the DDM workflows, to reduce load on the file catalog
- Under planning, probably utilising DFC (as discussed offline last year)

Clouds and Volunteer Computing

Belle II is using cloud resources via

- VMDIRAC
- Cloud Scheduler

(we have the expert of each)

Network issues

- Squid inside the clouds on static nodes
- Shoal to find the nearest squid
- Possible DIRAC services (eg.CS) within the clouds?
 - may not need this if CS accessible via http cache
- Some academic network connection to clouds...

BelleII@home

- Some development work for volunteer computing with BOINC
- To be presented at CHEP 2016

Works on DIRAC codes

SiteDirectors

- Quite some Belle II sites are non-grid
- Thus, some efforts have been, and are being, put into SiteDirector coding
 - SLURM : some bug fixes — to be reported as PRs
 - HTCondor : some fixes — to be reported as PRs?
 - IBM LoadLeveler : under development

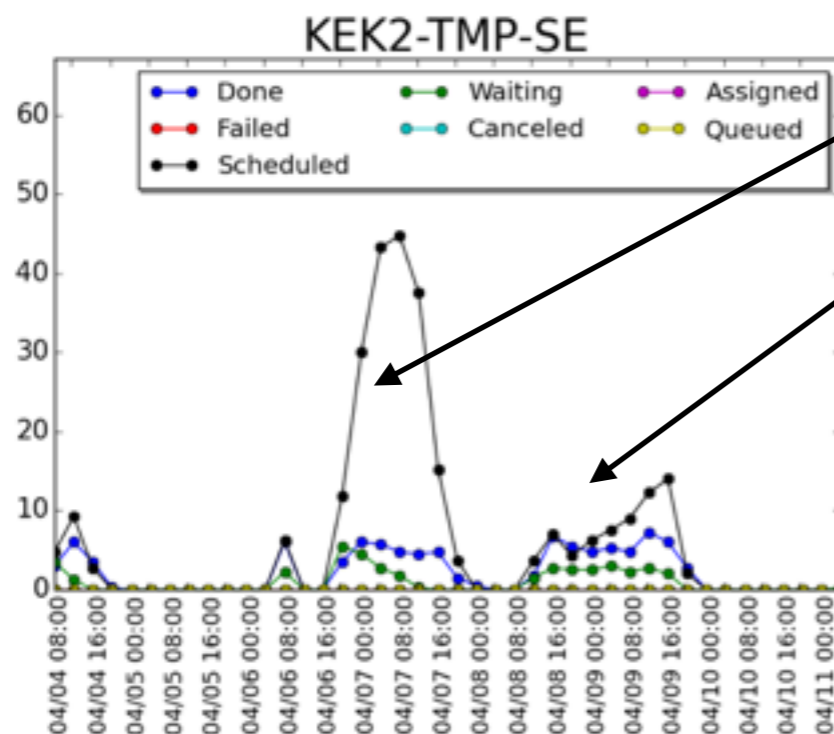
RSS CE Status?

- A person in Belle II was looking into this
 - following the communication in the diracgrid-forum ([DIRACGrid] “[Installing RSS](#)”)
 - but then he became busy in his new position before implementing any contribution
- Now we find new implementation <https://github.com/DIRACGrid/DIRAC/issues/2874>
- Is this only for v6r15? (or to be back-ported to r14?)

RMS/DMS

Replication requests will pile up


- with a higher request rate than the processing capacity
- The node running RMS+DMS gets fully loaded
- So, we try to “throttle” replication requests



Requests pile up with a too high request rate

Number of requests stays moderate with a lower request rate

The system fully loaded when requests pile up



BelleDIRAC RMS extensions

- ReplicateAndRemove for “move” operations
- Retrieve “destination SE” from replication requests (we monitor requests per dest SE)
- Retrieve a list of requests within a time range together with their status with a single call (rather than retrieving the status per request one-by-one)
- Probably to be pull-requested

DMS

GFAL2HTTP

- “Major to come” in the DMS presentation last year (C. Haen)
- Some on-going activities in Belle II on HTTP/DaviX data access (currently w/o DIRAC)
- We would be happy to try out GFAL2HTTP in DIRAC (when?)

Different protocols for different aspects?

- Following WLCG, Belle II will get rid of SRM, but there will be a transition period with multiple protocols to use (even for a single SE)
- SRM for write
 - Currently we rely on SRM for space accounting (space token)
- non-SRM for read (http/davix, root/xrootd)
 - We will need direct read by ROOT in some use cases (jobs read files on SE without downloading to WNs)
- Protocol to use in FTS transfers depends on the other SE
- Will this be possible in DIRAC?

Revisiting some issues from last year...

DIRAC config retrieval with http

- so that we can make use of http proxy cache

MatchingDelay with Site=ANY

- Pilot submission get affected
 - idle cycles without pilot submission even when there are available slots
- Pilots get wasted
 - increases the loads on the CE and the DIRAC server

PFN definition of DIRAC-SE (dips)

- used to be PFN = LFN, now PFN = hostname + LFN, solved?
- Should be solved now that PFN is “abandoned”?

Reference Slides

Issues: WMS



❧ **MatchingDelay with Site=ANY**

- ❧ MatchingDelay affects both site specific job (Site=AAA) and generic job (Site=ANY)
 - ❧ Due to behavior of TQDB.__generateNotSQL()
- ❧ Since SiteDirector refers Matcher, both pilot submission and payload matching are affected by the behavior
- ❧ As a result, number of "Site=ANY" job is quite unstable (sometimes SD submits pilot once per 6 hours)
- ❧ Practically it can be avoided if site specific queue (e.g. Site="AAA") exists in TQ
 - ❧ Dirty workaround to fill TQ by DIRC Agent

❧ **SiteDirector affected by Matcher**

- ❧ Related with the issue above, SD refers Matcher for pilot submission
- ❧ It means MatchingDelay is applied for SD
- ❧ Suggest to skip MatchingDelay for pilot submission

Issues: WMS



☞ MatchingDelay with Site=ANY

- ☞ Site = ANY = AAA or BBB or CCC
- ☞ $!(AAA \text{ or } BBB \text{ or } CCC) = !AAA \text{ and } !BBB \text{ and } !CCC$
- ☞ Imply Site=ANY is affected by all of negative conditions

- ☞ Even if a pilot is submitted at Site=BBB, it is vetoed if other pilot was submitted to Site=AAA just before
- ☞ Same for payload matching
- ☞ As a result, job execution for Site=ANY is quite opportunistic

- ☞ Workaround: to fill TQ by Site=BBB job (at least pilot can be submitted)
 - ☞ Payload is still affected (afak)

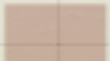
Matcher with Delayed Matching



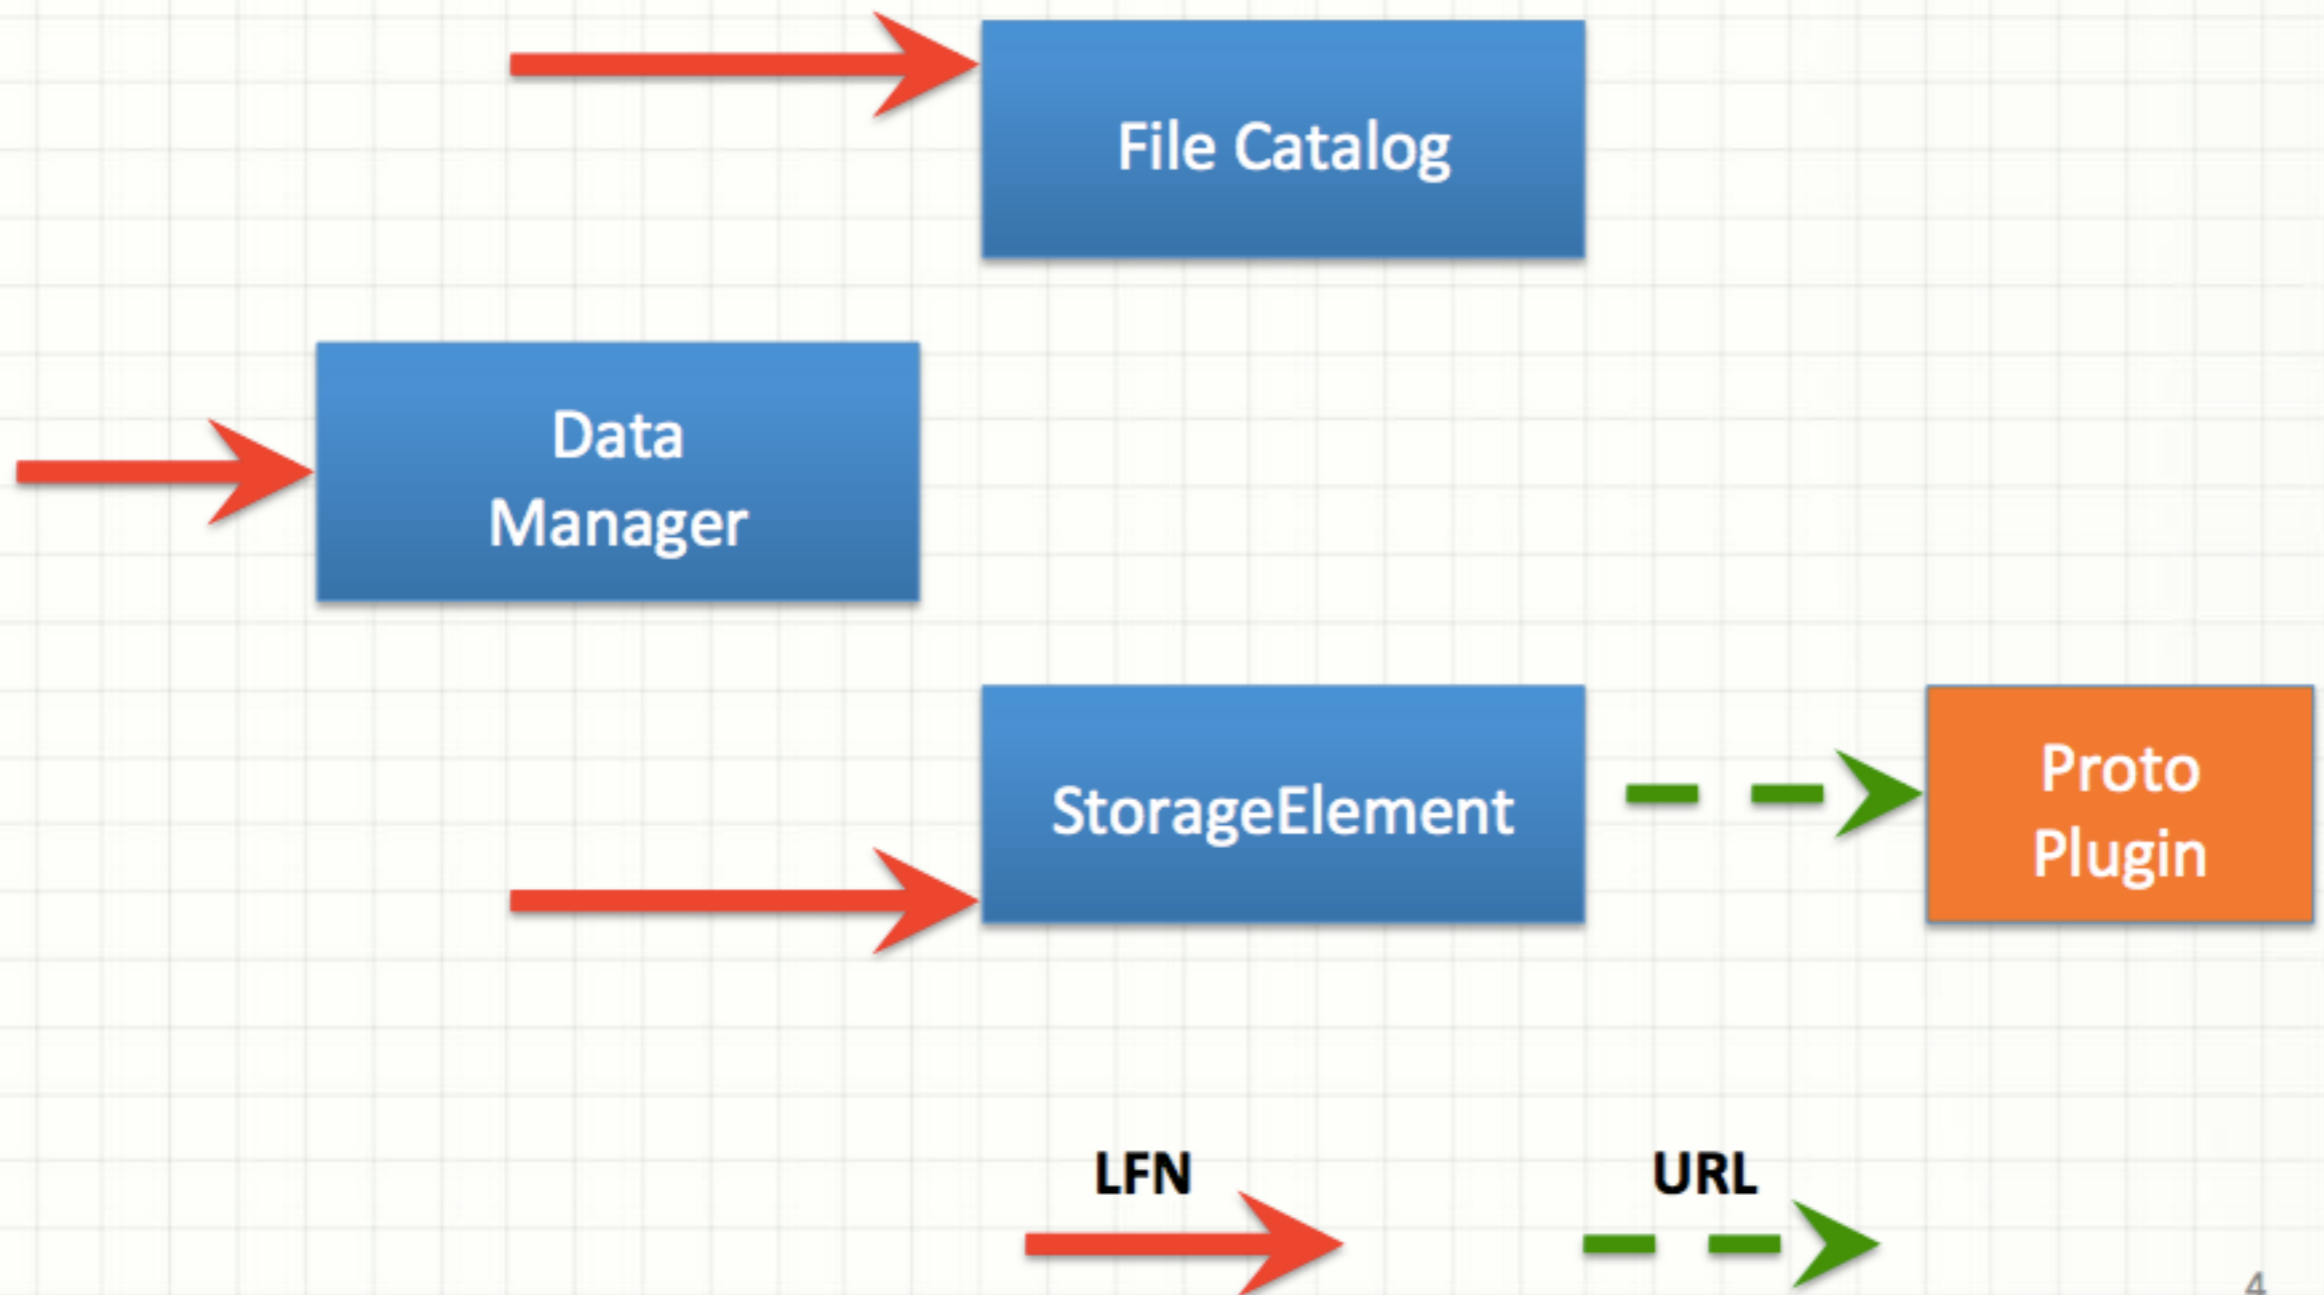
Delayed Matching

- Useful to prevent SE overload
- However many pilots are wasted after the delay
- It increases both site and server loads
- Do you have any idea to reduce wasted pilots under delayed matching configuration?

| PilotRefId | Status | Site | ComputingEl... | Broker | CurrentJobID | OwnerGroup | LastUpdateTime (s) |
|------------------|--------|---------------|------------------|---------------------------------|--------------|-------------|--------------------|
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:25 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14892866 | belle_pilot | 2015-03-19 05:04 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:25 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:27 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:30 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:26 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14900539 | belle_pilot | 2015-03-19 06:06 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14896288 | belle_pilot | 2015-03-19 05:52 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:26 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14888619 | belle_pilot | 2015-03-19 05:03 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:30 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14891168 | belle_pilot | 2015-03-19 05:04 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:30 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:26 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14895662 | belle_pilot | 2015-03-19 05:37 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14897336 | belle_pilot | 2015-03-19 05:51 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | 14891925 | belle_pilot | 2015-03-19 05:06 |
| https://recas... | Done | LCG.Napoli.it | recasna-ce01... | can66.cc.kek.jp | - | belle_pilot | 2015-03-19 04:26 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:24 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | 14891001 | belle_pilot | 2015-03-19 05:10 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:25 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:25 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:25 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | 14888582 | belle_pilot | 2015-03-19 06:00 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:25 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:24 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | 14891459 | belle_pilot | 2015-03-19 04:29 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:26 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:25 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:24 |
| sshondor/ib... | Done | DIRAC.Uvic.ca | bellecs.heprc... | host206-12-154-50.heprc.uvic.ca | - | belle_pilot | 2015-03-19 04:26 |

 = Waste pilots

Major change: Good bye PFN



Major to come: gfal2

