GridPP DIRAC

Daniela Bauer & Simon Fayer

Imperial College London

GridPP

The GridPP DIRAC instance

- The GridPP project is a collaboration of 19 UK universities
- The GridPP DIRAC instance is the main entry point for access to UK (and a limited number of non-UK) resources for non-LHC VOs with UK stakeholders
- Currently supports 17 VOs:
  - 5 main users: lz, pheno, na62, dune, snoplus
  - 2 small users: gridpp, vo.moedal.org
  - 2 VOs 'resting': lsst, comet.j-parc.edu
  - 2 VOs being weaned off the WMS: t2k.org, mice
  - 1 VO being commissioned: solid experiment
  - 5 (very small) VOs: Candidates for removal

What's new ?

- Since the DIRAC Workshop 2016
  - Improved upgrade procedure
  - First large scale use of our DIRAC instance: DUNE
  - LZ: GridPP DIRAC as part of the LZ UK data centre
  - EL7
  - Containers

Maintaining GridPP DIRAC

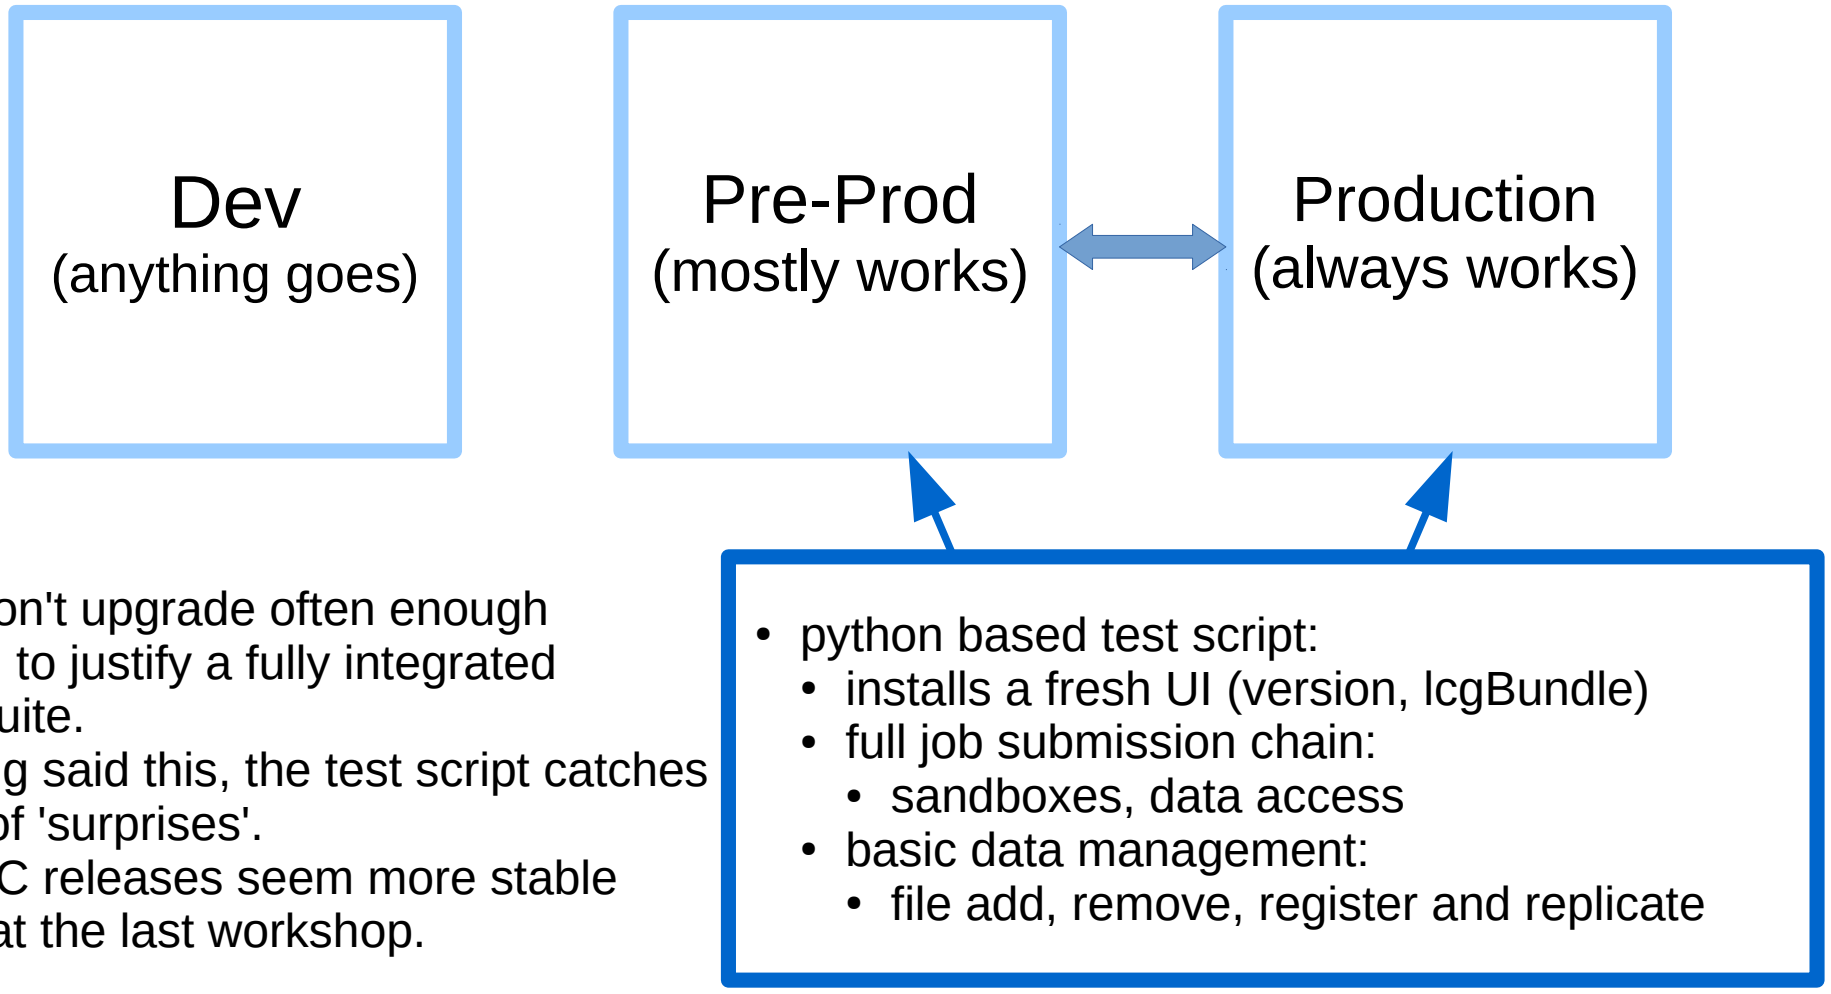

- We don't upgrade often enough (yet?) to justify a fully integrated test suite.
- Having said this, the test script catches a lot of 'surprises'.
- DIRAC releases seem more stable than at the last workshop.

Imperial College London

GridPP

GridPP DIRAC Usage

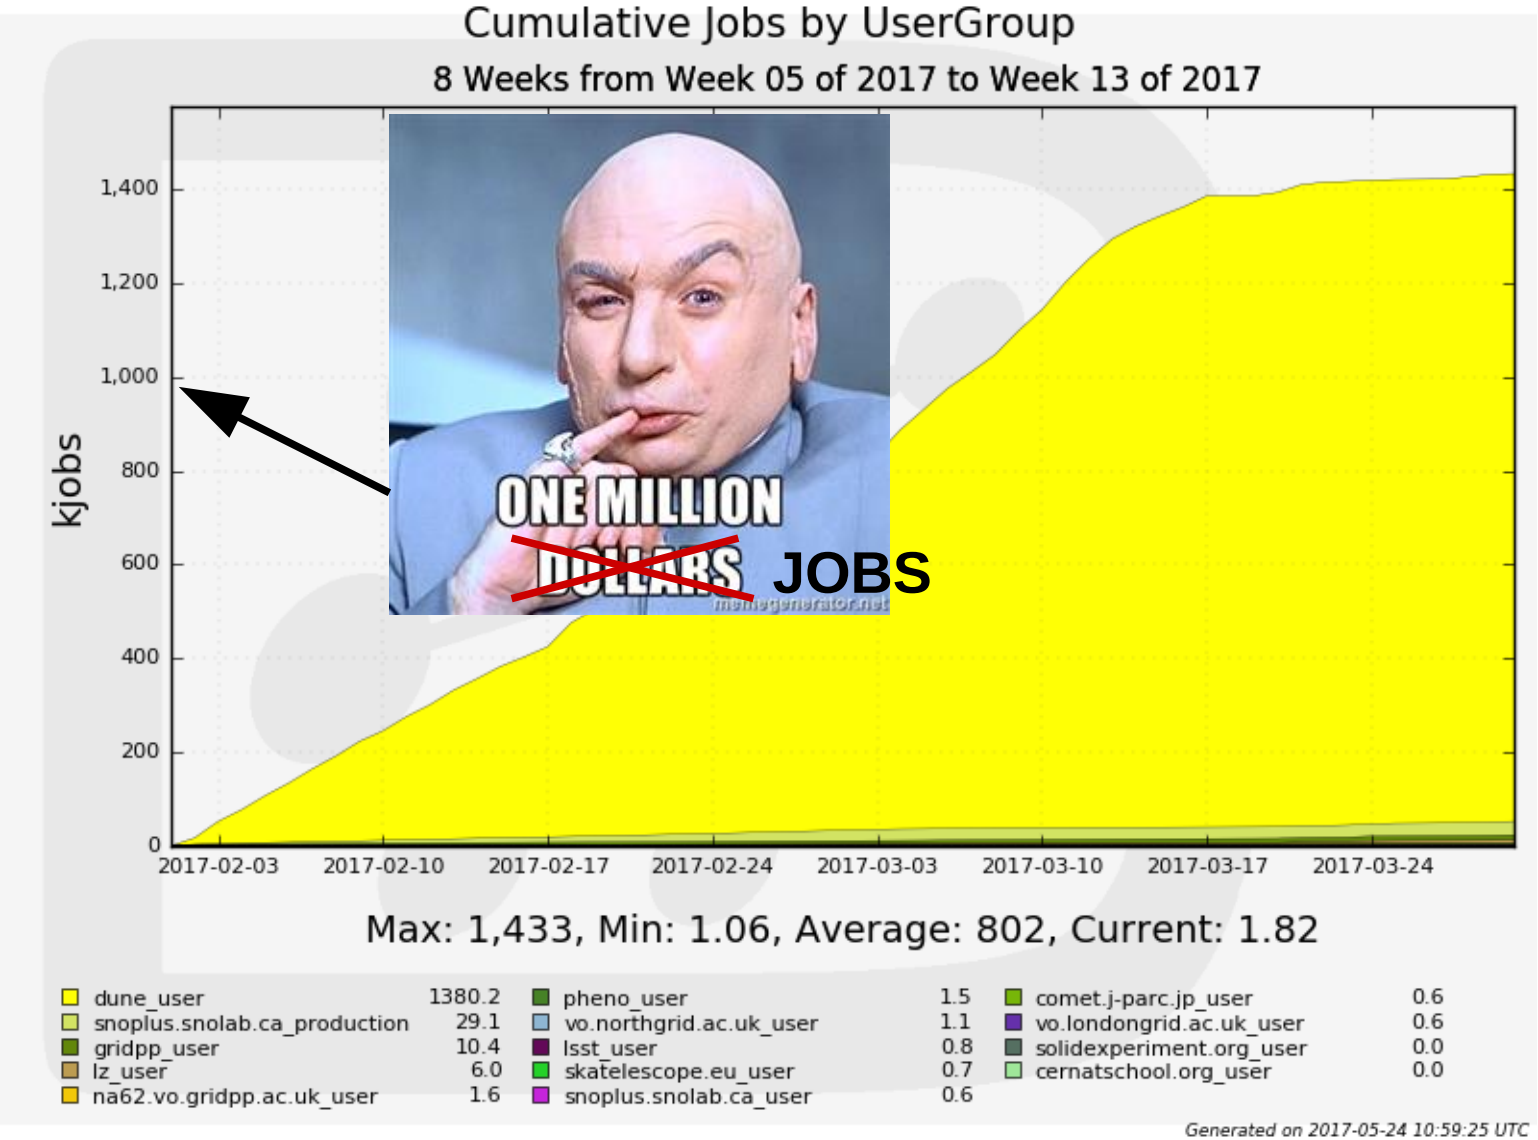

Imperial College London

GridPP

What have we learned from DUNE ?

- DUNE is a Neutrino experiment currently being built in South Dakota and Fermilab.
- Data stored at Sheffield, analysis run all over the UK.
- This was the first time we had 100000+ jobs waiting and ~4400 jobs simultaneously running.
- Our DIRAC instance held up, but we hit queue limits at sites.
- Networking was not a problem.

VOs using GridPP DIRAC

Wrt 2016:
- lsst
+ dune
+ na62

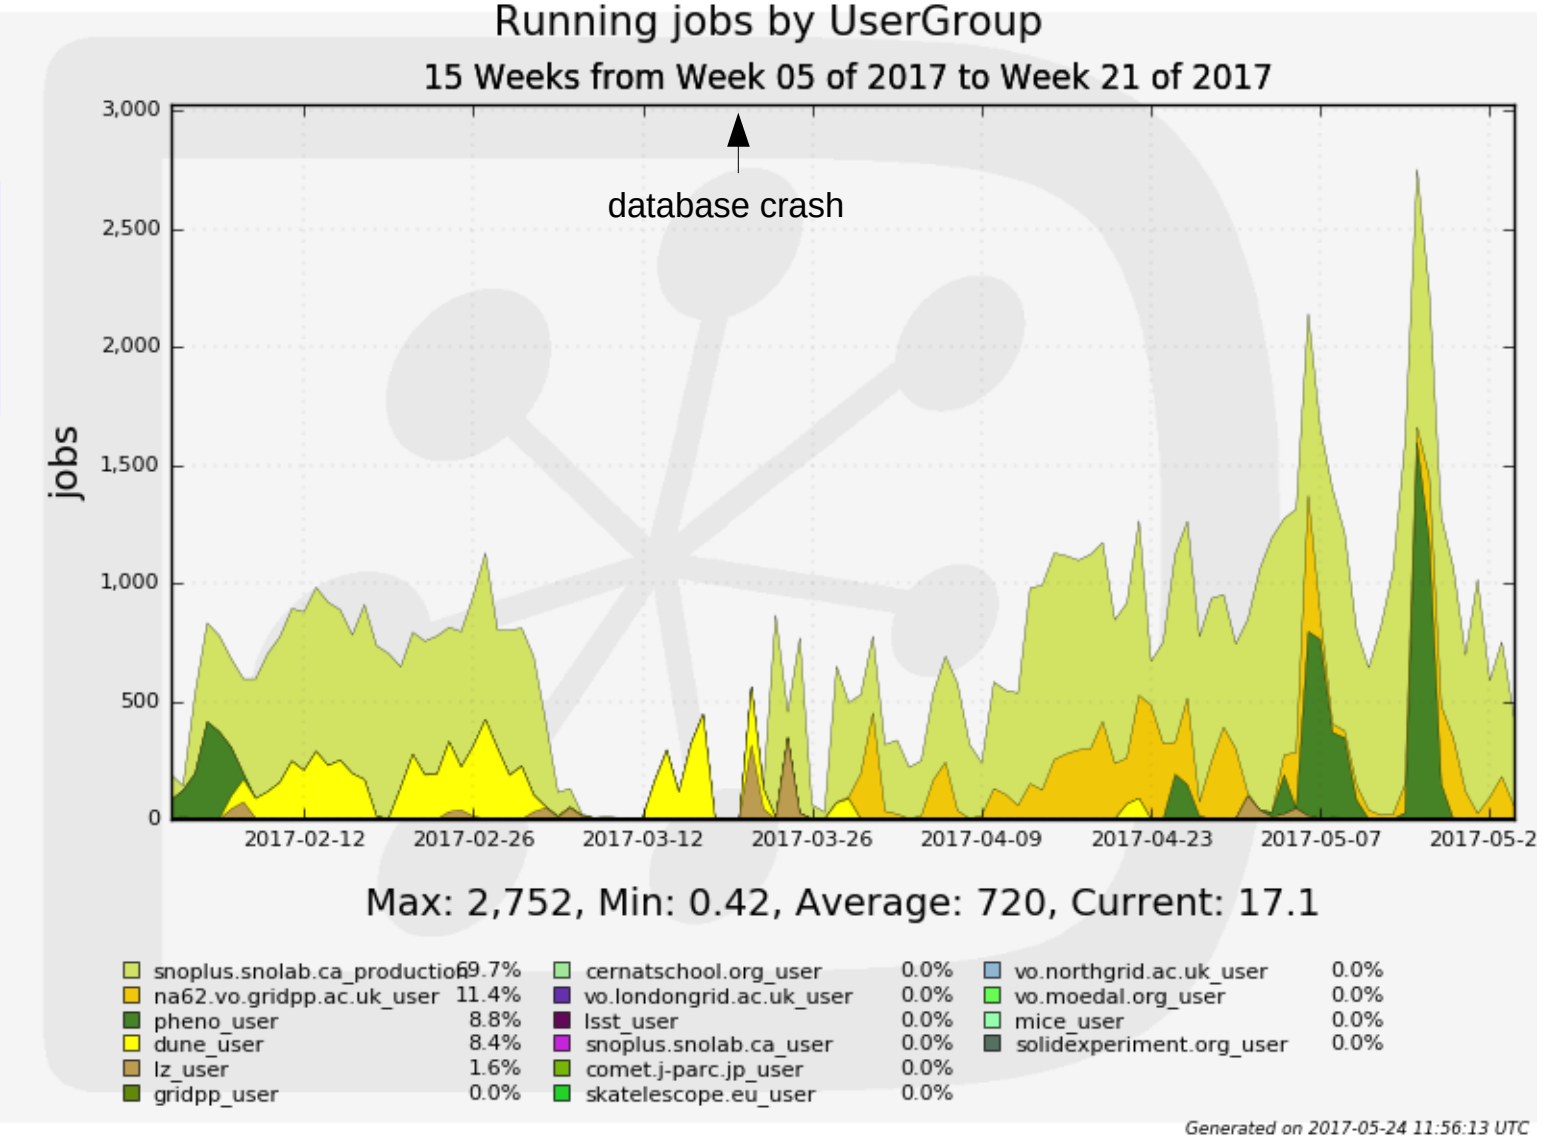

Imperial College London

GridPP

The LZ experiment

Dark Matter experiment looking for WIMPs located in a mine in South Dakota.

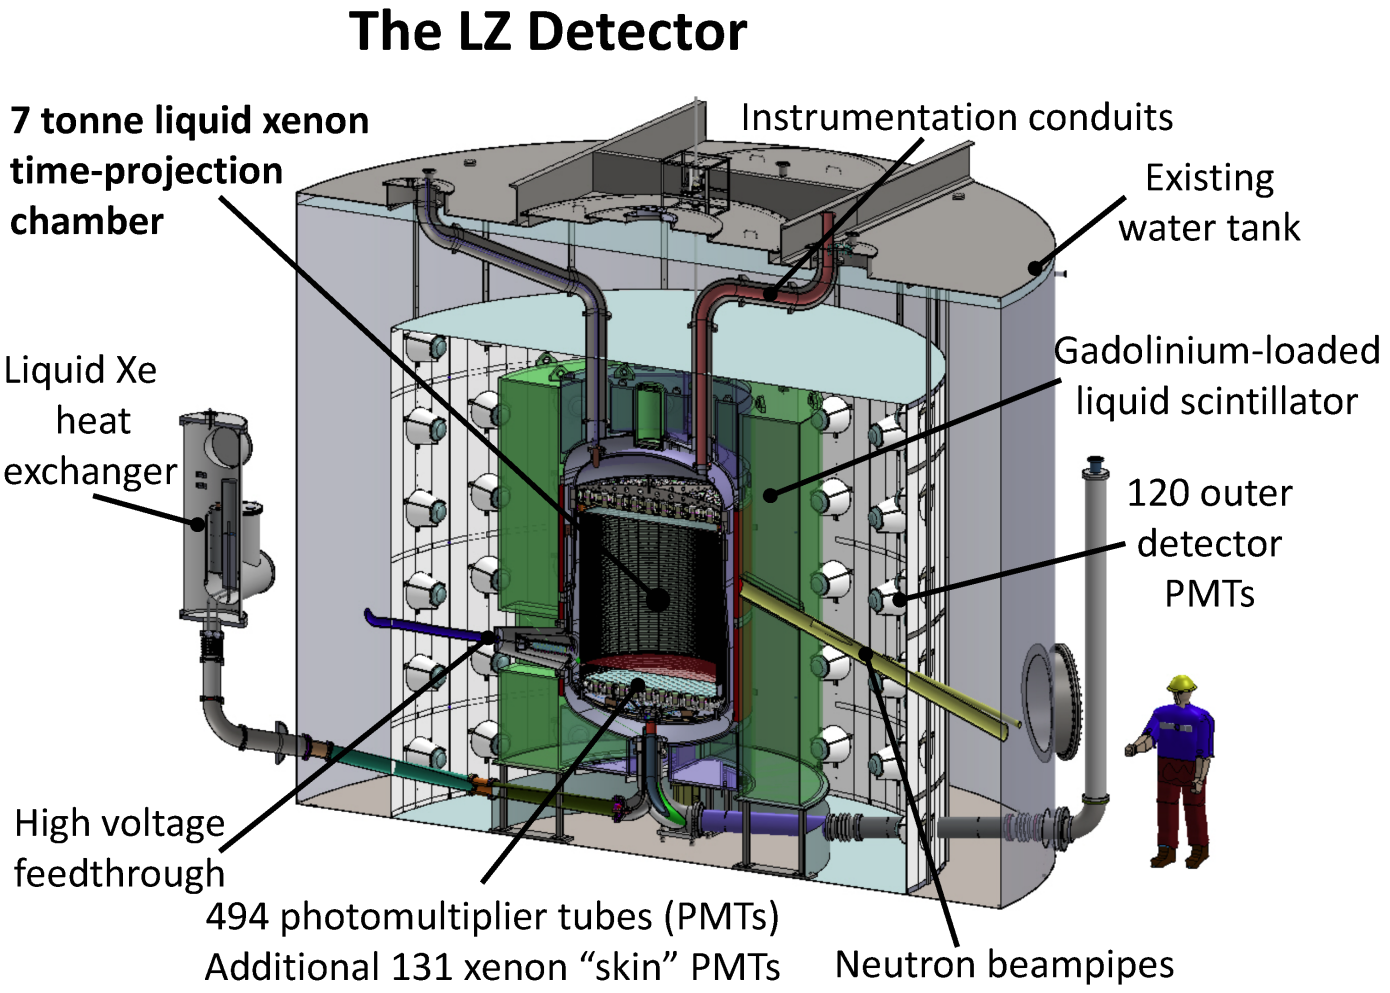

Imperial College London

GridPP

LZ: The UK Data Centre

- The UKDC based at Imperial College will hold a complete copy of the LZ data.
- CPU to analyse this data and for Monte Carlo production will be provided by GridPP sites.
- LZ has been successfully using GridPP resources via DIRAC for over a year.
- We expect LZ to leave the 'small VO' remit by 2020 at the latest.

Imperial College London

GridPP

LZ Submission System: Behind the scenes

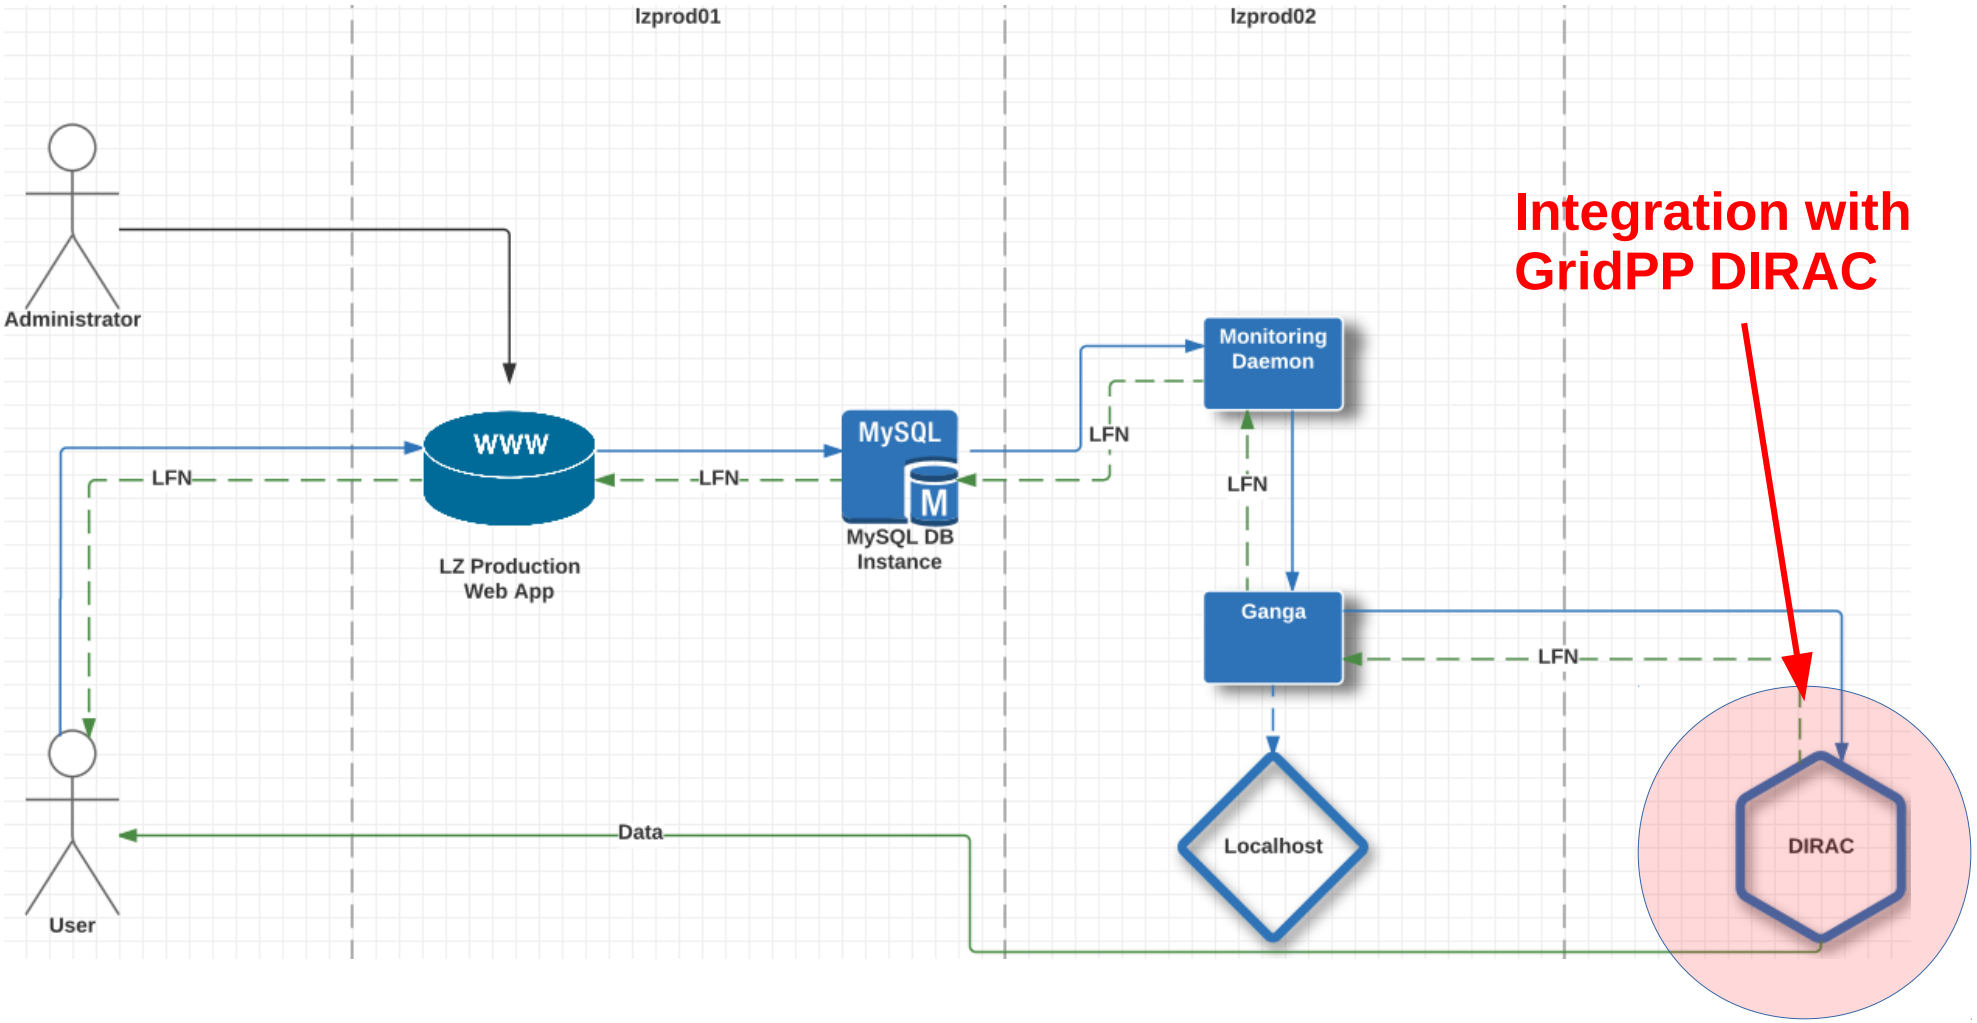

LZ: Production requests web interface

LZ Production Requests

gangad up DIRAC up Admins

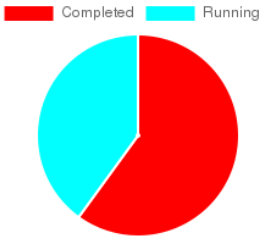

Show 10 entries Search:

	ID	Description	Sim Lead	Status	Request Date	Requester
⊕	41	BACCARAT Test 4	Rob	Running	10/02/2017	rob taylor
⊕	39	BACCARAT Test 3	Rob	Running	10/02/2017	rob taylor
⊕	37	BACCARAT Test 2	Rob	Running	08/02/2017	rob taylor
⊕	35	BACCARAT Test	Rob	Running	08/02/2017	rob taylor
⊕	40	Alex	Alex	Completed	10/02/2017	alexander richards
⊕	38	Alex	Alex	Completed	10/02/2017	alexander richards
⊕	36	Final pre merge LUXSim test	Alex	Completed	08/02/2017	alexander richards
⊕	34	test	Alex	Completed	08/02/2017	alexander richards
⊕	33	test3	Rob	Completed	07/02/2017	rob taylor
⊕	30	test	EK	Completed	07/02/2017	elena korolkova

Showing 1 to 10 of 10 entries Previous 1 Next

⊕ New Request

Imperial College London

GridPP

EL6 to EL7 transition

- Three sites in the UK currently have EL7 resources:
  - ECDF: EL7 only (shared cluster)
  - Imperial: 1/3 of all worker nodes are EL7
  - Liverpool: EL7 test queue
- It's not going well.

GridPP EL6/EL7 selection

- In the GridPP module we squash all operating systems variations to “EL6” or “EL7”.
- User can select either “EL6” or “EL7” or “AnyPlatform” by setting:
  - e.g. Platform = “EL6” in a JDL
- If no Platform is set the system defaults to EL6 to comply with current user expectations.
- We have not been able to run production on EL7 queues due to **gfal2 stage out problems** (nagios tests pass).

Containers

- Yes, we have some.
- CMS requires Singularity on EL7.
- Some VOs (e.g. LZ) require EL6 for the foreseeable future.
- LZ are planning to use Shifter with docker images at NERSC.

Containers ?

- Is there any plan for a container supporting DIRAC pilot ?
  - User picks a platform (e.g. EL6)
  - If a site supports containers, the pilot should start a container matching the requested platform
  - Currently there is no way for sites to advertise that they support containers (can we use software tags ?)
  - Bonus feature: Let users specify their own container image

Containers

- If there is any interest, we have written a proposal at:
  - https://github.com/ic-hep/DIRAC/issues/79
- We'd be happy to code (and document) it.

Conclusions

- It's getting better.
- We are now mainly new stuff work, rather than fixing old bugs.
- gfal2 is the bane of my existence.