Big Sites: Imperial College

Daniela Bauer

GridPP38

Imperial College London

GridPP

What makes a Big Site™ big ?

- There is more to a Big Site™ than lots of CPU and Storage:
  - number of VOs supported
  - number of experiments the group is actively involved with:
    - additional resources due to overlap with experiment specific software developments
    - additional pressures due to users
  - services to the community (GridPP *and beyond*)
    - GridPP DIRAC, voms, WMS, Cloud
  - community involvement
    - e.g. perfSonar, ipv6

Imperial College London

GridPP

Overview: VOs

- LHC VOs: CMS, LHCb, Atlas
- non-LHC VOs:
  - LZ, mice, solid, t2k.org, comet
  - biomed, ilc, lsst, snoplus, dune, na62, pheno

LHC VOs: CMS

- CMS currently uses of 70% of fair share
- Though not very efficiently: multicore jobs
- 2.4 Pb Storage (user data duplicated as backup)
- Accessible via the grid and locally

Big Site™ contributions:

- T3 CMS specific configuration
- Phedex server (data transfer)
- UK xrootd server
- VM images for HLT (Cloud running on Trigger hardware)
- WMAgent (not via GridPP)
- EL7 and Singularity (almost there)

Imperial College London

GridPP

LHC VOs efficiency

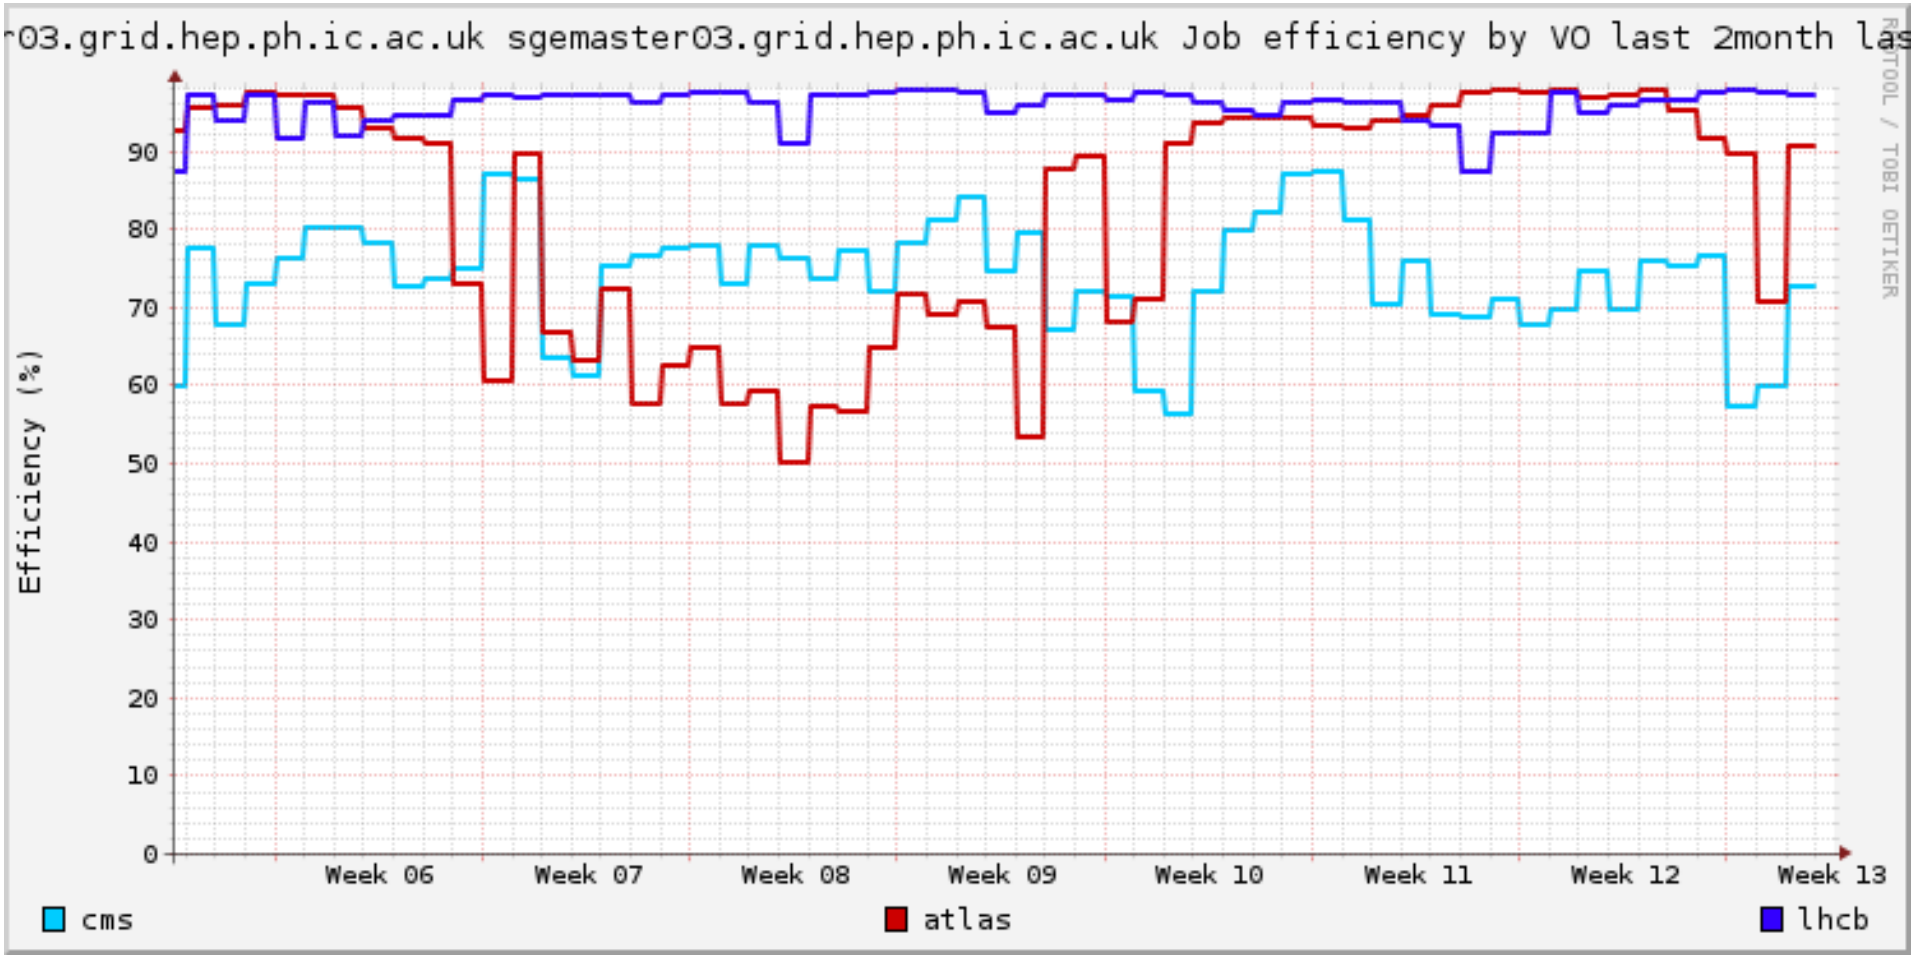

It would be funnier if it wasn't CPU delivered that counted.

Imperial College London

GridPP

And it's not us....

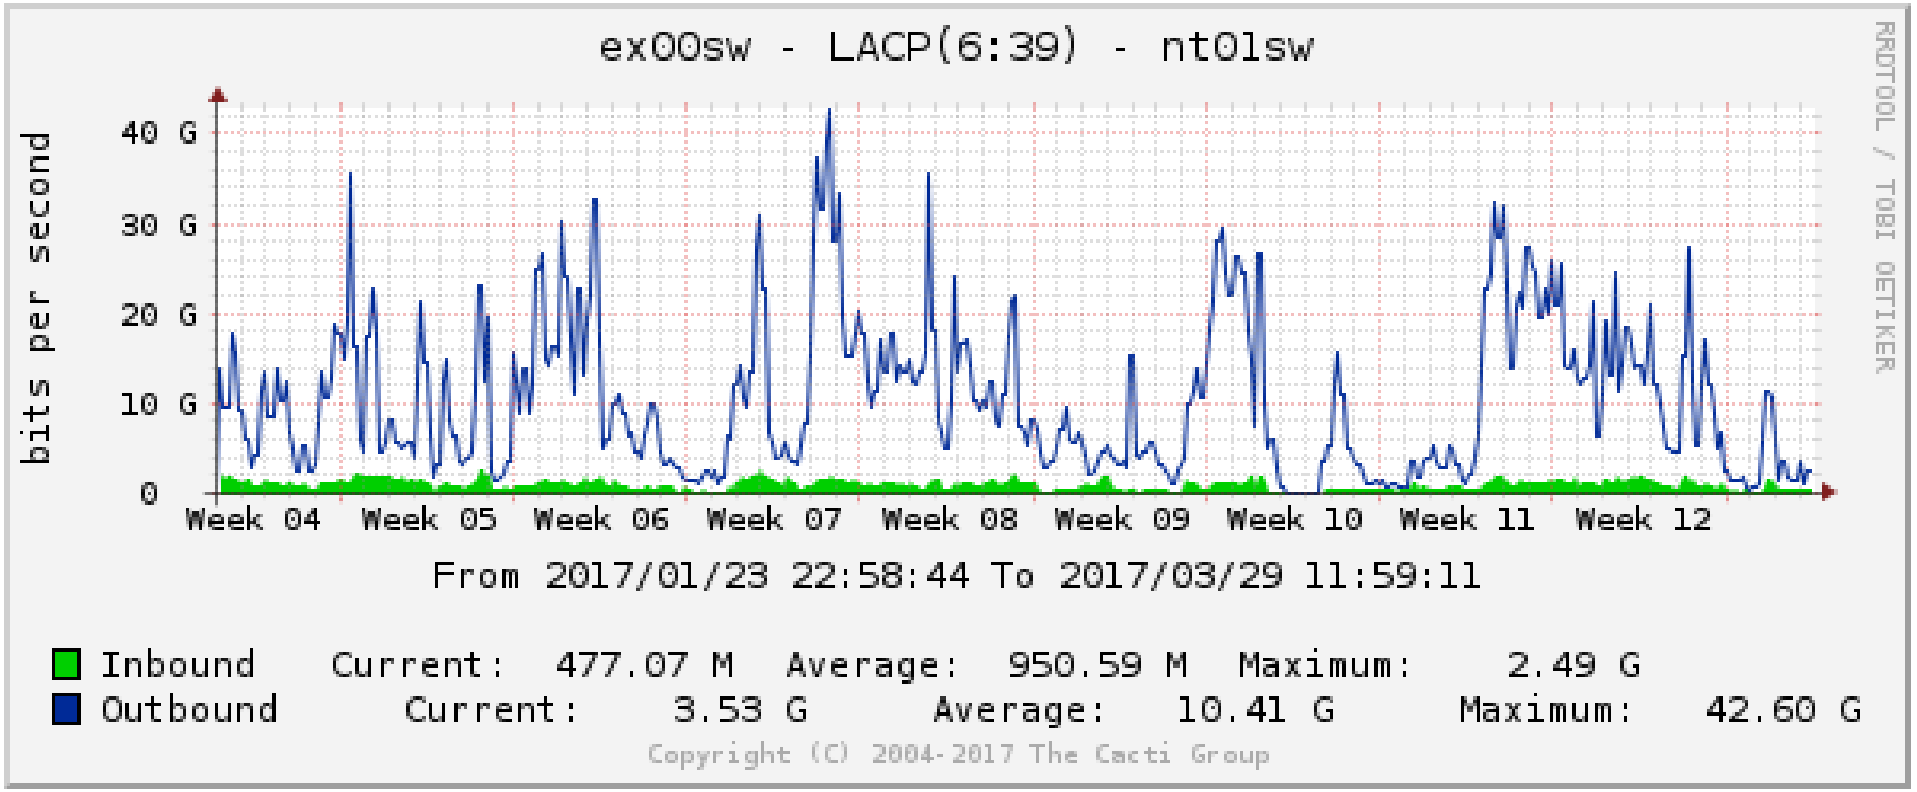

LHC VOs: LHCb

- 416 TB of storage, 15 % of fair share.
- Very efficient use of resources.
- Currently extending usage to cloud.
- Voted “most low maintenance LHC experiment” several years in a row.

non-LHC VOs

- Imperial College contributes significantly to the computing effort of several non-LHC VOs.
- Current work:
  - LZ: UK Data Centre
  - Mice: DataMover (experiment → T0), configuration database
  - Solid: Getting a new VO production ready

The LZ experiment

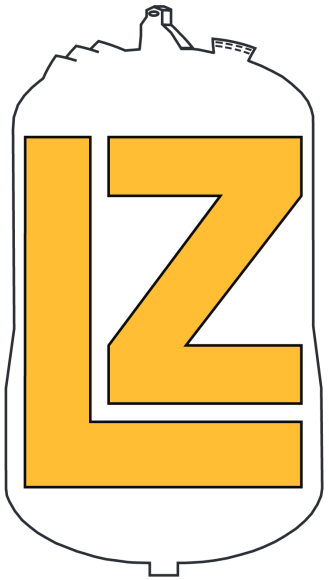

Dark Matter experiment looking for WIMPs located in a mine in South Dakota.
For details: Please see Antonin's talk.

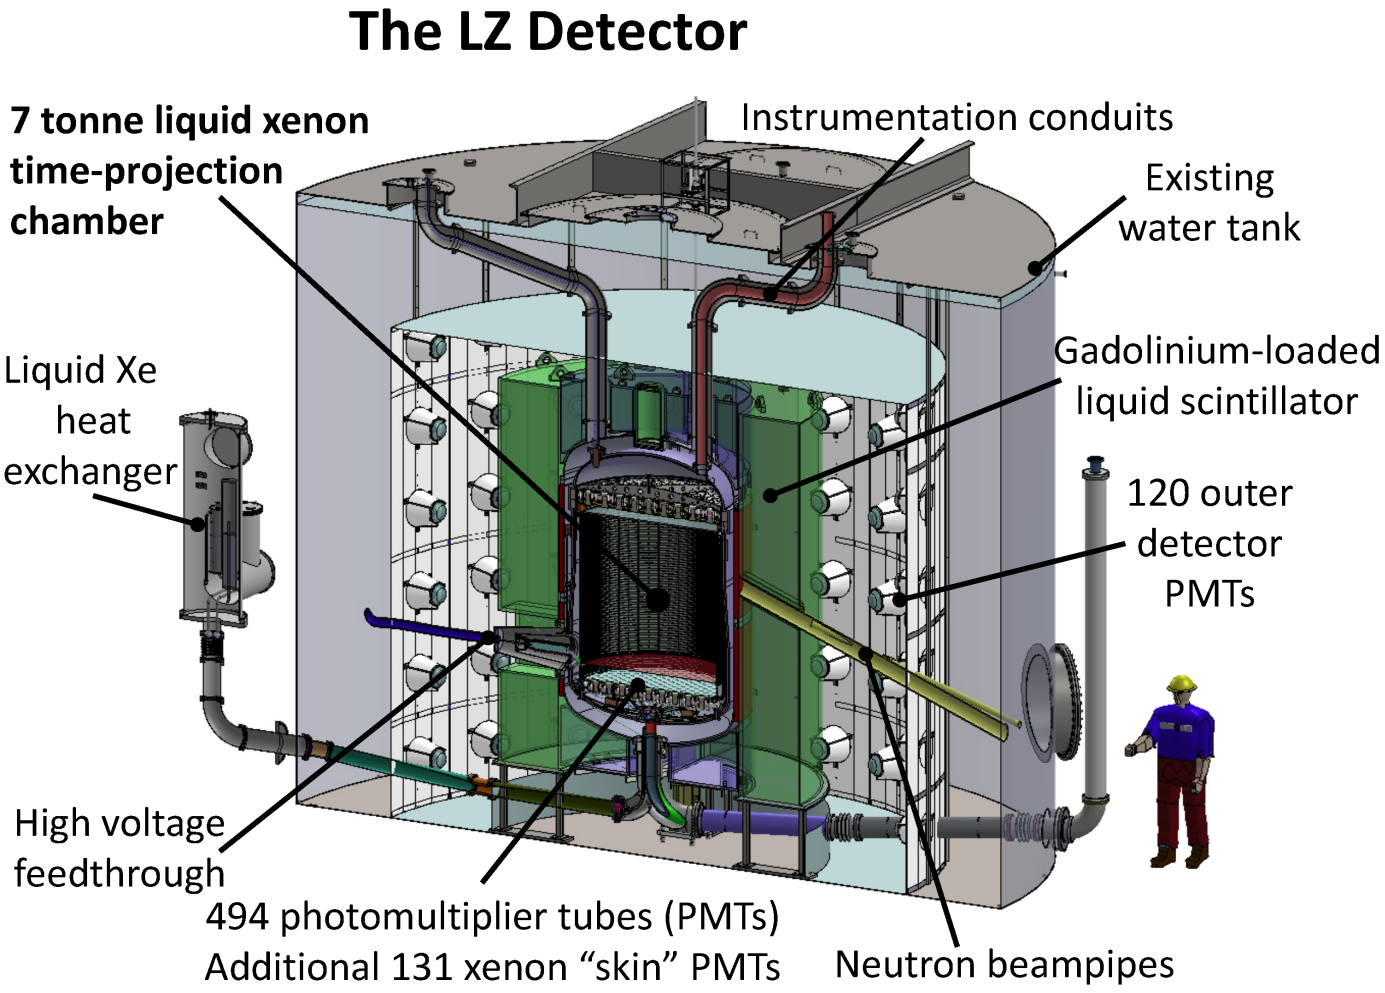

Imperial College London

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 (storage, CPU) via DIRAC for over a year.
- We expect LZ to leave the 'small VO' remit by 2020 at the latest.

LZ: Planned Storage and Processing Capacity (from TDR)

FY	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025
Raw data	–	–	–	–	–	560	1680	2800	3920	5040	6160
Calibration data	–	–	–	–	–	160	480	800	1120	1440	1760
Simulation data	40	80	80	100	100	200	200	200	200	200	200
Processed data	20	40	40	50	50	172	316	460	604	748	892
User data	20	40	40	50	50	55	134	213	292	371	451
Total data	80	160	160	200	200	1147	2810	4473	6136	7799	9463
USDC: Disk space	40	220	220	220	220	1360	3360	5360	7360	9360	11360
USDC: CPU cores	–	–	175	350	350	390	830	1270	1710	2150	2590
UKDC: Disk space*	150	220	220	270	650	1597	3260	4923	6586	8249	9913
UKDC: CPU cores	150	175	350	350	350	390	830	1270	1710	2150	2590

*in TB

LZ UKDC Final Design Review

From the committee report:

GENERAL FINDINGS

- The UKDC has provided essential CPU and storage resources during the CD[1] process and during the preparation of the CDR[2] and TDR. On two separate iterations, the LZ background model has been entirely simulated at the UKDC. The success of LZ in achieving its CD milestones would have been impossible without this essential contribution from the UKDC.
  [1]Critical Decision [2]Conceptual Design Report
- The UKDC team is highly qualified and have successfully supported many other HEP experiments, and LZ will be well-served to rely on the experience and ability as the experiment moves towards first data.
- TL;DR: GridPP is awesome.

Imperial College London

GridPP

LZ: GridPP overlap with experiment specific work

Because we have people working on experiment specific software on site, we can work on a seamless integration of this software with GridPP resources.

So we convert this*

* (otherwise known as the LZ MC production master google spread sheet)

MC production table

File Edit View Insert Format Data Tools Add-ons Help

imperial.richards@gmail.com

View only

Share

	A	Requester (not just initials)	Date of request (MM/DD/YY)	Description	Source types and stats	Detector components	LUXSim/ BACCARAT/... ver.	Sim lead (not just initials)	Job submission tab	Status	Date complete (MM/DD/YY)
2	25	Hugh L.	2/22/17	Energy depositions in the skin	Xe131m decays in skin?	LiquidXenonSkin comp	BACCARAT new release				
3	24	Hugh L.	2/22/17	Multiple scattering simulation for LZap developmer	200 keV gammas might be doable - 1000? on a FNAL machine it's 39060 seconds per event	LiquidXenon, PMT Pho	BACCARAT new release				
4	23	Cees C.	11/30/16	Kr-83 sims for DER validation	100 events made with the following mac	LiquidXenon, PMT Pho	POSTPONED UNTIL NEXT RELEASE OF DER. Do not use DER release-1.0.4 and LUXSim release-4.4.6_geant4.9.4				
5	22	Jim Dobson	10/09/16	Check of full optical sims for LZAP for BACCARAT	0-10 keV ER, 0-30 keV NR, uniform in t	LiquidXenon	BACCARAT/release-[0.X.X, POSTPONED TILL	Maria Elena	BACCARAT v0.2.1 ER and NR full sims for LZap	2) LUXSim/macros complete	
6	21	Maria Elena	10/26/16	Full optical sims (ER+NR) for LZap validations	0-10 keV ER, 0-30 keV NR, uniform in t	LiquidXenon	LUXSim/release-4.4.6 DER/release-1.0.4	Maria Elena	LUXSim v4.4.6 ER and NR full sims for LZap	6) Jobs completed	
7	20	Melih S.	10/07/16	Full neutron sims for the latest OD design+Water	Bckg U_early/late/Th neutrons + Water,	All components used i	Will use a forthcoming release of LUXSim after	Melih S.	LUXSim v4.X.X NR and ER with updated OD geome	5) Jobs submitted	
8	19	Maria Elena	09/27/26	Full stats for Xe131m, uniform in TPC (and RFR)	Xe131m decays uniform in LXe volume	LXe (TPC+RFR)	LUXSim/release-4.3.3 TDRAnalysis/release-3.3	Maria Elena	LUXSim v4.3.3 Xe131m e-deposit sims	6) Jobs completed	
9	18	Steve W.	09/23/16	Full optical sims for DER background deliverable	0-10 keV betas uniform in the LXe volun	LiquidXenon	LUXSim/release-4.4.4 reduction:LUXSim/releas	Maria Elena	LUXSim v4.4.4 ER full sims for DER	6) Jobs completed	10/28/16
10	17	Tyana S.	09/18/16	BG sims for BACCARAT validation using Geant4.9	U_early/late/Th neutrons + gammas	neutrons for TPC PMTs	BACCARAT/0.2.0 TDRAnalysis/release-3.13.1	Tyana S.	BACCARAT v0.2.0 BG sims for G4.9.5 validations	6) Jobs completed	
11	16	Henrique A.	30/08/16	BG sims for assessing alternative materials for gri	U/Th/60Co/34Sc gammas + Neutrons fr	All grids separately (bo	LUXSim/release-4.3.3 TDRAnalysis/release-3.3	Tyana S.	LUXSim v4.3.3 BG sims for Oct. review	6) Jobs completed	
12	15	Tyana S.	8/20/16	BACCARAT v0.2.0 validation sims for g4decay gar	U238 and Th232 gammas	Vessels and TPC PTFE	BACCARAT/0.2.0	Tyana S.	BACCARAT v0.2.0, g4decay gamma BG checks	6) Jobs completed	
13	14	Jim D.	08/18/16	BG sims for Oct. backgrounds review	U_early/late/Th neutrons + gammas	Initially limited set of components, will be increased once checked		Elena K.	LUXSim v4.3.3 BG sims for Oct. review	6) Jobs completed	
14	13	Peter S	08/12/16	Full optical for saturation study (~1000 events per	131mXe at radii in range (~50 cm -- ~75 cm) and z=0			Asher K./Jim D.	LUXSim v4.3.3 Xe131 sims for saturation study	4) Job tab filled	
15	12	Jim D.	07/27/16	BG sims of key components for collaboration mee	U_early/late/Th neutrons + gammas	TPC PMTs, PTFE walls	LUXSim/release-4.3.2 TDRAnalysis/release-3.3	Elena K.	LUXSim v4.3.2 key BG sims for testing	6) Jobs completed	
16	11	Matthew S.	06/21/16	ER and NR events for LZAP development	1e4 ER unif(2,5), 1e4 NR unif(5,10)	LiquidXenon	LUXSim/release-4.3.2 ElectonicsSimulation/rel	Maria Elena	LUXSim v4.3.2 ER and NR for DER	6) Jobs completed	08/16/16
17	10	Jim D.	06/14/16	Additional BG sims for components missed during	U_early/late/Th neutrons + gammas + H	1) OD tank, 2) OD Liqu	LUXSim/3.7.0	Elena K.	LUXSim v3.7.0	6) Jobs completed	
18	9	Melih S.	05/09/16	Repeat TDR neutron sims for candidate OD geometry	U_early/late/Th neutrons (TDR stats)	All components used f	LUXSim/MS-CaseL-6122cd76	Melih S.	LUXSim MS-CaseL-6122cd76, stats for OD-rate stu	5) Jobs submitted	

Imperial College London

GridPP

LZ: Production requests web interface

LZ Production Requests

gangad up DIRAC up Admins

To this:

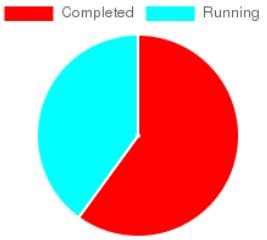

Show 10 entries Search:

	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

LZ Submission System: Behind the scenes

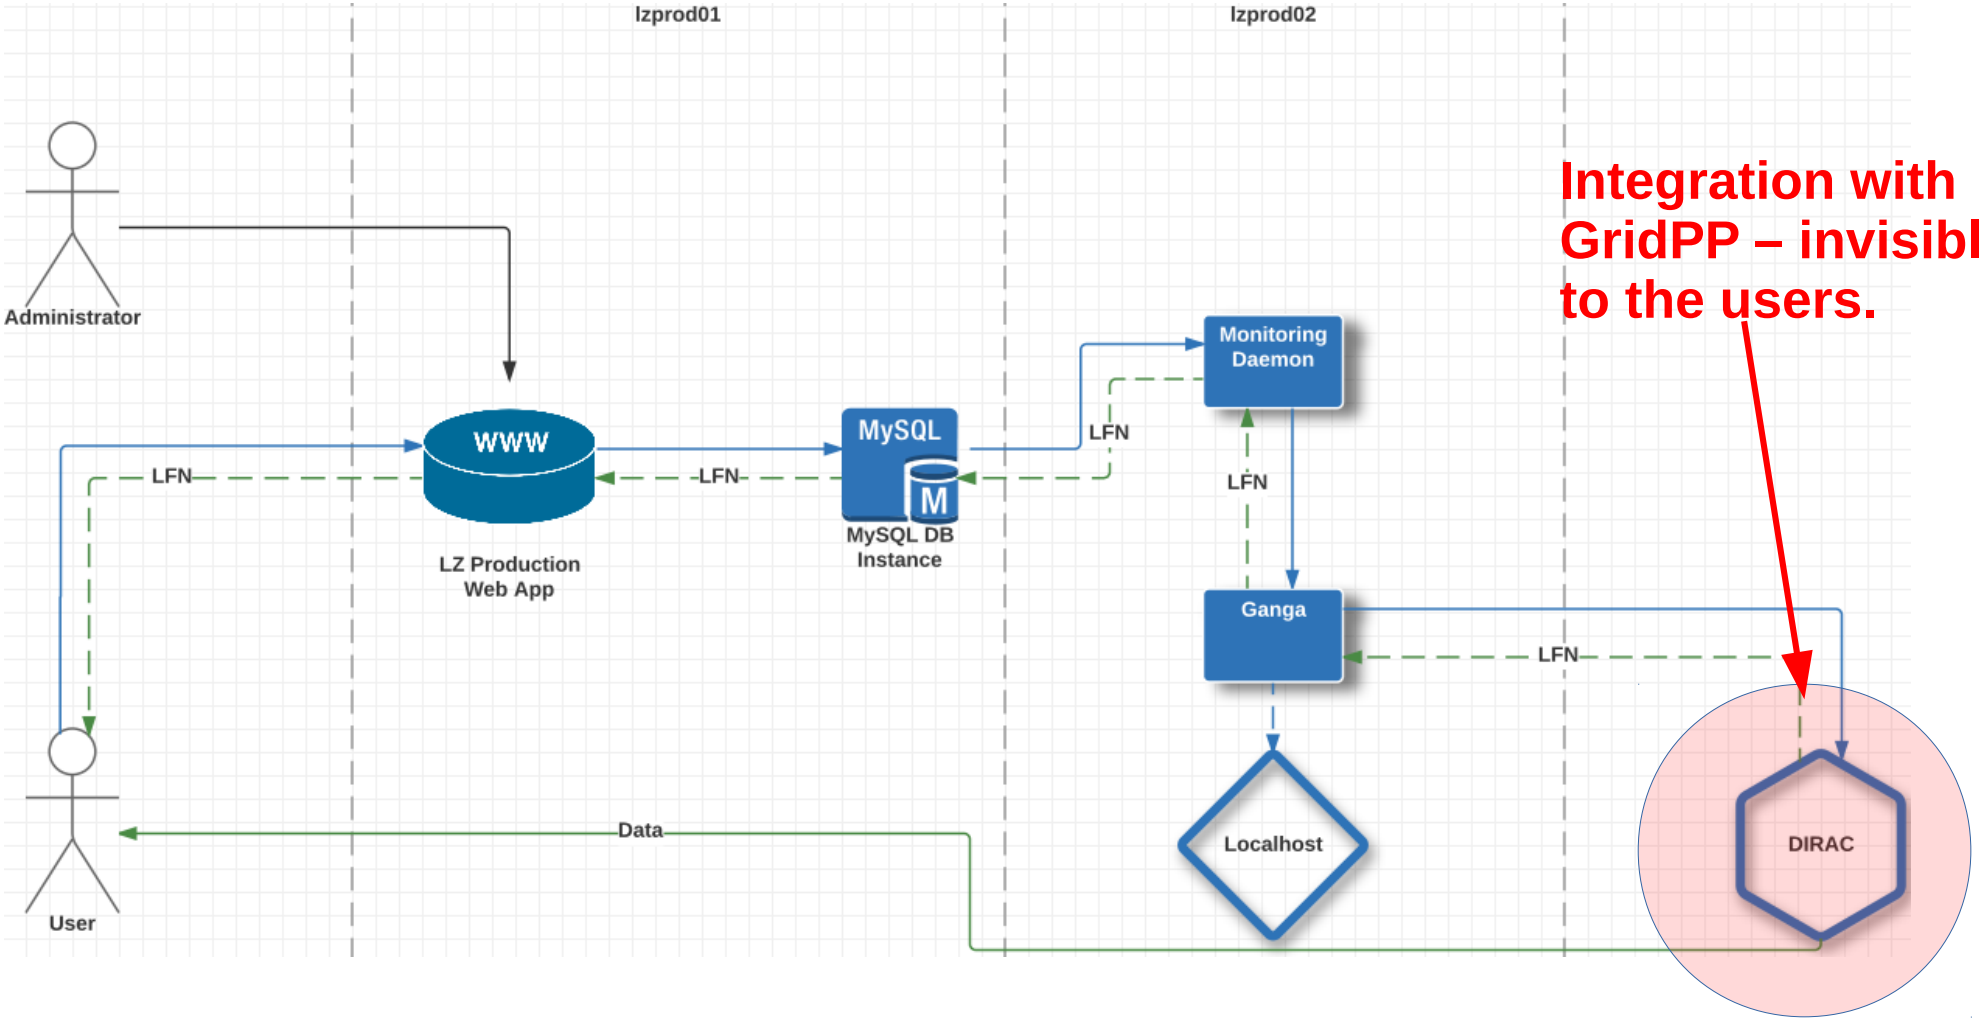

Services to the Community*

- GridPP DIRAC
- Cloud (finally...)
- voms, WMS

  * stuff that requires maintaining extra hardware

Community Involvement

- Networking

Services to GridPP: GridPP DIRAC

- **Overview in CHEP 2016 talk:** https://indico.cern.ch/event/505613/contributions/2230725/attachments/1340849/2038979/Oral-383.pdf
- Main conclusions:
  - **We have a fairly computer literate user base, so most problems and/or requests they have are non-trivial.**
  - The best venue to report dirac problems is the gridpp dirac help list.
  - Not everything that looks like a dirac problem is a dirac problem. But as users see their 'dirac'-jobs fail, Simon and me tend to be the first point of contact.
  - And often the last one as well (but we can't just fop our users off with 'site issue' or 'VO issue', though we try and channel their requests correctly)
  - Often figuring out whether it's a dirac/VO/site problem is the main part of the work.
  - We try to avoid expectation management, but occasionally it's unavoidable: It's a shared resource (just like the rest of the grid).

GridPP DIRAC - maintenance

- The GridPP dirac server is regularly upgraded, often in reponse to user/admin requests (xrootd, EL7, condor).
- We test our updates on a dedicated test server and it rarely works out of the box.
- We submit bugs and fixes to the dirac repo.
- Most of the technical issues are solvable, but we need time and money to do this.

DIRAC: GridPP *and beyond*

- Several VOs supported on DIRAC use sites outside GridPP:
  - na62: Italy, Belgium
  - solid: Belgium (different site to Na62)
  - snoplus: Canada, Portugal
- Usually cooperation with these sites works smoothly.
- Increases the visibility of GridPP.

Imperial College GridPP cloud (1)

- OpenStack (Ocata) in fully redundant tri-controller configuration.
- Image & Volume storage on CEPH (Jewel).
- Direct (provider/bridged) network connection for high-performance SE access.

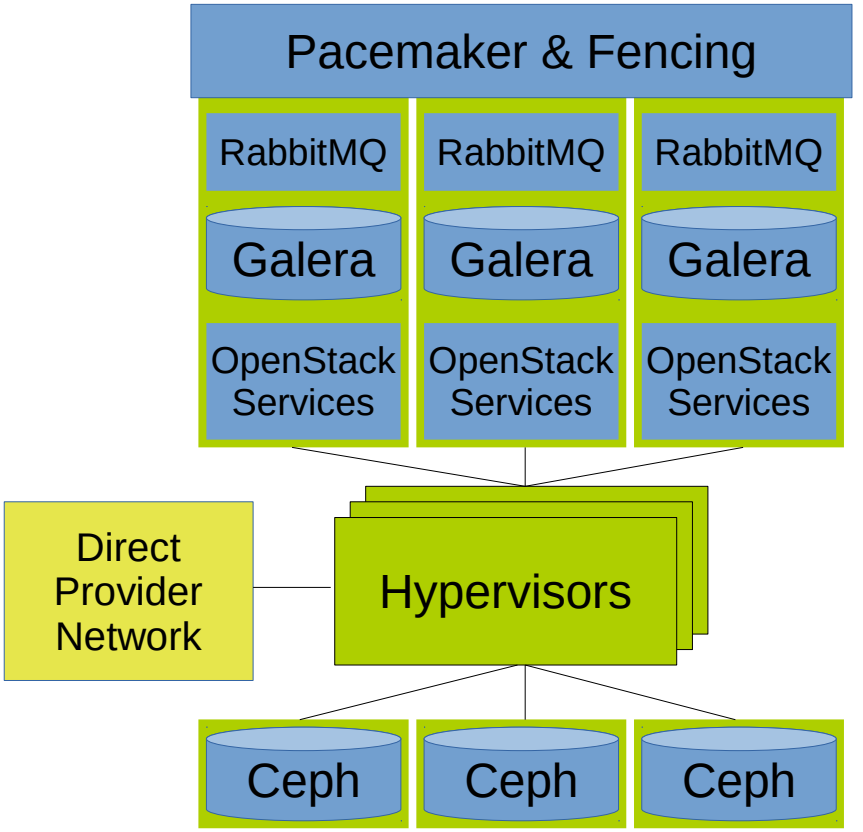

Imperial College London

GridPP

Imperial College GridPP cloud (2)

- Currently ~260 cores & ~30TB of storage.
- Now production ready (from 1st April).
- Available for development use:
  - Ensure you are a member of the GridPP VO and then e-mail *gridpp-cloud-admin@imperial.ac.uk* for access.
- Cloud has its own site-name, UKI-GridPP-Cloud-IC:
  - GGUS tickets can be opened against the site if there are problems.
  - We will soon run “grid” jobs on the idle cores

Imperial College London

GridPP

Networking

- IPv6
  - Active membership of HEPiX IPv6 working group
  - The site has been dual-stack for several years – allows us to be tested against
  - Running IPv6 FTS service for London CMS sites
  - Assisting in testing of the IPv6 WN at Brunel for LHC VOs
- perfSONAR
  - Membership of WLCG Network Throughput working group
  - Maintaining UK and dual-stack perfSONAR meshes
- Part of Jisc End-to-end Performance Initiative
  - sharing GridPP experiences of high-throughput networking, monitoring and IPv6 with other research domains and institutes

Non GridPP Funding

- College pays for electricity (incl. cooling) and machine room space
- Consolidated grant: (small) manpower bonus, some equipment (network switches)

Outlook

- As a Big Site™ we will continue to provide GridPP++ services.
- DIRAC has become public face of GridPP for non-LHC VOs.
- Medium term: CMS in stable mode right now, concentrate on non-LHC VOs until LHC ramps up.
- Long term: e-Infrastructure ?