High Energy Physics Data Popularity: ATLAS Datasets Popularity Case Study

Eugeny Tretyakov, Maria Grigorieva, , Alexei Klimentov, Dmitry Golubkov, Tatiana Korchuganova, Aleksandr Alekseev, Alexey Artamonov and Timofei Galkin

Ivannikov memorial workshop, Orel, September 25-26, 2020

Outline

- **Sources of initial information**
  - PanDA database
  - CRIC
  - CERN Phonebook directory
- **Data model**
- **ETL workflow**
- **Filtered data sample**
- **Data popularity**
  - Plots
  - Heatmaps
  - Cartography
- **Summary and Plans**

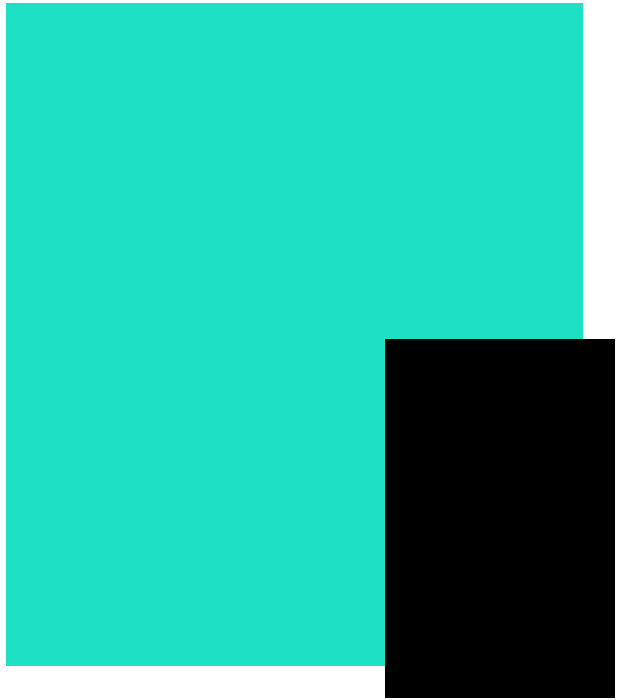

ATLAS Experiment at the LHC

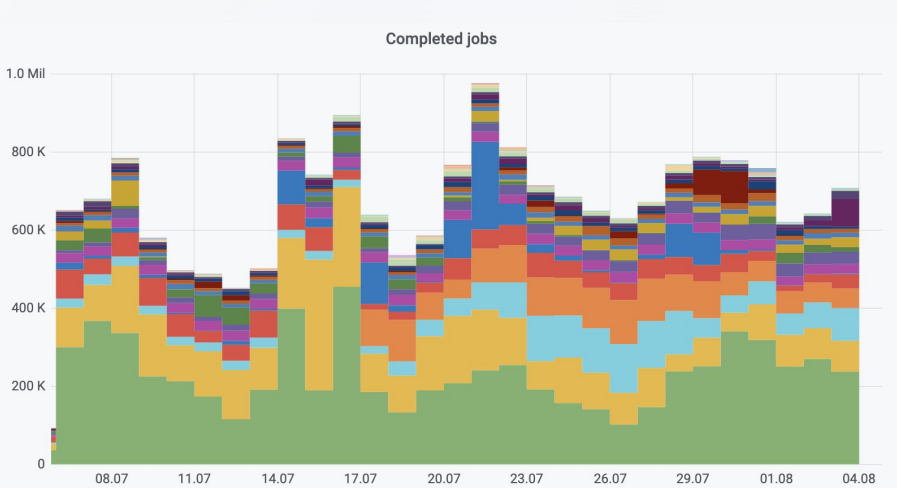

	min	max	avg	total
User Analysis	35.9 K	453.9 K	229.9 K	7.1261 Mil
MC 16 evgen	16.3 K	334.4 K	114.5 K	3.5493 Mil
MC 16 simul	3.7 K	125.0 K	54.1 K	1.6785 Mil
Reprocessing default	0	112.6 K	50.2 K	1.5567 Mil
MC Derivations	14.8 K	74.3 K	42.6 K	1.3220 Mil
Group Analysis	818	224.1 K	29.8 K	924.0 K
Test	4.1 K	27.0 K	24.2 K	751.6 K

ATLAS utilizes distributed resources involving about 160 computing centers spread around the world for analysis, data simulation and processing

Millions of computing jobs! are executed monthly within the distributed computing infrastructure

Types of data processing in ATLAS:

- ***Central production*** workflows *(transformation of the RAW detector/Monte-Carlo data to the formats applicable for physics analysis)* are planned in advance and organized by the experiment and physics groups
- ***User analysis*** (users can use the distributed infrastructure of ATLAS for data analysis) are more chaotic because bursts of task submissions might appear and there is also a correlation with major scientific events such as the main international conferences

Popularity of ATLAS Datasets

- ***ATLAS Central Production*** generates vast volumes of data files prepared for physics analysis
- ***Dataset*** = collection of files processed using the same version of software
- ***Dataset Popularity*** is the metric showing how often a dataset was used as an input for physics analysis within a certain time interval
  - Popularity can be estimated for a particular dataset or/and for a group of datasets

- The popularity of data among individual physicists and University groups has become one of the ***key factors of efficient data management and processing***

The popularity measurements can be divided into four metrics:

1. ***Tasks*** - the number of user analysis tasks that used datasets as an input (the number of accesses)
2. ***Users*** - the number of users which executed analysis tasks
3. ***Home institutes*** - the number of home institutes of the users
4. ***Number of countries*** - the number of home institutes countries

This particular combination of measurements allows to evaluate the popularity more accurate than just the number of accesses.

Data Sources

- ***Workload Management System (WMS) PanDA*** responsible for scheduling jobs on the grid based on available resources
  - PanDA Database provides the detailed information about datasets used in analysis tasks, and about these tasks (timings; execution process; user who executed the task)
- ***CRIC*** (Computing Resource Information Catalogue) provides APIs to obtain information about PanDA queues, ATLAS sites, storage systems, services and so on
- ***CERN Phonebook Directory*** provides anonymous read-access (from inside CERN only) to the user information in Active Directory

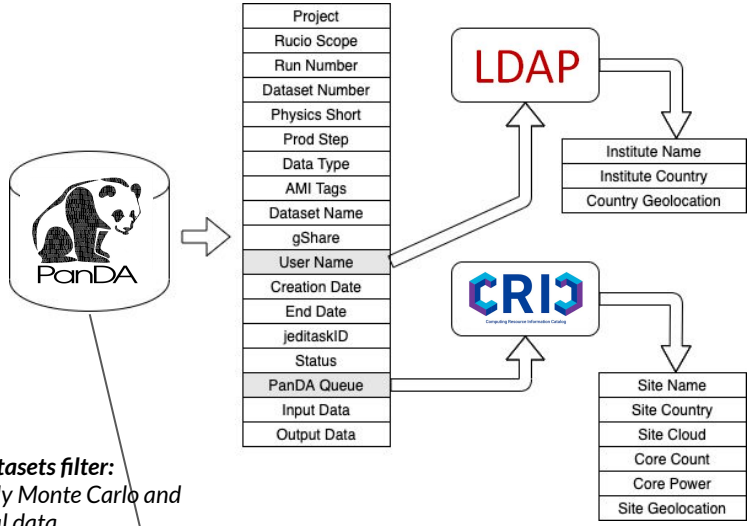

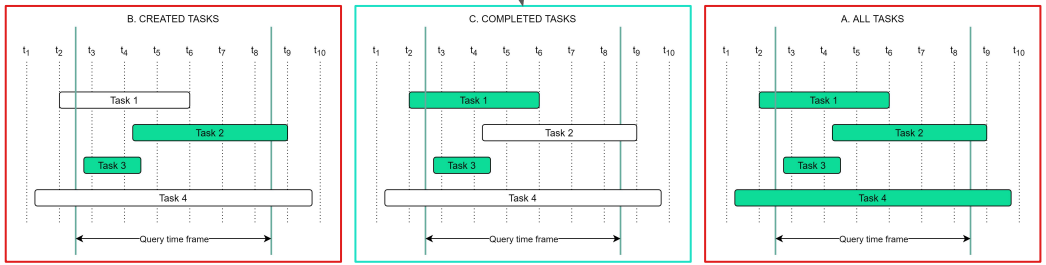

Dataset Naming Nomenclature:

Monte Carlo Datasets

mcNN_subProject.datasetNumber.physicsShort.prodStep.dataType.Version

mc15_13TeV.300402.Pythia8B_A14_CTEQ6L1_Bs_mu3p5mu3p5.evgen.EVNT.e4397

Real Data

DataNN_subProject.runNumber.streamName.prodStep.dataType.Version

data15_13TeV.00284484.physics_Main.merge.AOD.f644_m1518

Data model

Fields by sources

PanDA database	CRIC	LDAP
Dataset name	ATLAS site name	User home institute
Name of the job flow-group	ATLAS site geolocation	Home institute country
User name	Cloud name	Home institute country geolocation
ID of user analysis task	Number of cores at site	
Creation date	Corepower at site	
End date	Tier level	
Status of task	Resource type	
Type of resource		
Name of PanDA queue		
Size of input data		
Size of output data		

ETL workflow

Extract, transform, load

The ETL Workflow is described with Directed Acyclic Graph (DAG), where each node is a run unit.

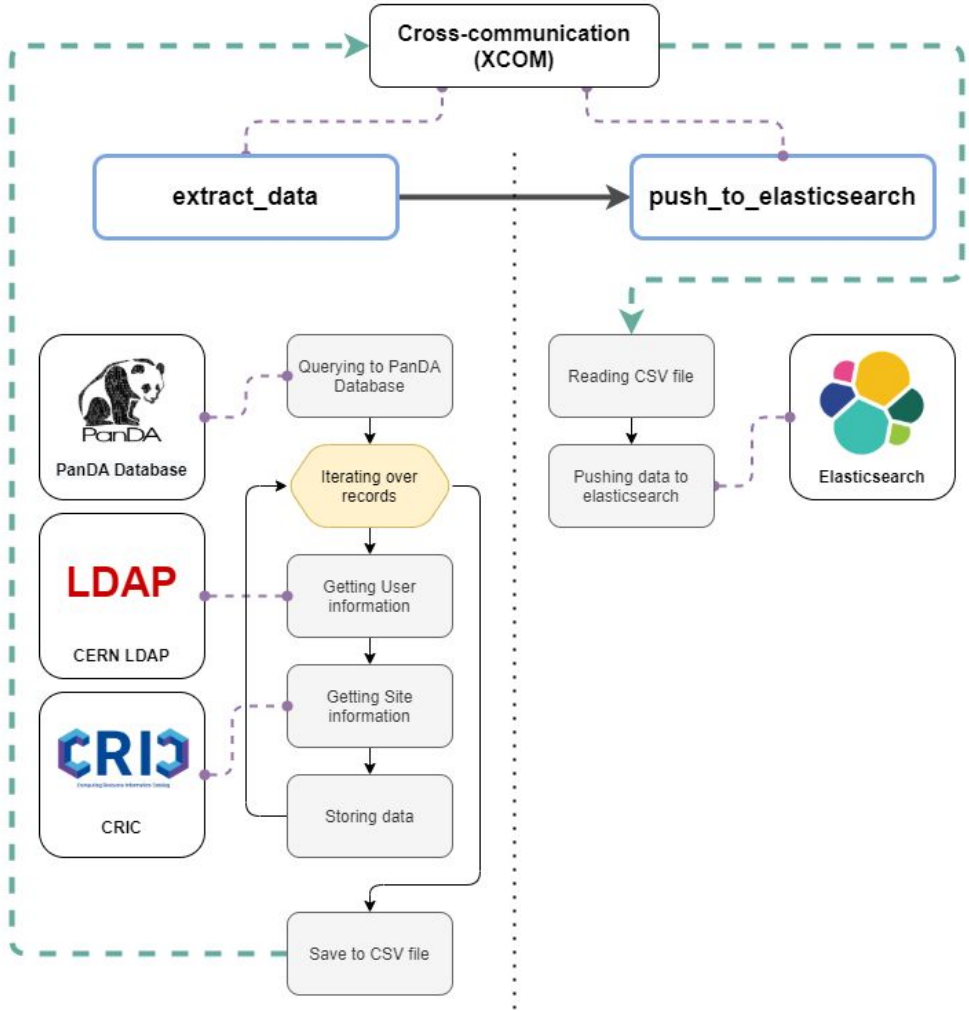

Selected Data Sample

Analysis tasks statistics

Time frame (by ENDTIME field):	**from 2020-07-02 to 2020-08-15**
Datasets:	**119,803**
JEDI tasks:	**138,176**
Users:	**622**
Home Institutes/Countries:	**155 /34**
Production Sites /Countries	**154/23**

Selected Data Sample. Sites

Top 10 sites where users tasks were executed

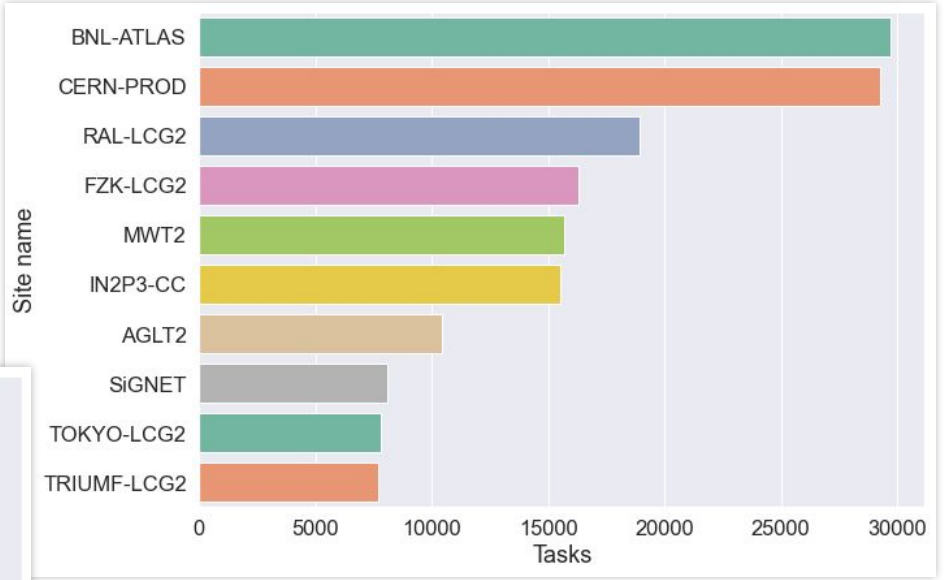

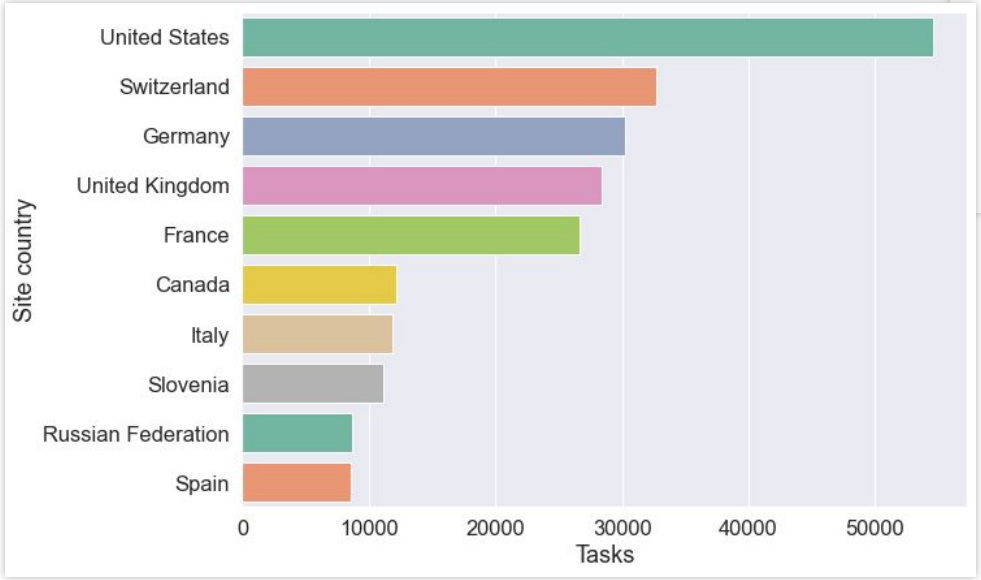

Selected Data Sample. Users Home Institutes

Top 10 Home Institutes for users tasks

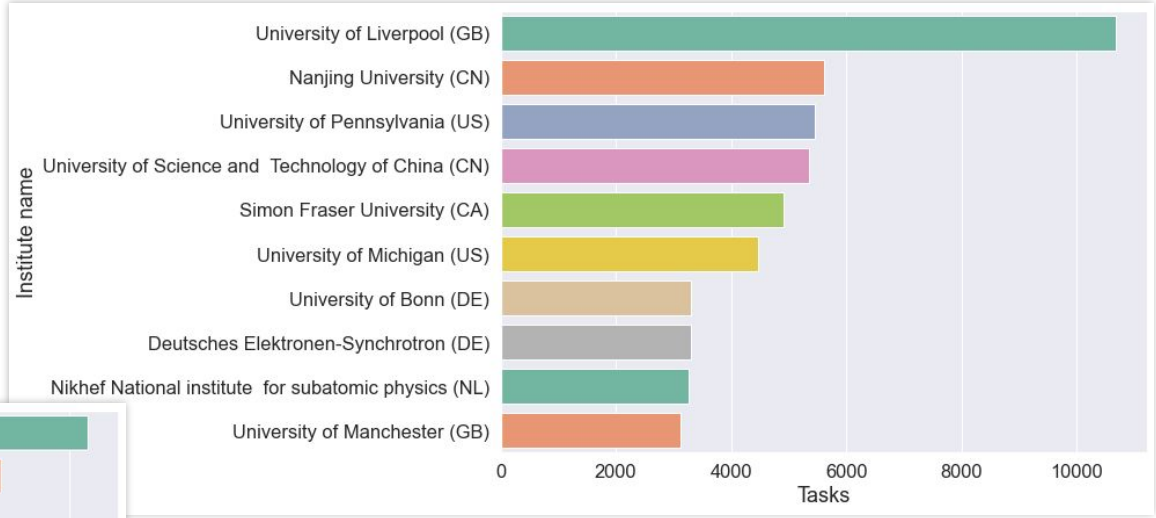

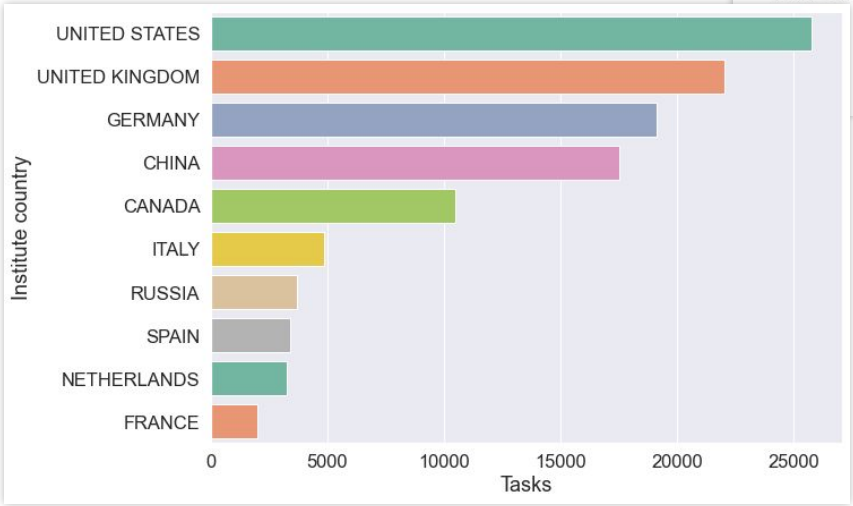

Data popularity (number of accesses by JEDI Tasks)

Data and Monte-Carlo datasets

Data

Datasets: 71,803
Tasks: 100,409

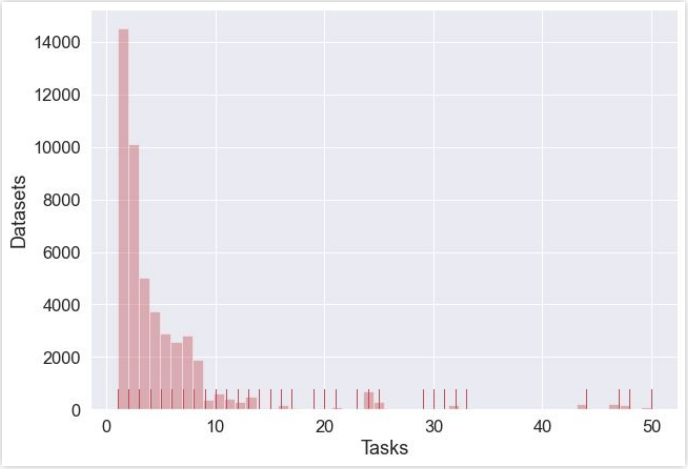

Monte-Carlo

Datasets: 48,000
Tasks: 37,767

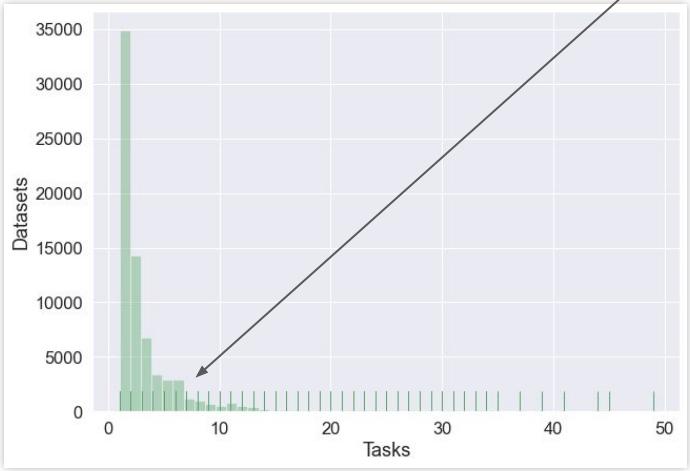

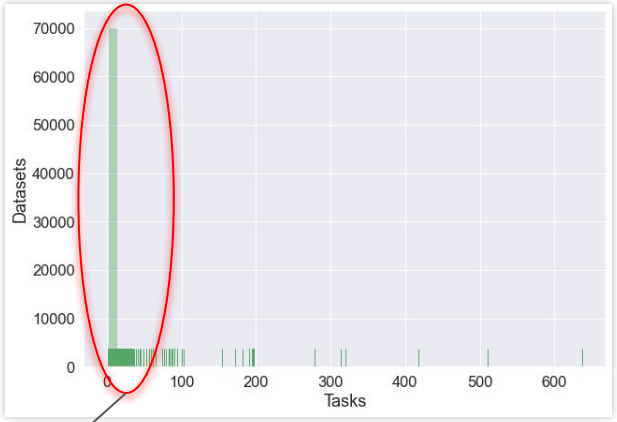

Data popularity (number of accesses by users)

Data and Monte-Carlo datasets

Data

Datasets: 71,803
Users: 289

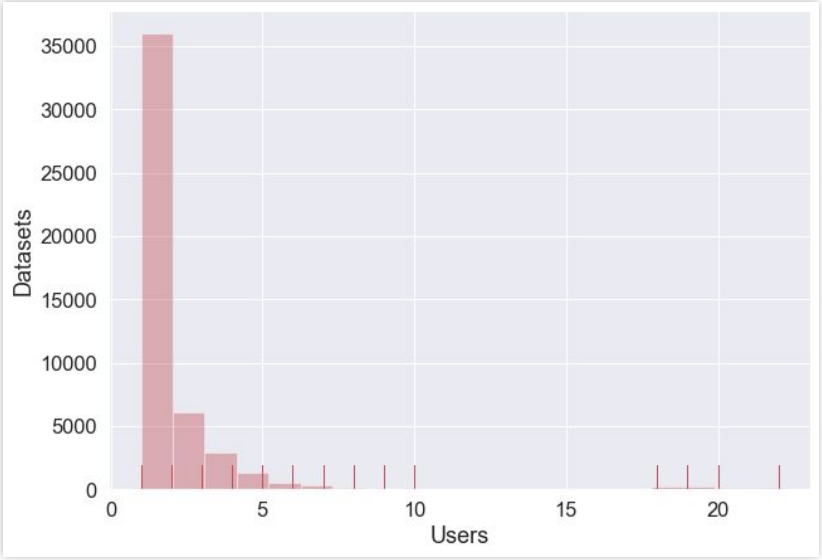

MONTE CARLO

Datasets: 48,000
Users: 558

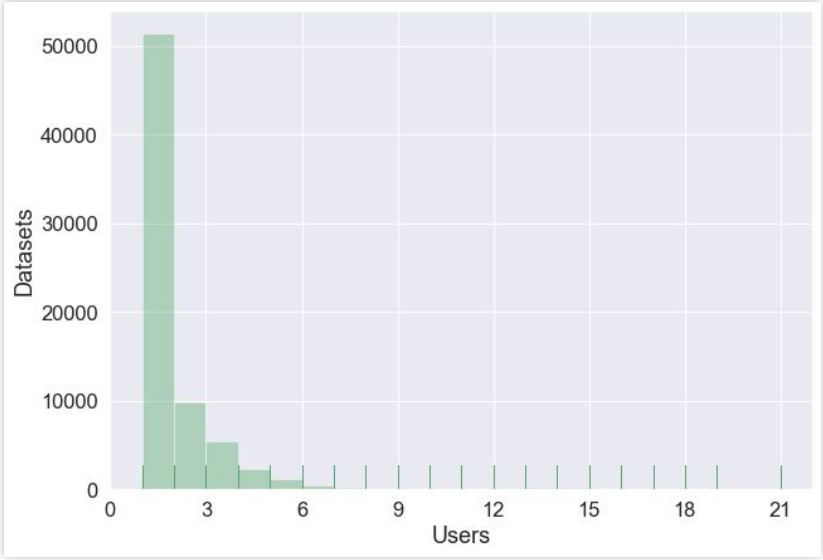

Projects and data types popularity (for analysis tasks)

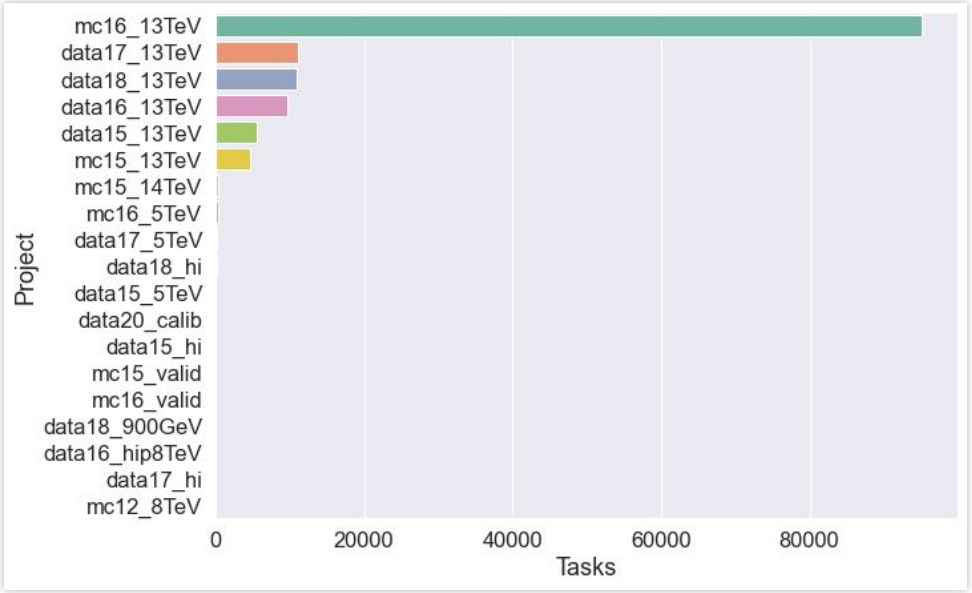

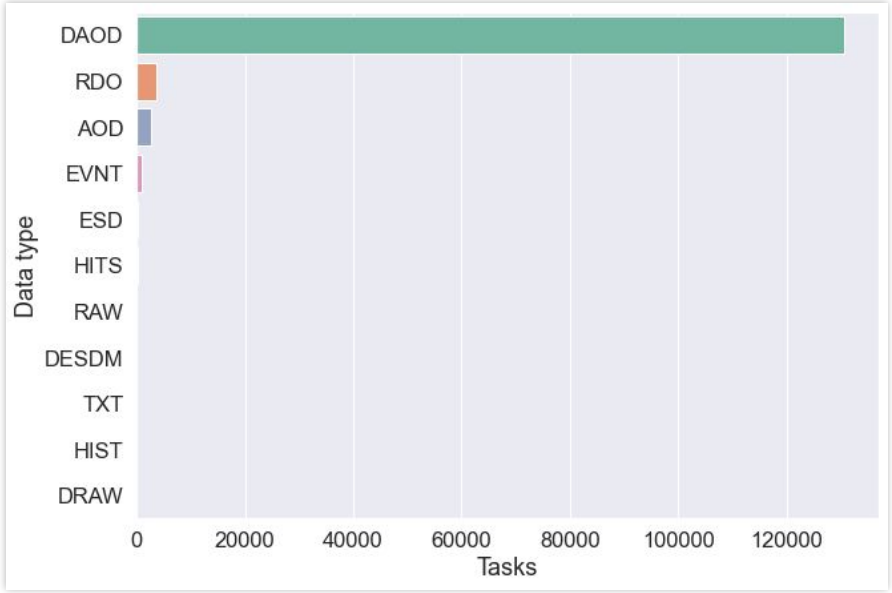

Users Home Institutes Countries

Heatmap plot

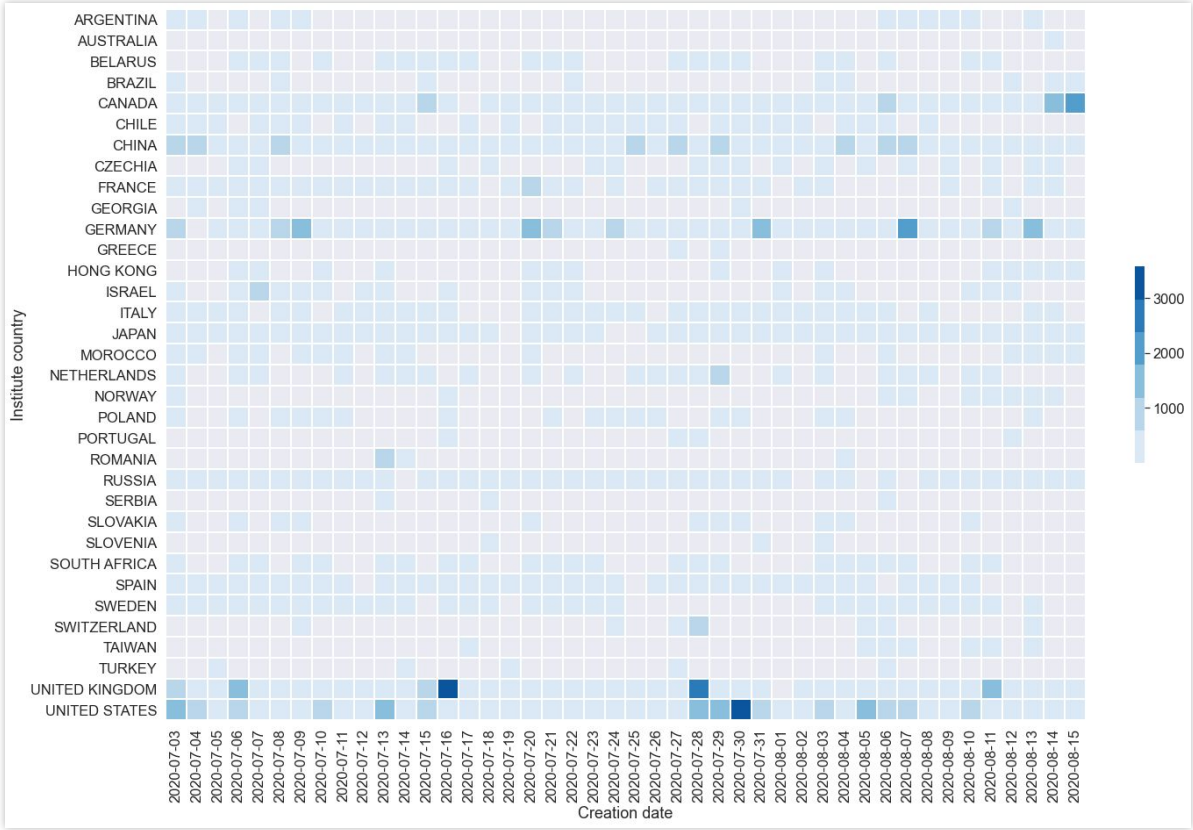

Analysis Tasks Execution Sites Countries

Heatmap plot

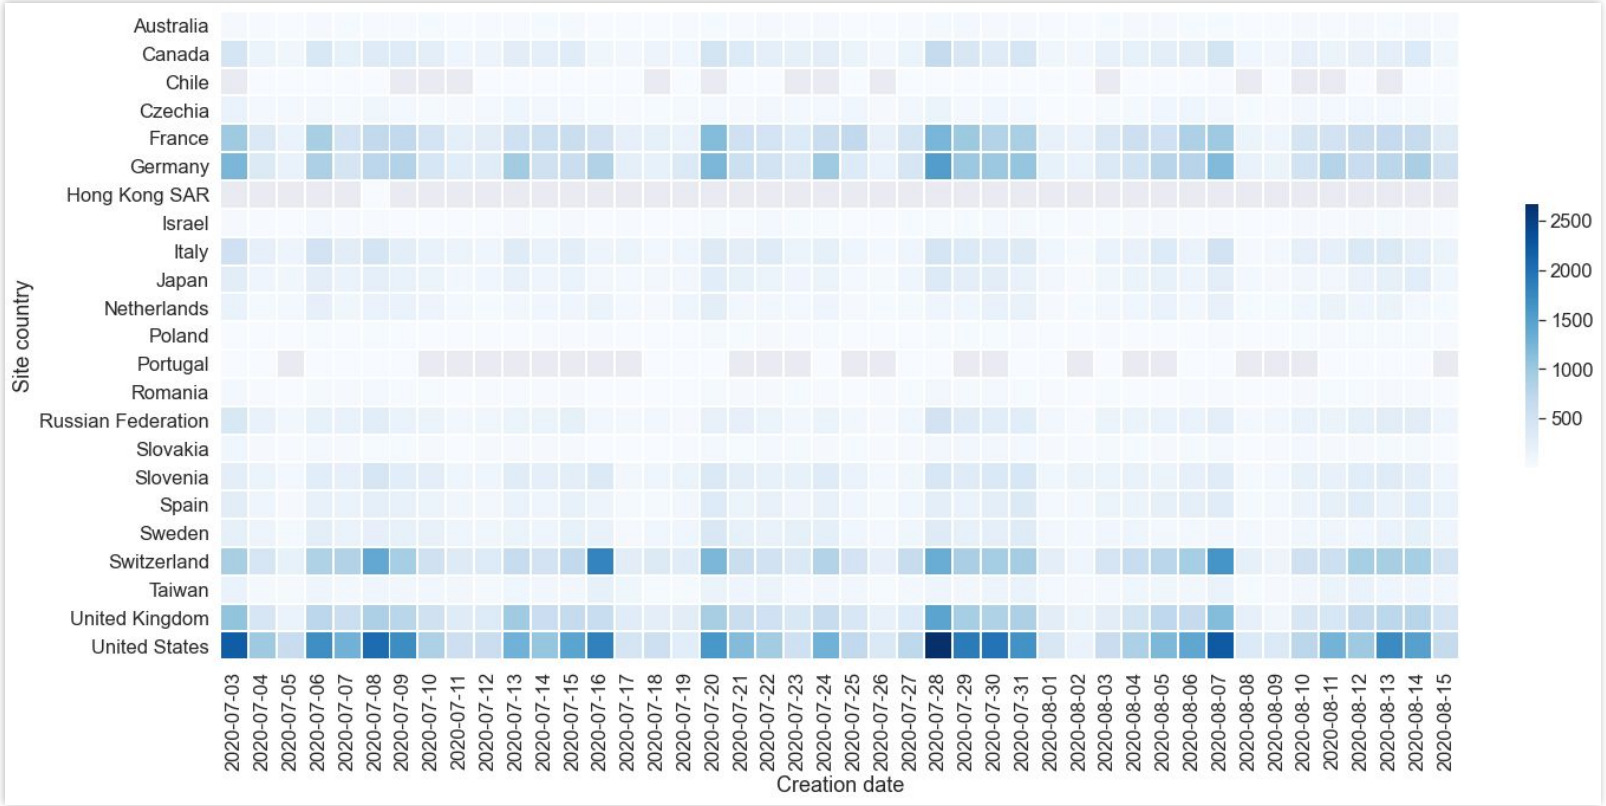

Data popularity. Users Tasks (Users from ATLAS Home Institutes)

Download map by link:
https://cernbox.cern.ch/index.php/s/LPqMqykakKnp504

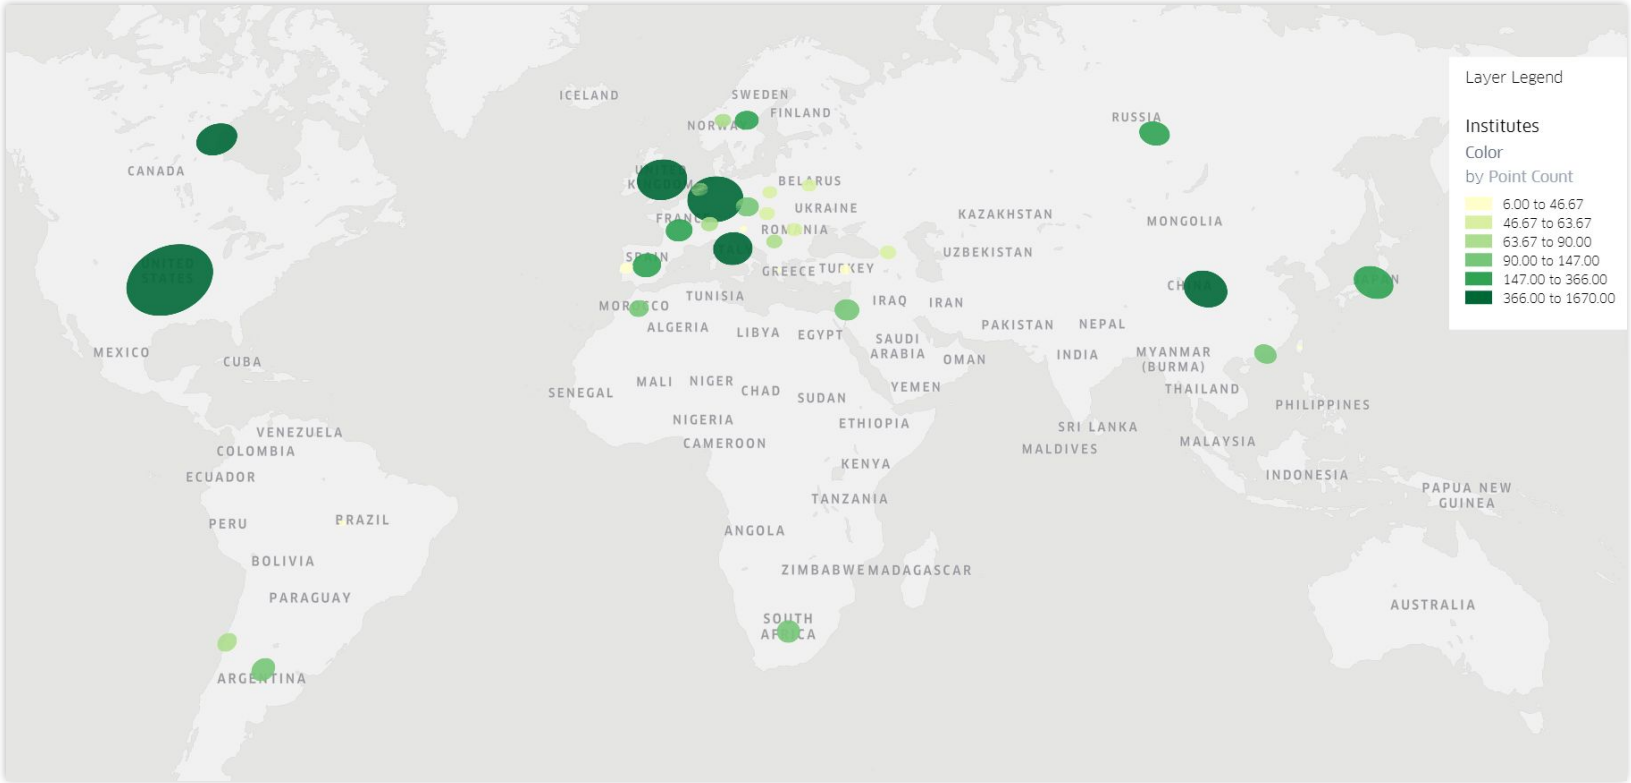

Data popularity. Users Tasks (Tasks execution sites)

Download map by link:
https://cernbox.cern.ch/index.php/s/LPqMqykakKnp504

All sites are in use by all users

Download map by link:
https://cernbox.cern.ch/index.php/s/LPqMqykakKnp504

All sites are in use by all users.
NRNU MEPhI ATLAS Collaborators

Download map by link:
https://cernbox.cern.ch/index.php/s/LPqMqykakKnp504

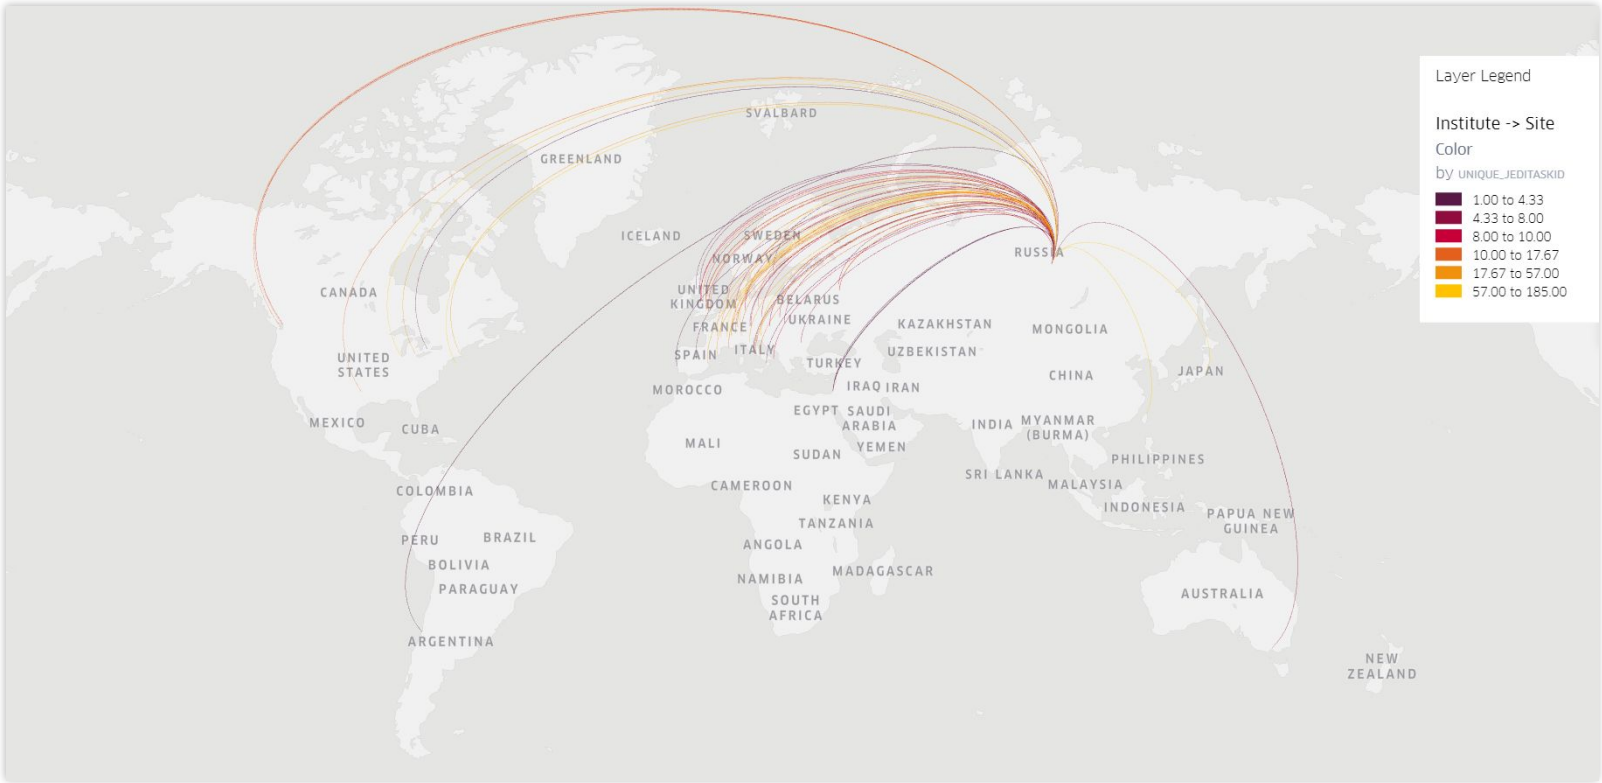

Summary and Plans

Results

Explored Data sources and developed Data model.

The ETL workflow was developed and brought in production. Aggregated data in CSV file format reached 3.3 Gb.

Explored Data popularity metrics. Suggested visualization techniques.

Explored and tested visualization engines for 3D maps.

Geolocations

Sites

Not all geolocations of computing sites were found using CRIC API. Among 90K of analysis tasks we did not find geolocations for about 2K of them.

Users

Sometimes user specify not standard username in task configuration (for example, user may add email address to the string with username). In this case now we don't find the information about this user in CERN Directory. But this issue can be solved partially by removing unnecessary emails from username parameter. Currently we didn't find users locations for 8K tasks.

Datasources

At the next stage we will add DDM data source(s) (Rucio transfers to analyze ATLAS datasets downloaded by users to their local resources. The number of such downloads will become one more metric of datasets popularity.) and PanDA/Rucio metrics correlation

Software

We plan to design web-application aimed at the visual analysis of data popularity, with the appliance of advanced interactive visualization techniques like parallel coordinates, world maps, heatmaps and treemaps.

Data storage

We have exhausted currently used hardware. A new request for new storage nodes is submitted

Thank you for your attention

This work is supported by the Russian Science Foundation grant №19-71-30008,
Moscow Center of Fundamental and Applied Mathematics

Backup slides

Data sources I

PanDA database

Filters

ATLAS data and
Monte-Carlo datasets

Completed tasks

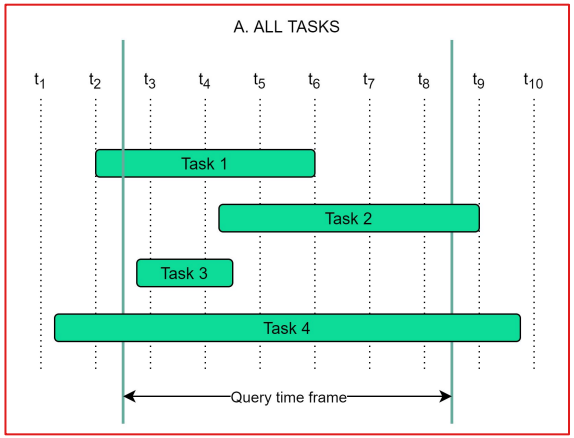

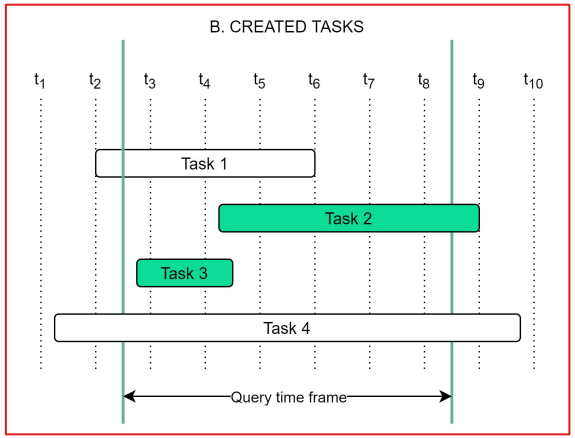

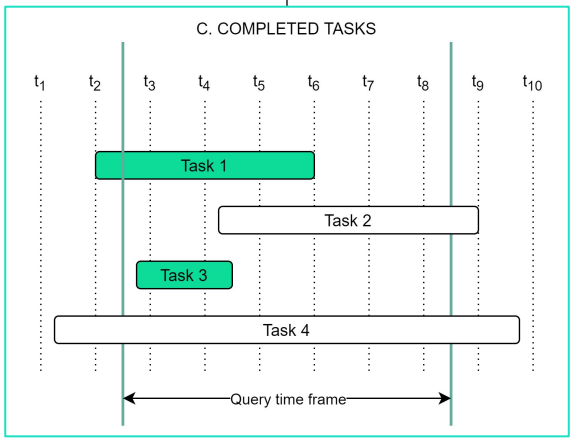

Data sources II

- Initially AGIS was used to get information about Computing sites but now we have migrated to CRIC.

- CRIC provides APIs to obtain information about PanDA Queues, ATLAS Sites, storage, services, etc.

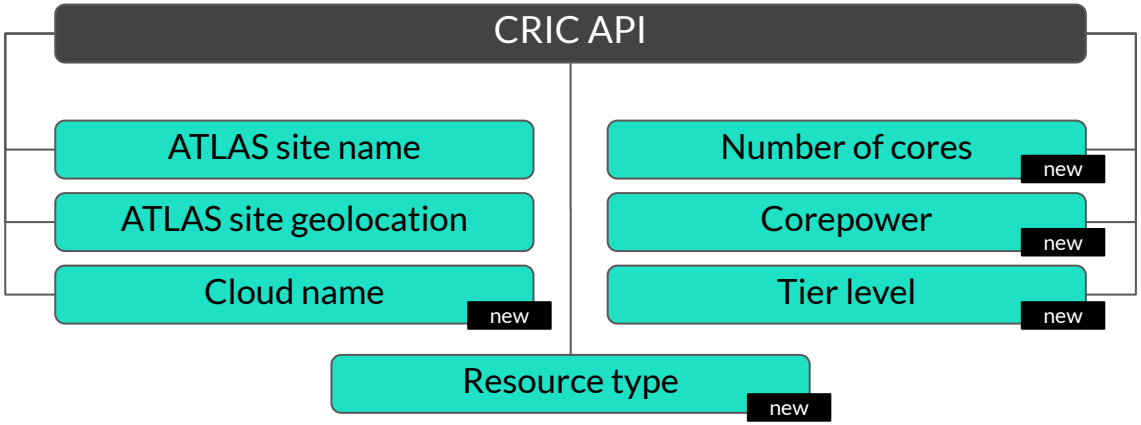

new - Fields added after migration from AGIS to CRIC

CRIC – computing resources information catalog

Data sources III

LDAP

The LDAP service allows anonymous read-access (from inside CERN only) to the user information in Active Directory.

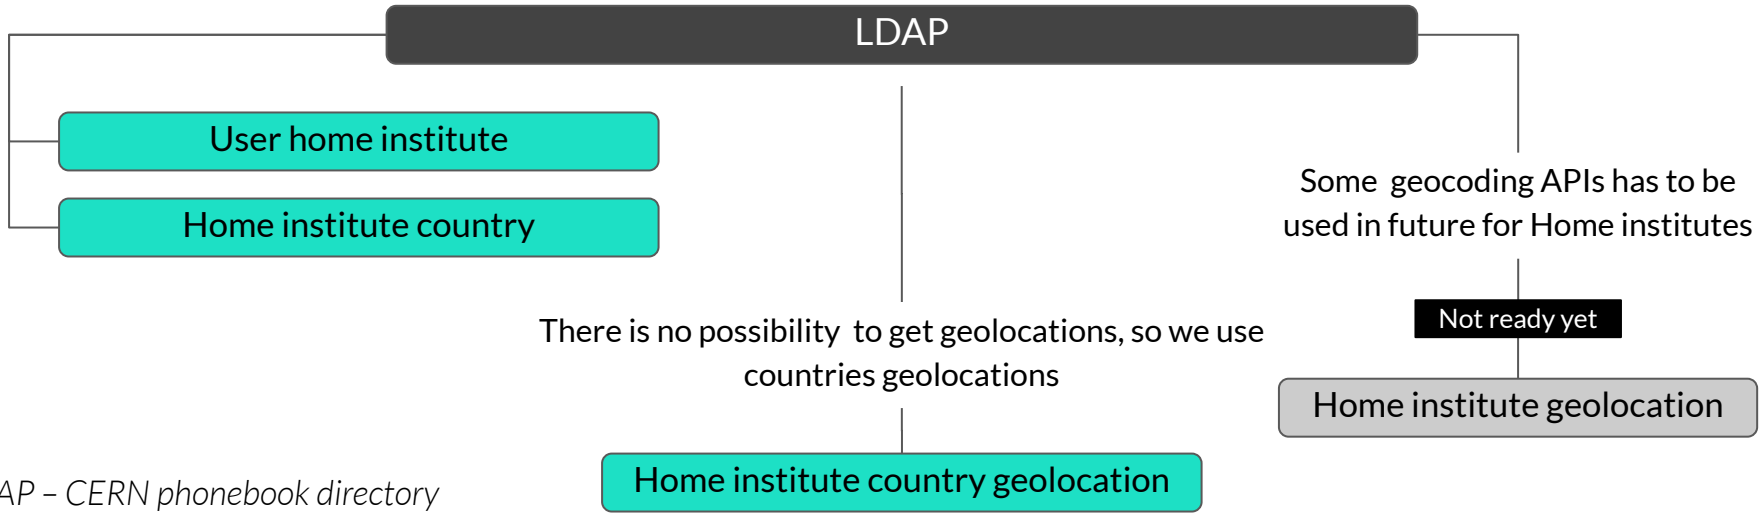

Data popularity (JEDI Tasks)

Monte-Carlo datasets left behind

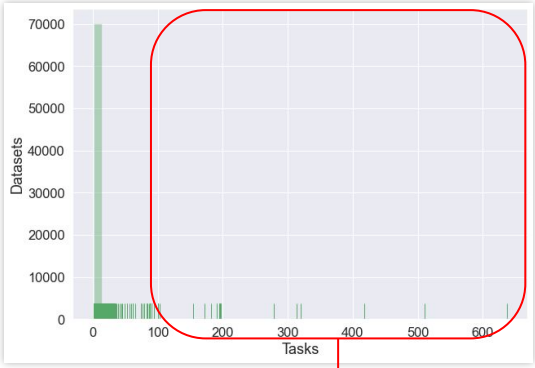

DATASETNAME	
mc15_13TeV:mc15_13TeV.410000.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_nonallhad.recon.RDO.e3698_s2608_s2183_r7193_tid06752771_00	637
mc15_13TeV:mc15_13TeV.410000.PowhegPythiaEvtGen_P2012_ttbar_hdamp172p5_nonallhad.recon.RDO.e3698_s2608_s2183_r7195_tid06752780_00	510
mc15_13TeV:mc15_13TeV.361107.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zmumu.recon.RDO.e3601_s2576_s2132_r7143_tid06718162_00	418
mc15_13TeV:mc15_13TeV.361106.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zee.recon.RDO.e3601_s2665_s2183_r7191_tid06753544_00	320
mc16_13TeV:mc16_13TeV.361108.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Ztautau.recon.RDO.e3601_s3126_r9546_tid11373415_00	314
mc15_13TeV:mc15_13TeV.423202.Pythia8B_A14_CTEQ6L1_Jpsie3e13.merge.AOD.e3869_s2608_s2183_r6630_r6264_tid05536542_00	278
mc16_13TeV:mc16_13TeV.364248.Sherpa_222_NNPDF30NNLO_ZZZ_4l2v_EW6.deriv.DAOD_HIGG2D1.e5887_e5984_s3126_r10724_r10726_p3870_tid18151531_00	197
mc16_13TeV:mc16_13TeV.364243.Sherpa_222_NNPDF30NNLO_WWZ_4l2v_EW6.deriv.DAOD_HIGG2D1.e5887_e5984_s3126_r10724_r10726_p3870_tid18151553_00	197
mc16_13TeV:mc16_13TeV.364283.Sherpa_222_NNPDF30NNLO_lllljj_EW6.deriv.DAOD_HIGG2D1.e6055_e5984_s3126_r10201_r10210_p3870_tid18270588_00	196
mc16_13TeV:mc16_13TeV.364247.Sherpa_222_NNPDF30NNLO_ZZZ_6l0v_EW6.deriv.DAOD_HIGG2D1.e5887_s3126_r9364_r9315_p3870_tid18151979_00	196
mc16_13TeV:mc16_13TeV.364247.Sherpa_222_NNPDF30NNLO_ZZZ_6l0v_EW6.deriv.DAOD_HIGG2D1.e5887_e5984_s3126_r10201_r10210_p3870_tid18151789_00	196
mc16_13TeV:mc16_13TeV.364245.Sherpa_222_NNPDF30NNLO_WZZ_5l1v_EW6.deriv.DAOD_HIGG2D1.e5887_s3126_r9364_r9315_p3870_tid18152076_00	196
mc16_13TeV:mc16_13TeV.364245.Sherpa_222_NNPDF30NNLO_WZZ_5l1v_EW6.deriv.DAOD_HIGG2D1.e5887_e5984_s3126_r10201_r10210_p3870_tid18151742_00	196
mc16_13TeV:mc16_13TeV.364243.Sherpa_222_NNPDF30NNLO_WWZ_4l2v_EW6.deriv.DAOD_HIGG2D1.e5887_s3126_r9364_r9315_p3870_tid18151933_00	196
mc16_13TeV:mc16_13TeV.364243.Sherpa_222_NNPDF30NNLO_WWZ_4l2v_EW6.deriv.DAOD_HIGG2D1.e5887_e5984_s3126_r10201_r10210_p3870_tid18151756_00	196
mc16_13TeV:mc16_13TeV.364248.Sherpa_222_NNPDF30NNLO_ZZZ_4l2v_EW6.deriv.DAOD_HIGG2D1.e5887_s3126_r9364_r9315_p3870_tid18152097_00	196
mc16_13TeV:mc16_13TeV.364250.Sherpa_222_NNPDF30NNLO_llll.deriv.DAOD_HIGG2D1.e5894_s3126_r9364_r9315_p3872_tid18151883_00	195
mc16_13TeV:mc16_13TeV.345706.Sherpa_222_NNPDF30NNLO_gglllll_130M4l.deriv.DAOD_HIGG2D1.e6213_e5984_s3126_r10724_r10726_p3872_tid18269926_00	194
mc16_13TeV:mc16_13TeV.364250.Sherpa_222_NNPDF30NNLO_llll.deriv.DAOD_HIGG2D1.e5894_e5984_s3126_r10724_r10726_p3872_tid18151559_00	191
mc16_13TeV:mc16_13TeV.364250.Sherpa_222_NNPDF30NNLO_llll.deriv.DAOD_HIGG2D1.e5894_e5984_s3126_r10724_r10726_p3872_tid18151550_00	191
mc16_13TeV:mc16_13TeV.345706.Sherpa_222_NNPDF30NNLO_gglllll_130M4l.deriv.DAOD_HIGG2D1.e6213_e5984_s3126_r10201_r10210_p3872_tid18269798_00	182
mc16_13TeV:mc16_13TeV.364283.Sherpa_222_NNPDF30NNLO_lllljj_EW6.deriv.DAOD_HIGG2D1.e6055_e5984_s3126_r9364_r9315_p3870_tid18270295_00	172
mc15_13TeV:mc15_13TeV.361036.Pythia8_A2MSTW2008LO_minbias_inelastic.recon.RDO.e3580_s2726_r7011_tid07344311_00	154
mc16_13TeV:mc16_13TeV.364283.Sherpa_222_NNPDF30NNLO_lllljj_EW6.deriv.DAOD_HIGG2D1.e6055_e5984_s3126_r10724_r10726_p3870_tid18270920_00	103
mc16_13TeV:mc16_13TeV.361239.Pythia8EvtGen_A3NNPDF23LO_minbias_inelastic_high.simul.HITS.e4981_s3087_s3111_tid10701335_00	100
mc16_13TeV:mc16_13TeV.361238.Pythia8EvtGen_A3NNPDF23LO_minbias_inelastic_low.simul.HITS.e4981_s3087_s3111_tid10701323_00	100
mc16_13TeV:mc16_13TeV.423001.ParticleGun_single_photon_egammaET.recon.RDO.e3566_s3113_r10470_tid14016886_00	100
mc16_13TeV:mc16_13TeV.423000.ParticleGun_single_electron_egammaET.recon.RDO.e3566_s3113_r10470_tid14016860_00	100
mc16_13TeV:mc16_13TeV.361106.PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zee.deriv.DAOD_JETM3.e3601_e5984_s3126_r10724_r10726_p3978_tid19330142_00	100
mc16_13TeV:mc16_13TeV.300012.Pythia8EvtGen_A14NNPDF23_W_JpsiDs.merge.AOD.e7194_e5984_a875_r9364_r9315_tid16767672_00	94