The ATLAS Event Picking Service and its evolution

E.I. Alexandrov[1,a], I.N. Alexandrov[1], D. Barberis[2], L. Canali[3], E. Cherepanova[4], E.J. Gallas[5], S. Gonzalez de la Hoz[6], F.V. Prokoshin[1], G. Rybkin[7], J. Salt Cairols[6], J. Sanchez[6], M. Villaplana Perez[6], A.V. Yakovlev[1]

[1]Joint Institute for Nuclear Research (Russia)
[2]University and INFN Genova (Italy)
[3]CERN (Switzerland)
[4]Nikhef National Institute for Subatomic Physics and University of Amsterdam (Netherlands)
[5]University of Oxford (UK)
[6]Instituto de Física Corpuscular IFIC (Spain)
[7]IJCLab (France)

GRID Conference at JINR, 04 July 2023

The ATLAS EventIndex

- The EventIndex is the global catalogue of all ATLAS events
- For each event, each data format and each processing version, it contains:
  - Event identifiers (run and event number)
  - Location (GUID of the file containing it) and provenance
  - Trigger and other useful metadata
- Main use case is event picking for detailed analysis and/or displays
  - Also production checks and overlap counts

EventIndex for Run 3 (1)

- The first EventIndex version was designed in 2013, implemented in 2014 and started operations in 2015
  - Software tools evolved considerably since then!
- The partitioned architecture allowed the replacement of individual components during the years
  - In particular, the Data Collections component, gathering together all indexing information produced by Grid jobs for each dataset, was revised and re-implemented since 2016 with the addition of a controller process: the Supervisor

EventIndex for Run 3 (2)

- The core data storage system was reimplemented during 2021 and deployed in 2022 for the start of LHC Run 3
  - HBase for the dataset and event tables
  - Phoenix interface for SQL queries
  - New client query service CLI also implemented for optimal performance

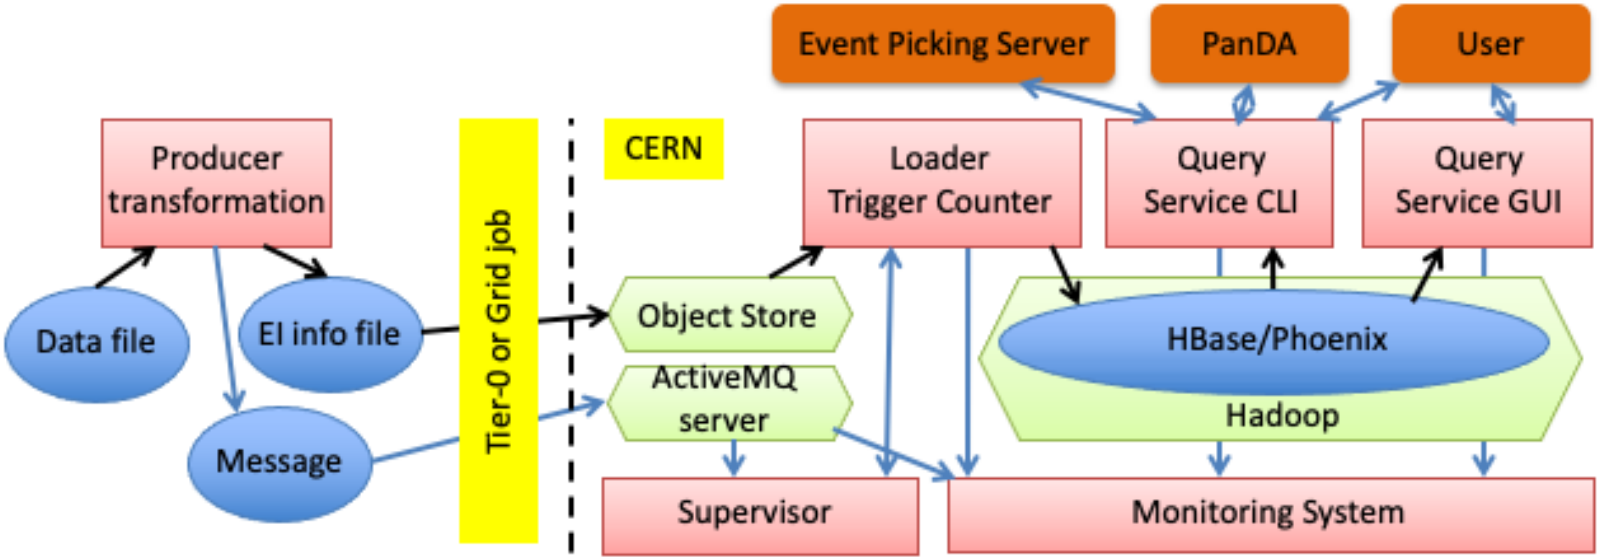

EventIndex for Run 3 (3)

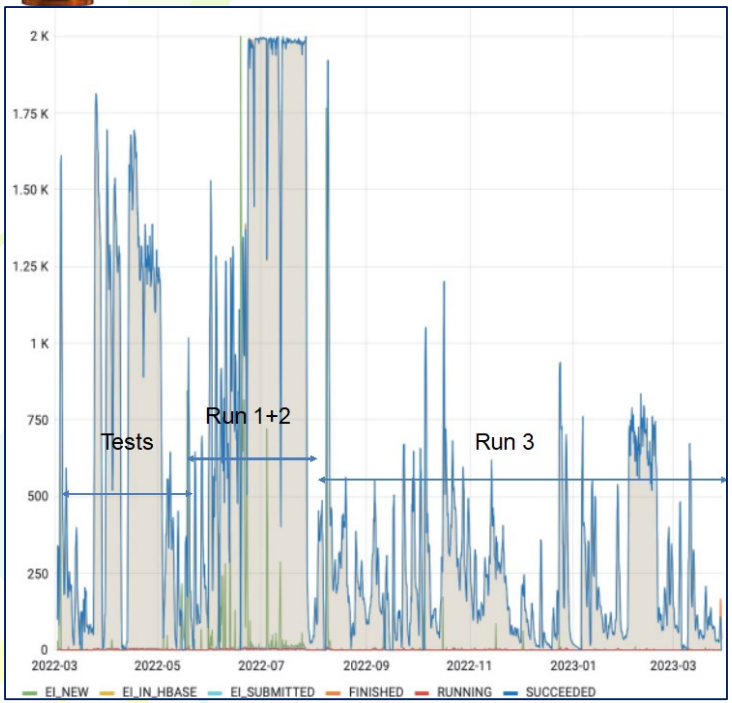

- After initial tests, HBase was loaded with all Run 1+2 data and then received Run 3 data in real time
  - Ingestion performance amply sufficient as shown in the figure (between March 2022 and March 2023)

- Search and retrieve performance is constantly monitored
  - The figure shows the response times to queries retrieving 10, 100, 1000 and 10000 events in seconds (between May 2022 and February 2023)

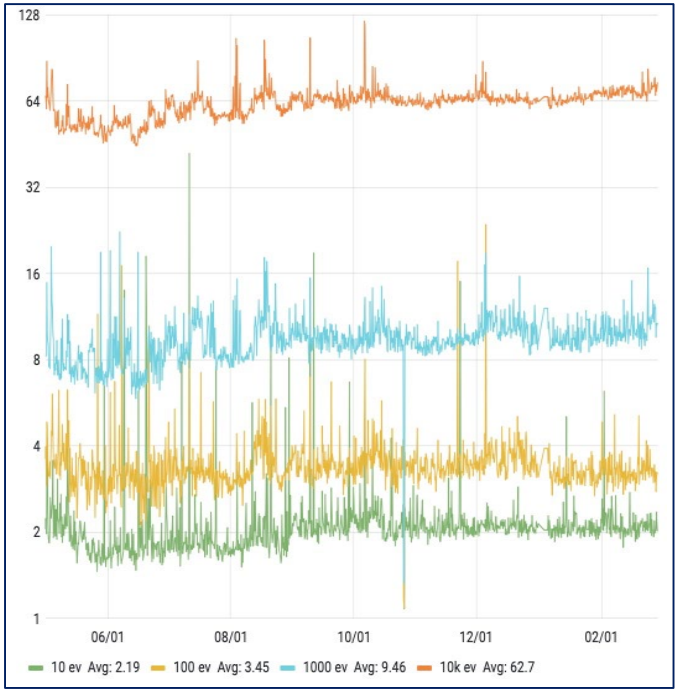

Event Picking Service

- Some physics analyses need to extract many events in order to process them with enhanced algorithms
  - γγ -> WW analysis:
    - The first round 50k events (2019)
    - The second round 136k events (2021)
  - B_c^* -> B_c gamma analysis:
    - 16K events (2023)
  - Z -> TauTau selections:
    - 11K events (2023)
- An automatic system to extract the requested events and deliver them to the requestors is therefore needed: the Event Picking Service

Tasks and problems (1)

- Split by Run: required for correct work with minimal time
  - The number of input data can be large
  - Input data may not be in order
  - 1 run 1 file
- Get GUIDs: this information should be added to panda and is required to get the dataset name template
  - Possible error answer from EI
- Get Dataset Name: require for panda job
  - Rucio has no information about some GUID

Tasks and problems (2)

- Start Panda Job: make a real copy of events
  - Long working time
  - The result may be an error (even for valid inputs)
- Validate : should validate the output data
  - Possible duplicate events
  - Possible skip events
- Set Metadata in RUCIO: panda does not set event count of events
  - Possible big number of output files
- Restart: events that were skipped or have errors
  - Possible big number such events in different runs

Request Workflows

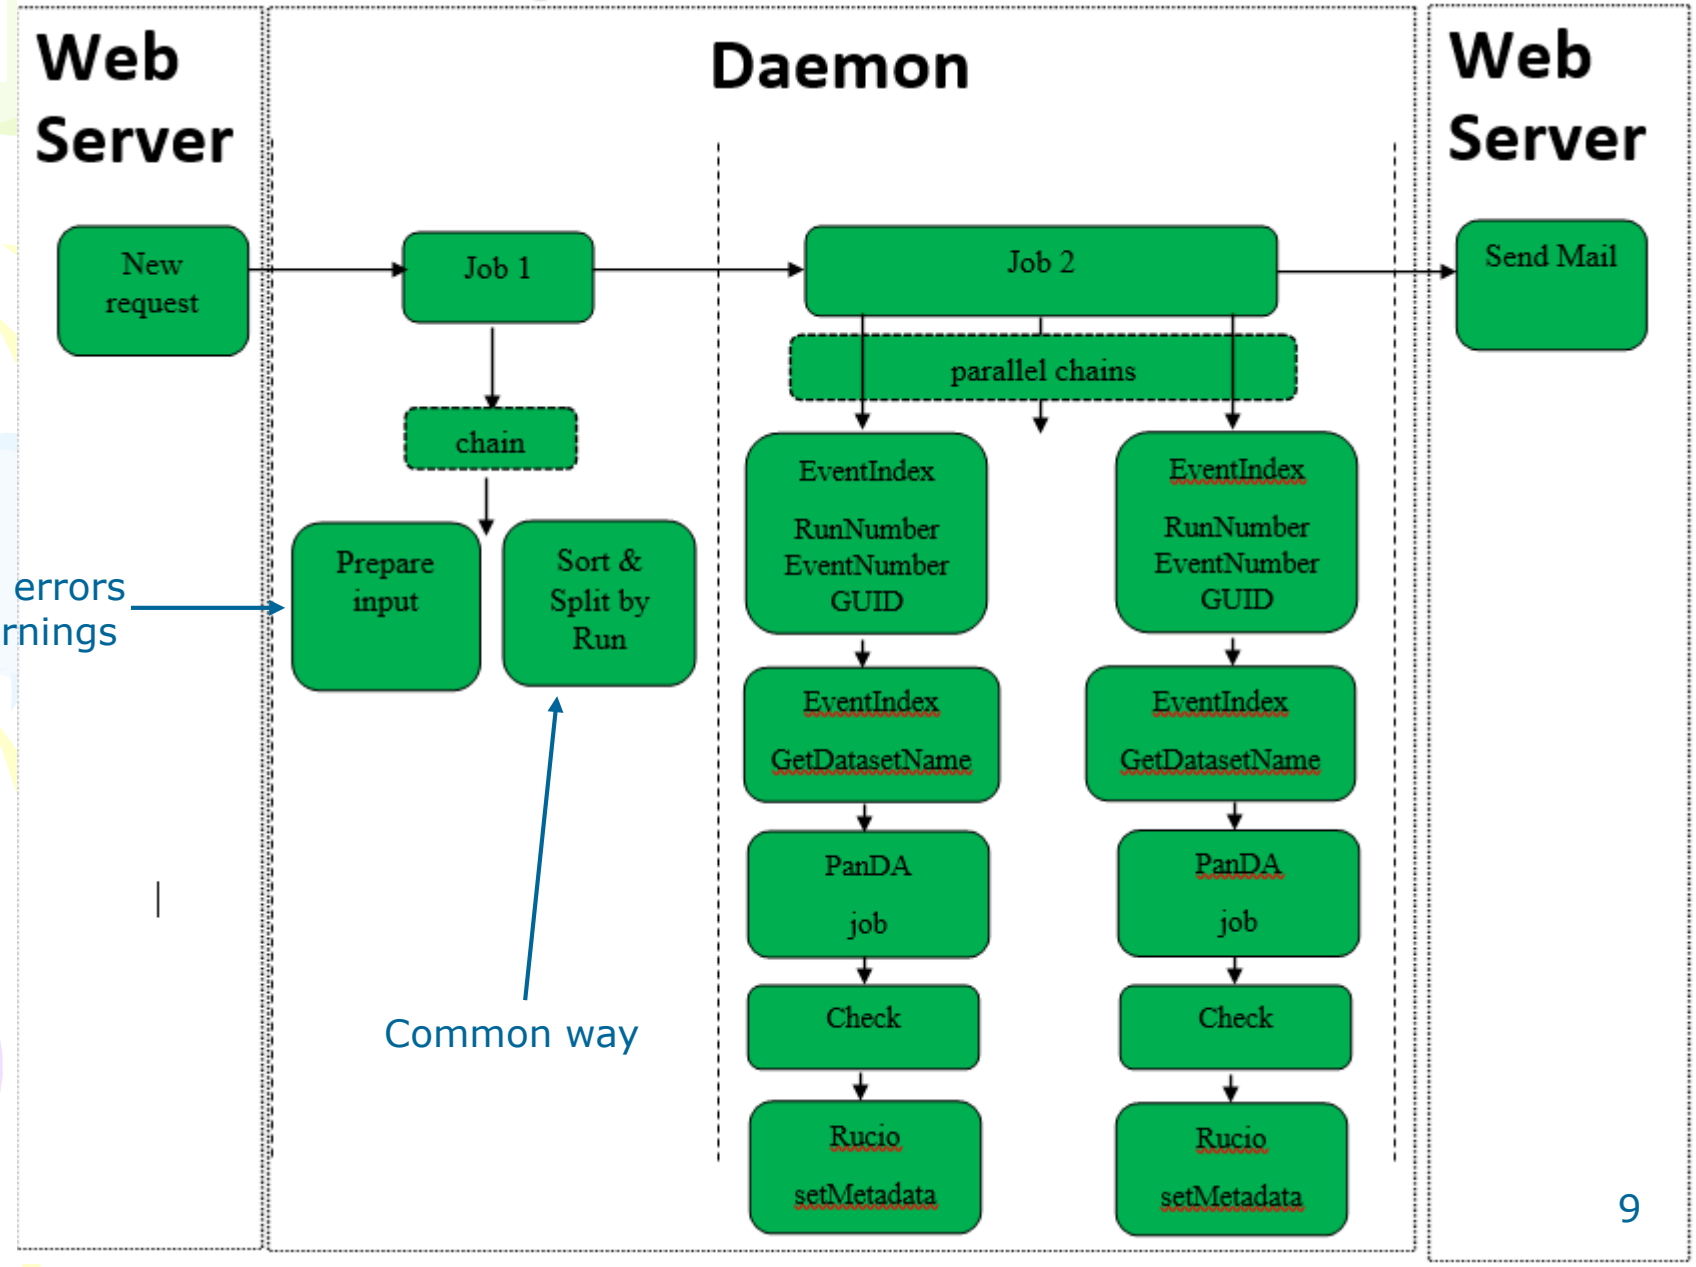

Architecture of the Event Picking service

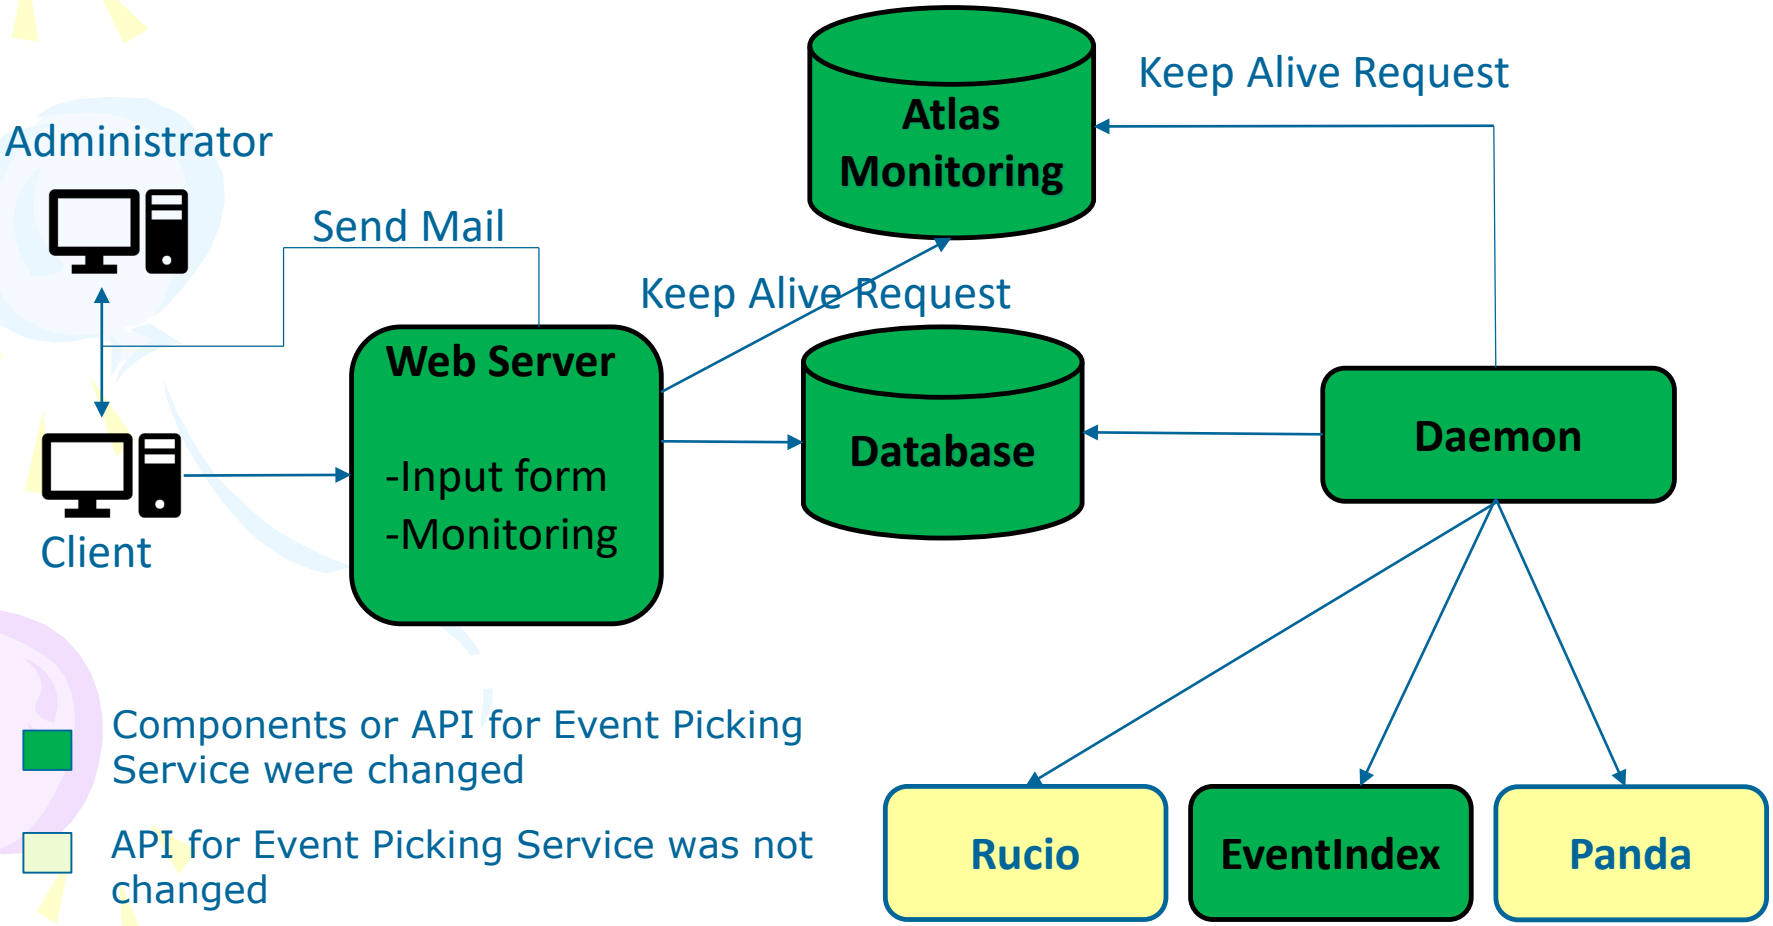

Web interface

Basic technologies:

- PostgreSQL;
- Apache Tomcat;
- JAVA;
- WALT*.

Version 1.1.2 available

*Web Application Lego Toolkit (Developer Sergey Kunyaev from JINR)

Start page :

https://atlas-event-picking.cern.ch/eventpicking/

Web service is available outside of CERN (need CERN SSO authorization)

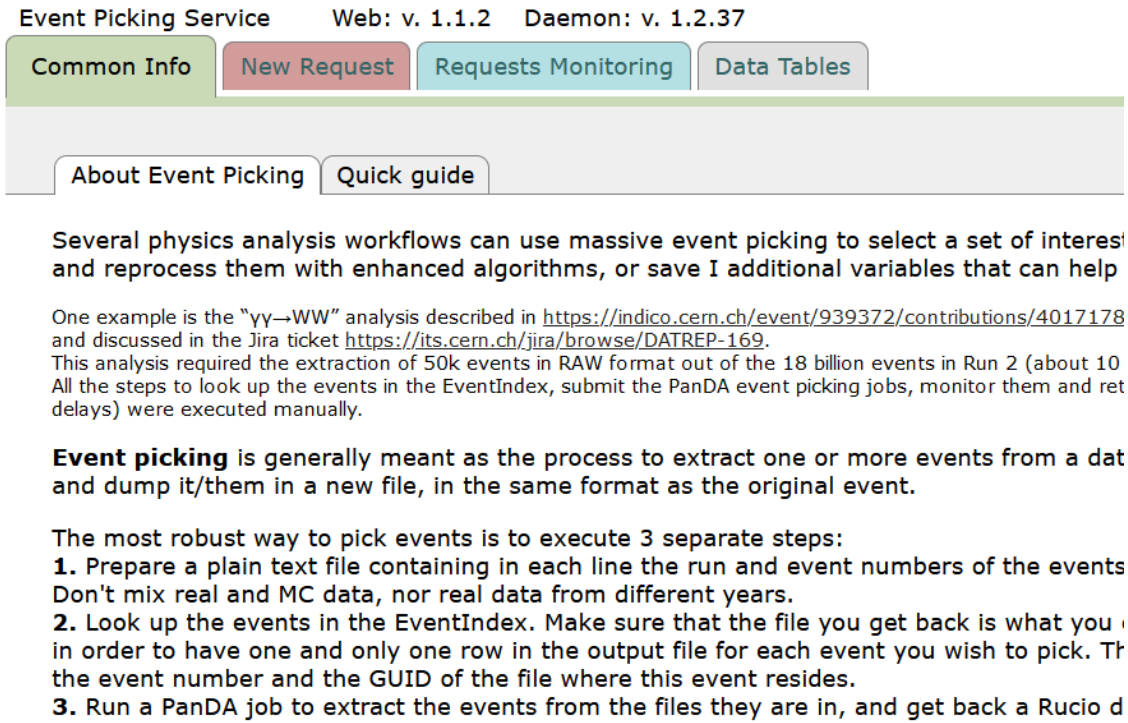

Create new request

- all required fields must be filled
- text file must contain only strings like XXX YYYYY

(first column – run number, second column – event number)

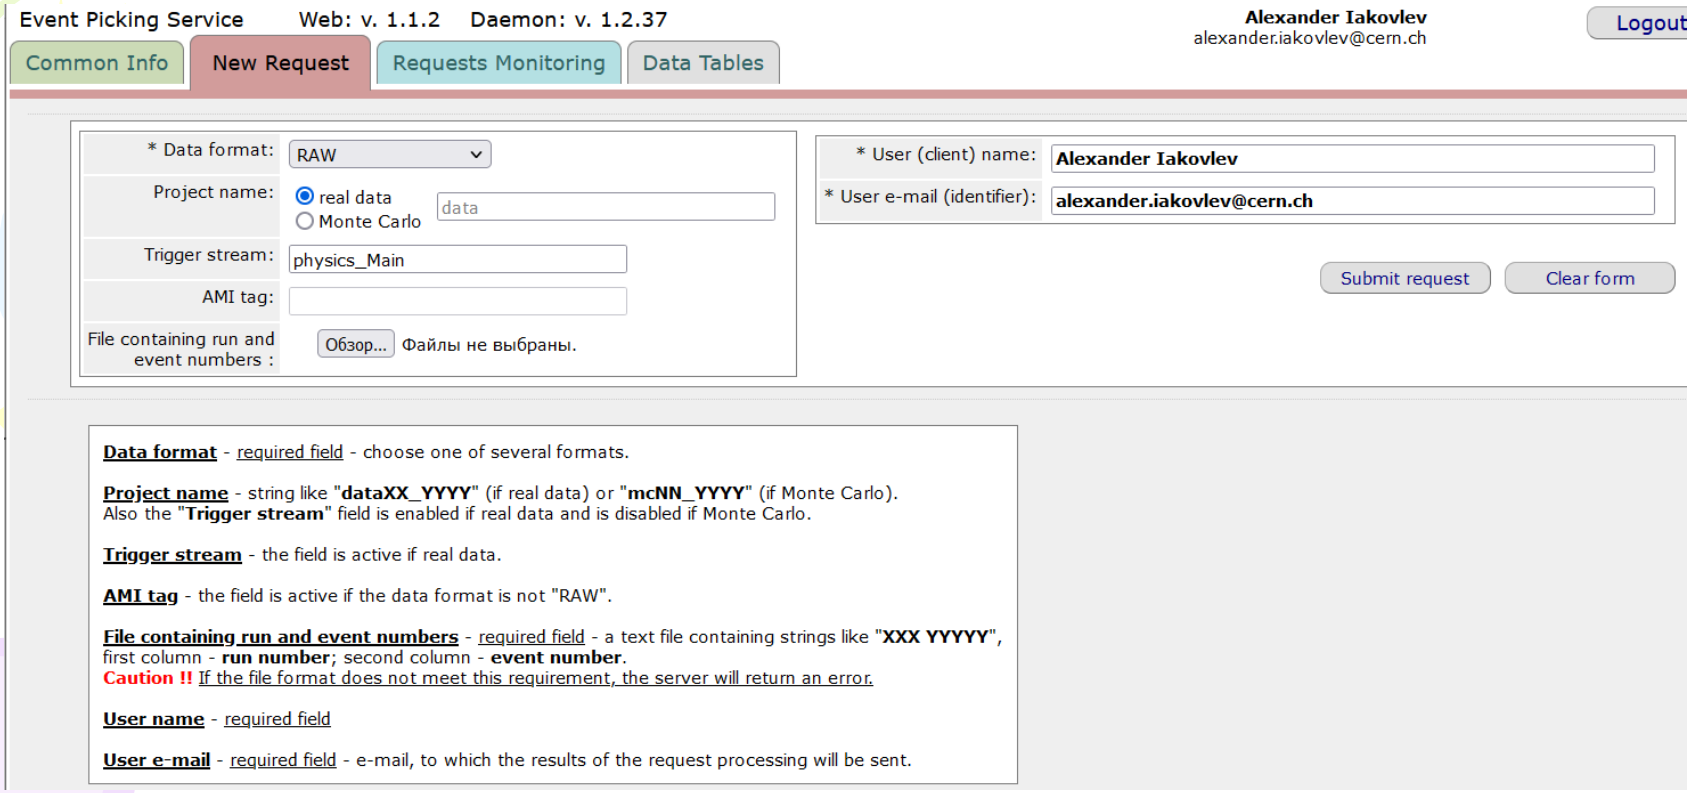

Requests monitoring

list of requests

search filter by different fields

current status of request

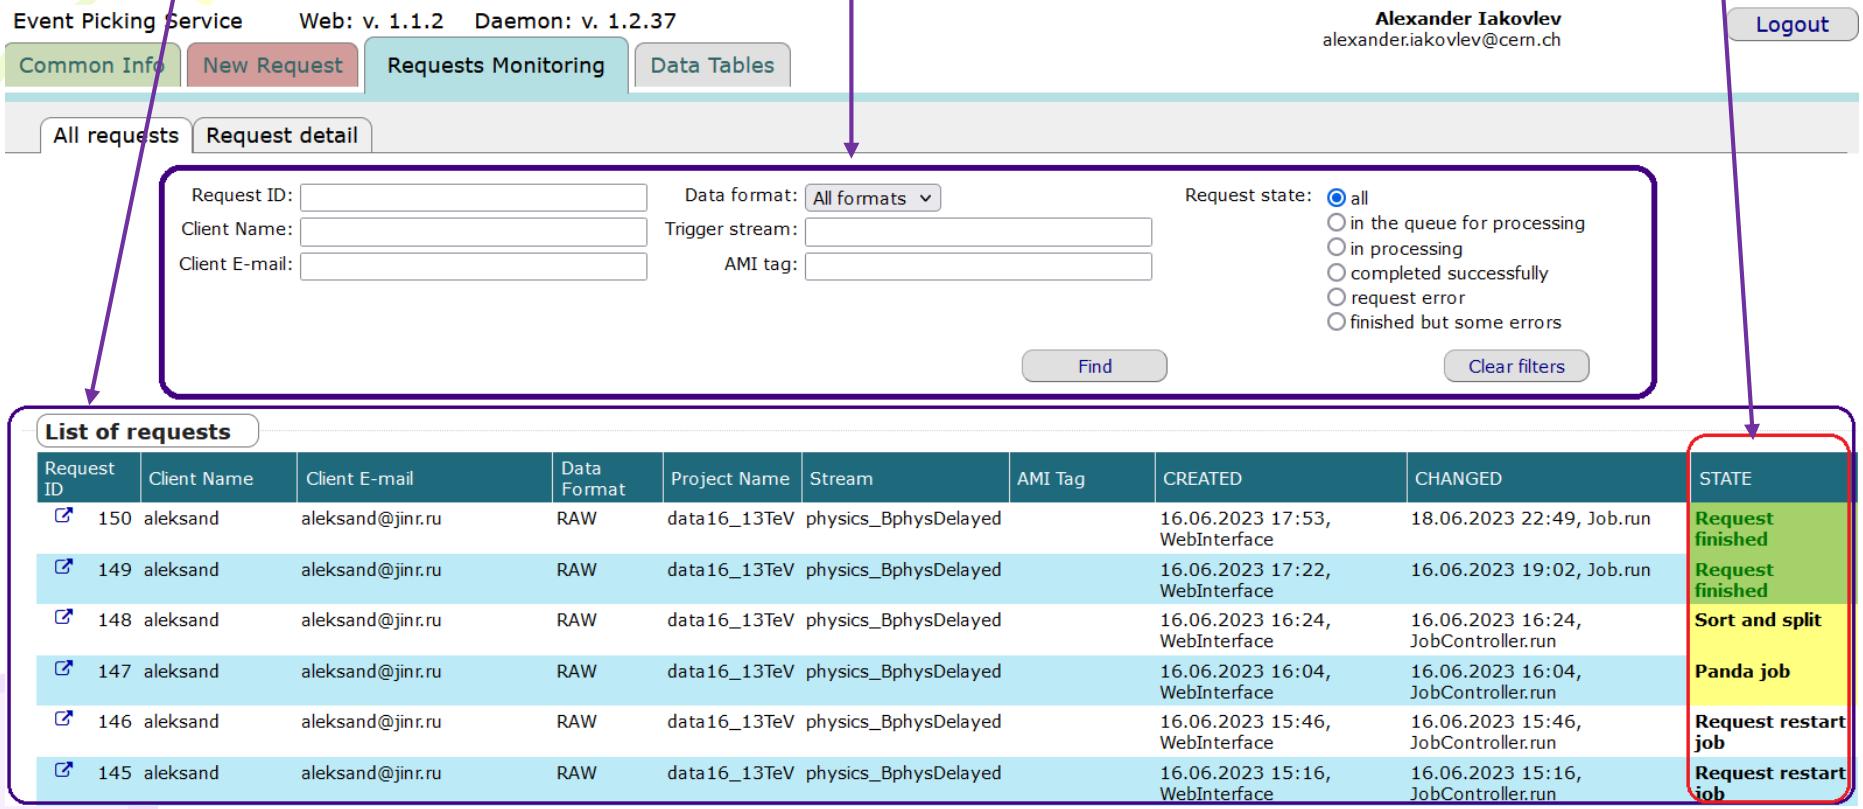

Request ID	Client Name	Client E-mail	Data Format	Project Name	Stream	AMI Tag	CREATED	CHANGED	STATE
150	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 17:53, WebInterface	18.06.2023 22:49, Job.run	Request finished
149	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 17:22, WebInterface	16.06.2023 19:02, Job.run	Request finished
148	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 16:24, WebInterface	16.06.2023 16:24, JobController.run	Sort and split
147	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 16:04, WebInterface	16.06.2023 16:04, JobController.run	Panda job
146	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 15:46, WebInterface	16.06.2023 15:46, JobController.run	Request restart job
145	aleksand	aleksand@jinr.ru	RAW	data16_13TeV	physics_BphysDelayed		16.06.2023 15:16, WebInterface	16.06.2023 15:16, JobController.run	Request restart job

Requests monitoring

Request detail information:

start and stop time for jobs, chains and tasks

link to Panda page for Panda task

output result in JSON format

current statuses of jobs, chains and tasks

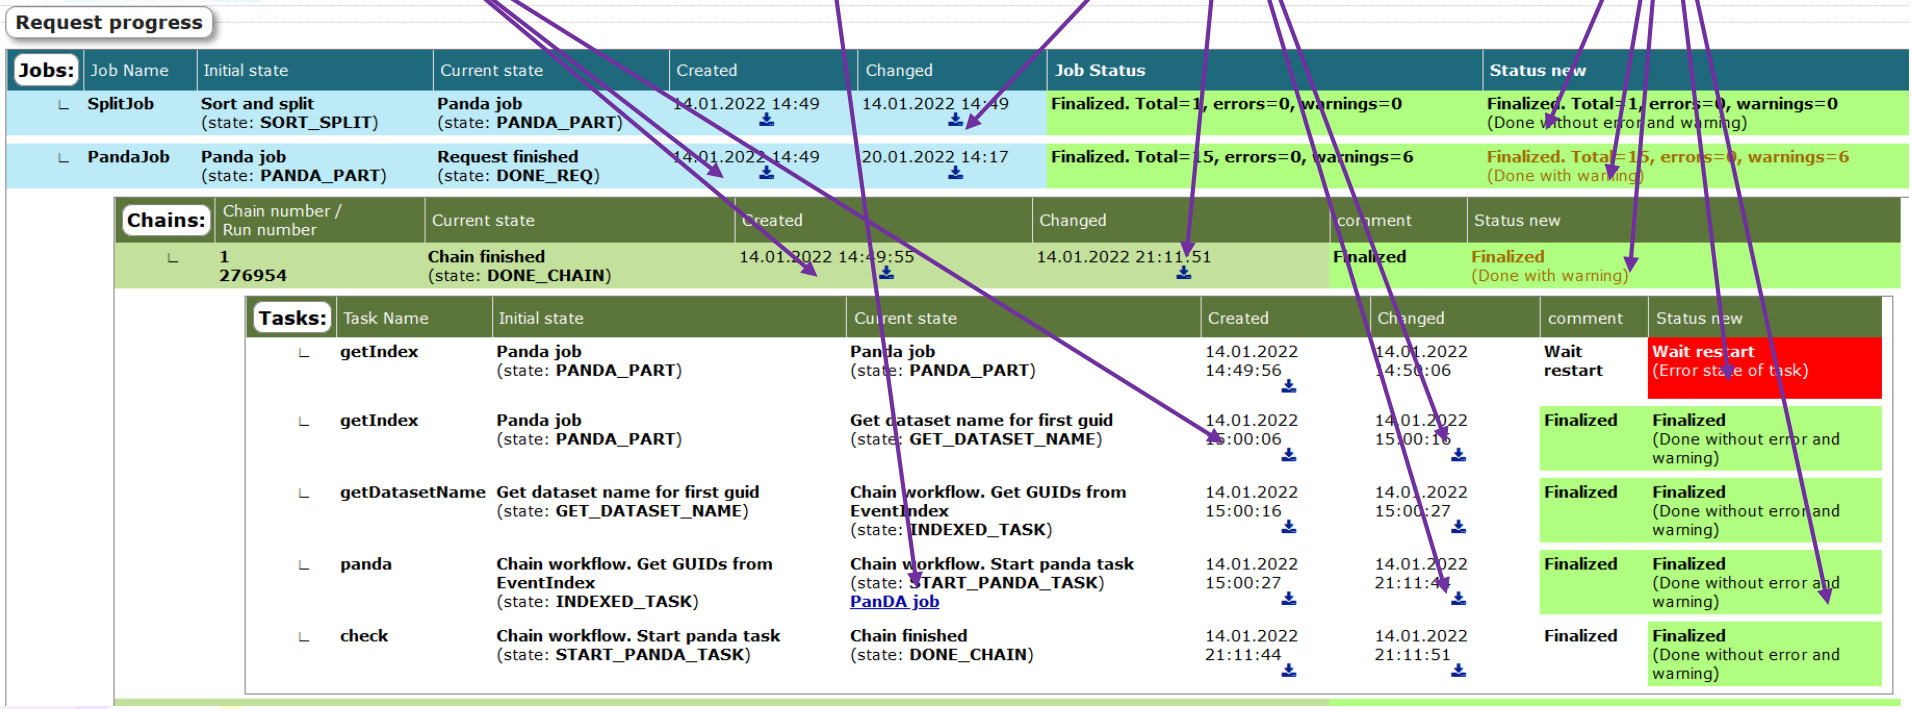

Request progress

Jobs:	Job Name	Initial state	Current state	Created	Changed	Job Status	Status new
	SplitJob	Sort and split (state: SORT_SPLIT)	Panda job (state: PANDA_PART)	14.01.2022 14:49	14.01.2022 14:49	Finalized. Total=1, errors=0, warnings=0	Finalized. Total=1, errors=0, warnings=0 (Done without error and warning)
	PandaJob	Panda job (state: PANDA_PART)	Request finished (state: DONE_REQ)	14.01.2022 14:49	20.01.2022 14:17	Finalized. Total=15, errors=0, warnings=6	Finalized. Total=15, errors=0, warnings=6 (Done with warning)

Chains:	Chain number / Run number	Current state	Created	Changed	comment	Status new
	1 276954	Chain finished (state: DONE_CHAIN)	14.01.2022 14:49:55	14.01.2022 21:11:51	Finalized	Finalized (Done with warning)

Tasks:	Task Name	Initial state	Current state	Created	Changed	comment	Status new
	getIndex	Panda job (state: PANDA_PART)	Panda job (state: PANDA_PART)	14.01.2022 14:49:56	14.01.2022 14:50:06	Wait restart	Wait restart (Error state of task)
	getIndex	Panda job (state: PANDA_PART)	Get dataset name for first guid (state: GET_DATASET_NAME)	14.01.2022 15:00:06	14.01.2022 15:00:16	Finalized	Finalized (Done without error and warning)
	getDatasetName	Get dataset name for first guid (state: GET_DATASET_NAME)	Chain workflow. Get GUIDs from EventIndex (state: INDEXED_TASK)	14.01.2022 15:00:16	14.01.2022 15:00:27	Finalized	Finalized (Done without error and warning)
	panda	Chain workflow. Get GUIDs from EventIndex (state: INDEXED_TASK)	Chain workflow. Start panda task (state: START_PANDA_TASK) PanDA job	14.01.2022 15:00:27	14.01.2022 21:11:44	Finalized	Finalized (Done without error and warning)
	check	Chain workflow. Start panda task (state: START_PANDA_TASK)	Chain finished (state: DONE_CHAIN)	14.01.2022 21:11:44	14.01.2022 21:11:51	Finalized	Finalized (Done without error and warning)

Requests monitoring

Request results panel:

Request view :

Data format:	**RAW**	Request ID:	**95**	Created:	**23.04.2022 23:21** **WebInterface**
Stream:	**physics_Main**	Client name:	**aleksand**	Last change:	**24.04.2022 15:50** **Job.run**
AMI tag:		Client e-mail:	**aleksand@jinr.ru**	Current status:	**Request finished**
				Detail status:	**Request finished** (Done with restart task)

Request result :

Download Summary Report file (+)

Event Processing Progress Log (+)

Result files (+)

typical request result (for completed successfully)

Request result

DatasetName	file_path
group.proj-evind.data15_13TeV.00276954.physics_Main.evtpick.DRAW_EVTPICK.r60t22992	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/67/c2/group.proj-evind.27840942._000001.event.dat
group.proj-evind.data15_13TeV.00276689.physics_Main.evtpick.DRAW_EVTPICK.r60t23003	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/5a/69/group.proj-evind.27840962._000001.event.dat
group.proj-evind.data15_13TeV.00276952.physics_Main.evtpick.DRAW_EVTPICK.r60t23002	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/a1/f7/group.proj-evind.27840961._000001.event.dat
group.proj-evind.data15_13TeV.00279345.physics_Main.evtpick.DRAW_EVTPICK.r60t23013	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/
group.proj-evind.data15_13TeV.00276790.physics_Main.evtpick.DRAW_EVTPICK.r60t23001	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/36/d1/group.proj-evind.27840958._000002.event.dat
group.proj-evind.data15_13TeV.00278880.physics_Main.evtpick.DRAW_EVTPICK.r60t23016	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/1c/78/group.proj-evind.27840979._000002.event.dat /eo /group/proj-evind/aa/a4/group.proj-evind.27840979._000001.event.dat /eos/atlas/atlasgroupdisk/proj-evind/rucio/gr evind.27840979._000003.event.dat /eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/ed/54/group.proj-evi
group.proj-evind.data15_13TeV.00279764.physics_Main.evtpick.DRAW_EVTPICK.r60t22996	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/19/99/group.proj-evind.27840952._000002.event.dat /ec /group/proj-evind/6c/a3/group.proj-evind.27840952._000001.event.dat
group.proj-evind.data15_13TeV.00279685.physics_Main.evtpick.DRAW_EVTPICK.r60t23017	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/
group.proj-evind.data15_13TeV.00279259.physics_Main.evtpick.DRAW_EVTPICK.r60t22994	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/43/05/group.proj-evind.27840946._000001.event.dat /ec /group/proj-evind/ba/aa/group.proj-evind.27840946._000002.event.dat
group.proj-evind.data15_13TeV.00279284.physics_Main.evtpick.DRAW_EVTPICK.r60t23010	/eos/atlas/atlasgroupdisk/proj-evind/rucio/group/proj-evind/a7/80/group.proj-evind.27840969._000002.event.dat /eo /group/proj-evind/d1/0e/group.proj-evind.27840969._000001.event.dat

Restart request

"Restart request" - In terms of the Event Picking Service, it means creating a new request using the input from the parent request.

Request view :

Data format:	**RAW**	Request ID:	**149**	Created:	**16.06.2023 17:22** **WebInterface**
Project name:	**data16_13TeV**	Client name:	**aleksand**	Last change:	**16.06.2023 19:02** **Job.run**
Stream:	**physics_BphysDelayed**	Client e-mail:	**aleksand@jinr.ru**	Current status:	**Request finished**
AMI tag:				Detail status:	**Request finished** (Done without error and warning)

Restart request

Restart request : Request ID: 149

"Restart request" - In terms of the Event Picking Service, it means **creating a new request** using the input from the parent request. A new request will be launched with **your credentials**.

Attention !! There are several options for restarting a request with different options.

Parent request

Request ID:	**149**
Data format:	**RAW**
Stream:	**physics_BphysDelayed**
AMI tag:	
Created:	**16.06.2023 17:22** **WebInterface**
Last change:	**16.06.2023 19:02** **Job.run**
Current status:	**Request finished**
Detail status:	**Request finished** (Done without error and warning)
Client name:	**aleksand**
Client e-mail:	**aleksand@jinr.ru**

* User (client) name:

* User e-mail (identifier):

Full Restart

Restart Warning & Error branches

Restart request | Clear form

Time results

Request	Number of events	Version	Time
γγ -> WW	50k	1.0.0	2 weeks
		manual	3 months
γγ -> WW	136k	Beta version	3 months
B_c^* -> B_c	16K	1.2.37	84h
Z -> TauTau	11K	1.2.37	40h

Conclusion

- The new implementation of the Event Index is working without problems.
- The Event Picking Service update has been completed and is running on the production server.
- Error handling and automatic fixes have been improved, which is why the speed of the new version has been increased.
- The number of users of the service is growing.