

DIRAC - a global workload management system

Daniela Bauer

Overview

- What is DIRAC and what is it being used for ?
- Providing a DIRAC server: Admin perspective
- DIRAC from a user perspective:
 - How do different communities use DIRAC

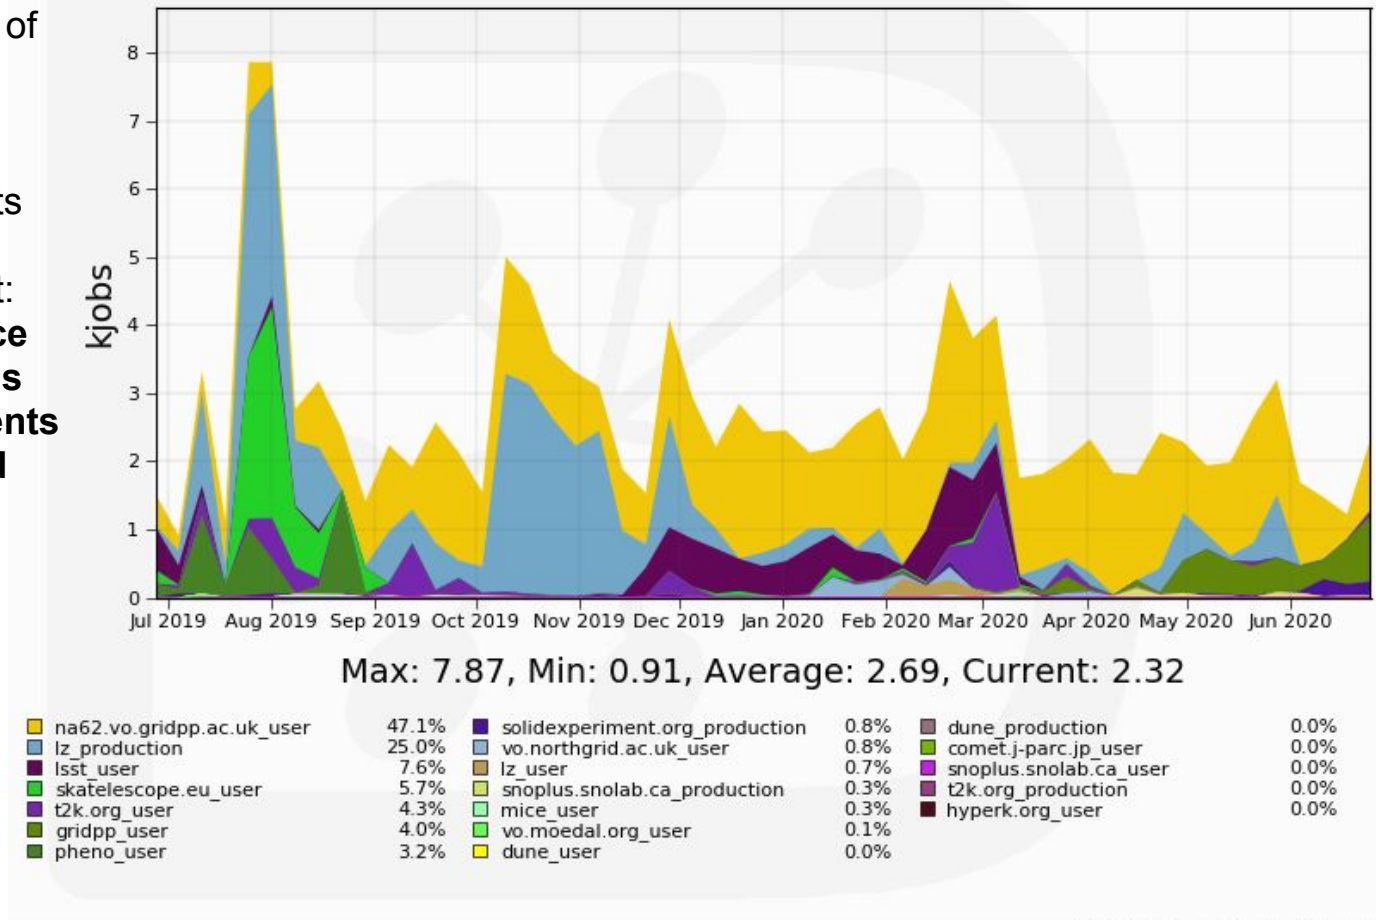
Overview

- DIRAC is a software originally developed by LHCb that comprises of:
 - **Workload Management System** (“Global batch system”)
 - File Catalog
 - Workflow Management System
 - Documentation: <https://dirac.readthedocs.io/en/latest/>
- Provides a standardized user interface to multiple compute and storage resources
- Written in Python (for Linux)
 - OpenSource: <https://github.com/DIRACGrid/DIRAC>
- Apart from LHCb it is used by a number of communities to manage the various aspects of their data processing:
 - Experiment specific: e.g. Belle2, ILC, Cherenkov Telescope Array, NICA (JINR), BES (Beijing), biomed
 - **Multi-Community**: e.g. GridPP, France-Grilles, EGI

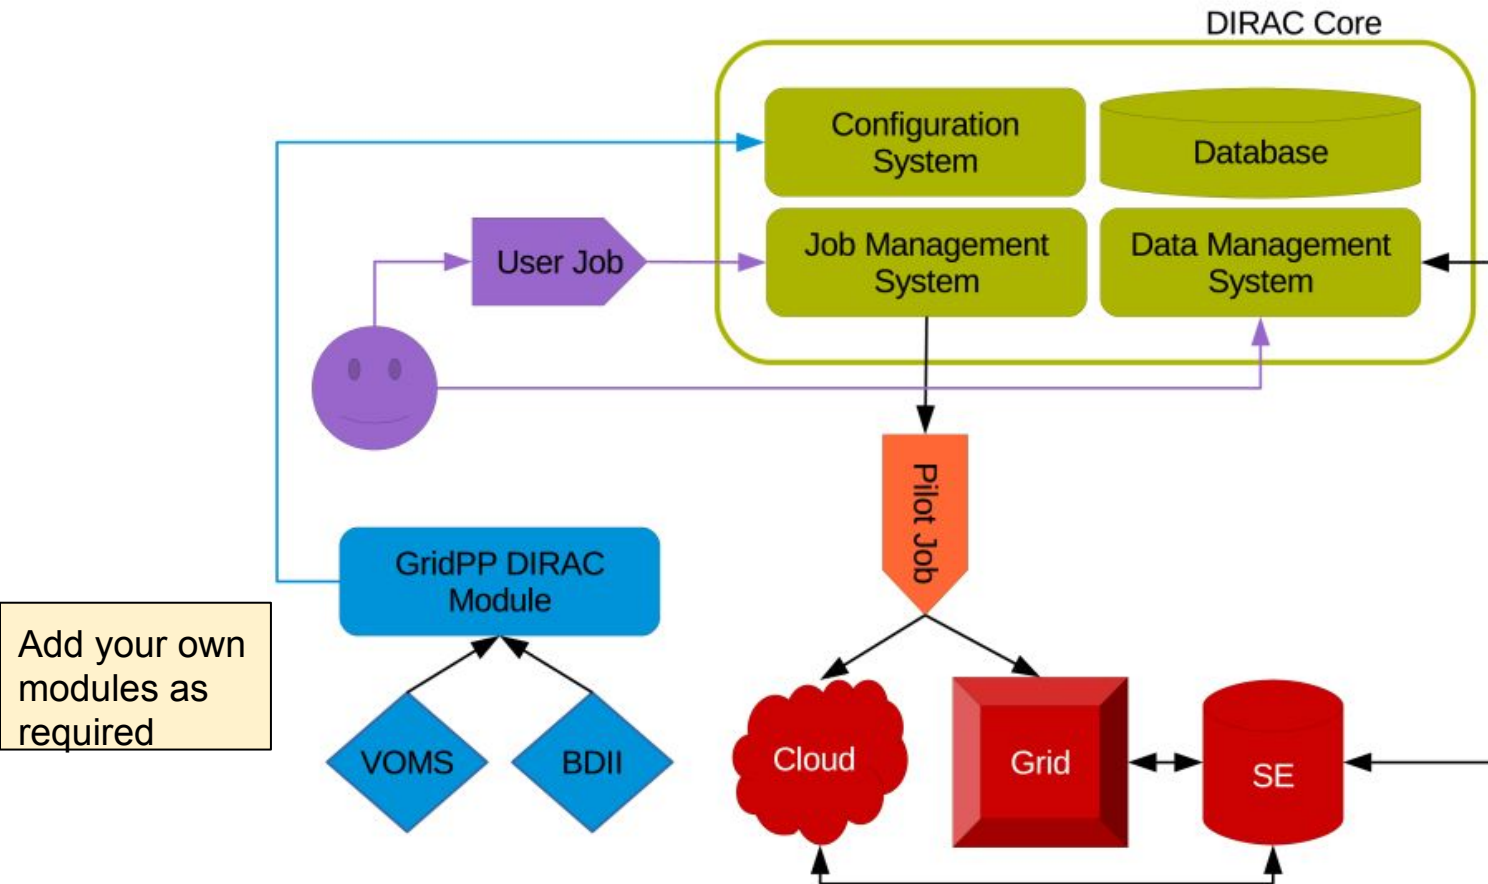
Background: GridPP/DIRAC

- GridPP: A collaboration of 19 UK universities
- Focus on providing computing resources to the big LHC experiments
- Dedicated “other experiment” component: **GridPP DIRAC instance is provided as a means for non-LHC experiments to access GridPP (and other compatible) resources**

Running jobs by UserGroup
52 Weeks from Week 25 of 2019 to Week 25 of 2020



Example: The GridPP DIRAC implementation



DIRAC: Admin view

GridPP - DIRAC - Mozilla Firefox

System Administration [Untitled 2]

Overall System Information

Restart Start Stop

System	Name	Module	Status	Uptime	PID	CPU(%)	MEM(%)	RSS(MB)	VSZ(MB)	
Type: Services (24 Items)										
<input type="checkbox"/>	DataManagement	FTS3Manager	FTS3Manager	Run	265152	3241	0	0.5	82.578125	704.69921875
<input type="checkbox"/>	DataManagement	StorageElement	StorageElement	Run	265152	2749	0	0.3	53.9921875	600.15234375
<input type="checkbox"/>	DataManagement	FileCatalog	FileCatalog	Run	265152	3166	1.4	0.7	126.08984375	1135.5078125
<input type="checkbox"/>	Framework	SystemLogging	SystemLogging	Run	265153	2641	0	0.3	58.34375	602.6953125
<input type="checkbox"/>	Framework	Monitoring	Monitoring	Run	265152	2672	1.1	0.5	83.80859375	1521.3671875
<input type="checkbox"/>	Framework	Notification	Notification	Run	265152	3087	0	0.3	54.44921875	602.67578125
<input type="checkbox"/>	Framework	SecurityLogging	SecurityLogging	Run	265152	2926	0.6	0.4	74.75	813.859375
<input type="checkbox"/>	Framework	SystemLogging...	SystemLogging...	Run	265152	3089	0	0.3	56.58984375	602.6953125
<input type="checkbox"/>	Framework	Plotting	Plotting	Run	265153	2541	0	0.5	84.91796875	1035.62109375
<input type="checkbox"/>	Framework	ComponentMo...	ComponentMo...	Run	265152	2899	0.1	0.5	85.5859375	697.15625
<input type="checkbox"/>	Framework	UserProfileMan...	UserProfileMan...	Run	265153	2508	0.1	0.4	73.7109375	605.1953125
<input type="checkbox"/>	Framework	ProxyManager	ProxyManager	Run	265152	2830	1.1	2.8	458.78125	2249.2578125
<input type="checkbox"/>	Framework	SystemAdminis...	SystemAdminis...	Run	265152	2920	0.2	0.3	62.5703125	704.09375
<input type="checkbox"/>	Framework	BundleDelivery	BundleDelivery	Run	265153	2525	1.7	0.4	79.34765625	1060.31640625
<input type="checkbox"/>	WorkloadMana...	Matcher	Matcher	Run	265152	2969	4.4	0.9	148.359375	1663.24609375
<input type="checkbox"/>	WorkloadMana...	JobMonitoring	JobMonitoring	Run	265152	2816	8.3	0.5	94.98828125	1031.4921875
<input type="checkbox"/>	WorkloadMana...	JobManager	JobManager	Run	265152	2733	0.6	0.6	101.63671875	1070.69140625
<input type="checkbox"/>	WorkloadMana...	JobStateUpdate	JobStateUpdate	Run	265153	2601	3.6	0.6	109.22265625	1915.984375
<input type="checkbox"/>	WorkloadMana...	WMSAdminstr...	WMSAdminstr...	Run	265152	3236	1.9	0.6	111.05859375	1384.56640625
<input type="checkbox"/>	WorkloadMana...	OptimizationMind	OptimizationMind	Run	265152	3091	0.5	0.5	92.90234375	1547.2285625
<input type="checkbox"/>	RequestManag...	ReqManager	ReqManager	Run	265152	3156	0.1	0.5	94.04296875	827.078125

Auto: Disabled Updated: 2018-10-08 16:05 [UTC]

Reduced Overview

Default

Displaying 1 - 64 of 64

No such thing as a free lunch: Support effort

- Disclaimer: This is the GridPP experience, supporting ~6 active communities at any one time; one or two new communities a year.
- Core software development is driven by LHCb
 - But: Active developer group (~10 people) drawing from all communities
- Two research software engineers at approximately ½ FTE each to provide a production level service:
 - Maintenance/upgrades
 - Bug fixes/Coding new features (first) needed by local communities
 - DIRAC↔Resource provider debugging
 - User assistance

User interactions with DIRAC

DIRAC as a 'direct' user interface:

- Command line tools from DIRAC UI:
 - Once a DIRAC server is set up, it can be used with minimal setup (on the users' side)
 - Job submission
 - Data access
- DIRAC Python API:
 - Because “by hand submission/shell scripts” don't scale
- Ganga:
 - <https://ganga.readthedocs.io/en/latest/>
- Web interface to monitor jobs status etc

Restricted/production usage:

- Community developed submission interfaces build on the DIRAC Python API

The LZ job submission interface

Key features:

- Only valid combinations of software versions can be selected.
- Requests require approval by a production manager before submission.
- Database keeps track of status of individual requests that might consist of hundreds of jobs.

The screenshot shows a web-based form titled "New Request Form". It is divided into two main sections: "Basic Information" and "Application Setup".

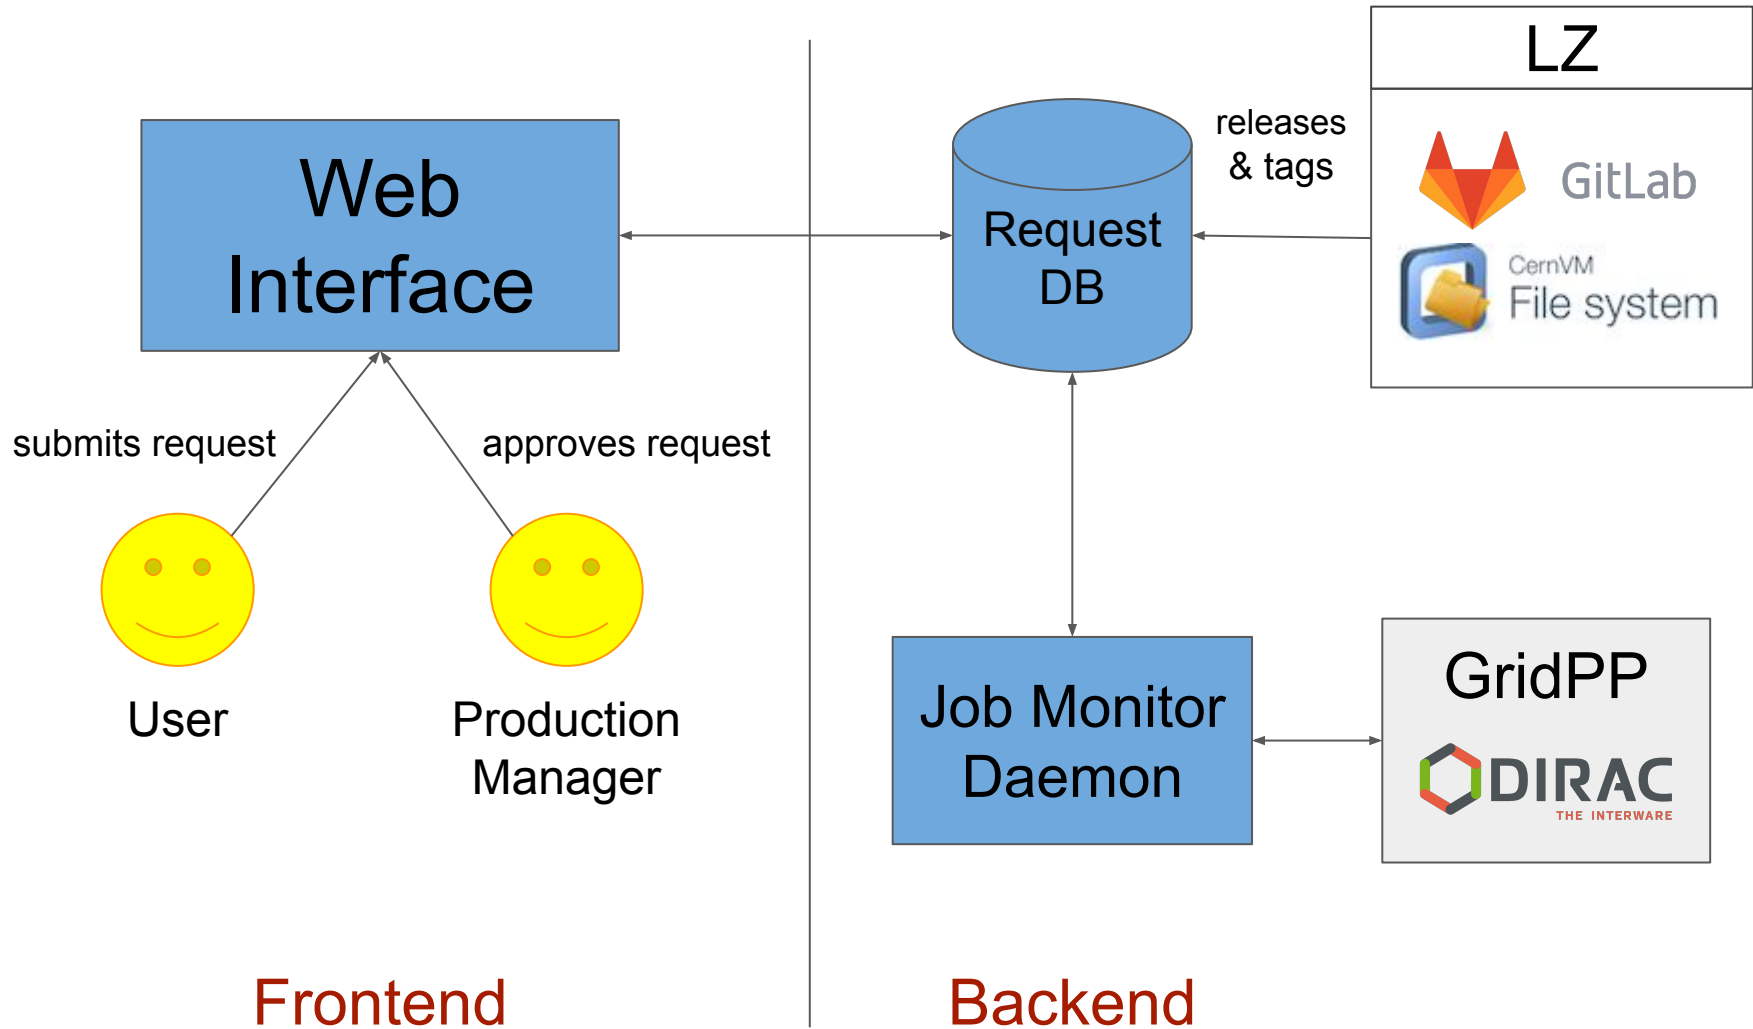
Basic Information:

- Site:** A dropdown menu with "ANY" selected.
- Priority:** A dropdown menu with "3" selected.
- Sim Lead:** A text input field containing "Sim Lead".
- Source type:** An empty text input field.
- Detector components:** An empty text input field.
- Description:** A large text area with the placeholder text "Enter a description of the request here."

Application Setup:

- Four tabs are visible: "Simulation" (active), "Reduction", "DER", and "LZap".
- A toggle switch is currently turned on, with a power icon and the text "Off/On" next to it.
- Select Application Type:** A dropdown menu with "Application" selected.
- Select Application Version:** A dropdown menu with "Version" selected.
- Simulation output LFN dir:** An empty text input field.
- MCTruth output LFN dir:** An empty text input field.

LZ: Job submission interface





LZ Production Requests

monitoringd


Up

DIRAC

Unknown

Admins

■ Completed
 ■ Failed
 ■ Submitted

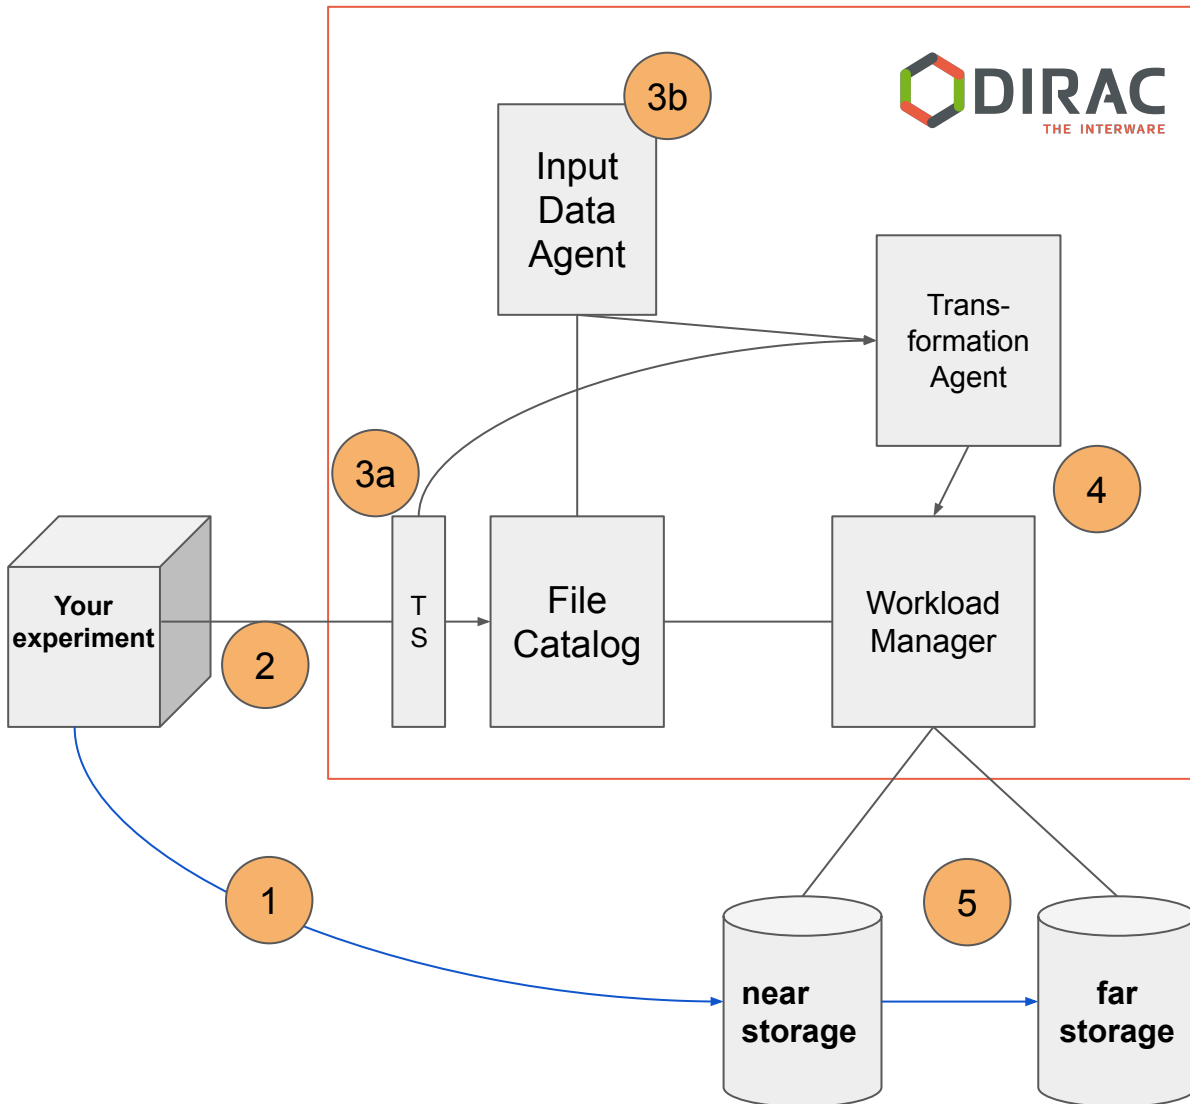


Show 10 entries

Search:

ID	Description	Sim Lead	Status	Request Date	Requester	
360	MDC2: background: job_config_1083000000....	EK	Completed	18/06/2018	elena korolkova	
Macro	NJobs	NEvents	Seed	Output	Status	Progress
MDC2Macros/lz_mdc2_background.mac	1	-1	1083003203	/lz/data/MDC2/background/BACCARAT-3.16.2_DER-7.4.7/20170830/DER/	Completed	<div style="width: 100%; background-color: green;"></div> 1
MDC2Macros/lz_mdc2_background.mac	1	-1	1083003418	/lz/data/MDC2/background/BACCARAT-3.16.2_DER-7.4.7/20170830/DER/	Completed	<div style="width: 100%; background-color: green;"></div> 1
MDC2Macros/lz_mdc2_background.mac	1	-1	1083003543	/lz/data/MDC2/background/BACCARAT-3.16.2_DER-7.4.7/20170830/DER/	Completed	<div style="width: 100%; background-color: green;"></div> 1
359	MDC2: background: job_config_1082400000....	EK	Completed	18/06/2018	elena korolkova	
358	MDC2: background: job_config_1082000000....	EK	Completed	18/06/2018	elena korolkova	
357	MDC2: background: job_config_1082500000....	EK	Completed	18/06/2018	elena korolkova	

Advanced DIRAC features: Workflow management



Workflow: Any raw data uploaded to the 'near' storage element is automatically processed and copied to the 'far' storage element. Must account for temporary unavailability of target storage.

(1) Raw data uploaded from DAQ to near storage

(2) File is registered in catalog.

(3) Transformation System detects new file:

(a) TS Client intercepts request
-or-

(b) InputDataAgent runs regular metadata query.

(Combination of the two is possible.)

(4) Transformation Agent applies rules and creates processing tasks.

(5) Workload manager submit jobs to process file.

Thanks for listening

- Questions ?
- Further reading:
 - DIRAC: <https://dirac.readthedocs.io/en/latest/>
 - Multi-VO DIRAC: Bauer D, Fayer S, 2017, *GridPP DIRAC: Supporting non-LHC VOs on LHC centric resources*, J. Phys.: Conf. Ser. 898 052003 ([Link](#))
 - [The LZ UK Data Centre](#) (CHEP 2018)

Backup slides

What the users see: “raw” data

The screenshot displays a web-based monitoring interface with a sidebar menu on the left and a main content area on the right. The sidebar menu includes categories like Desktops&Applications, Tools, Applications, and Help. The main content area is divided into a 'Selectors' panel and a 'Job Monitor' table.

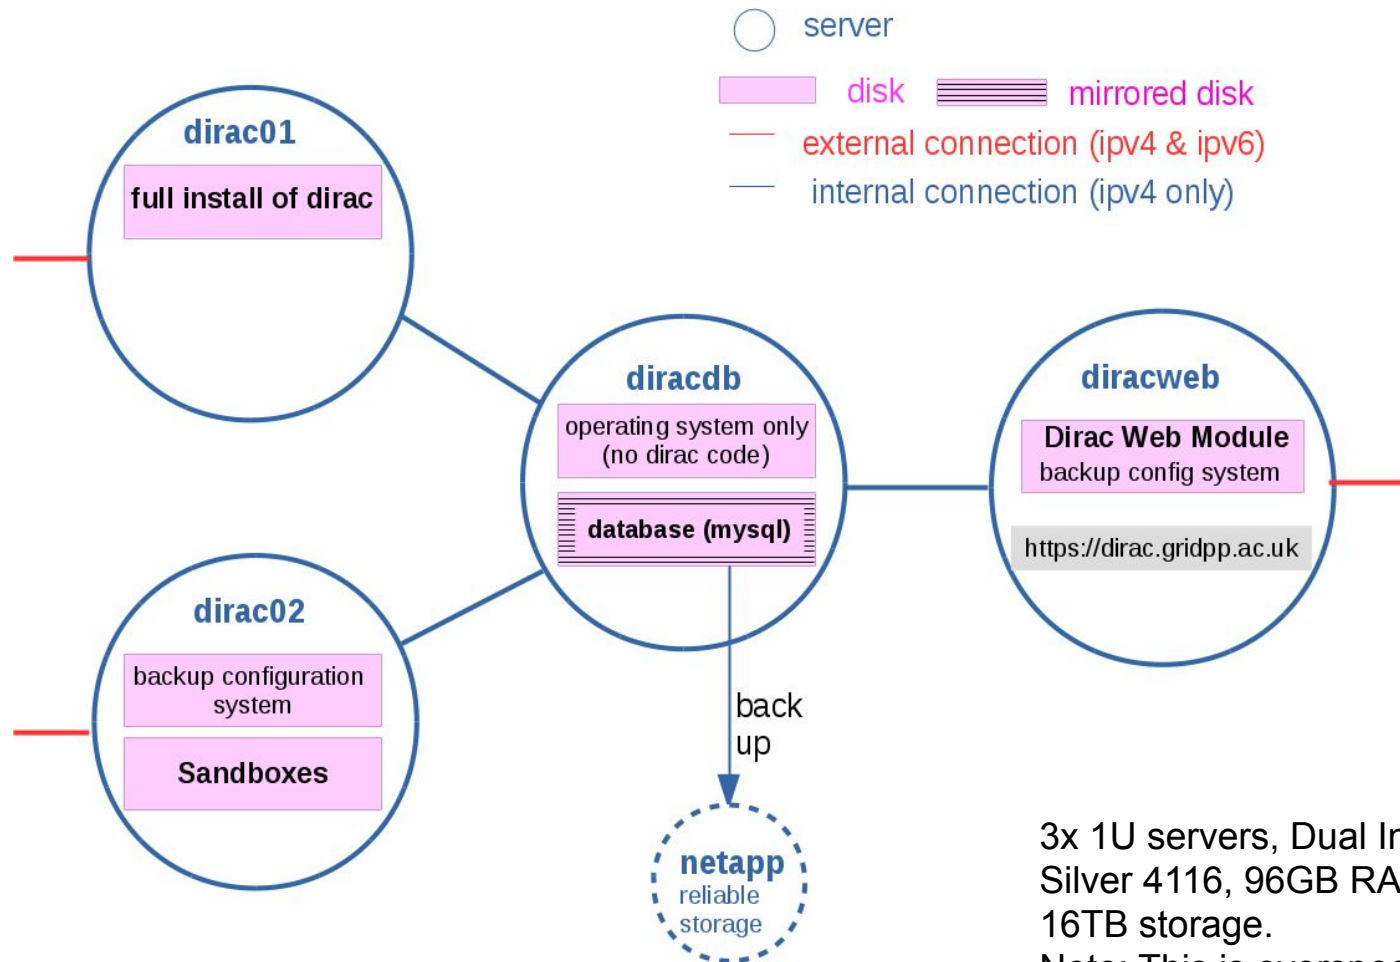
Selectors:

- Site: [Empty]
- Status: Running
- Minor Status: [Empty]
- Application Status: [Empty]
- Owner: [Empty]
- OwnerGroup: lz_production
- Job Group: [Empty]
- Job Type: [Empty]
- Time Span: [Empty]
- From: [Empty]
- To: [Empty]
- Reset Time Panel
- JobID(s): [Empty]

Job Monitor Table:

JobID	Status	MinorS	ApplicationStatus	Site	JobNa	LastUpdate[UTC]
13153971	Running	Ap...	Unknown	VAC.UKI-SOUT...	Iz_...	2018-10-03 12:48:51
13153954	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:42:40
13153944	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:42:31
13153943	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:44:01
13153927	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:43:04
13153924	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:41:14
13153905	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:37:45
13153893	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:39:10
13153880	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:45:07
13153879	Running	Ap...	Unknown	VAC.UKI-SOUT...	Iz_...	2018-10-03 12:48:50
13153874	Running	Ap...	Unknown	LCG.UKI-SOU...	Iz_...	2018-10-03 12:35:07
13153873	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:42:55
13153861	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:43:37
13153859	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:39:32
13153857	Running	Ap...	Unknown	LCG.UKI-SOU...	Iz_...	2018-10-03 12:32:51
13153853	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:43:55
13153851	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:44:51
13153842	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:43:29
13153841	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:31:20
13153839	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:43:51
13153836	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:42:14
13153832	Running	Ap...	Unknown	LCG.UKI-NOR...	Iz_...	2018-10-03 12:32:09
13153831	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:37:16
13153826	Running	Ap...	Unknown	VAC.UKI-SOUT...	Iz_...	2018-10-03 12:48:33
13153822	Running	Ap...	Unknown	VAC.UKI-SCOT...	Iz_...	2018-10-03 12:42:17

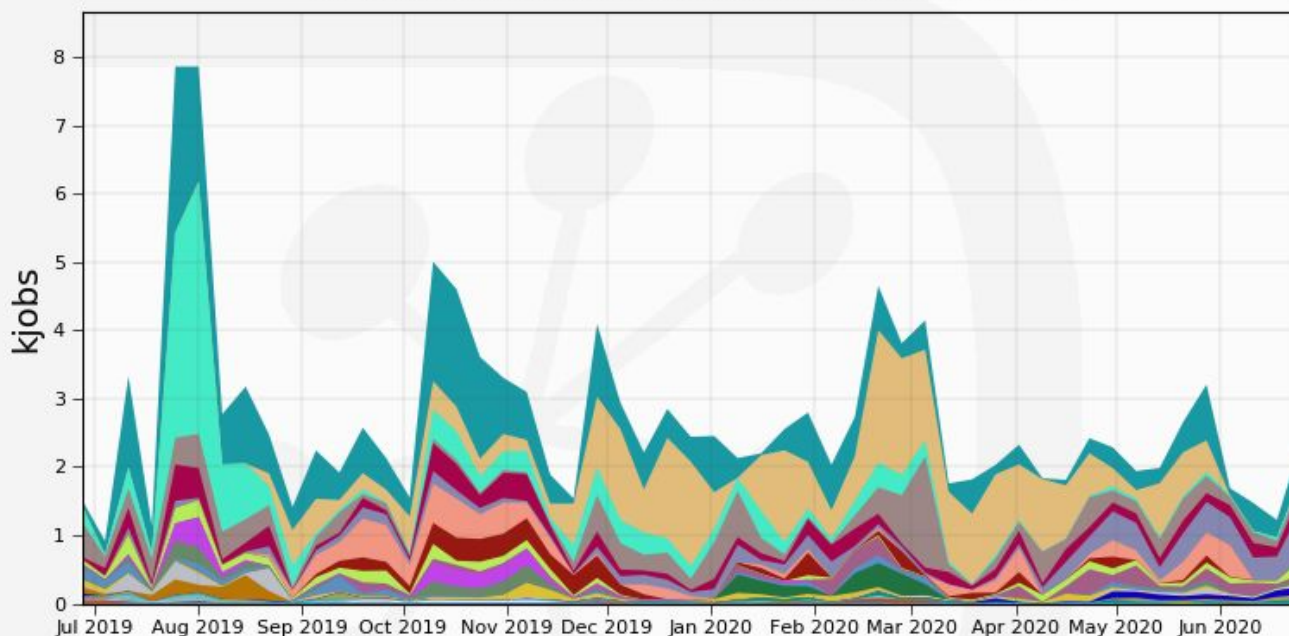
Hardware: The GridPP DIRAC server setup



3x 1U servers, Dual Intel Xeon Silver 4116, 96GB RAM each, 16TB storage.
Note: This is overspec'd.

Jobs by site

Running jobs by Site
52 Weeks from Week 25 of 2019 to Week 25 of 2020



Max: 7.87, Min: 0.91, Average: 2.69, Current: 2.32

■ LCG.UKI-LT2-IC-HEP.uk	21.4%	■ LCG.BelGrid-UCL.be	2.2%
■ LCG.CERN-NA62.ch	20.2%	■ VAC.UKI-NORTHGRID-MAN-HEP.uk	1.7%
■ LCG.UKI-NORTHGRID-MAN-HEP.uk	11.4%	■ LCG.UKI-SCOTGRID-ECDF.uk	1.7%
■ LCG.RAL-LCG2.uk	9.6%	■ LCG.UKI-LT2-QMUL.uk	1.4%
■ LCG.UKI-NORTHGRID-LIV-HEP.uk	6.0%	■ LCG.IN2P3-CC.fr	1.3%
■ LCG.UKI-SOUTHGRID-OX-HEP.uk	5.5%	■ VAC.UKI-SCOTGRID-GLASGOW.uk	1.3%
■ LCG.INFN-T1.it	5.2%	■ LCG.UKI-LT2-Brunel.uk	1.2%
■ CLOUD.UKI-LT2-IC-HEP-Iz.uk	3.3%	■ LCG.UKI-SCOTGRID-GLASGOW.uk	0.9%
■ LCG.UKI-SOUTHGRID-RALPP.uk	2.8%	... plus 21 more	

NA62 Workflow Management Interface

NA62 Simulations ³ Reconstruction Analysis Jobs ³³⁴ Files Transfers Stats Wiki Tools

Jobs - 324 running, 59 started, 4 waiting, 6 submitted

Show more | Default view | Expand

✓	Job number	Decay type	Number of events	NA62 s/w version	Submission Date/Time	Submitted by	Site	----- Archived Files -----					-- Grid job --		Output files*
								JDL	Mac	Exe	Out	Err	status	log	
<input type="checkbox"/>	1372047	31 ▼	2000	v22M1.0.1	03 Oct 13:16:49	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372046	31 ▼	2000	v22M1.0.1	03 Oct 13:16:38	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	
<input type="checkbox"/>	1372045	31 ▼	2000	v22M1.0.1	03 Oct 13:16:27	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372044	31 ▼	2000	v22M1.0.1	03 Oct 13:16:16	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372043	31 ▼	2000	v22M1.0.1	03 Oct 13:16:06	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372042	31 ▼	2000	v22M1.0.1	03 Oct 13:15:54	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	
<input type="checkbox"/>	1372041	31 ▼	2000	v22M1.0.1	03 Oct 13:15:43	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	
<input type="checkbox"/>	1372040	31 ▼	2000	v22M1.0.1	03 Oct 13:15:32	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372039	31 ▼	2000	v22M1.0.1	03 Oct 13:15:22	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372038	31 ▼	2000	v22M1.0.1	03 Oct 13:15:11	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372037	31 ▼	2000	v22M1.0.1	03 Oct 13:14:55	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	
<input type="checkbox"/>	1372036	31 ▼	2000	v22M1.0.1	03 Oct 13:14:45	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372035	31 ▼	2000	v22M1.0.1	03 Oct 13:14:35	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372034	31 ▼	2000	v22M1.0.1	03 Oct 13:14:24	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372033	31 ▼	2000	v22M1.0.1	03 Oct 13:14:13	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372032	31 ▼	2000	v22M1.0.1	03 Oct 13:14:01	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372031	31 ▼	2000	v22M1.0.1	03 Oct 13:13:49	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	
<input type="checkbox"/>	1372030	31 ▼	2000	v22M1.0.1	03 Oct 13:13:38	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372029	31 ▼	2000	v22M1.0.1	03 Oct 13:13:27	protopop	CERN	🔗	🔗	🔗			STARTED	-	
<input type="checkbox"/>	1372028	31 ▼	2000	v22M1.0.1	03 Oct 13:13:16	protopop	CERN	🔗	🔗	🔗			SUBMITTED	-	

Total: 20 runs, 40000 events.

1 2 3 4 5 6 7 8 9 ... >

DIRAC Speak

From the DIRAC Overview page:

<https://dirac.readthedocs.io/en/stable/AdministratorGuide/SetupStructure/index.html>

Databases:

Keep the persistent state of a *System*. They are accessed by *Services* and *Agents* as a kind of shared memory. The GridPP DIRAC instance currently has 15 *Databases* and counting.

Services:

Are passive components listening to incoming client requests and reacting accordingly by serving requested information from the *Database* backend or inserting requests on the *Database* backend. *Services* themselves can be clients of other *Services* from the same DIRAC *System* or from other *Systems*.

Agents:

Are the active components which are running continuously invoking periodically their execution methods. *Agents* are animating the whole system by executing actions, sending requests to the DIRAC or third party services.

System:

Is delivering a complex functionality to the rest of DIRAC, providing a solution for a given class of tasks. Examples of *Systems* are Workload Management System or Configuration System.