Hyperloop Train System

O^2 Analysis Tutorial

November 6th, 2023

Raluca Cruceru

Jan Fiete Grosse-Oetringhaus

Nicolas Poffley

Hyperloop Train System

- Train system for $\mathbf{O^2}$, replacing the LEGO train system
- Allows organized analysis on the **GRID** and **Analysis Facilities**
- Fully integrated with $\mathbf{O^2}$, allowing task configuration
- Individual workflows - known as wagons - are combined into trains
- **Skimmed / Derived data** stored for further processing in subsequent trains
- **Data available**: converted Run 2 data, Run 3 data and MC, Derived data
- Dedicated views for regular **users** and **operators**

- **Train support**
  - 24/5 Operation (different timezones)
  - Institutes: 1 in Americas, 2 in Europe,1 in Asia
  - Shift-type support during working hours
  - Organized feedback sessions

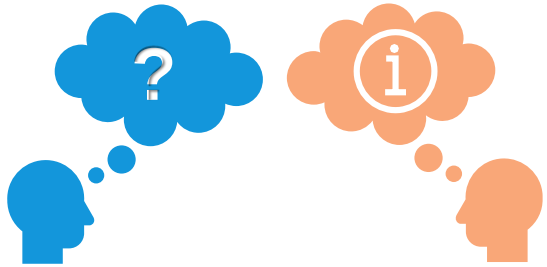

USER	OPERATOR
MyAnalyses	Train Submission
AllAnalyses	Train Runs
Dashboard	Datasets
	Derived Data
	DPG runlists
	Trains with issues

Hyperloop Train System

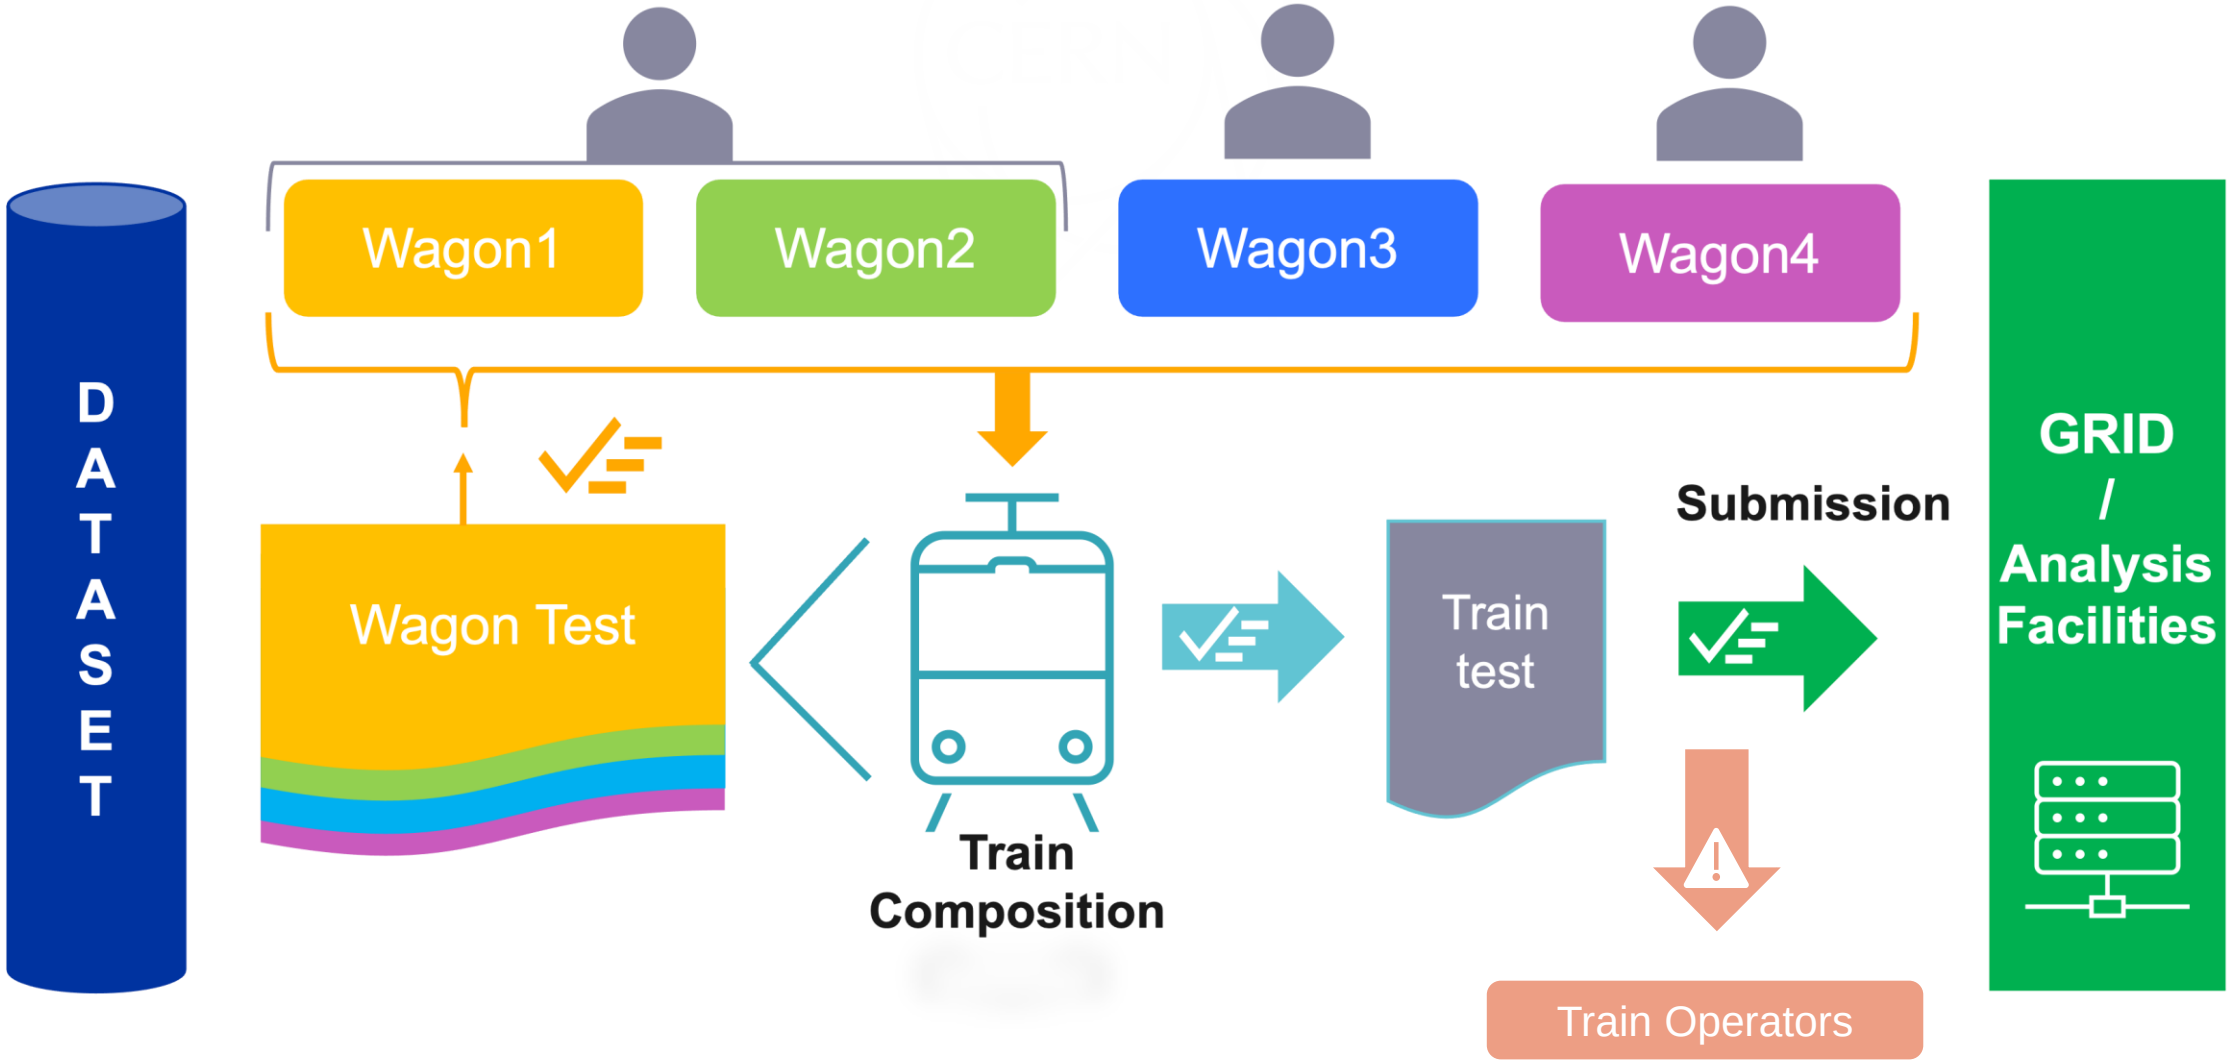

Train runs

There are 4 types of trains that can run on Hyperloop:

1. Analysis train
2. Slim derived data train
3. Derived data train
4. Linked derived data train

Let’s see what are the differences between these train types

Analysis Train

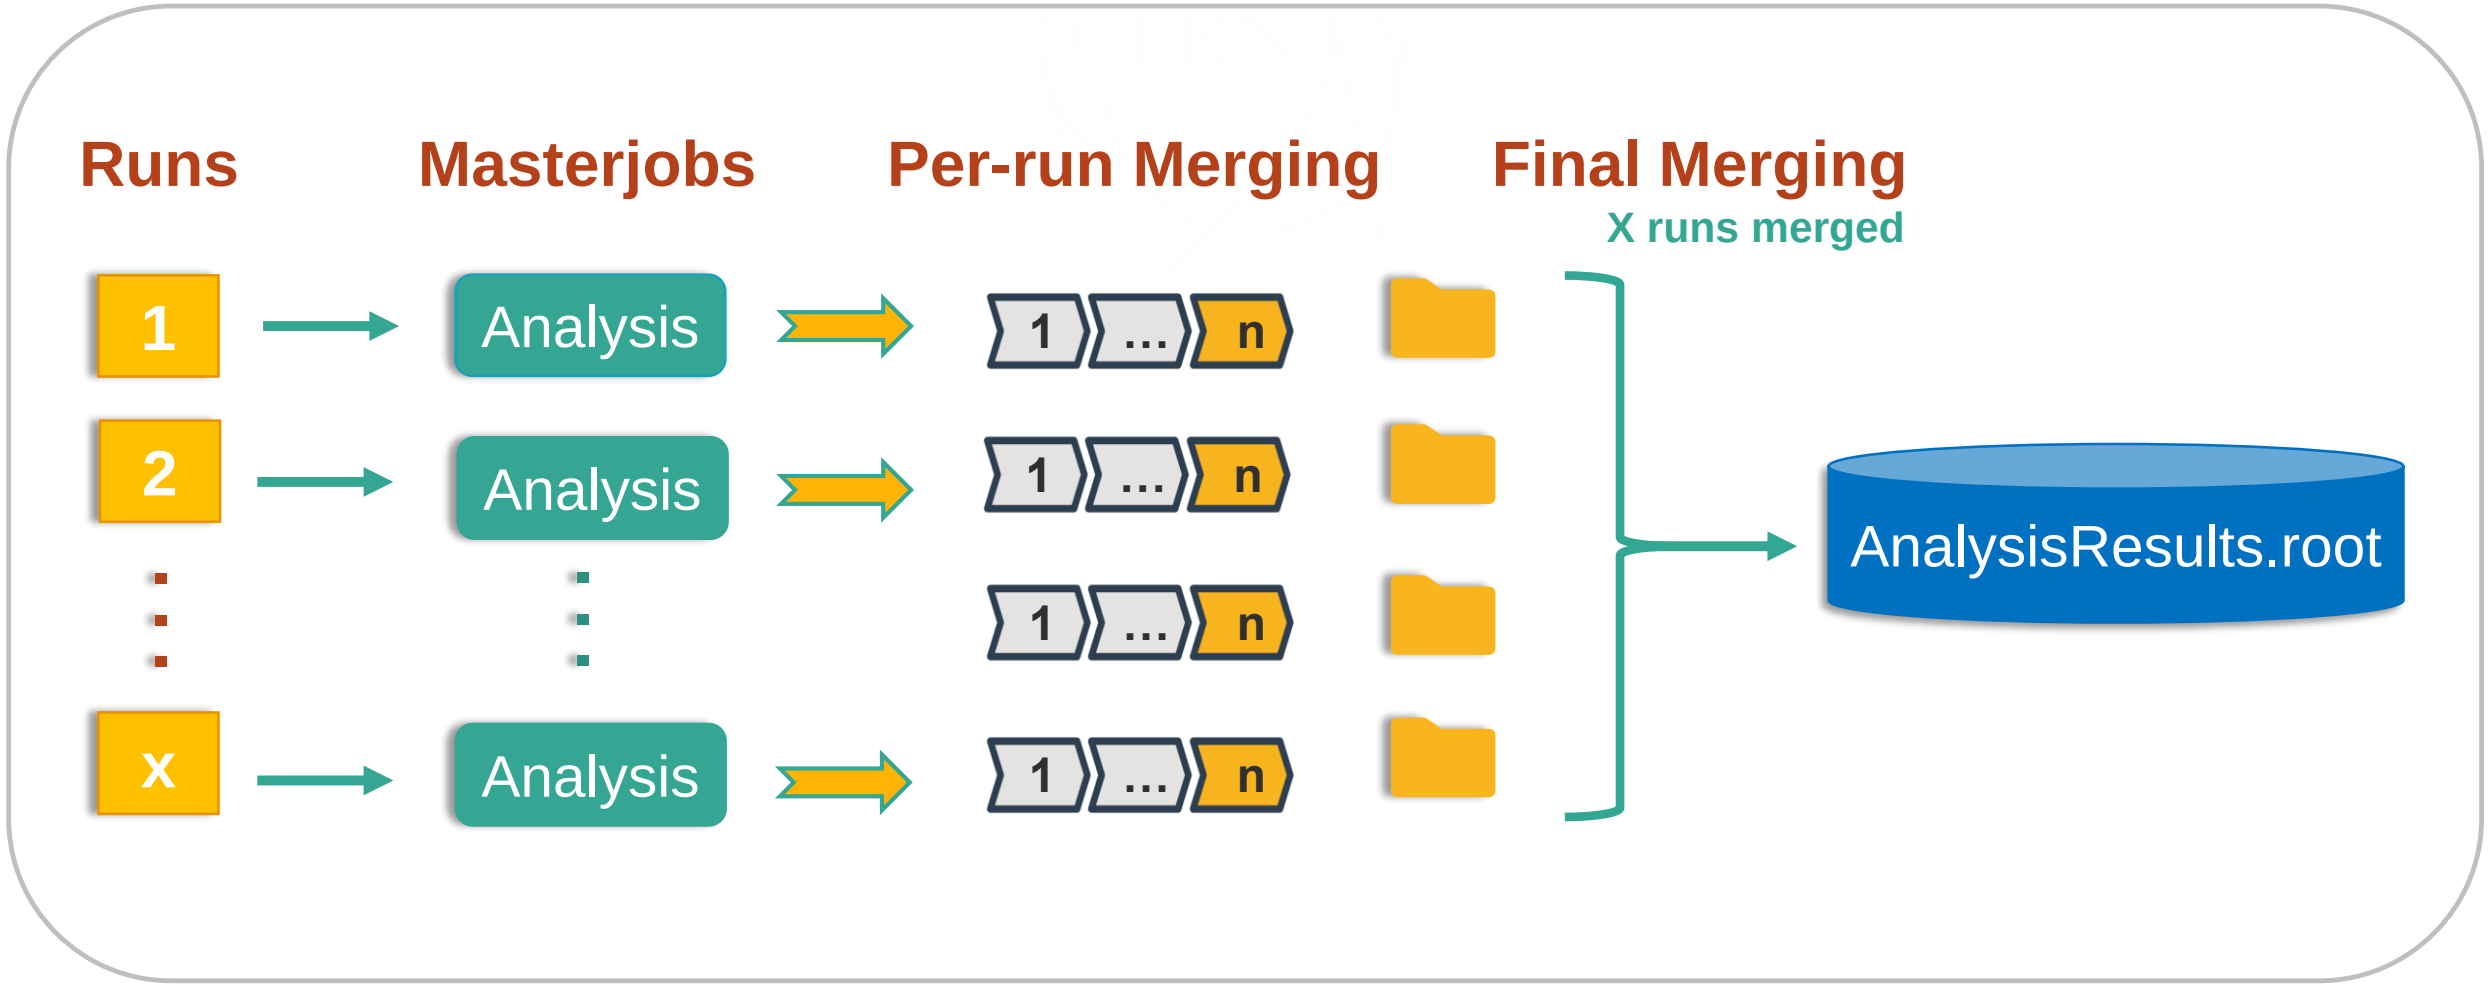

Analysis Train (mergelists)

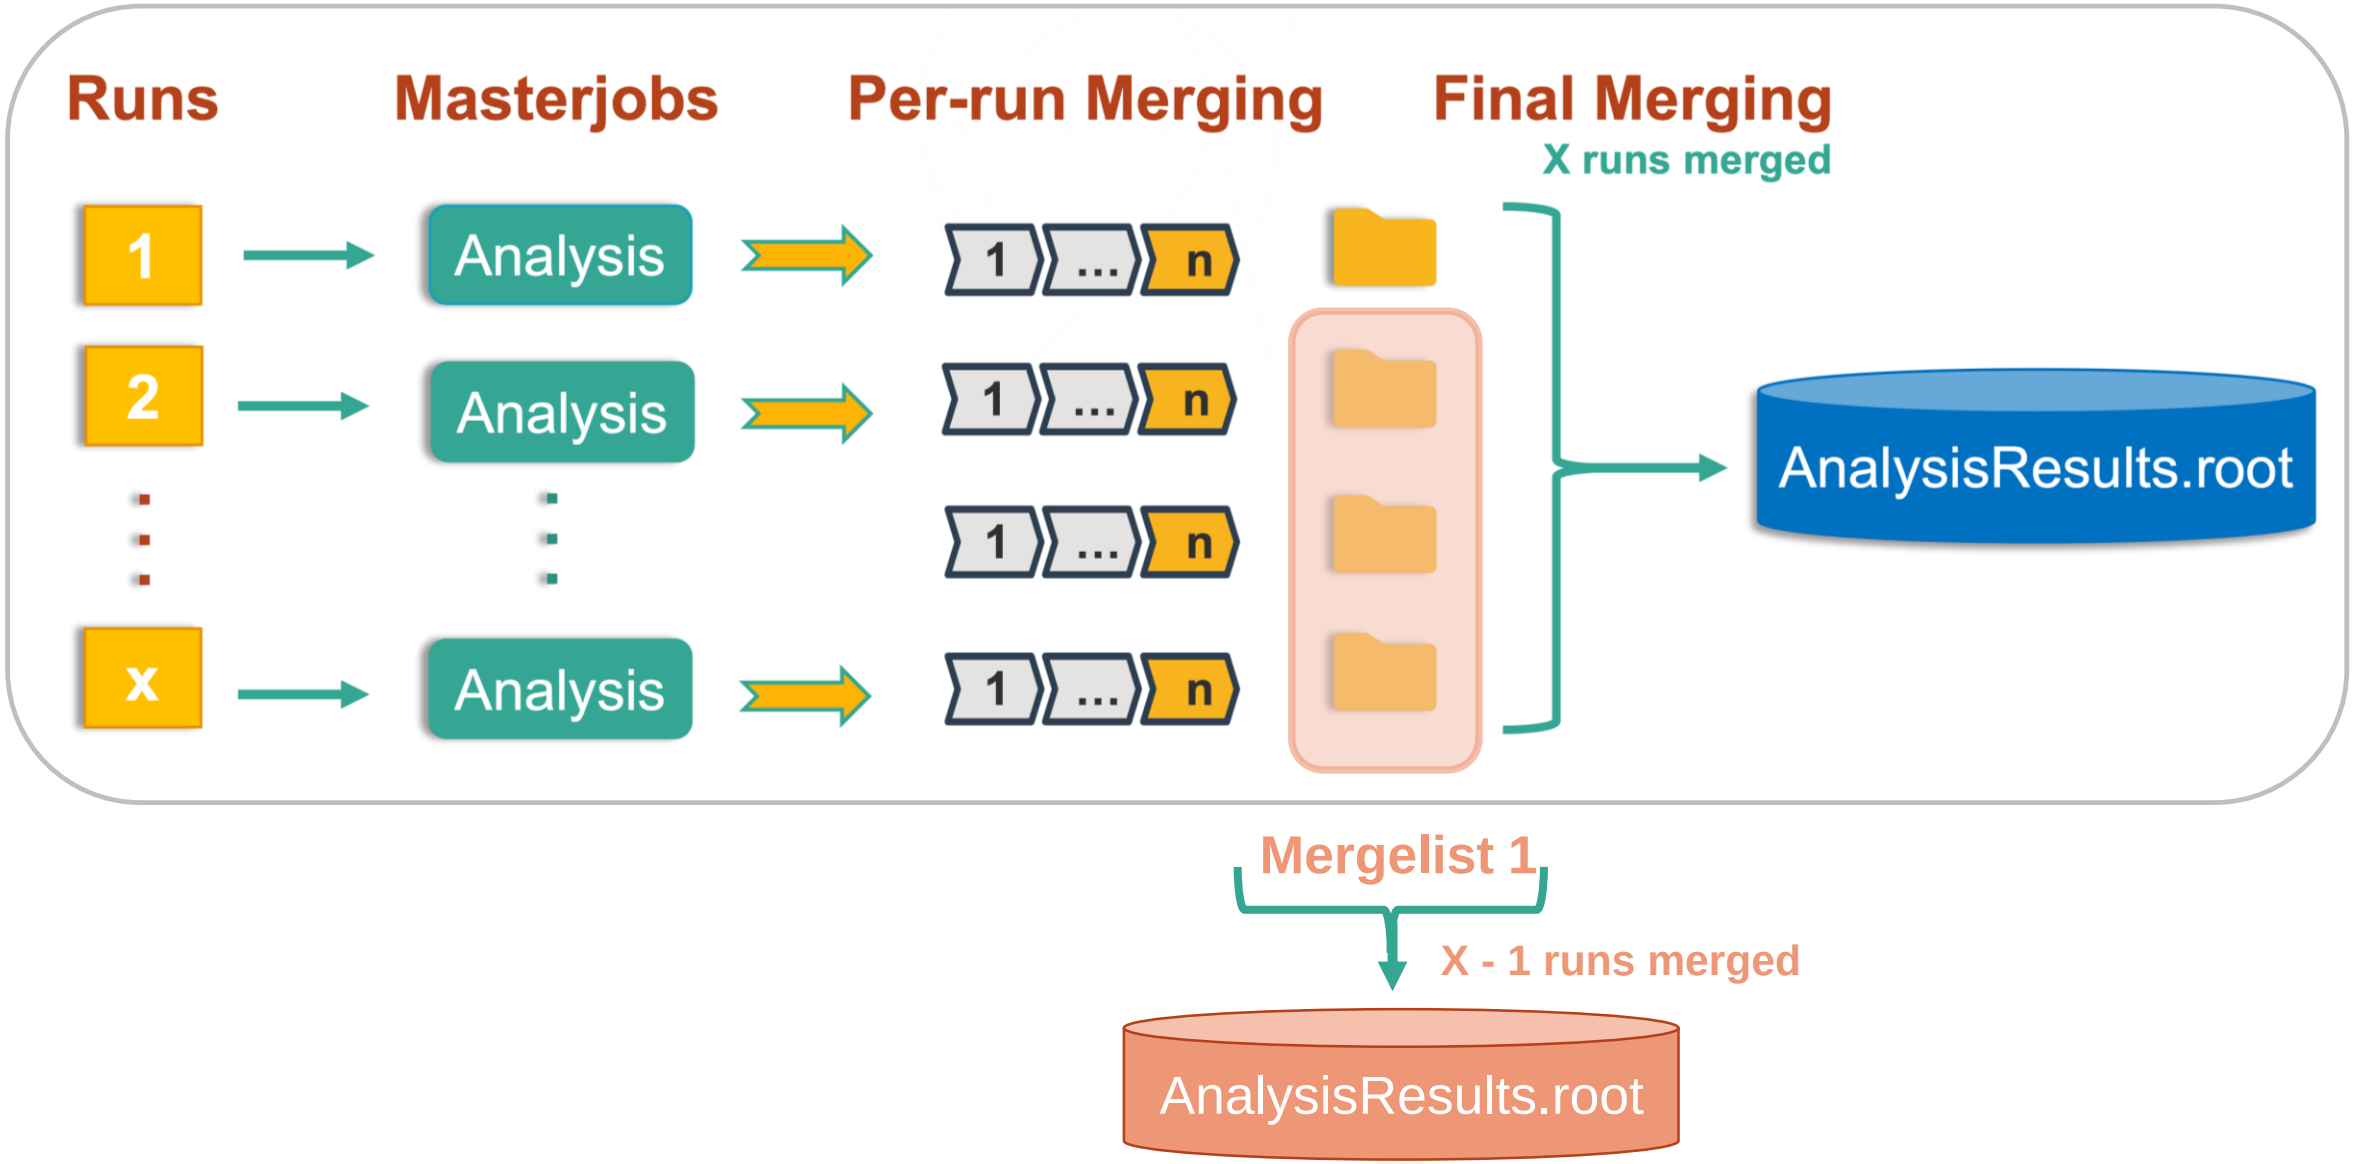

Slim Derived Data Train

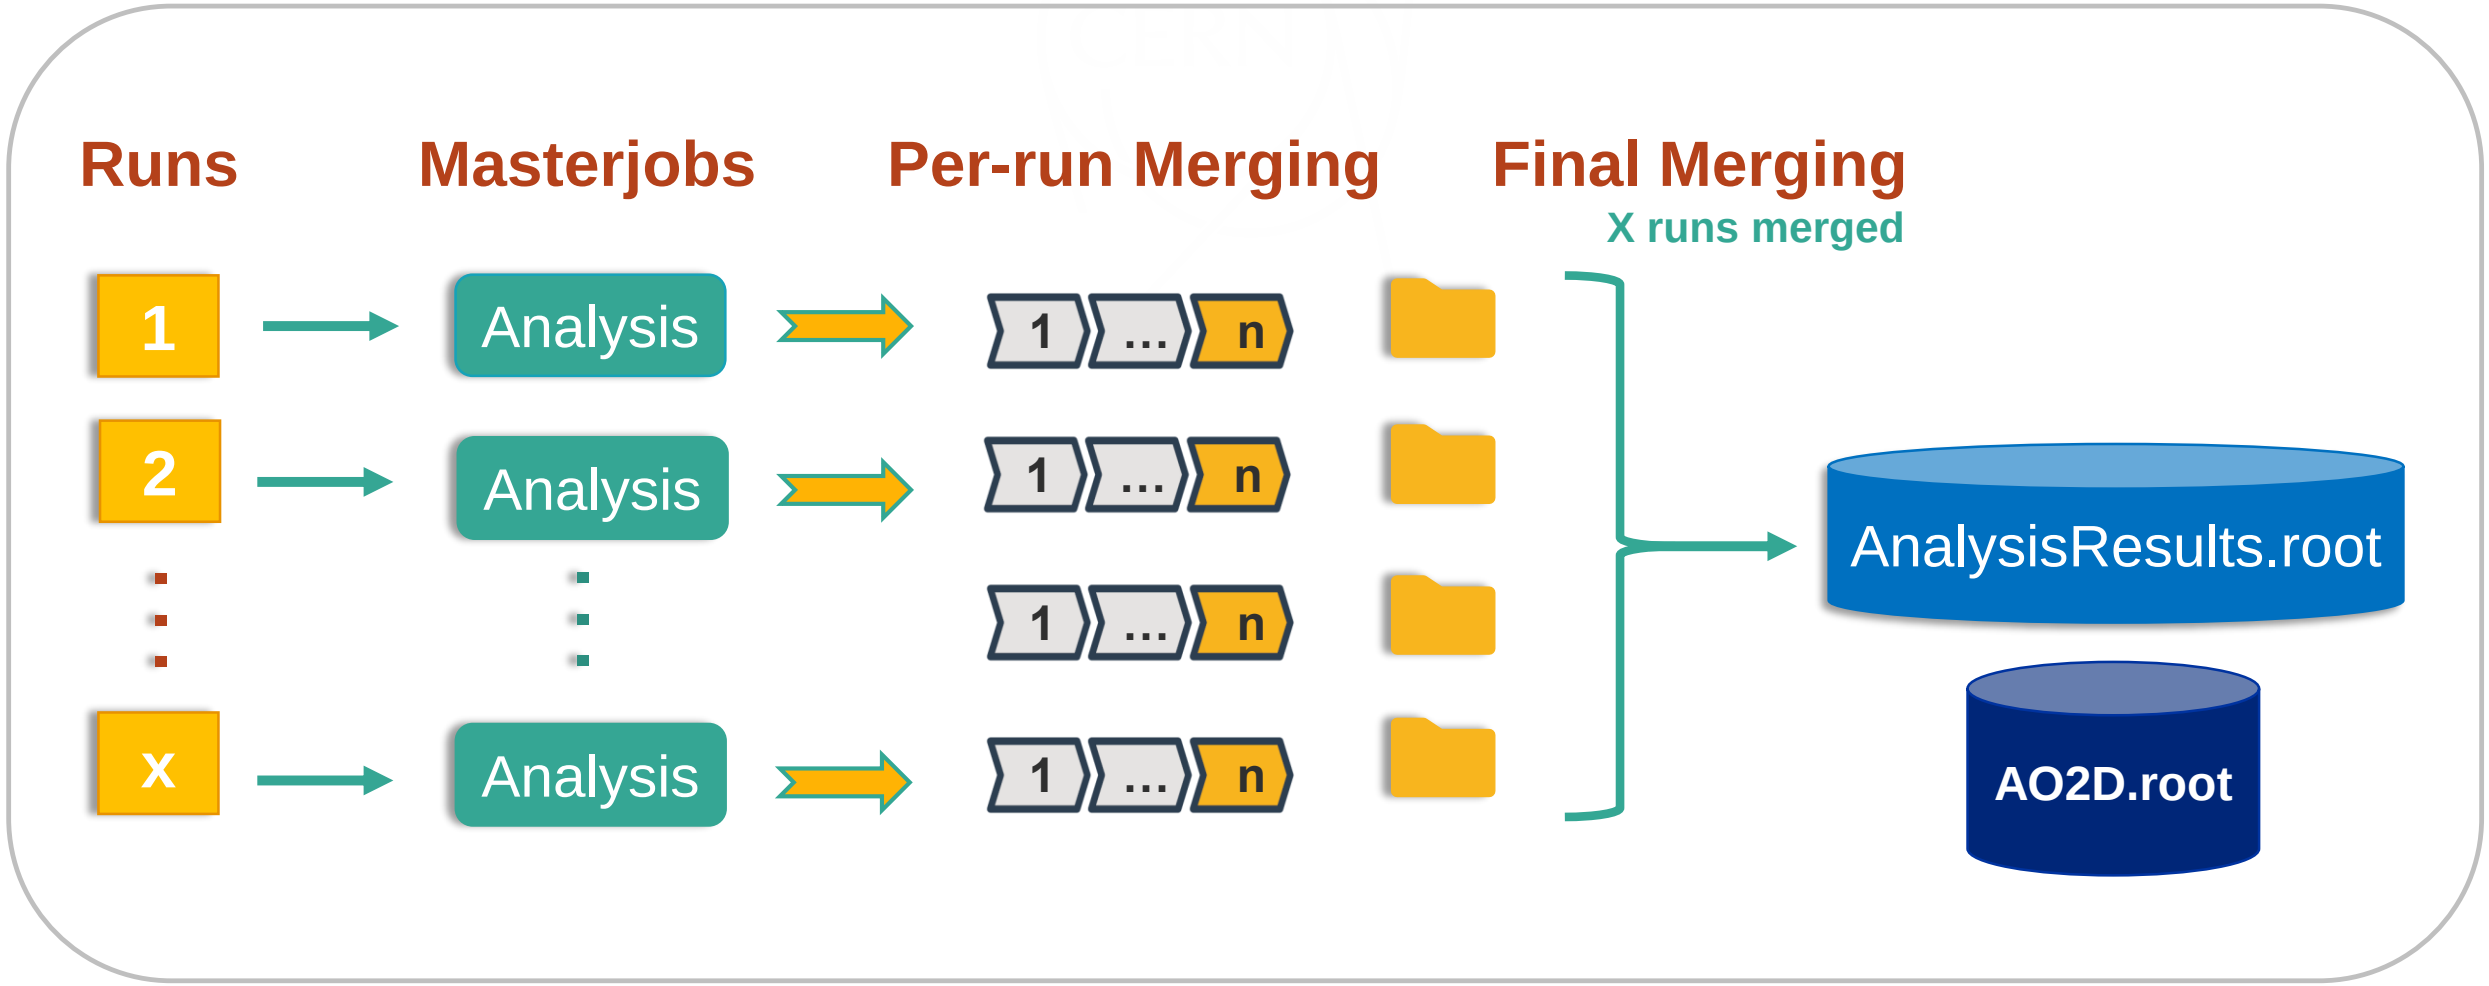

Derived Data Train

Significant differences!

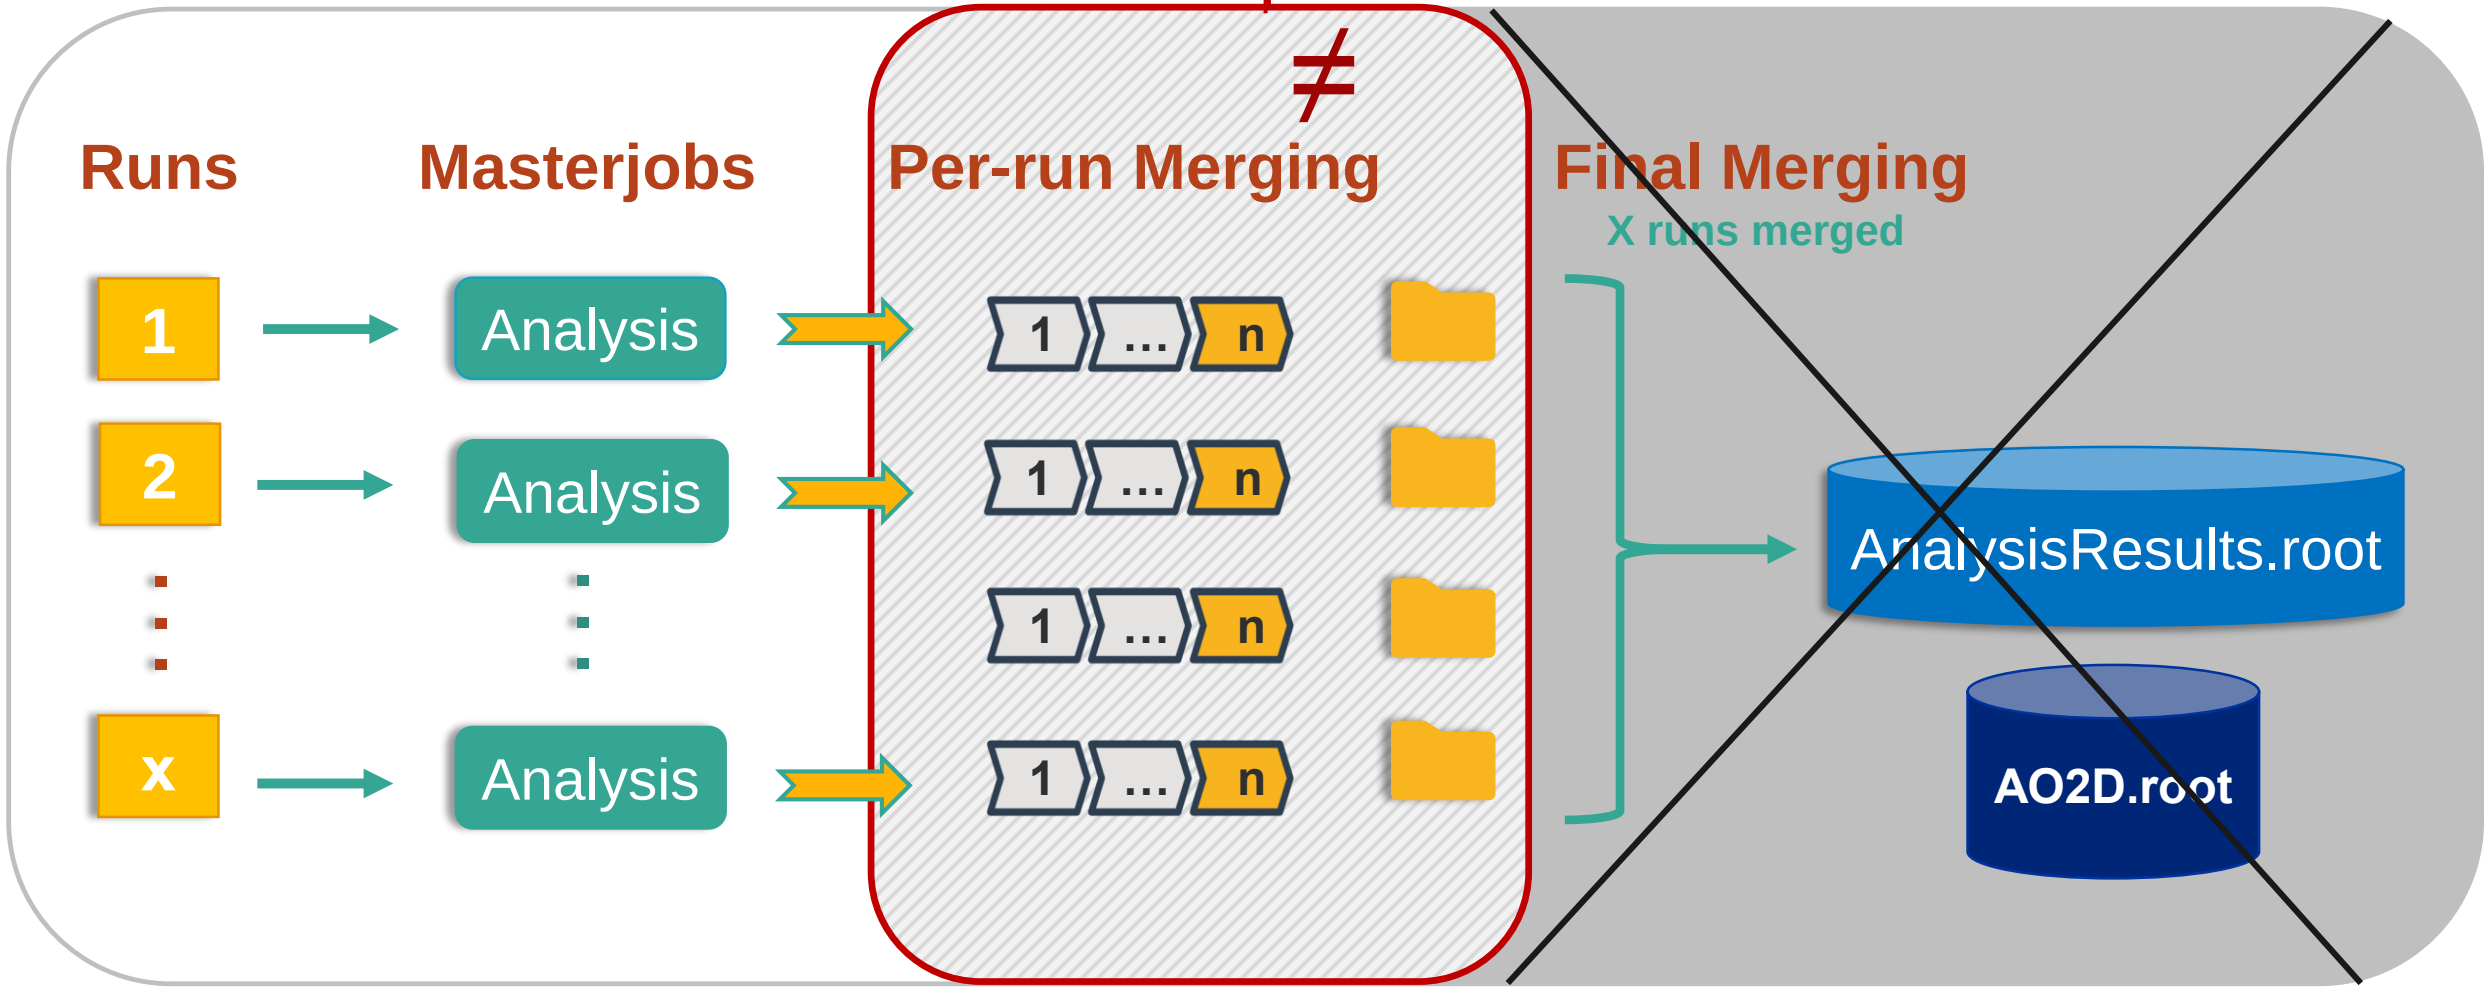

Derived Data Train

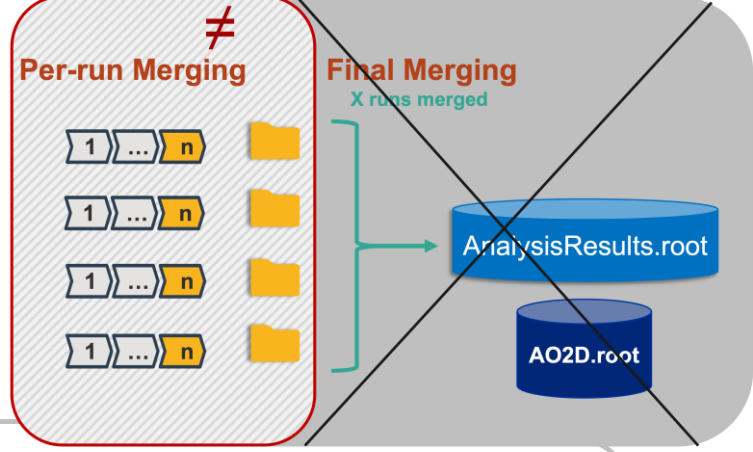

Runs **Masterjobs** **Per-run Merging** (1 level of merging per run)

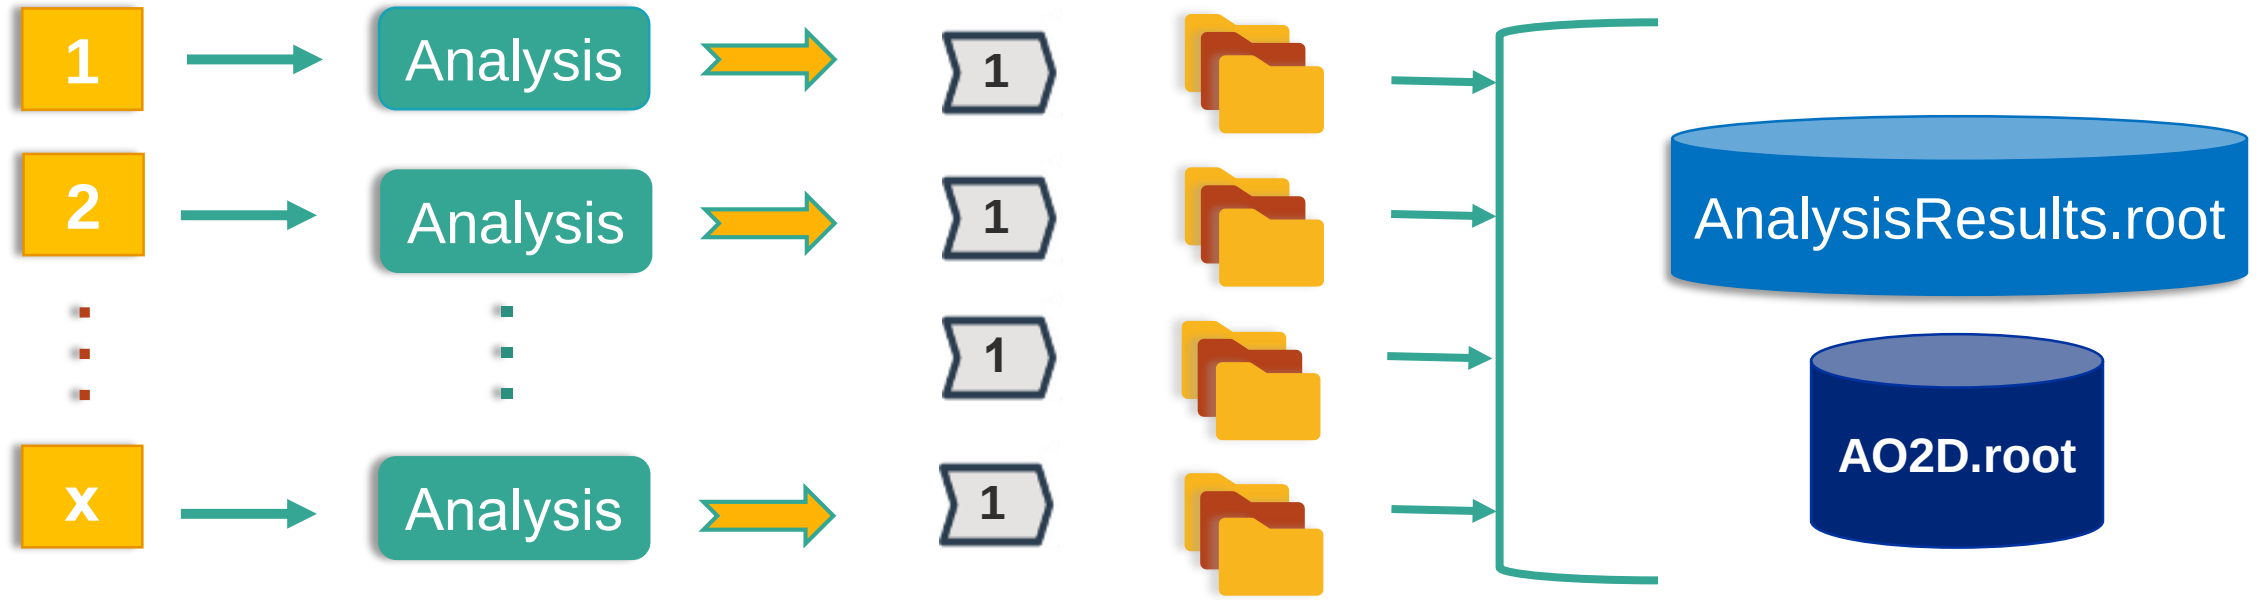

Linked Derived Data Train

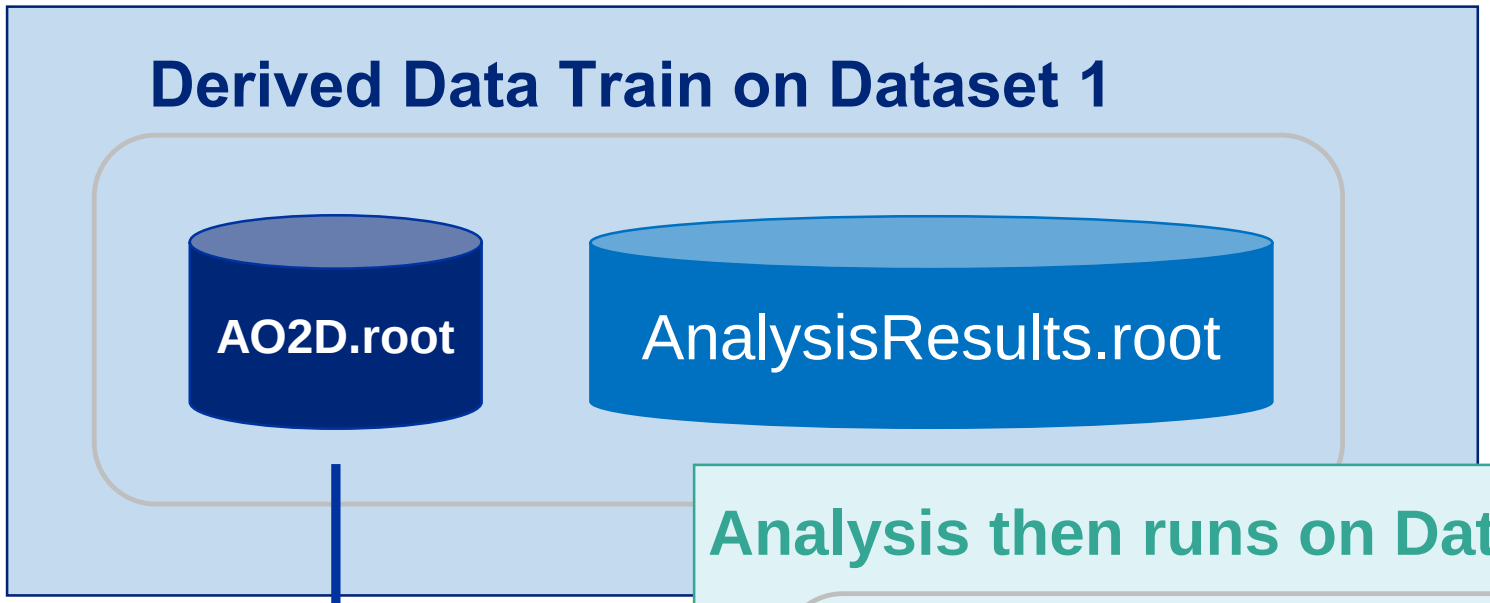

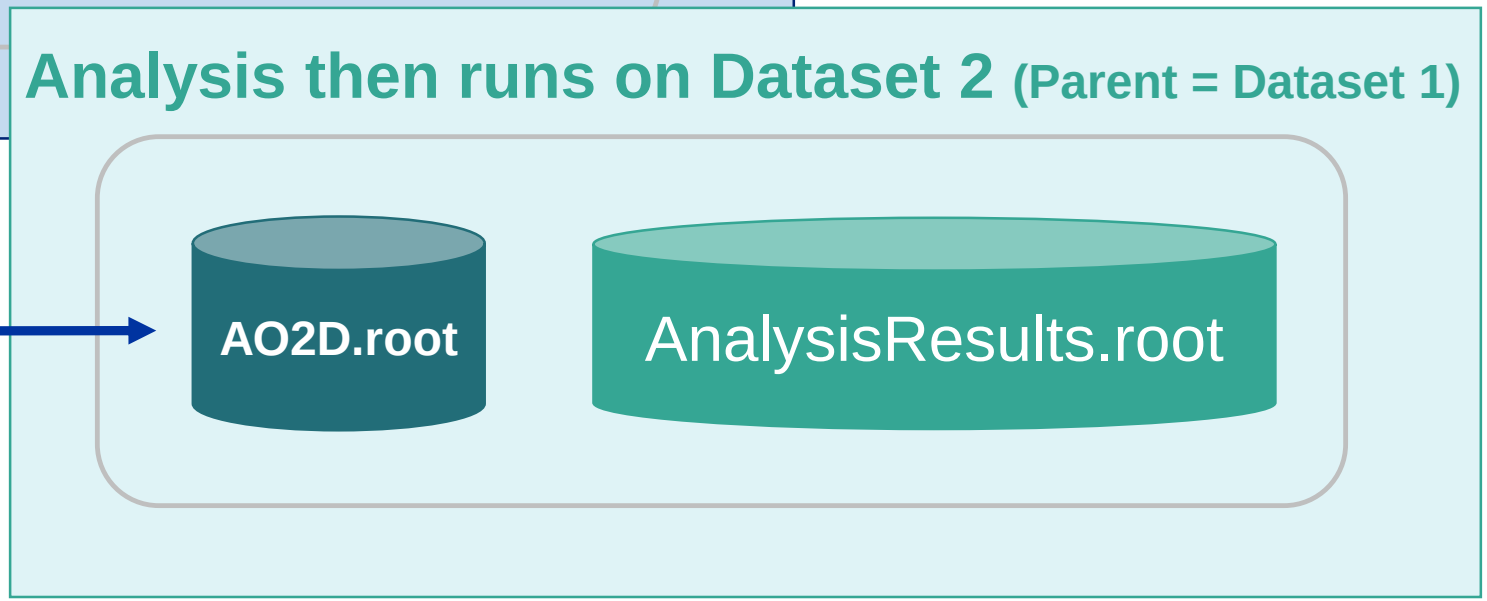

Train runs

The train type is decided by operators at composition in the Train Submission page

1. **Analysis Train** - is a standard analysis train and no derived data will be produced
2. **Slim Derived Data** - reserved for derived data of **small output size**
   - Similarly to the standard derived data case, this train **will produce derived data** to be used for further analysis
   - The **results will be merged across runs** and are **not available to use in future train runs**
   - The data will be **automatically deleted** after a pre-set period of time
3. **Standard Derived data** - **will produce derived data** to be used for further analysis
   - The **results will not be merged across runs** and can be used **as input for future train runs**
4. **Linked Derived data** - this option is for **derived data which needs to access its parent file when it is processed**
   - The derived data file produced will remember its parent files, inheriting also their storage location
   - The **results will not be merged across runs** and can be **used as input for future train runs**
   - Datasets composed from this train need to have parent access level activated

Train runs

	Specifics	Results merged across runs	Produces derived data	Can be used as input for future train runs
Analysis train		✓		
Standard derived data			✓	✓
Slim derived data	• For derived data of small output size • Data will be automatically deleted	✓	✓	
Linked derived data	• For derived data which needs to access its parent file when it is processed • Derived data file produced will remember its parent files, inheriting also their storage location • Datasets composed from this train need to have parent access level activated		✓	✓

New to Hyperloop? Take the tour!

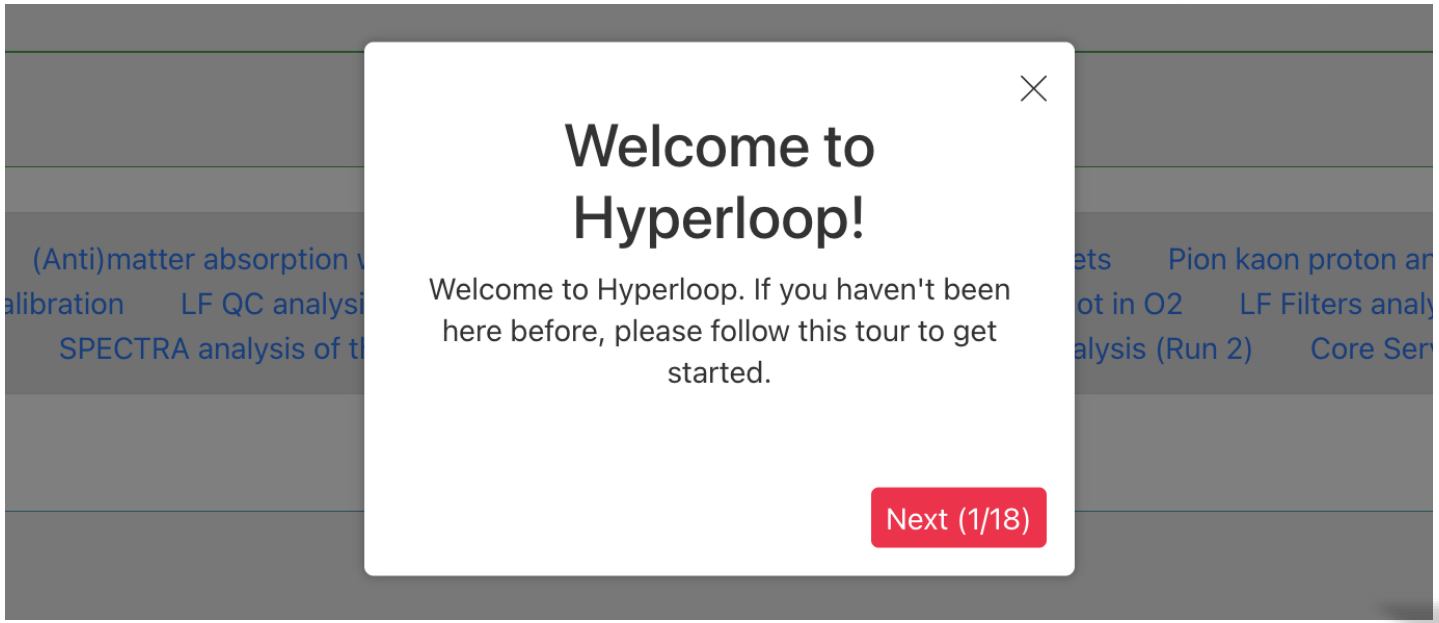

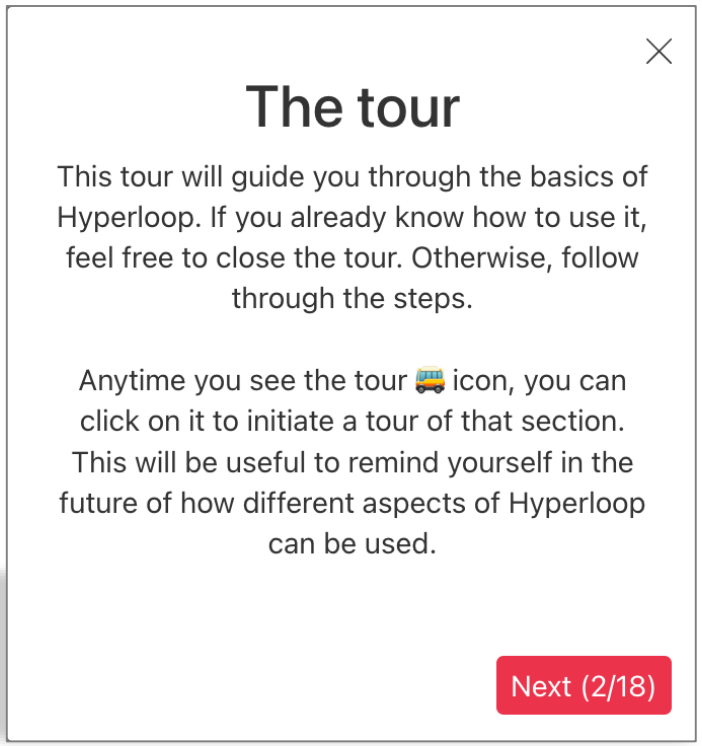

- When opening a page in Hyperloop that you did not visit before, a guided tour will explain key concepts
- Ideal for beginners, but also for refreshing knowledge ***(many tours offer direct links to relevant sections of the documentation)***
- Where appropriate, when one tour ends, the next will begin to explain the next section of Hyperloop
- Tours can be exited at any time. Once closed, they will not automatically begin on future page visits
- A tour can be retaken anytime by clicking

New to Hyperloop? Take the tour!

Currently available tours:

- My Analyses page
- Wagon view
- Wagon test view
- Train run view

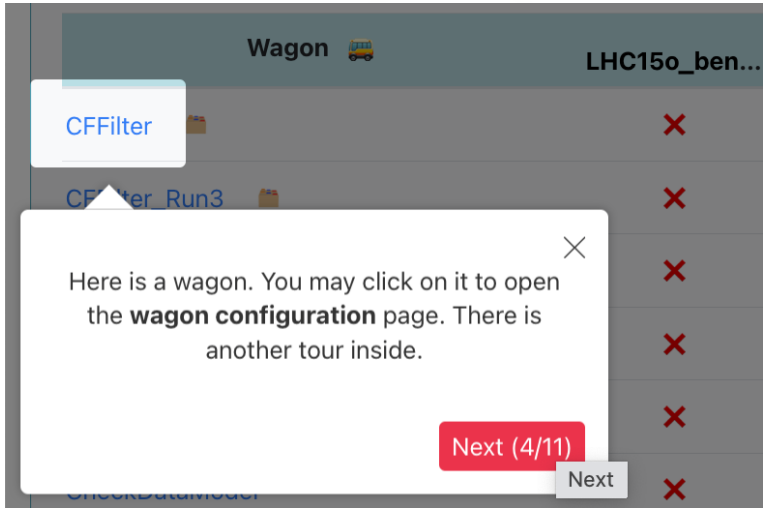

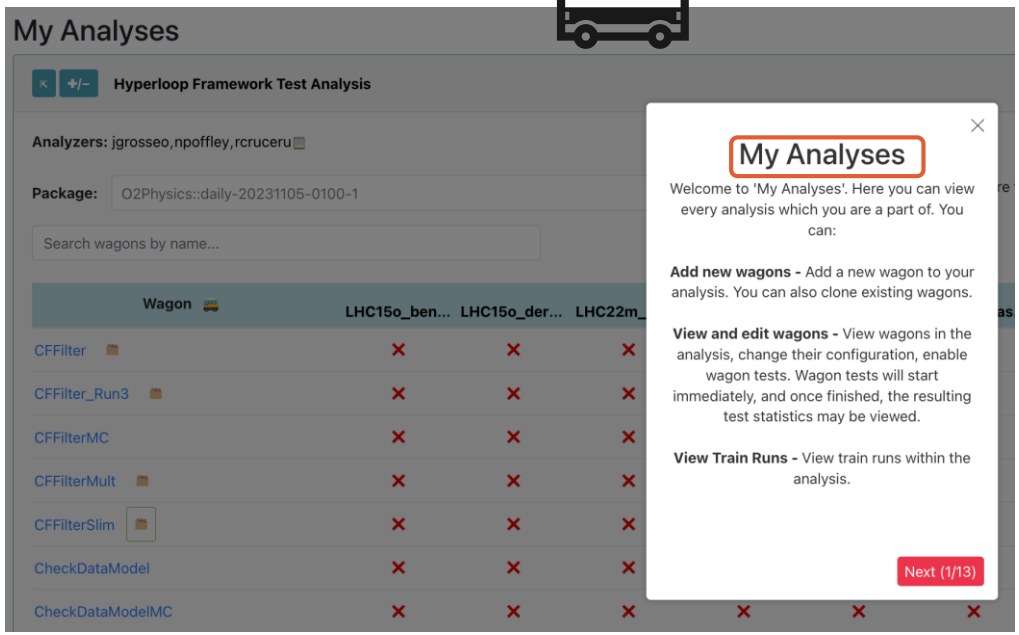

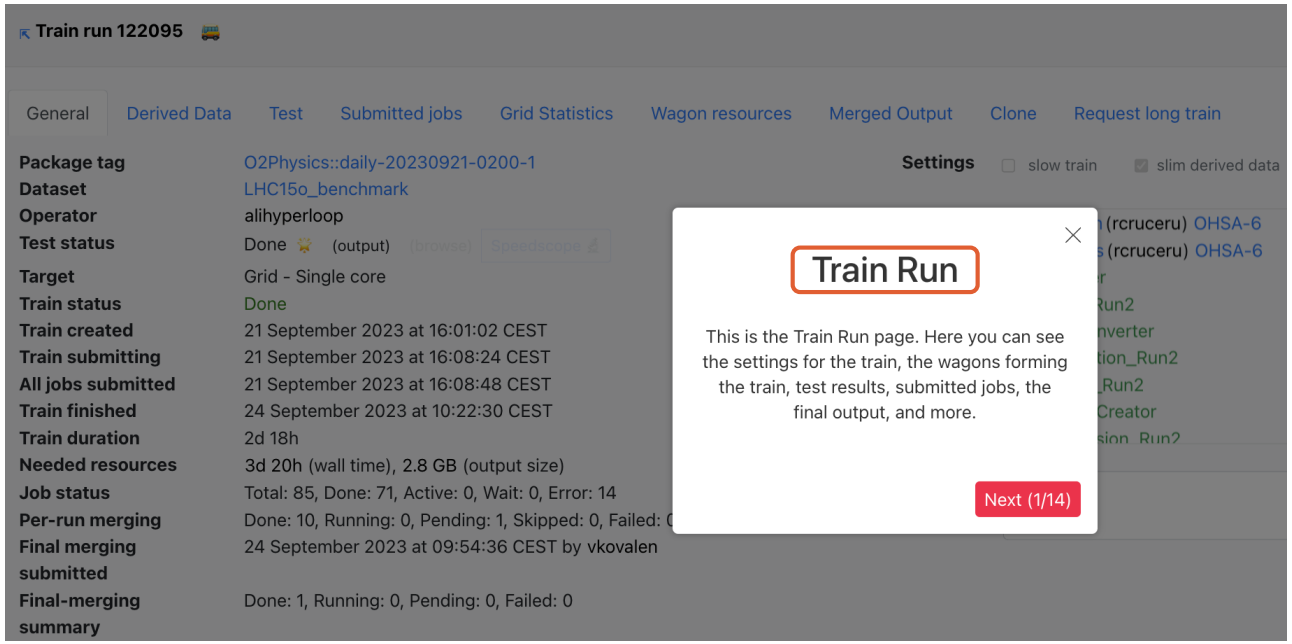

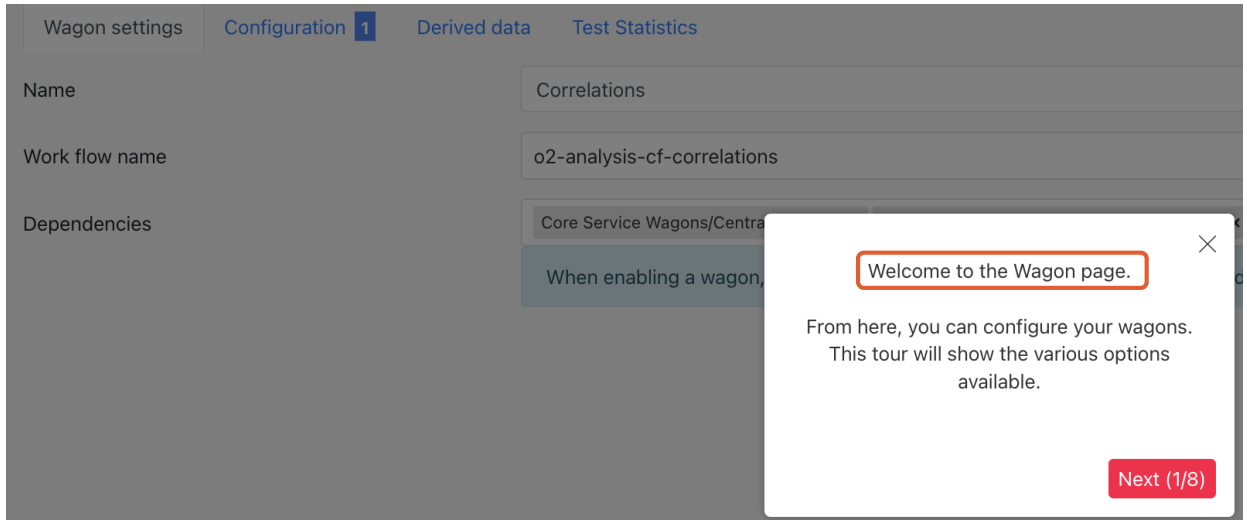

My Analyses

- Displays all the analyses the user belongs to
- The analyses are defined in JIRA of the respective PWG, along with the users that will be part of them
- Several analyzers can share an analysis
- Datasets are enabled per analysis
- The user decides the package tag

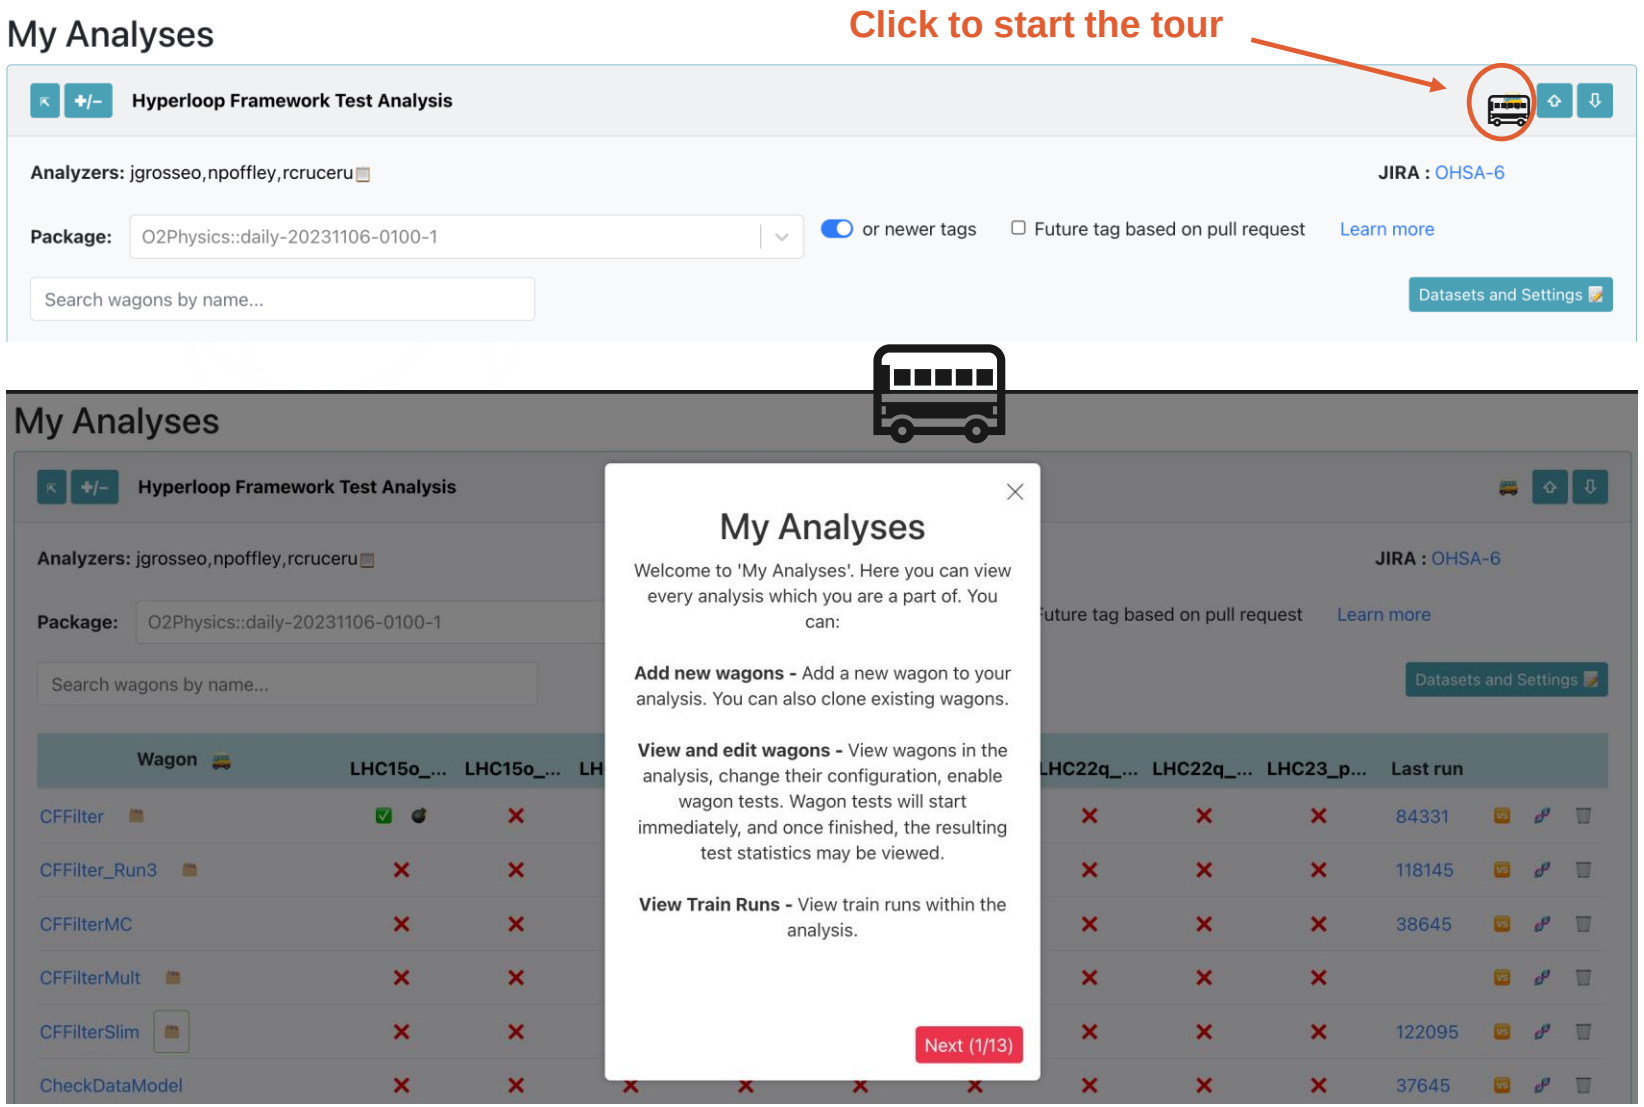

My Analyses

- **Displays all the analyses the user belongs to**
- **The analyses are defined in JIRA of the respective PWG, along with the users that will be part of them**
- **Several analyzers can share an analysis**
- Datasets are enabled per analysis
- The user decides the package tag

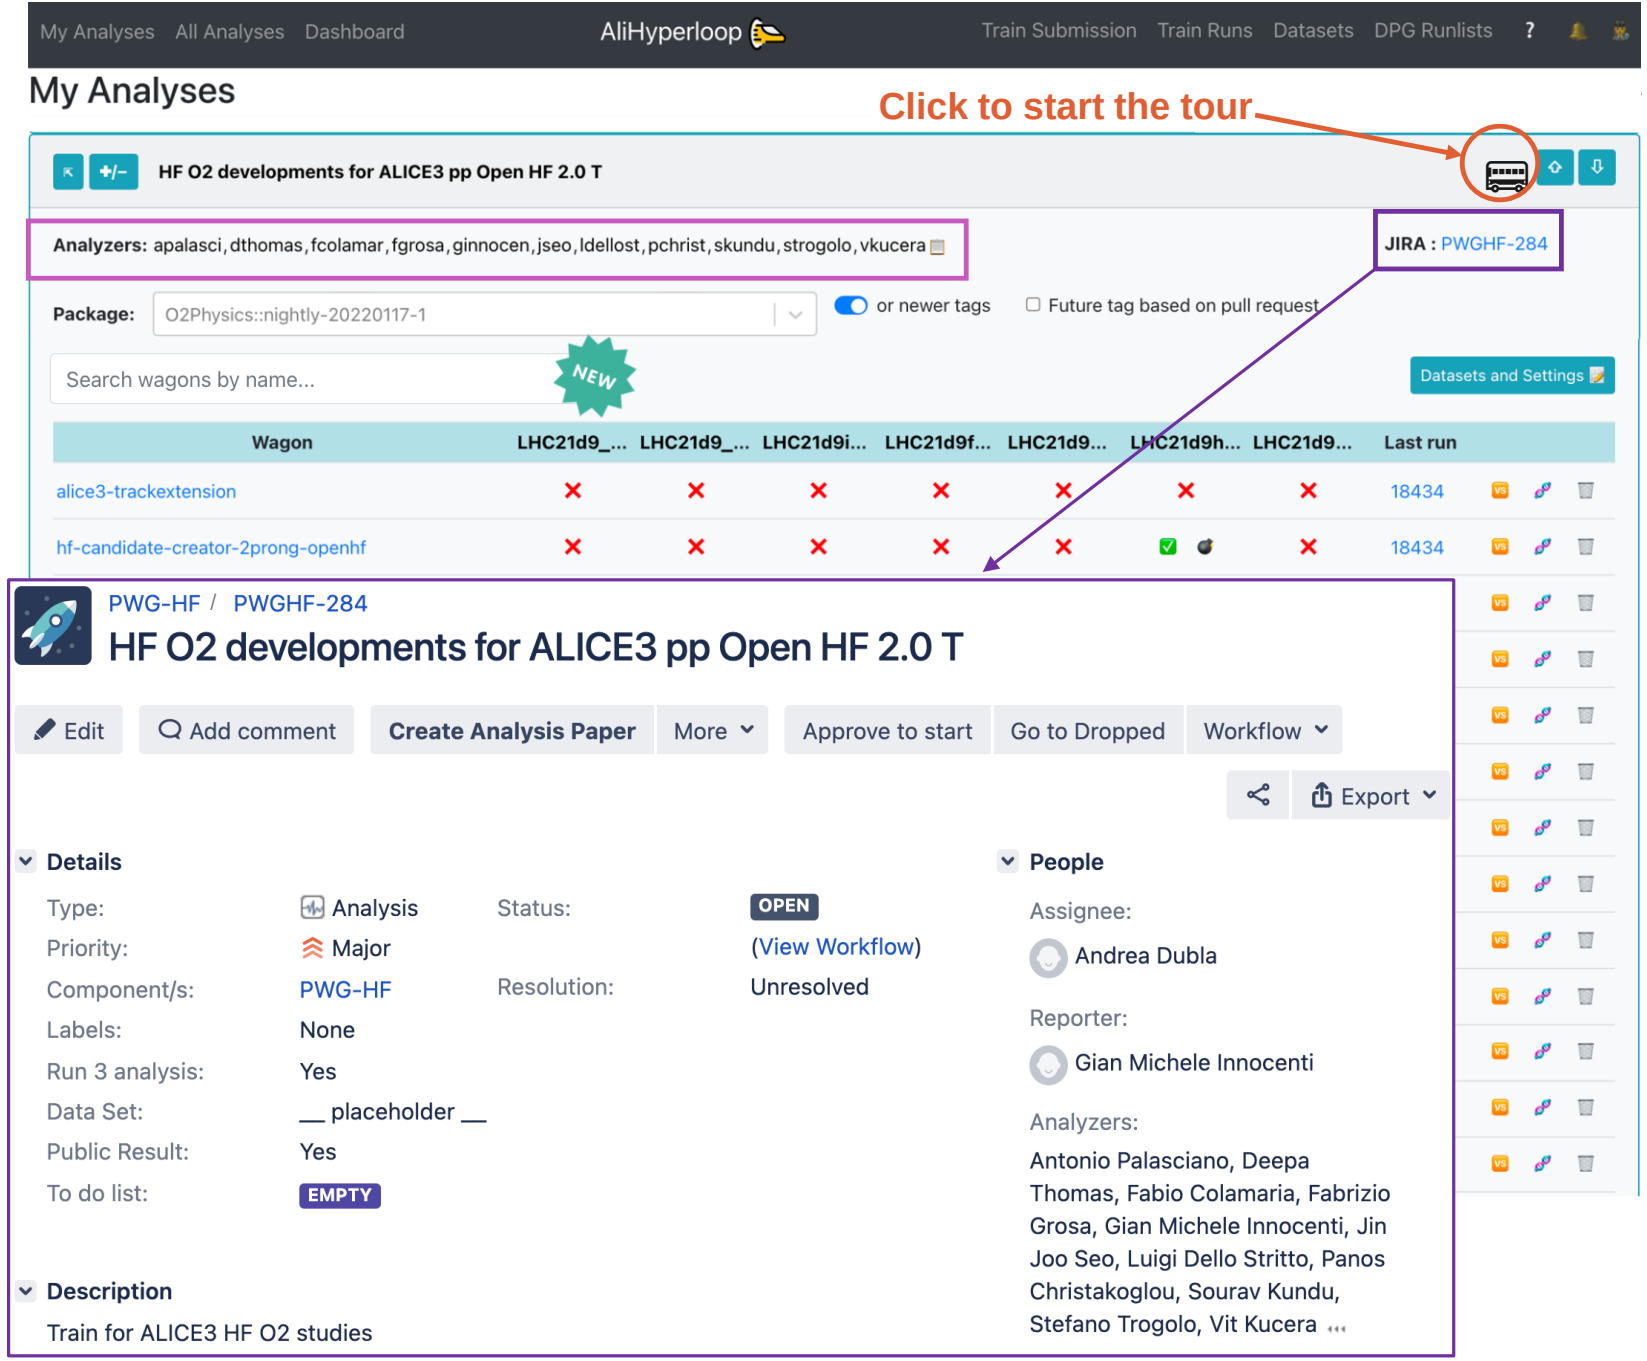

My Analyses

- Displays all the analyses the user belongs to
- The analyses are defined in JIRA, along with the users that will be part of them
- Several analyzers can share an analysis
- **Datasets are enabled per analysis**
- The user decides the package tag

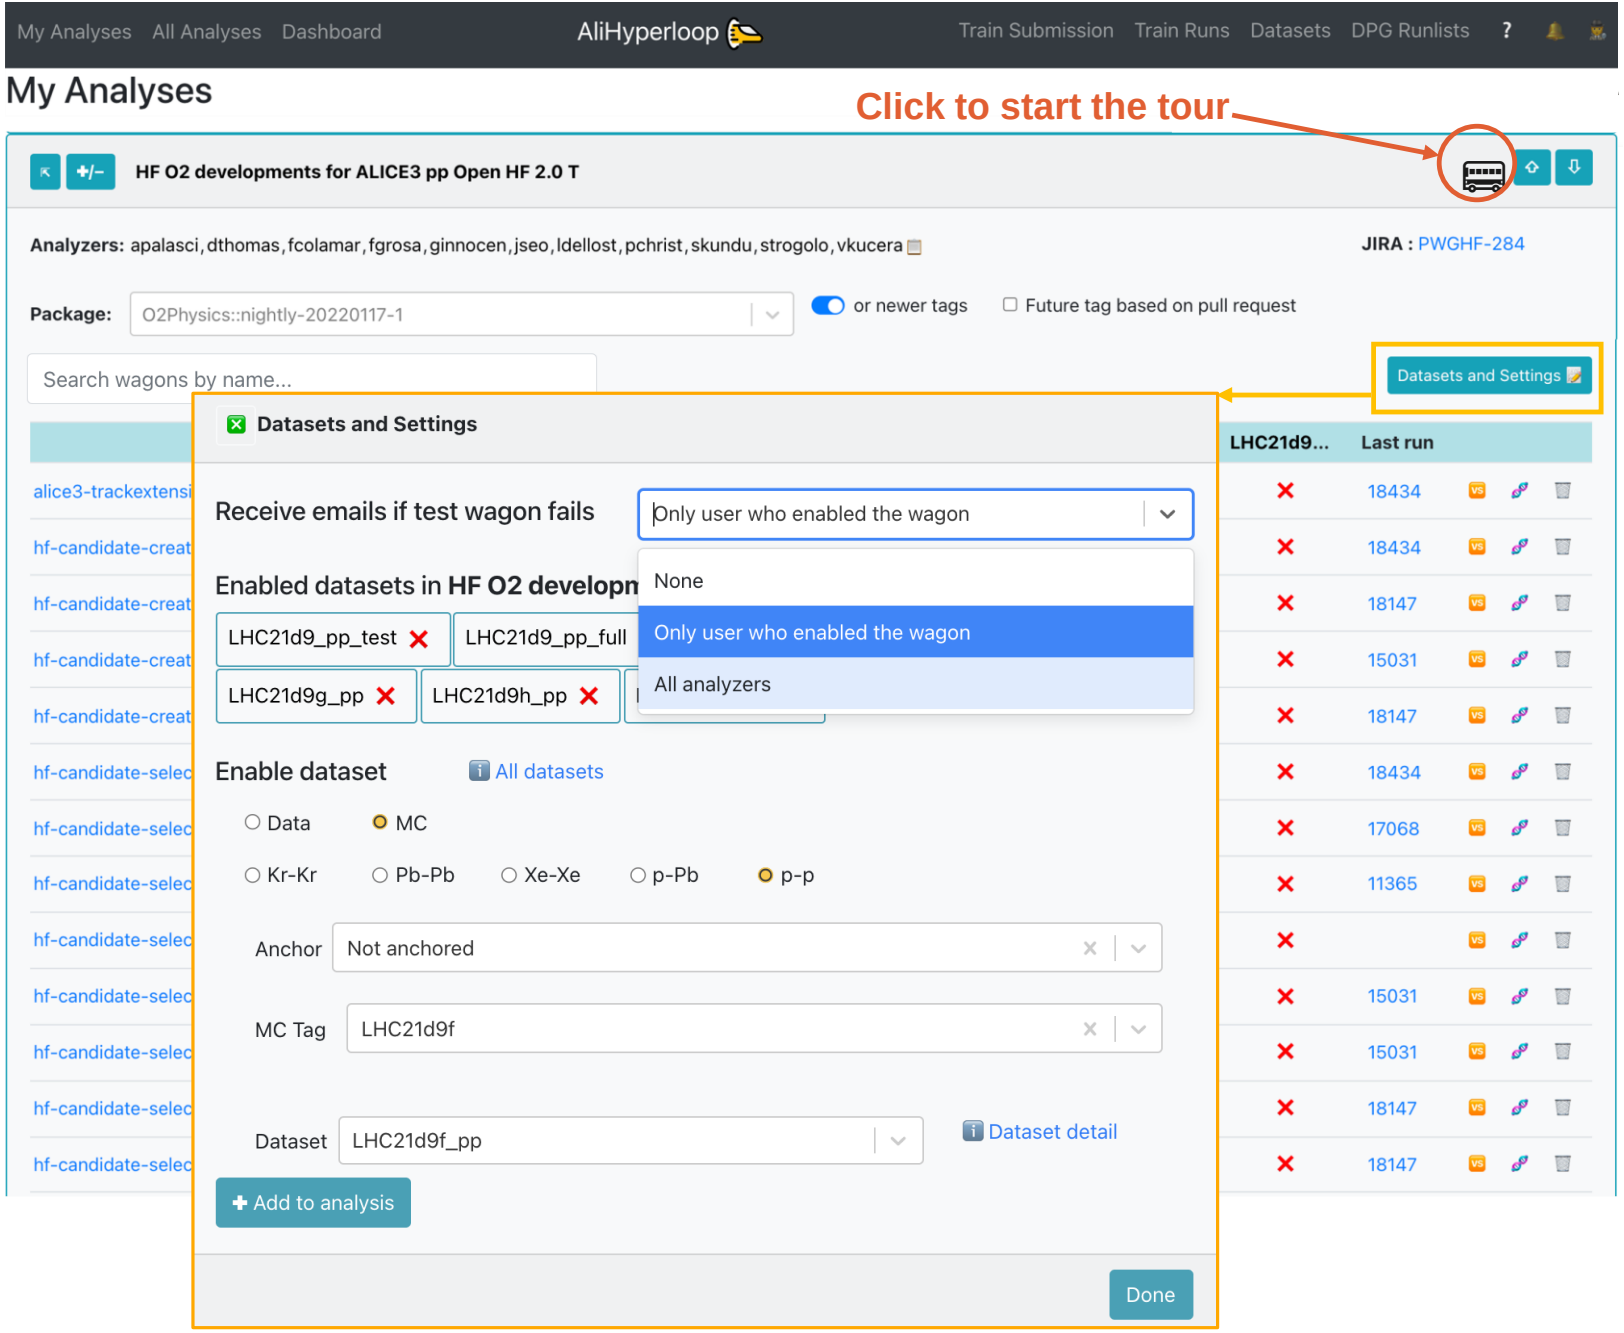

My Analyses

- Displays all the analyses the user belongs to
- The analyses are defined in JIRA, along with the users that will be part of them
- Several analyzers can share an analysis
- Datasets are enabled per analysis
- **The user decides the package tag**

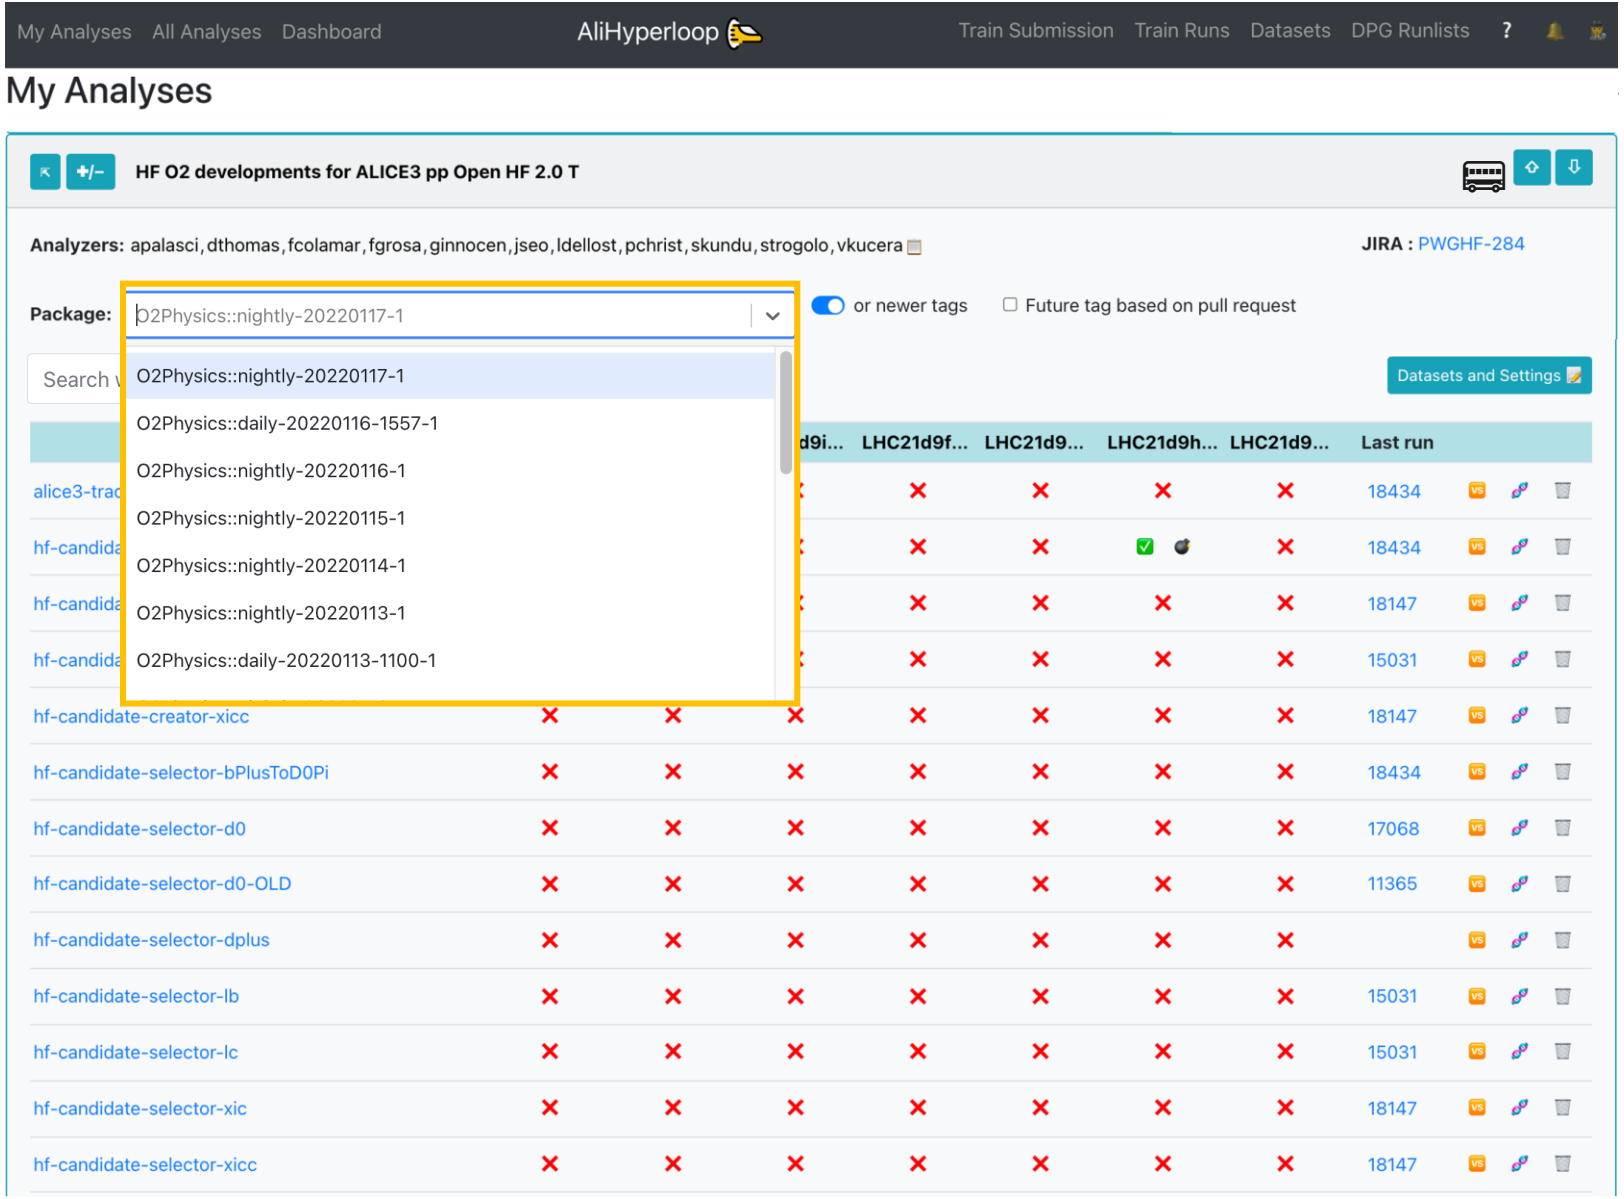

My Analyses

Deciding on package tag:

- Choose a tag from the ***dropdown*** and the wagon will be tested only with that specific tag
- Selecting ***or newer tags*** will allow the wagon to be included in a train run that has the previously selected tag or a **newer tag** *(this is useful to group analysis together in a train)*

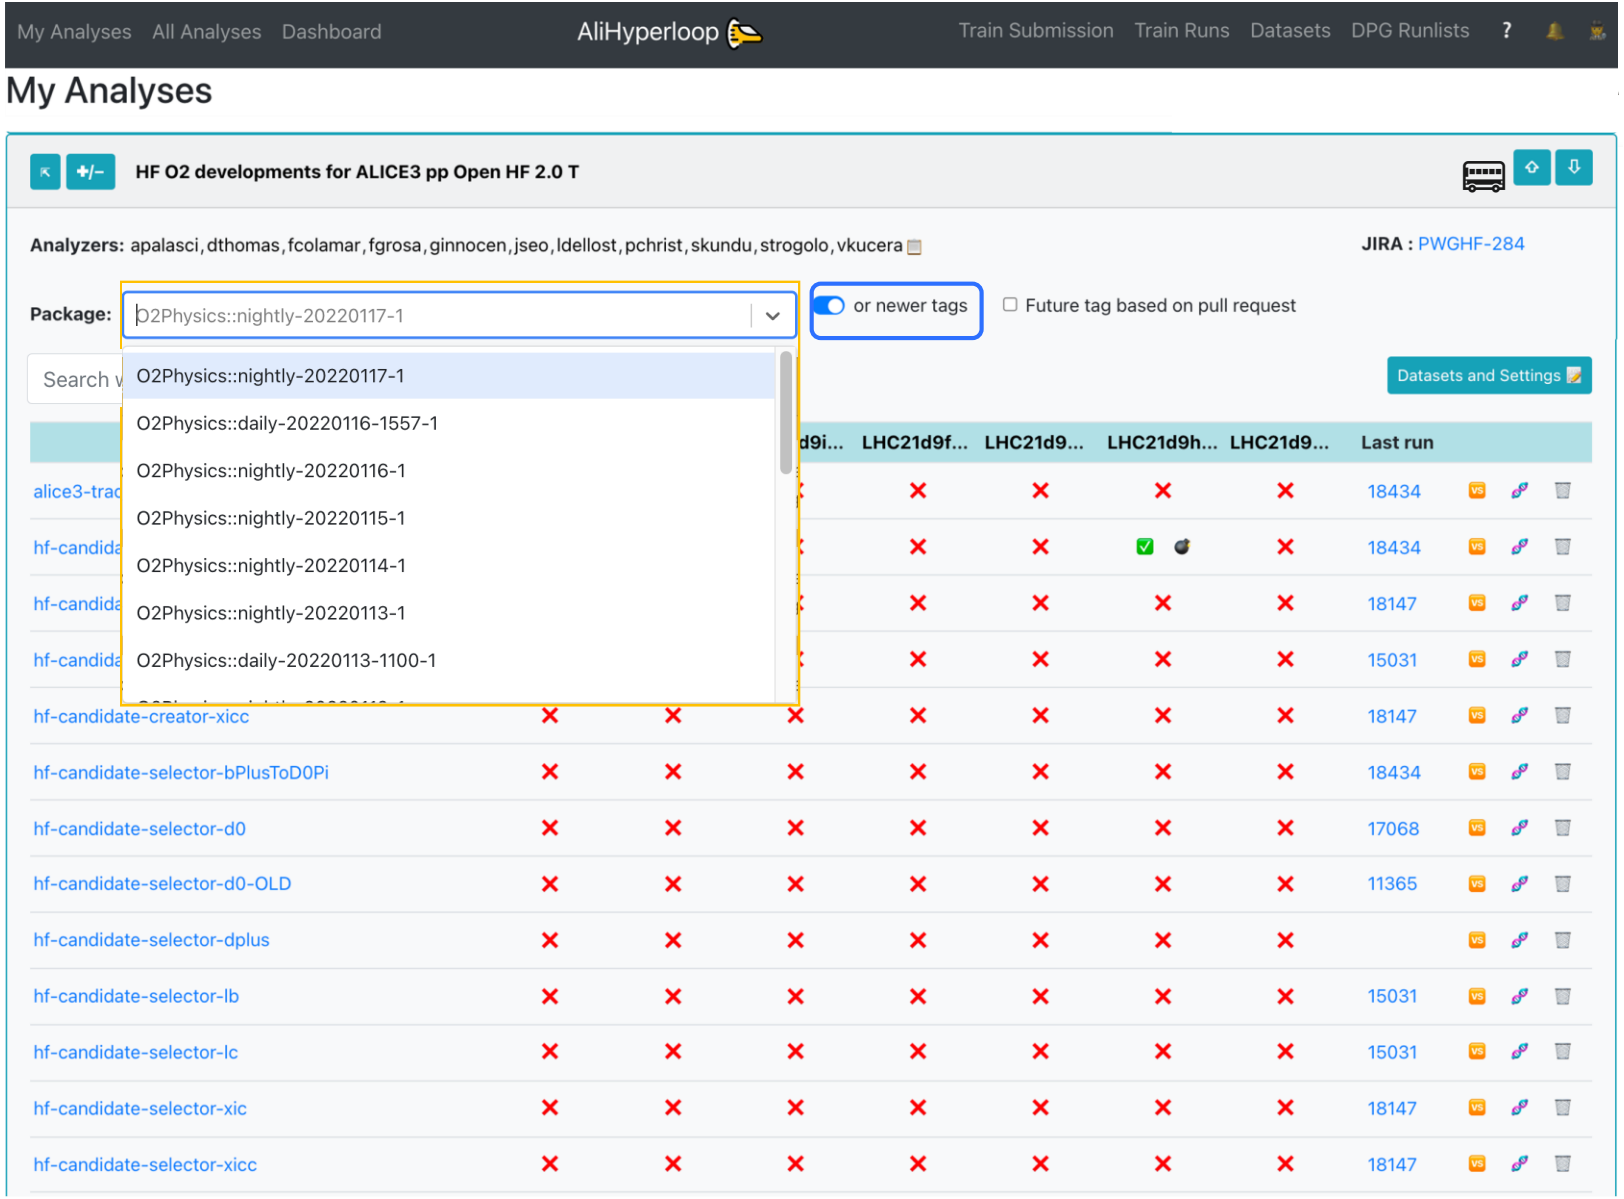

My Analyses

Deciding on package tag:

- Choose a tag from the ***dropdown*** and the wagon will be tested only with that specific tag
- Selecting ***or newer tags*** will allow the wagon to be included in a train run that has the previously selected tag or a **newer tag** *(this is useful to group analysis together in a train)*
- ***Future tag based on PR*** **:**
  - The wagon is tested and runs once the *pull request is included in a tag*
  - Once the PR is part of the tag, the wagon can be tested on that tag or a future tag

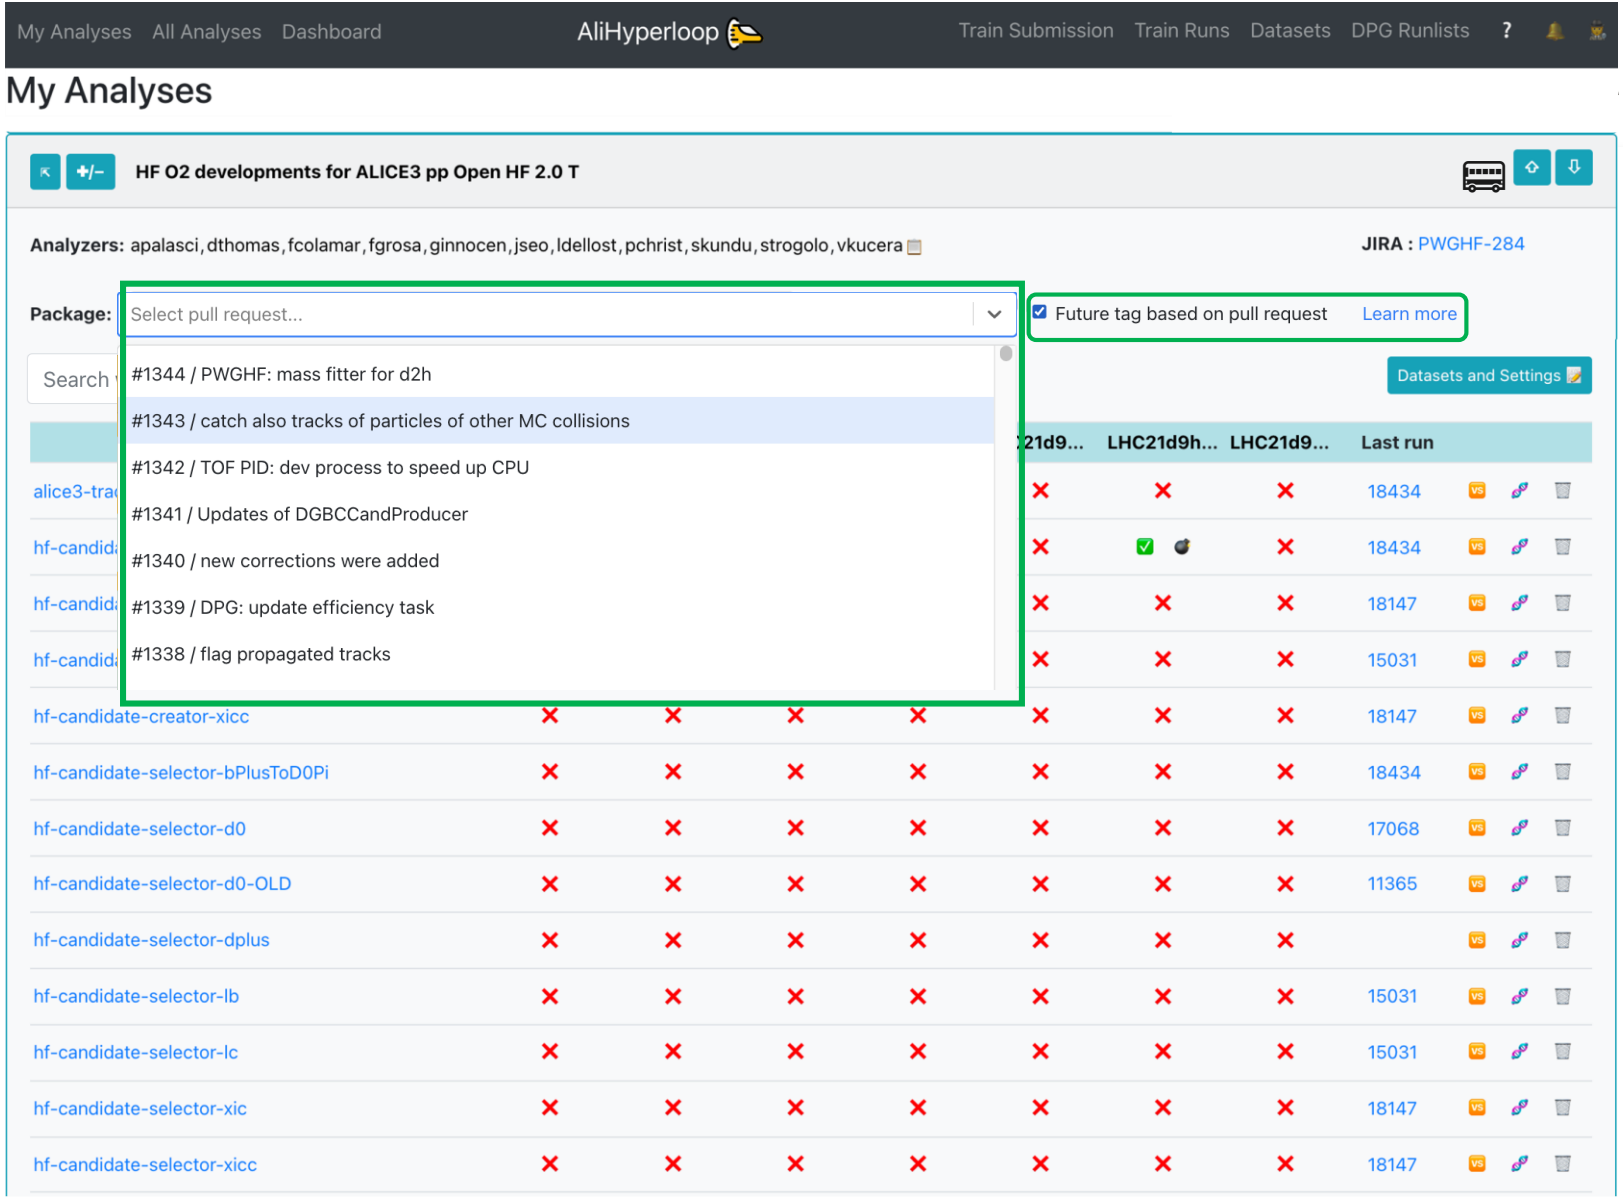

My Analyses

- Wagons defined from O^2 workflows
- Displays the list of train runs per analysis

- The user can:
  - Add/remove wagons
  - Enable/disable wagon > immediate wagon test
  - Clone/compare wagons

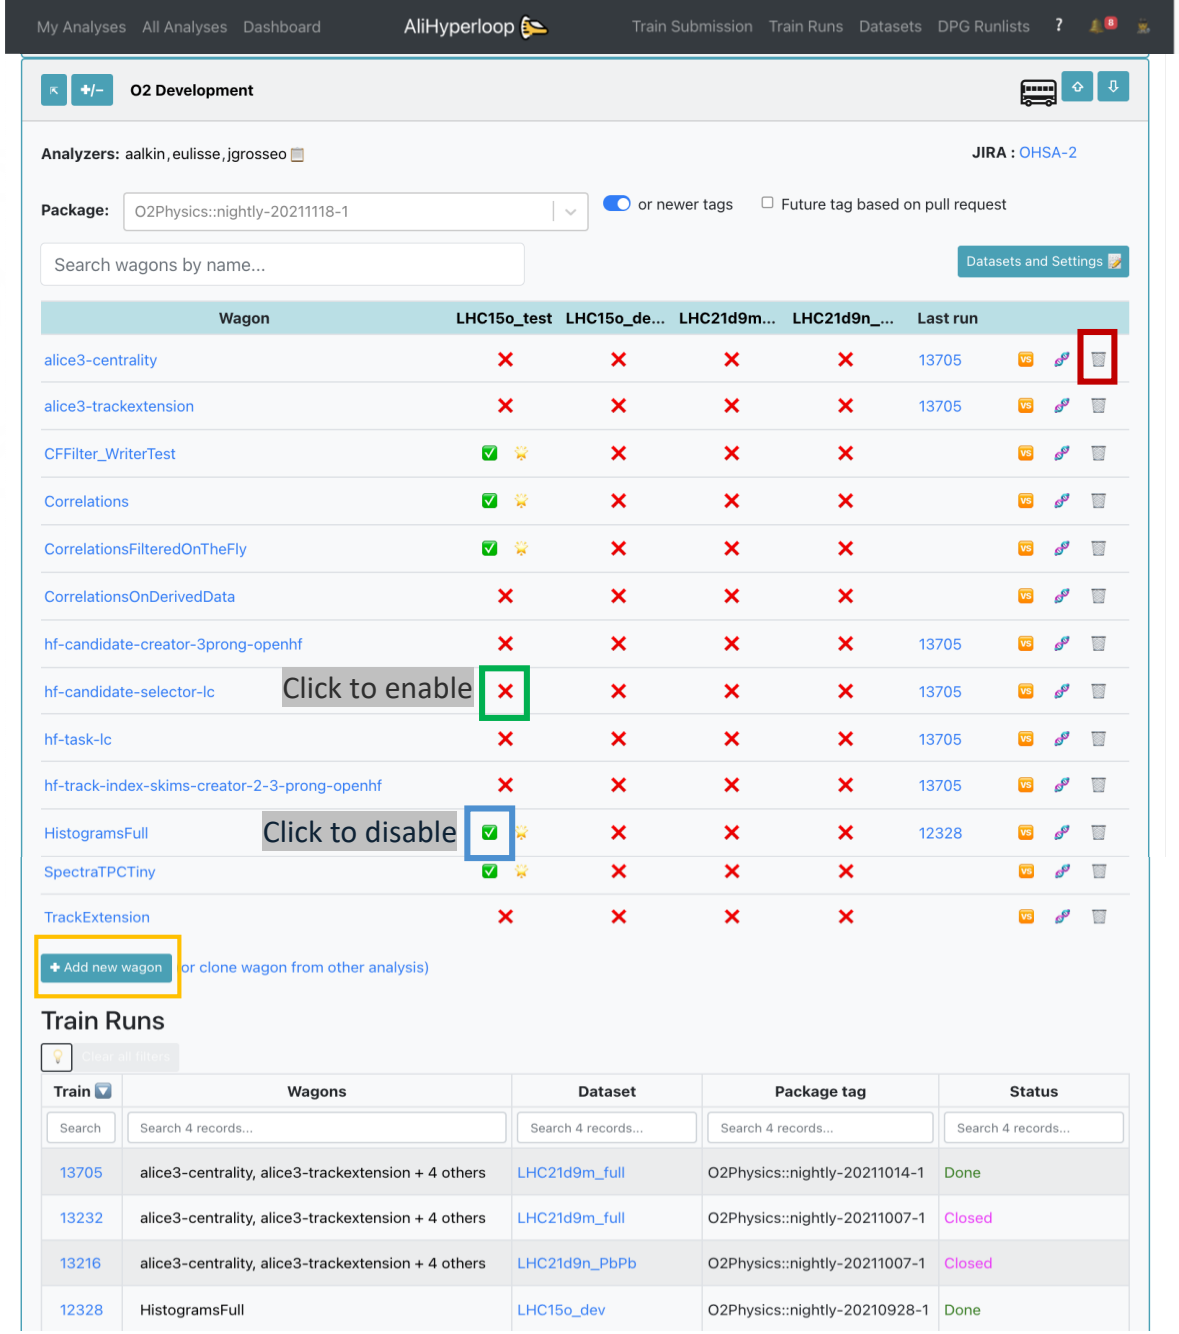

My Analyses

- Wagons defined from O^2 workflows
- Displays the list of train runs per analysis

- The user can:
  - Add/remove wagons
  - Enable/disable wagon
  - Clone/compare wagons

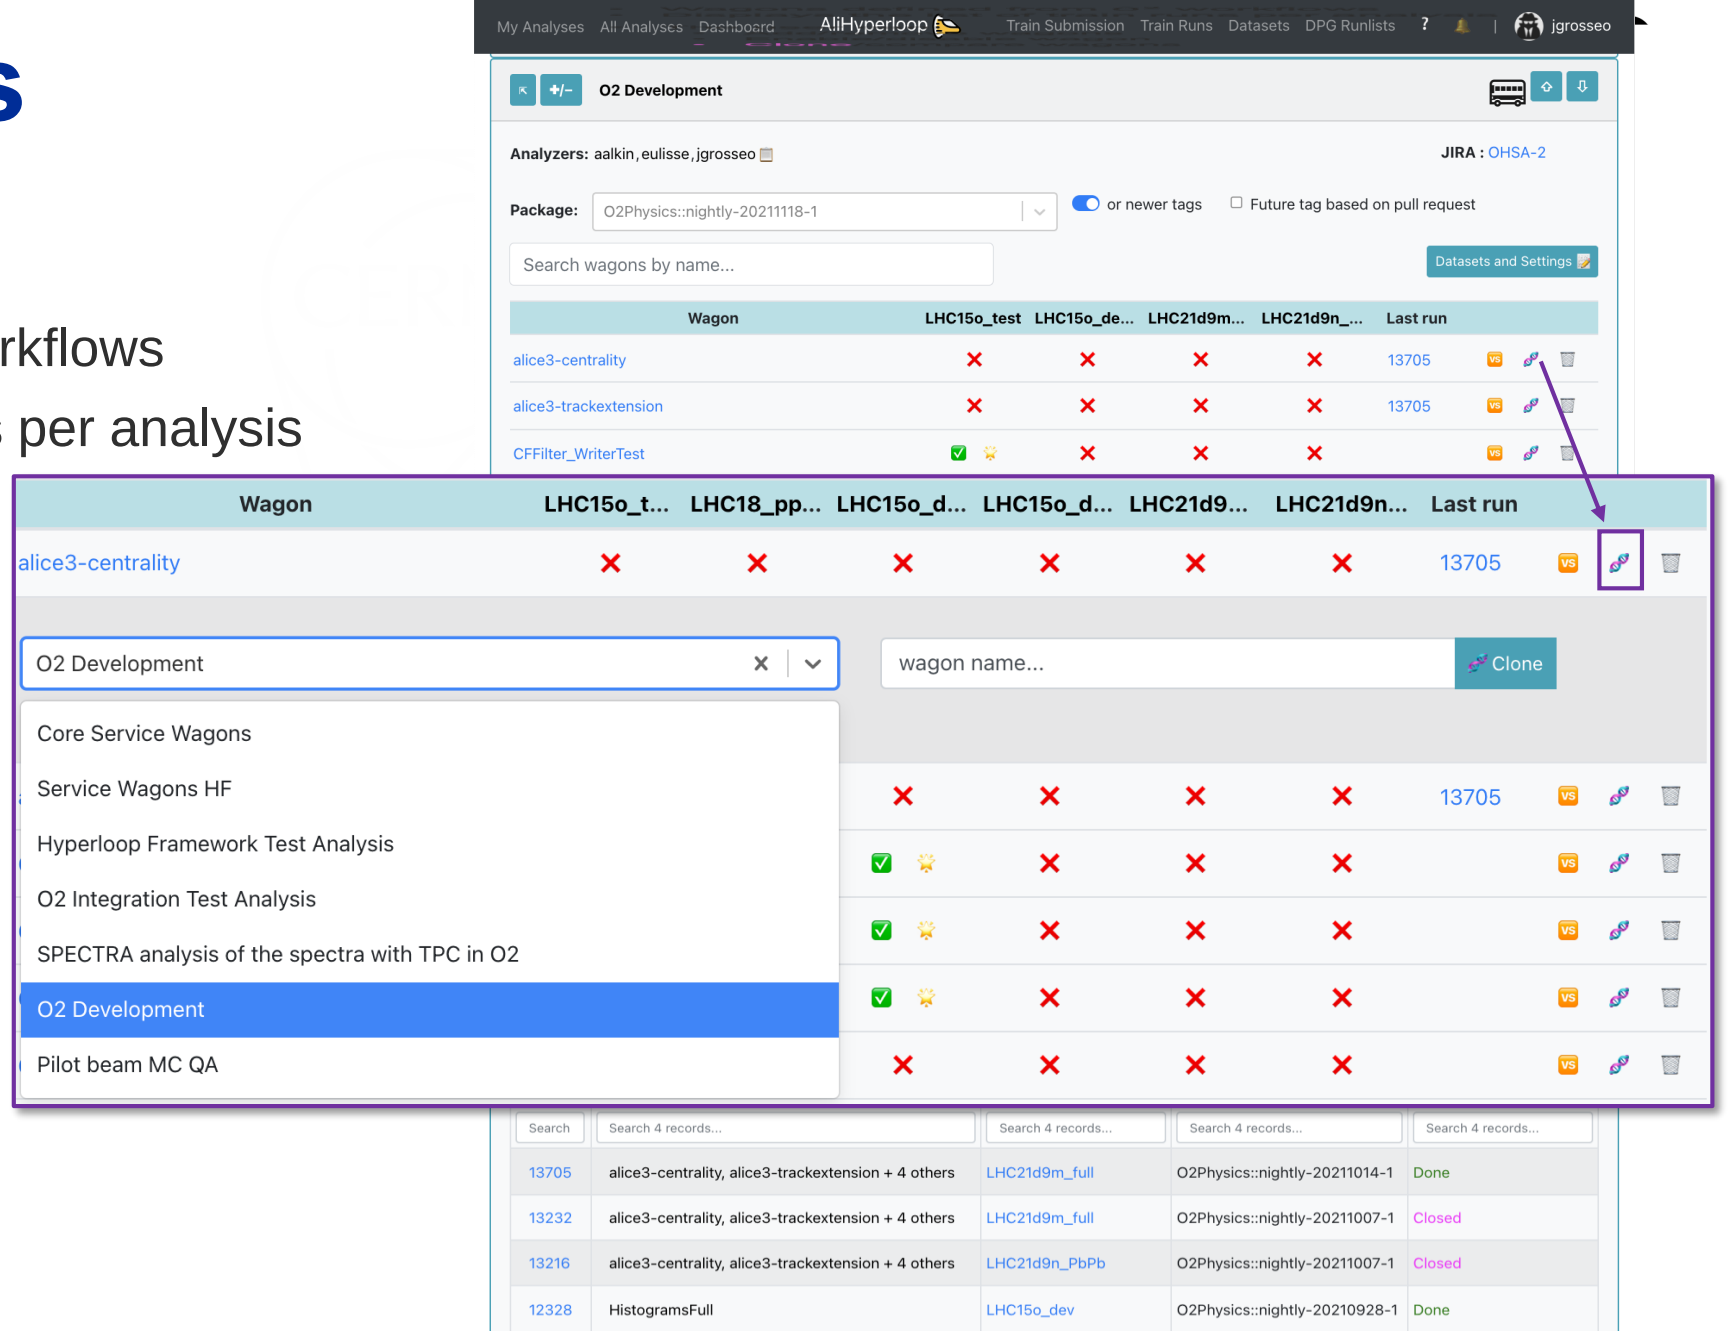

My Analyses

- Wagons defined from O^2 workflows
- Displays the list of train runs per analysis

- The user can:
  - Add/remove wagons
  - Enable/disable wagon
  - Clone/compare wagons

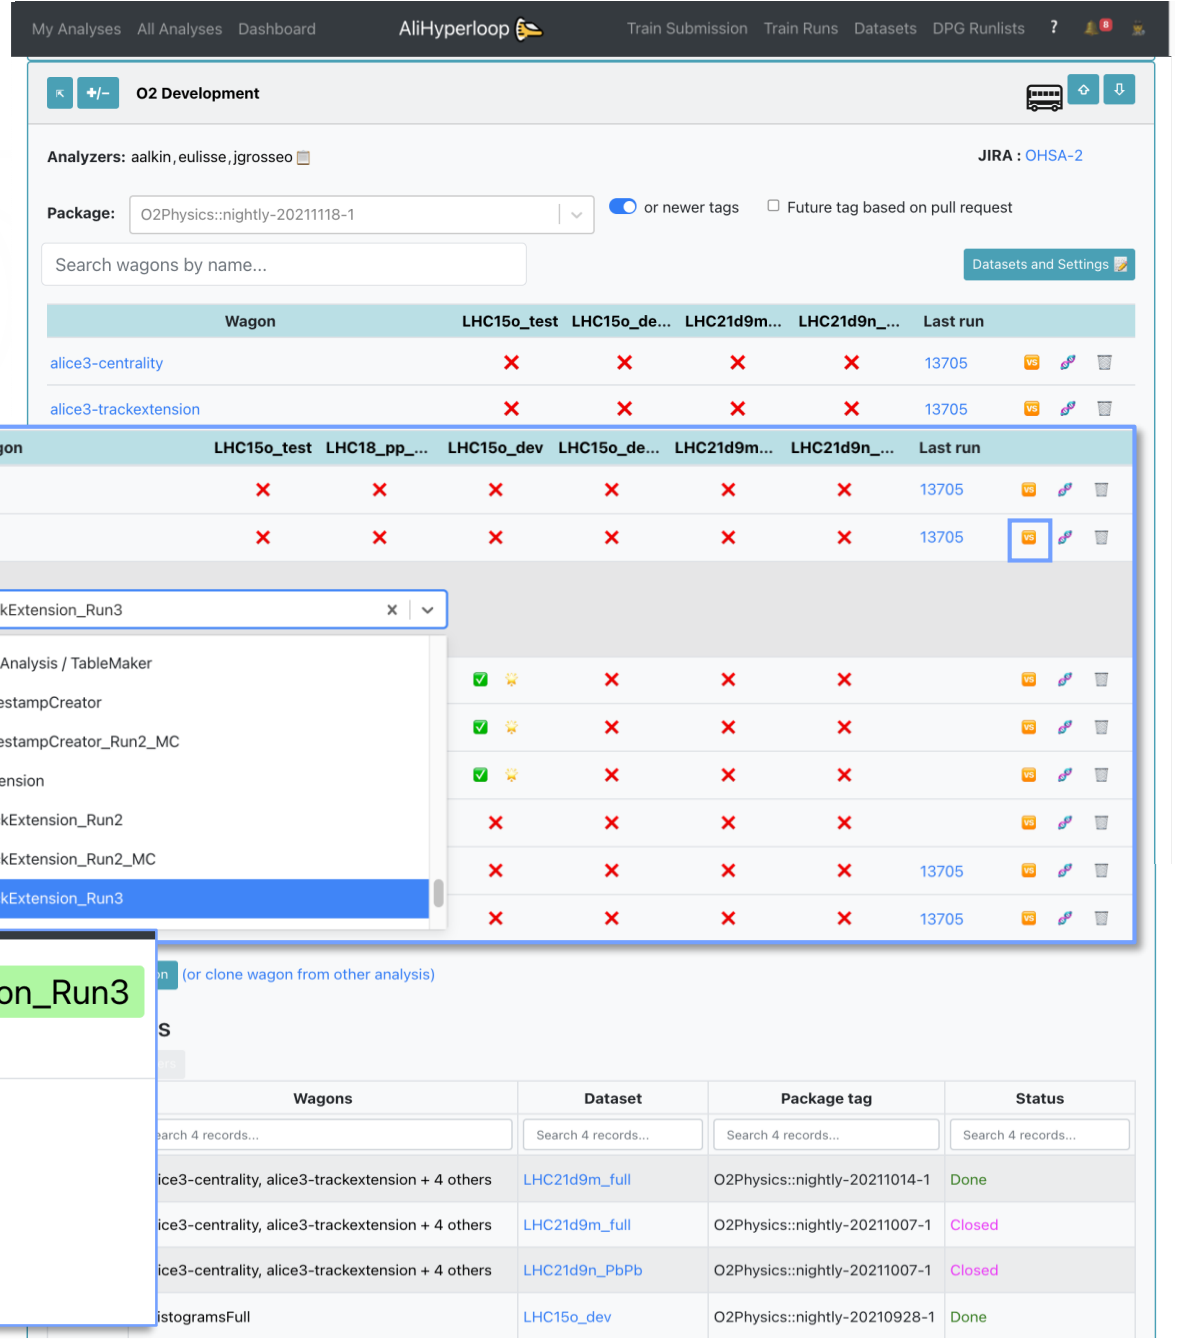

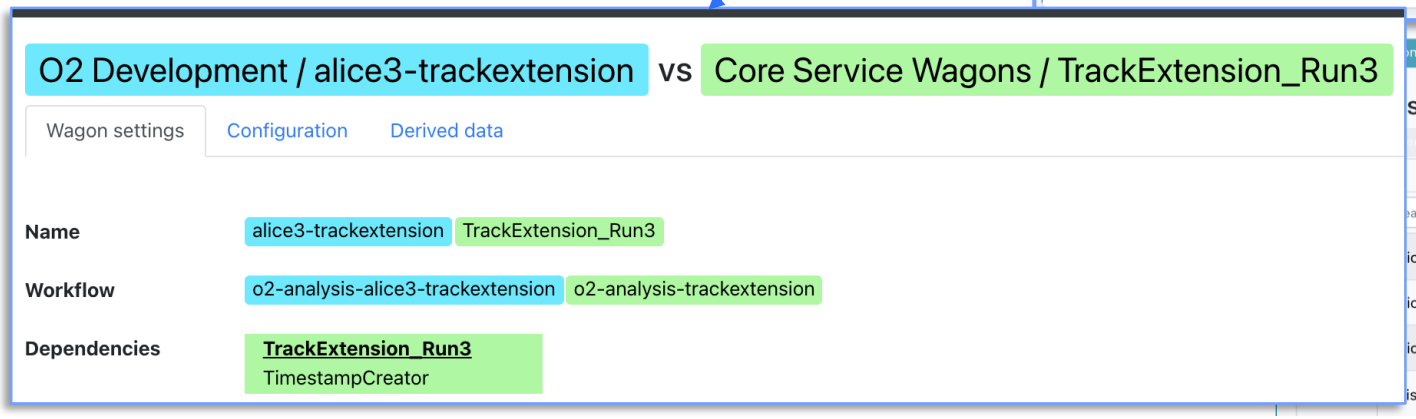

How to create a wagon

- Go to **My Analyses** page
- Click **Add new wagon** within your analysis
- This will open a pop-up window:
  - Insert wagon name
  - Choose Package tag
  - Select Workflow name
  - Click Save to create your new wagon

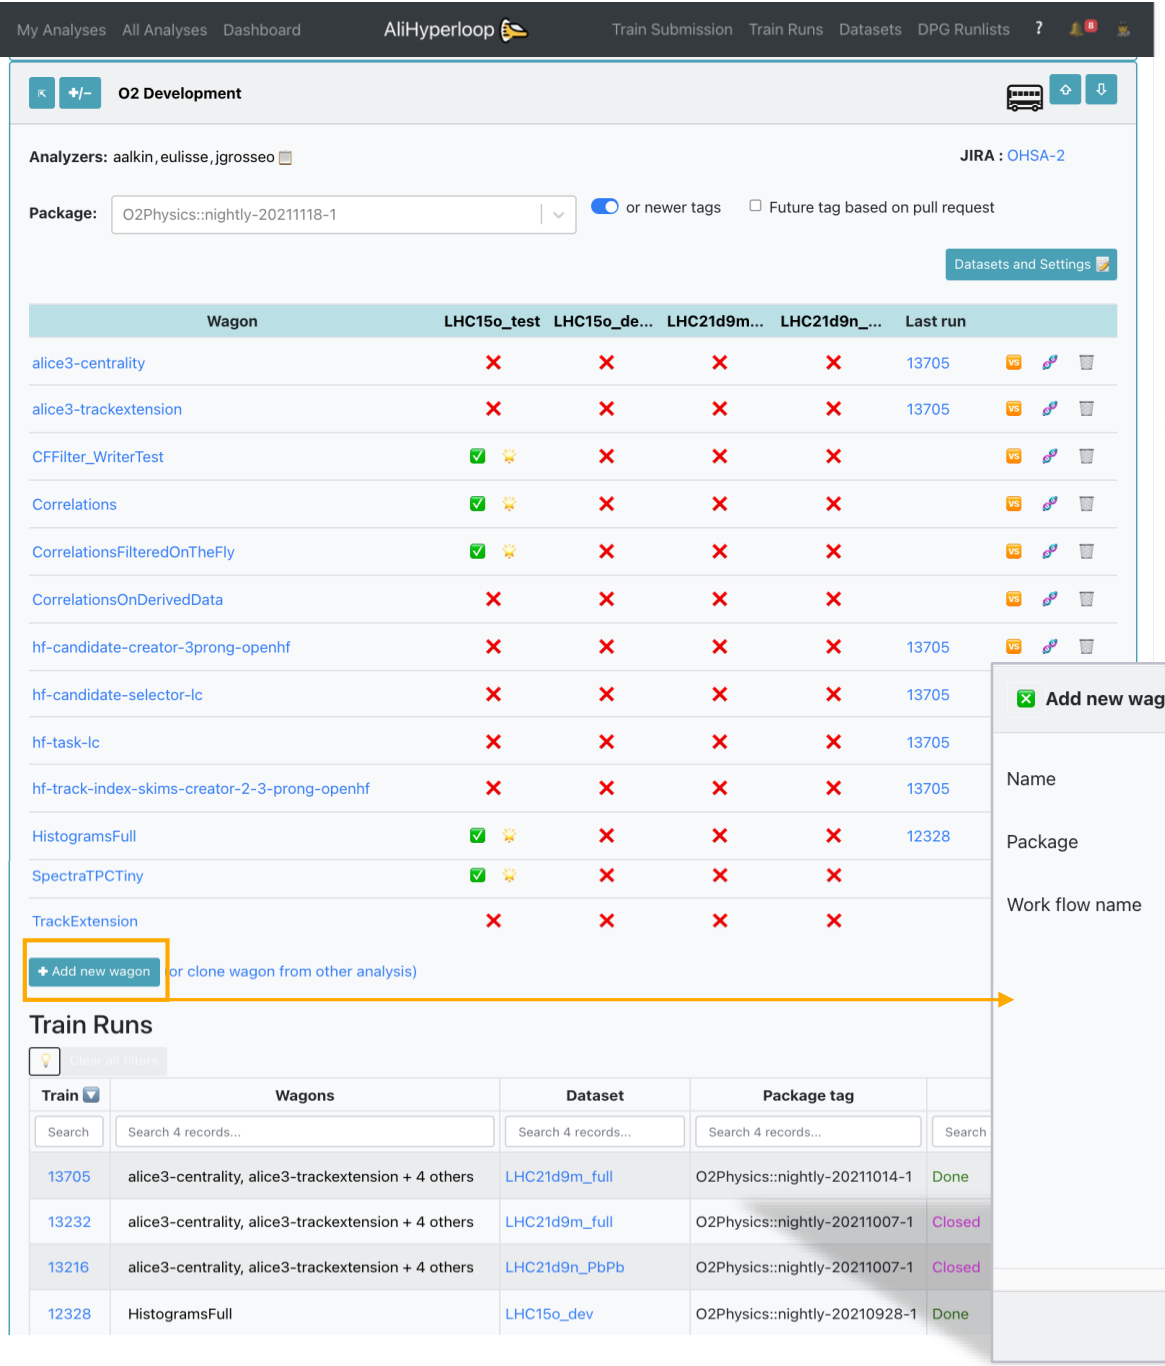

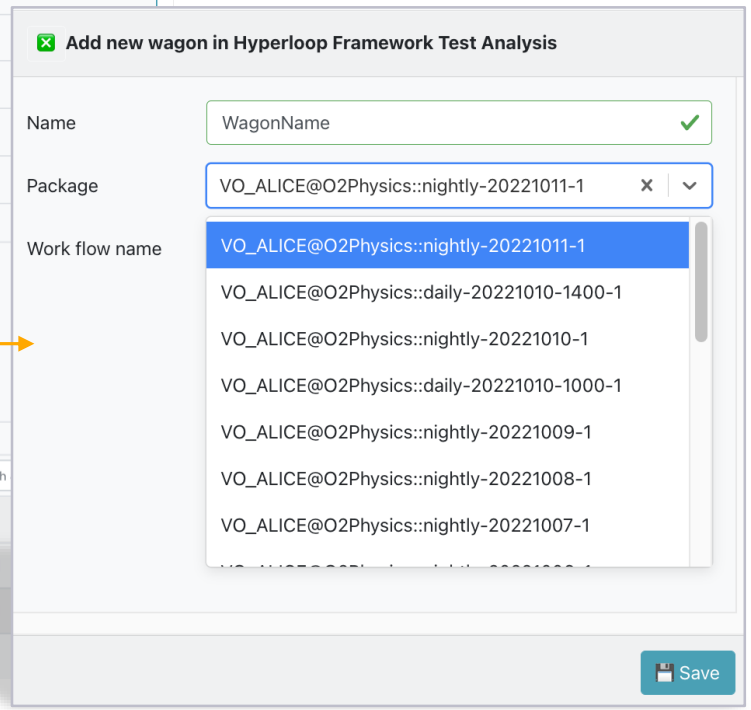

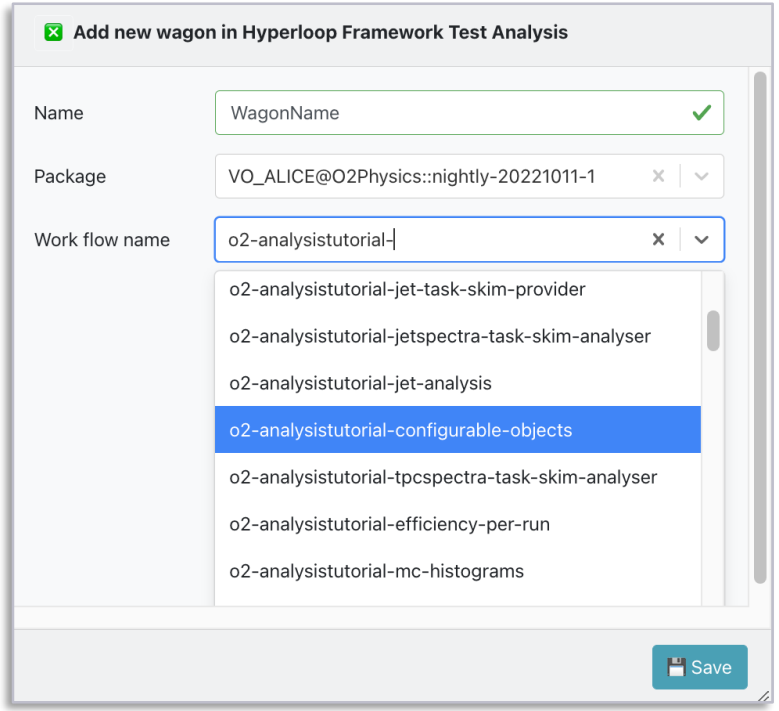

Edit Wagon

Click to start the tour

Documentation

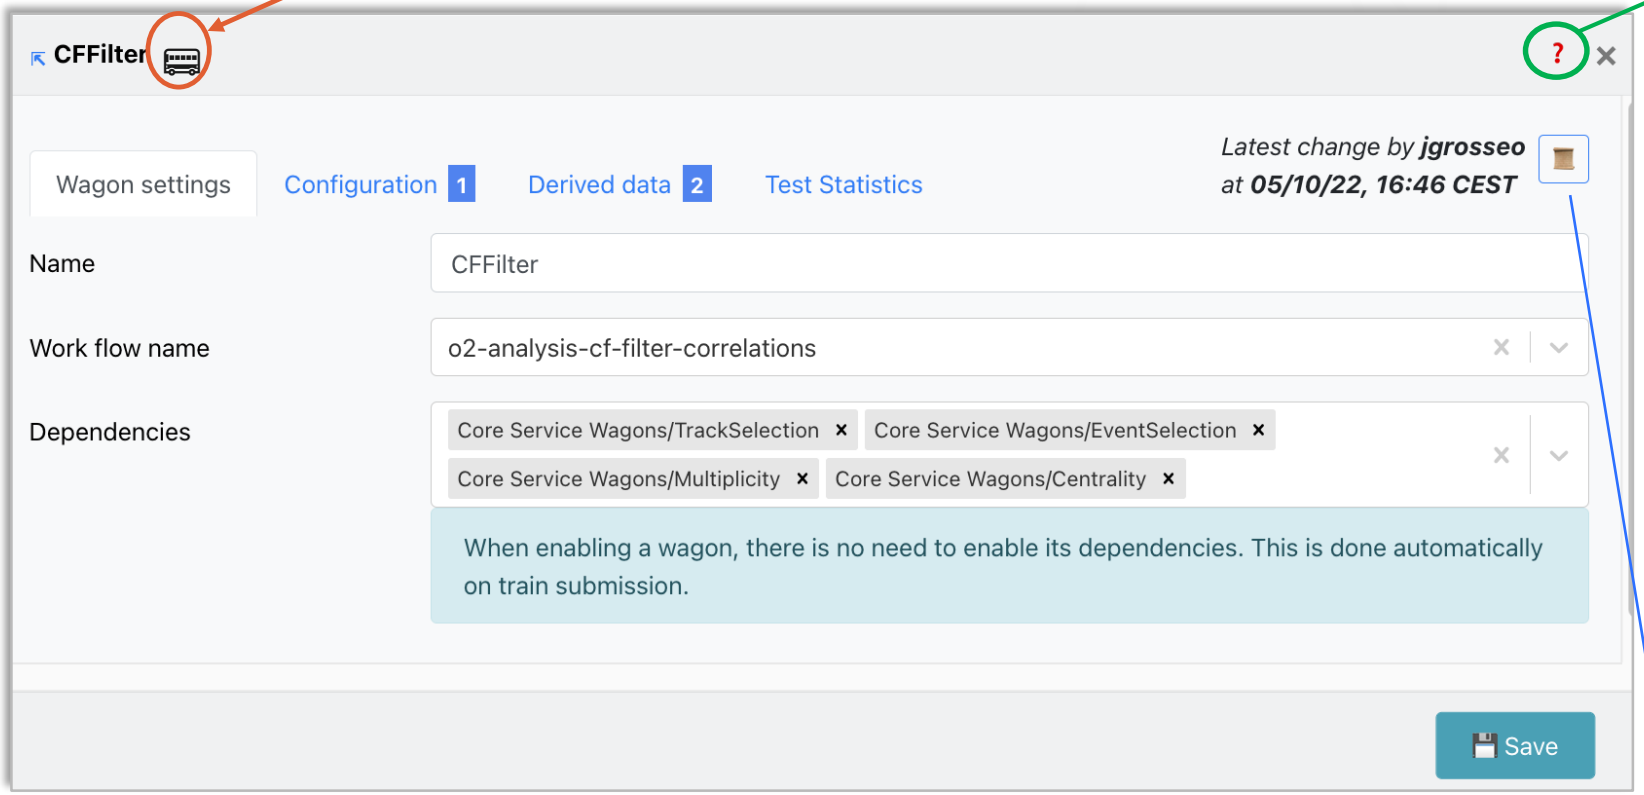

You can edit a wagon by clicking on the *Wagon* name in *My Analyses* page.

- Analyzers who are part of the analysis can add/edit/enable a wagon
- Add dependencies from service wagons or wagons from the same analysis
  - **Dependencies are automatically enabled when enabling a wagon** (this is different from the LEGO trains)
- Can open the detailed **wagon changelog** in a new tab (allows detailed comparison of wagon configuration for different timestamps)

- Direct link to a read-only view in a new tab: https://alimonitor.cern.ch/hyperloop/view-wagon/55
- Used to send to colleagues or to the support list when needed

- Available for wagons, datasets and DPG runlists, this always leads to the history page of the current element

Service Wagons

- These are wagons which are dependencies to other wagons
- They are created by experts, and can be added as a dependency to a wagon in the Wagon Settings
  - Each PWG can have a service wagon analysis
- They are displayed in a read-only view in the My Analyses page
- The service wagons of the same type (e.g. EventSelection) will be reduced to a single wagon once AOD metadata is read by $\mathbf{O^2}$

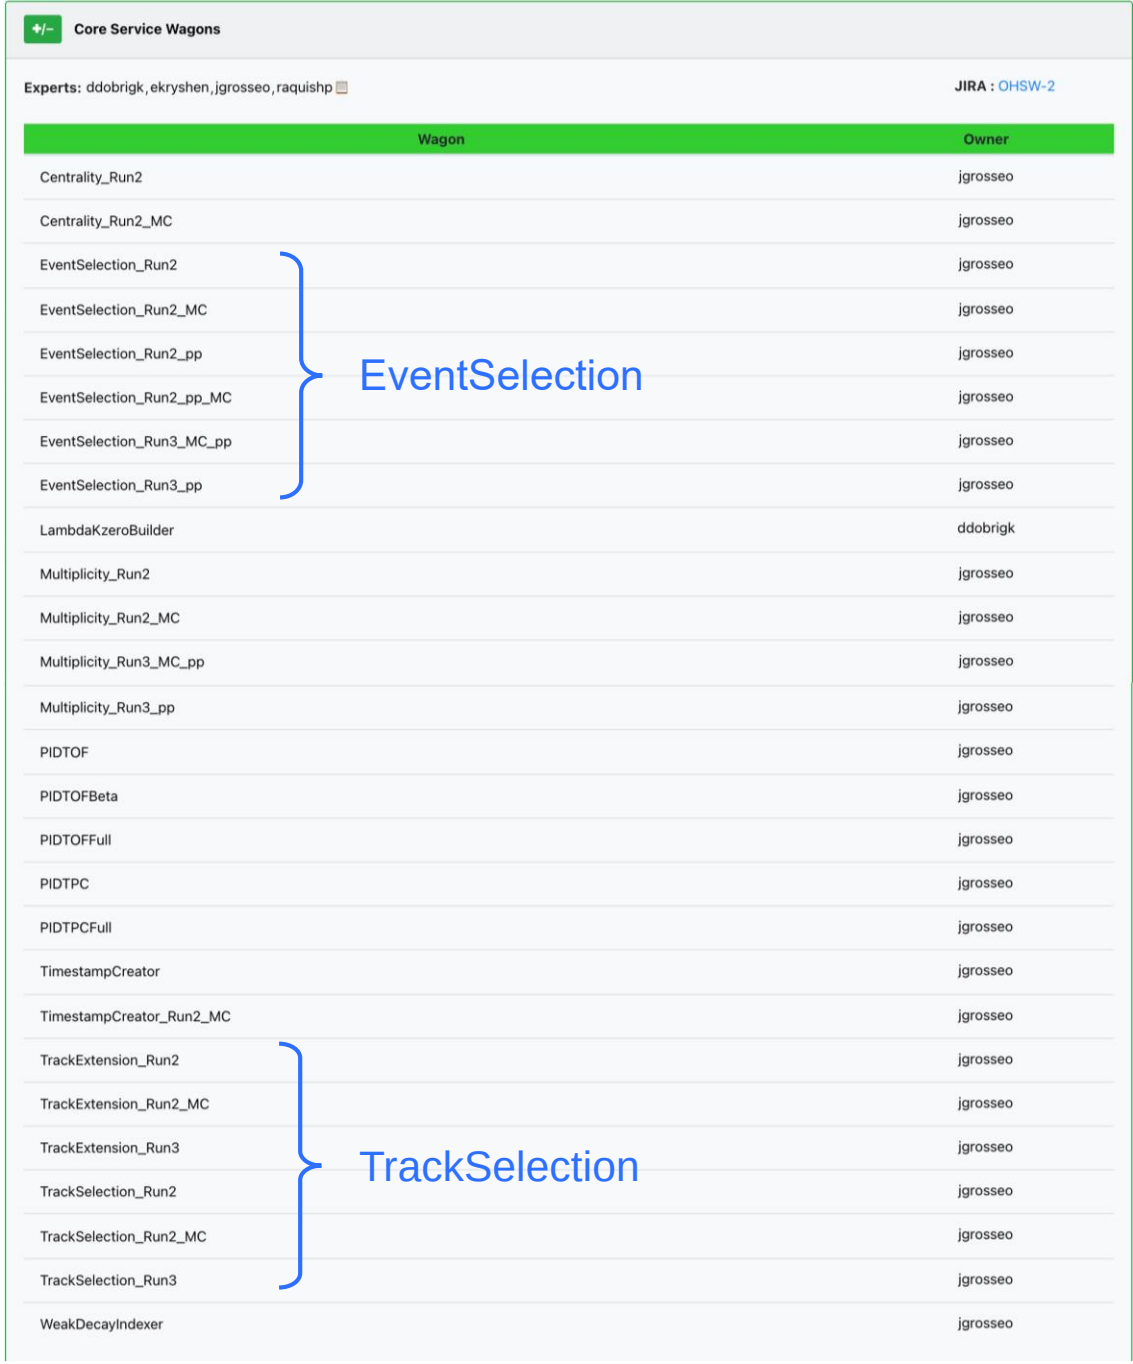

Core Service Wagons

Experts: ddobrigk, ekryshen, jgrosseo, raquishp

JIRA : OHSW-2

Wagon	Owner
Centrality_Run2	jgrosseo
Centrality_Run2_MC	jgrosseo
EventSelection_Run2	jgrosseo
EventSelection_Run2_MC	jgrosseo
EventSelection_Run2_pp	jgrosseo
EventSelection_Run2_pp_MC	jgrosseo
EventSelection_Run3_MC_pp	jgrosseo
EventSelection_Run3_pp	jgrosseo
LambdaKzeroBuilder	ddobrigk
Multiplicity_Run2	jgrosseo
Multiplicity_Run2_MC	jgrosseo
Multiplicity_Run3_MC_pp	jgrosseo
Multiplicity_Run3_pp	jgrosseo
PIDTOF	jgrosseo
PIDTOFBeta	jgrosseo
PIDTOFFull	jgrosseo
PIDTPC	jgrosseo
PIDTPCFull	jgrosseo
TimestampCreator	jgrosseo
TimestampCreator_Run2_MC	jgrosseo
TrackExtension_Run2	jgrosseo
TrackExtension_Run2_MC	jgrosseo
TrackExtension_Run3	jgrosseo
TrackSelection_Run2	jgrosseo
TrackSelection_Run2_MC	jgrosseo
TrackSelection_Run3	jgrosseo
WeakDecayIndexer	jgrosseo

Wagon Configuration

- Hyperloop imports all the defined configurables of your analysis workflow (defined by C++ code)
- The wagon configuration corresponding to the workflow will be available in the ***Base***
- Supports a variety of parameter types defined in task as ***Configurable*** including arrays, matrices, labelled matrices and histogram binning
- Subwagons run the same task with different parameter values. Click **+ New subwagon** to add a new subwagon
- In order to update the base and subwagon configuration with the latest O2Physics version of the workflow, click on the ↻ **sync** button
- By **synchronizing** the configuration, the parameters which no longer belong to the workflow will be removed, and the values of the wagon's *Base* will be updated as well if they have not been modified by the user

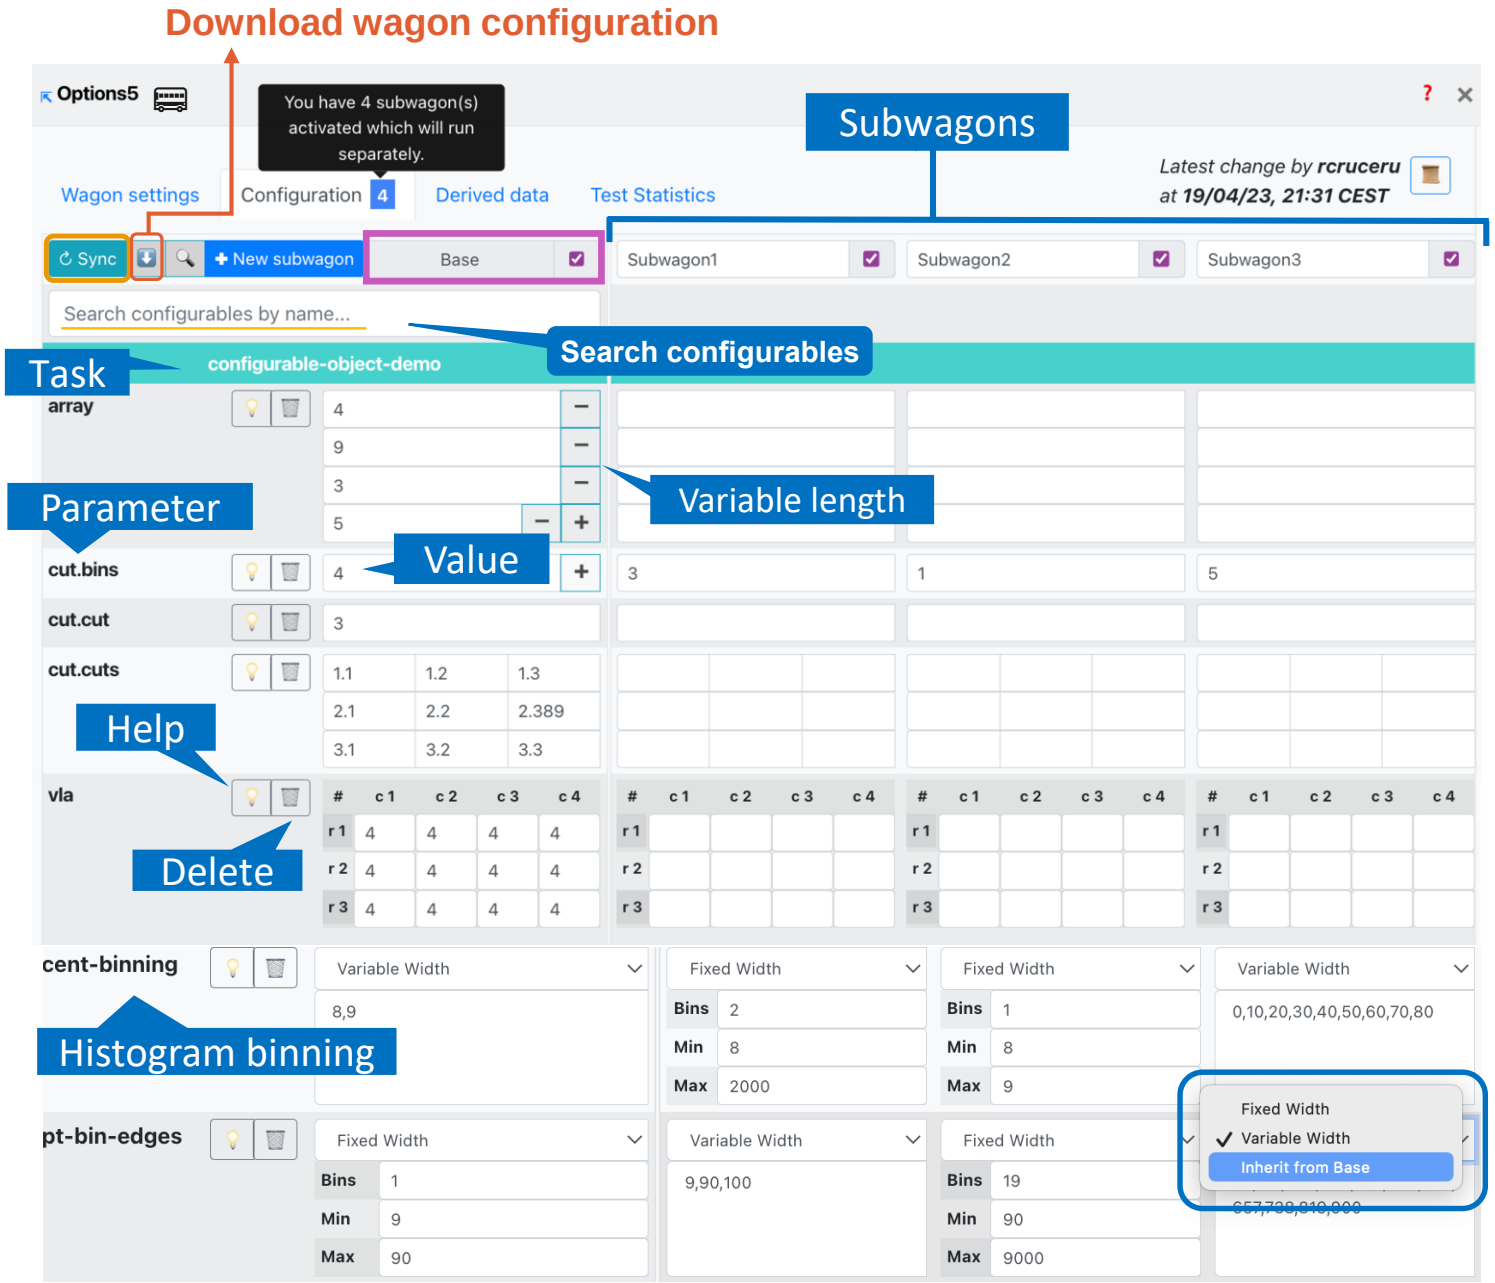

How to create a Subwagon?

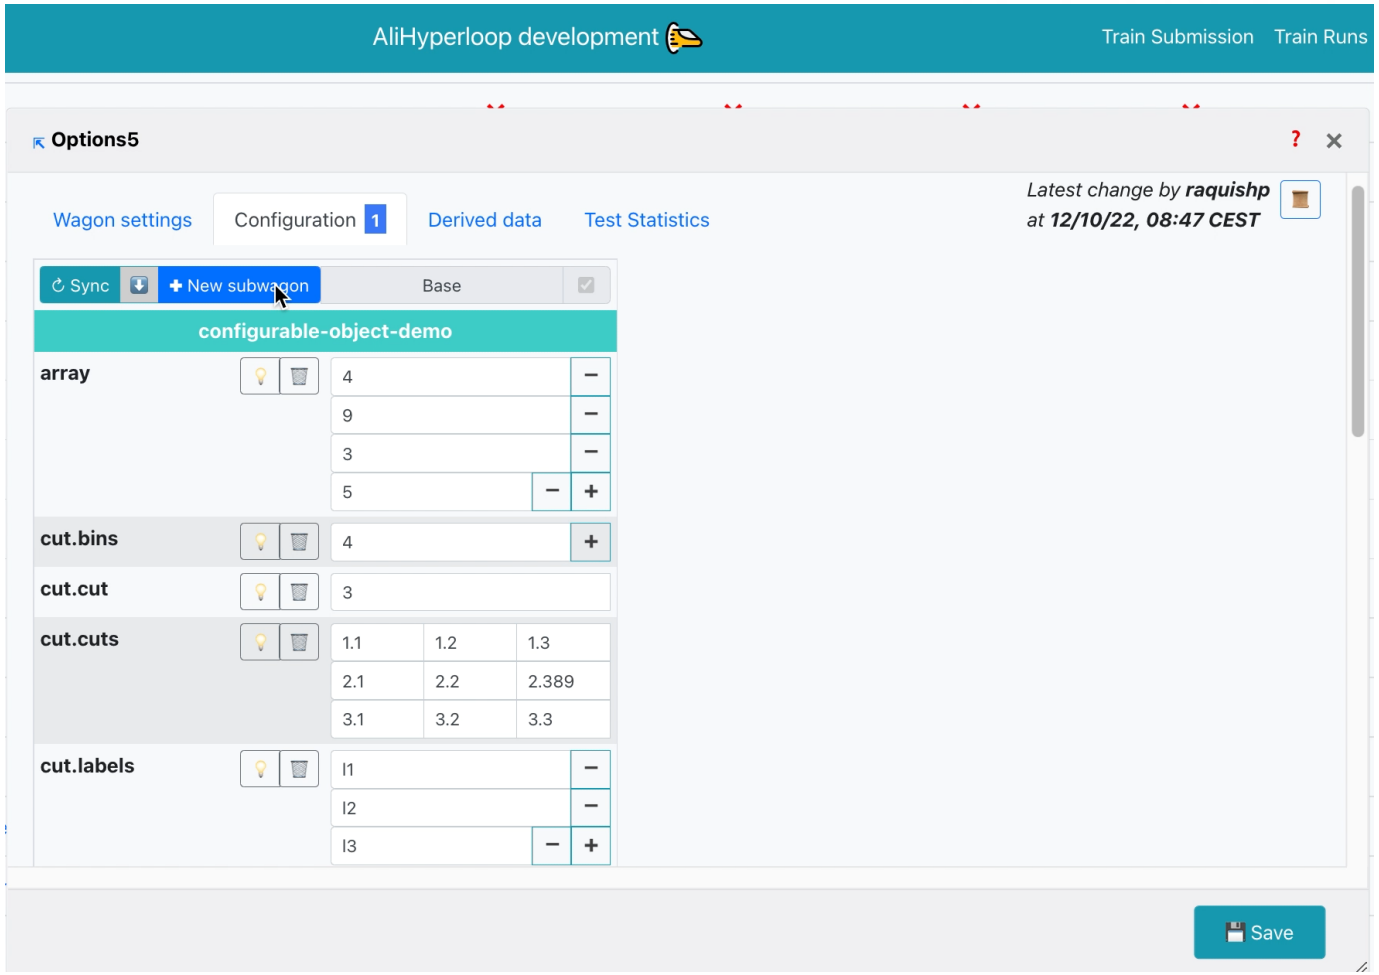

1. Open the wagon edit view and navigate to the **Configuration** tab
2. Click **+New subwagon**
3. Choose a name
4. Modify the variables of interest
5. Click the **+** button
6. Click Save

Note:

- ✓ At least one variable must have a different value from the base
- ✓ A subwagon can't be empty
- ✓ Subwagons can't be identical

Derived data settings

- Displays the tables which are produced by the task
- Here you can enable tables which should be saved into an **AO2D.root output file**
- ***This requires a derived data train which, unless 'Ready for slim' is checked, does not submit automatically and may need additional approval***
- ***If you just need the information in these tables in a subsequent wagon in the same train, there is no need to enable the tables***
- ***For derived data of small output size, you can enable the slim derived data option***

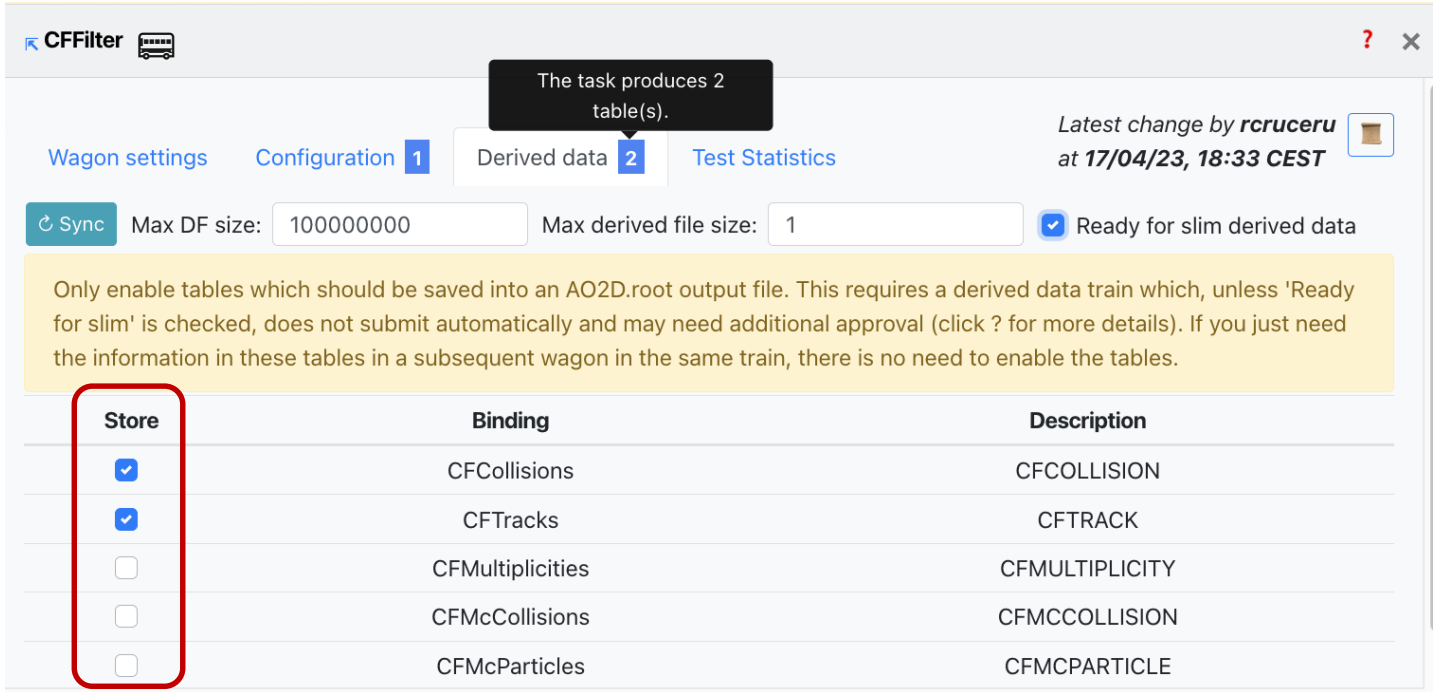

- In order to ***update*** the derived data configuration with the latest O2Physics version of the workflow, click on the ↻**sync button**
- By synchronizing the derived data, the tables which no longer belong to the workflow will be removed, and the values of the tables will be updated

Slim derived data settings

- Reserved for derived data of small output size
- Settings in Derived data tab
  - **Max DF size** - sets the maximal dataframe size in the merging step
    - Has to be 0 for not-self contained derived data (which need parent file access)
  - **Max derived file size** - Limit for the output file size of the derived data file
    - This is an expert parameter which usually does not have to be changed
    - Only change this value if the processing in subsequent trains takes so long that the jobs fail
    - If set to 0 a good value will be automatically determined

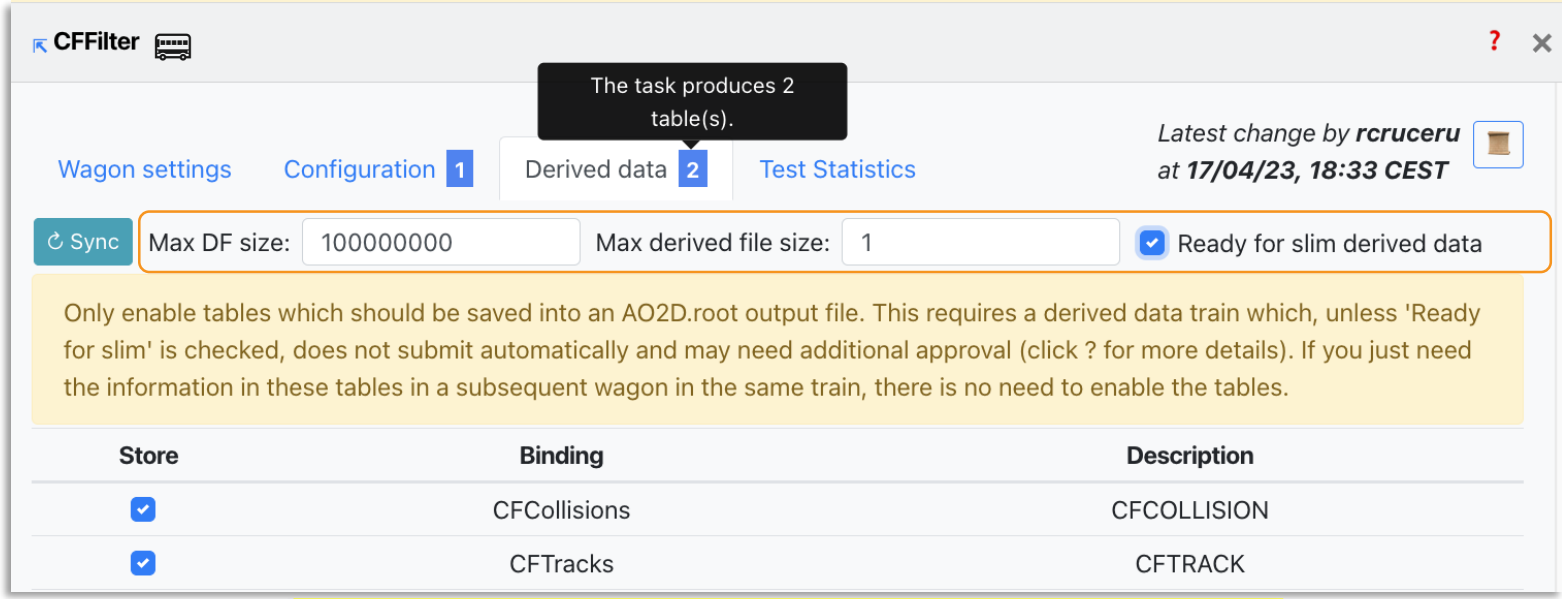

If unsure, please do not edit **max DF size** and **max derived file size**. Contact experts if needed.

- Both max DF size and max derived file size need to be the same for the wagon and its dependencies (which have derived data enabled)
- Check "**Ready for slim derived data**" to enable slim derived data
  - The test will run as slim derived data
  - Wagons (with successful tests) will be included in automatic train composition (you no longer need to request a train on the Operation channel)
  - All wagons (with derived data activated) included in the train need to have the same max DF size & derived file size

Wagon – Test statistics

- Long term statistics (per Test id) for enabled datasets per analysis
  - PSS Memory
  - CPU/Wall Time
  - Throughput
  - Output size
- Direct links to a read-only view of wagon test output by clicking into the graph
https://alimonitor.cern.ch/hyperloop/wagon-test/*test-id*

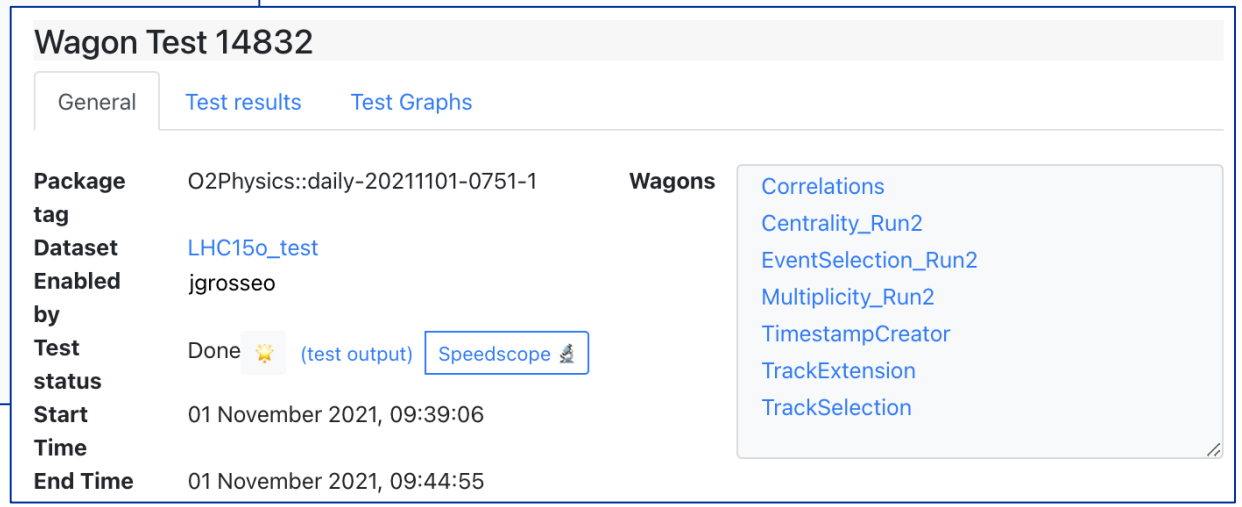

Wagon tests

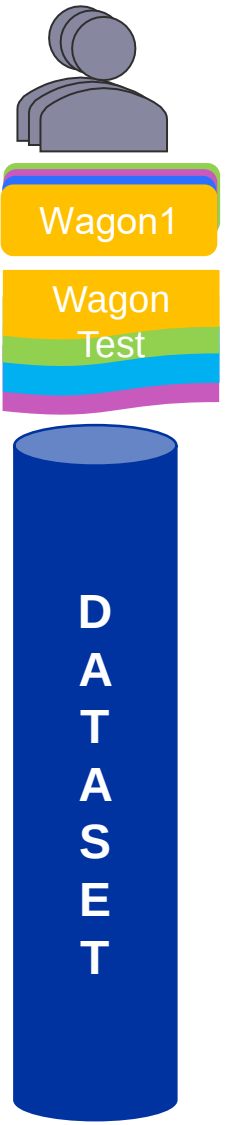

- When enabling a wagon, the test will start immediately

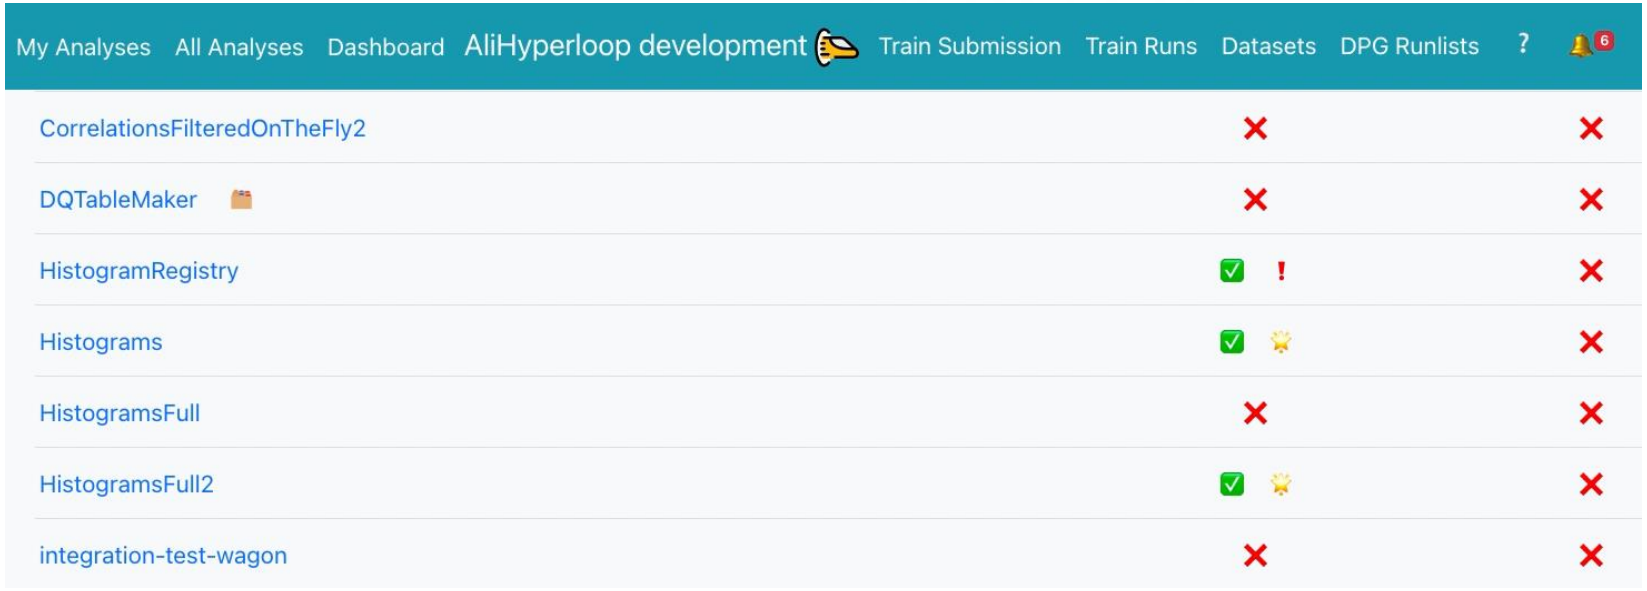

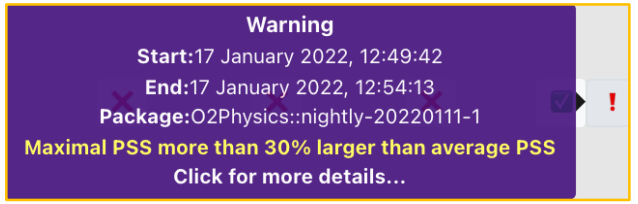

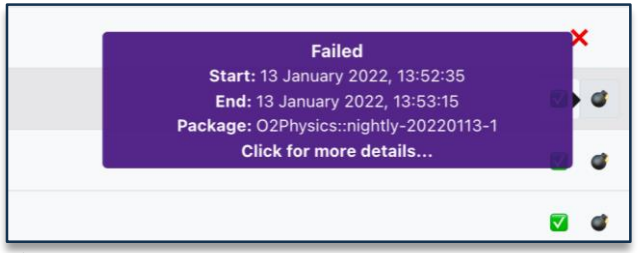

- **Successful** wagon tests will be marked with ✨. **Failed** wagon tests will be marked with 💣
- Wagon tests finished in **Warning** will be marked with ❗.
- The types of warnings that can occur (details here):
  - CPU usage too large
  - Memory consumption too large
  - Reduction factor too small (occurs when < 50)
  - Maximal PSS more than 30% larger than average PSS
  - Log output is too large
  - Too many CCDB queries
  - Derived output too large for slim train
  - Unbound indices detected NEW

Click the icon to open the wagon test result

Wagon test

- You can open the test output in a new tab by clicking (test output)

- Here you can analyze the:
  - Full configuration
  - Performance metrics
  - Status
  - Standard error
  - Standard output

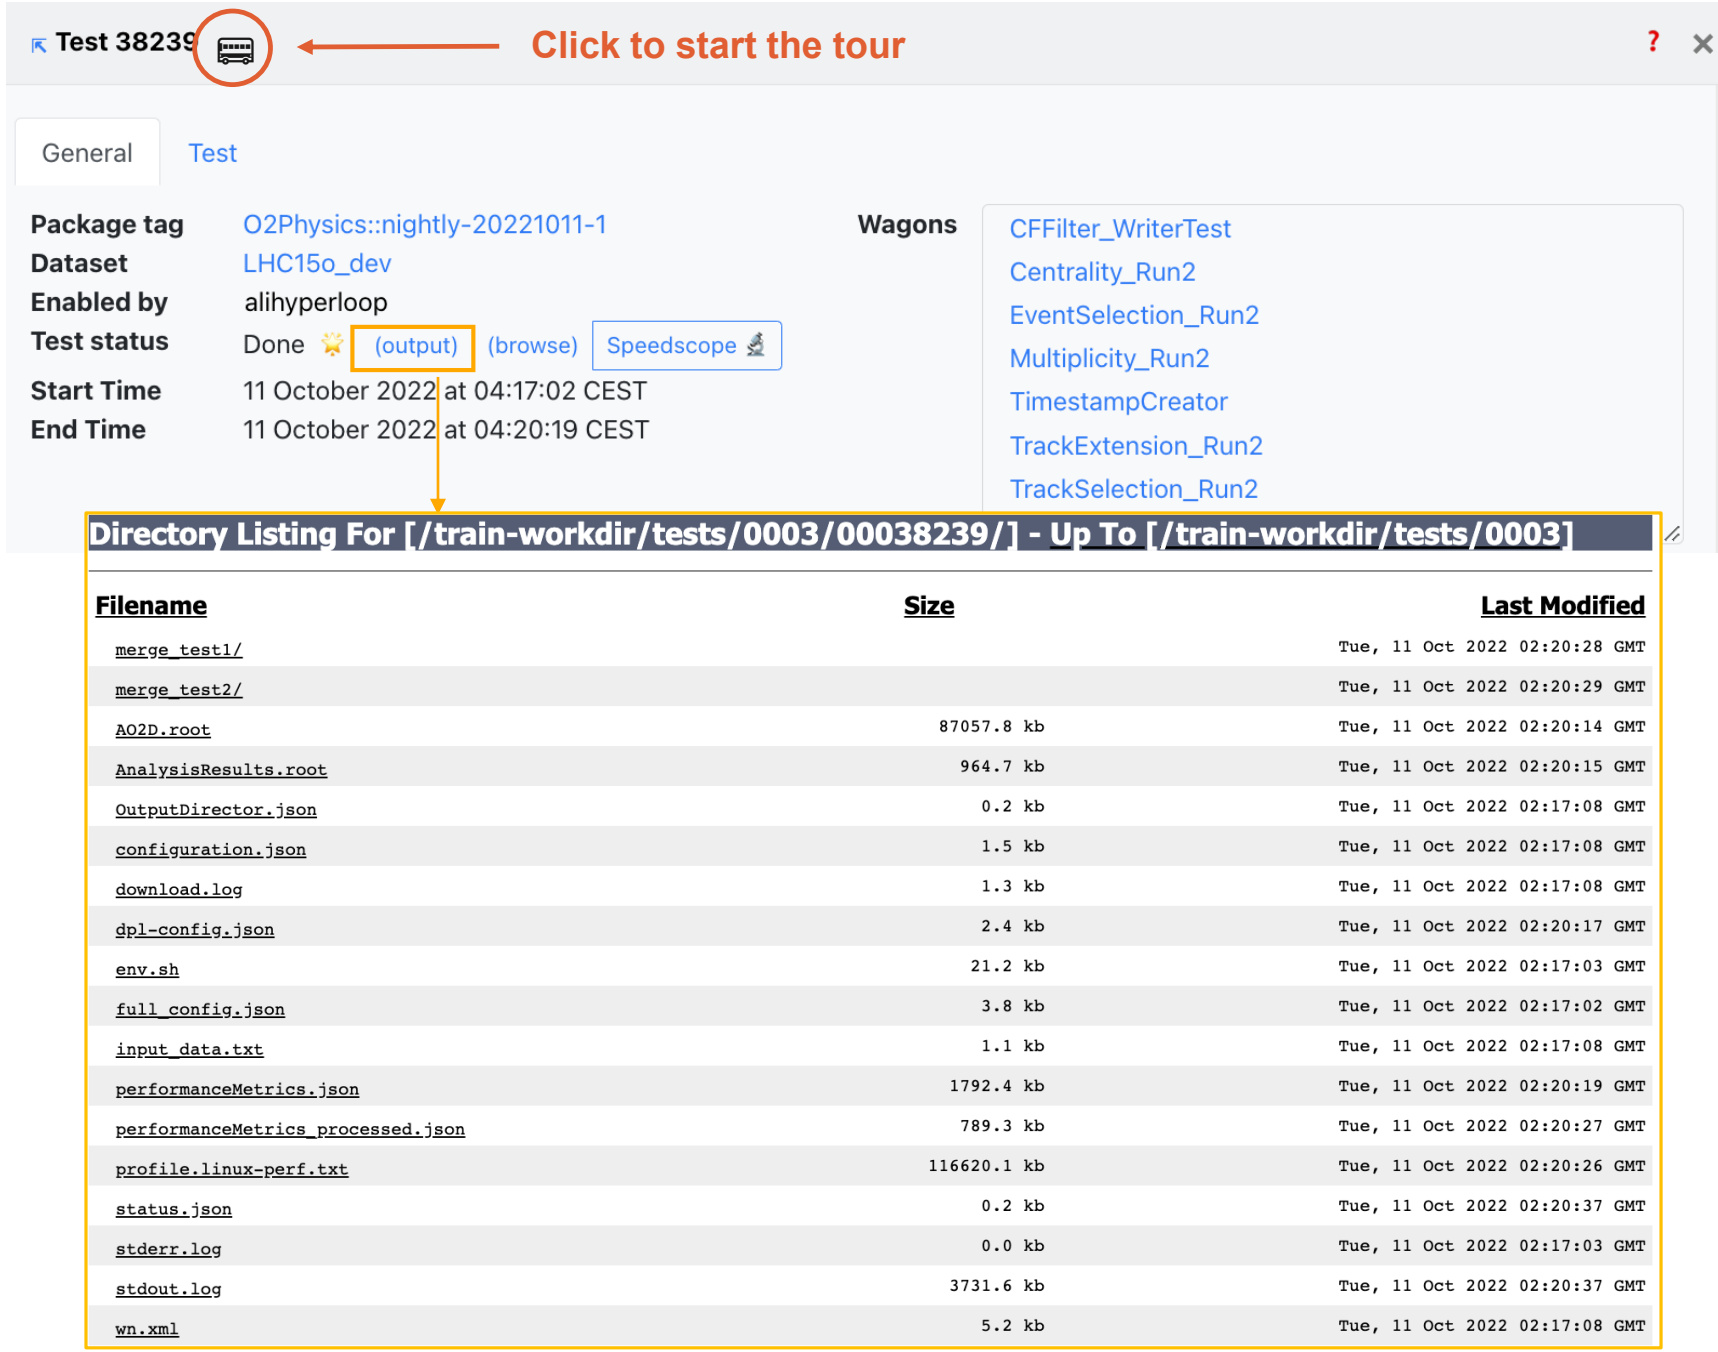

Wagon test – General

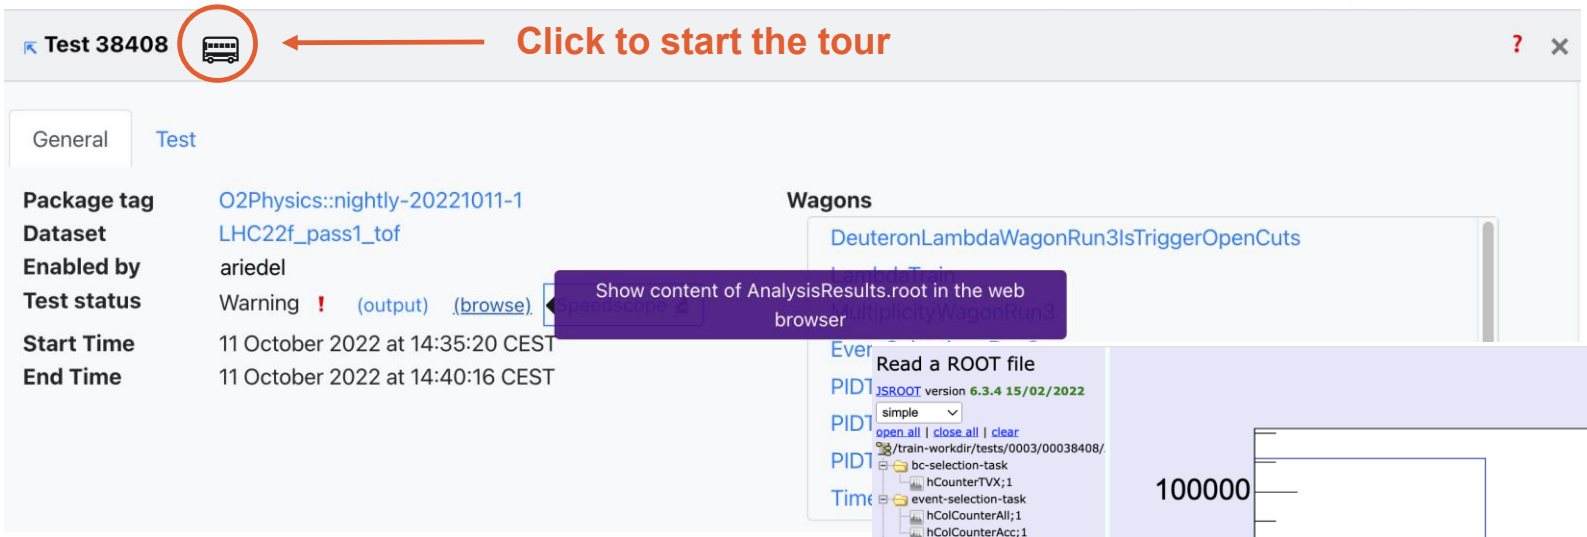

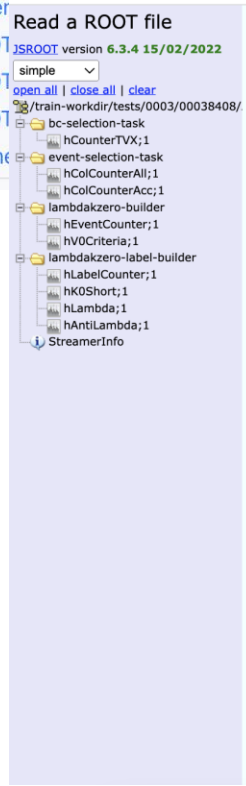

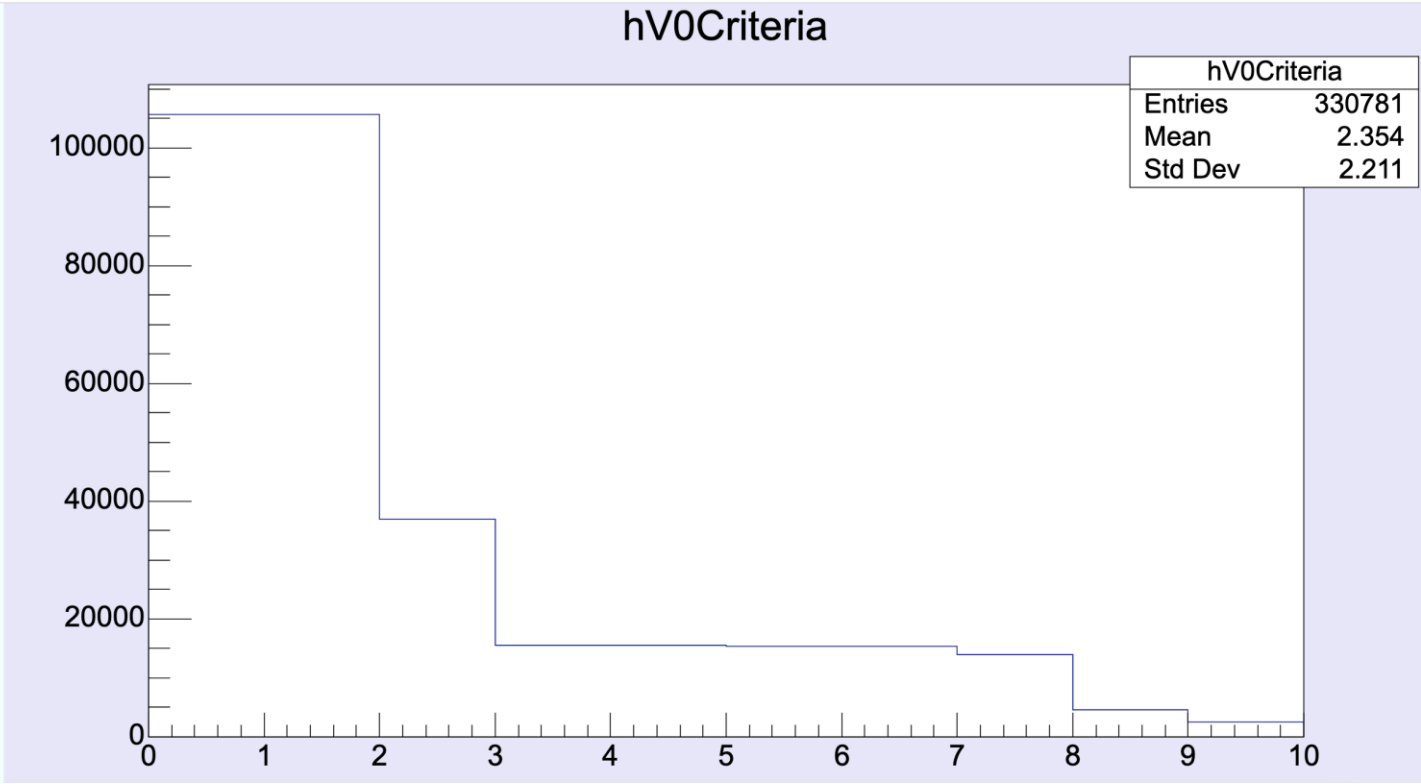

Link “browse” in wagon tests to show content of AnalysisResults.root in the web browser

Wagon test – General

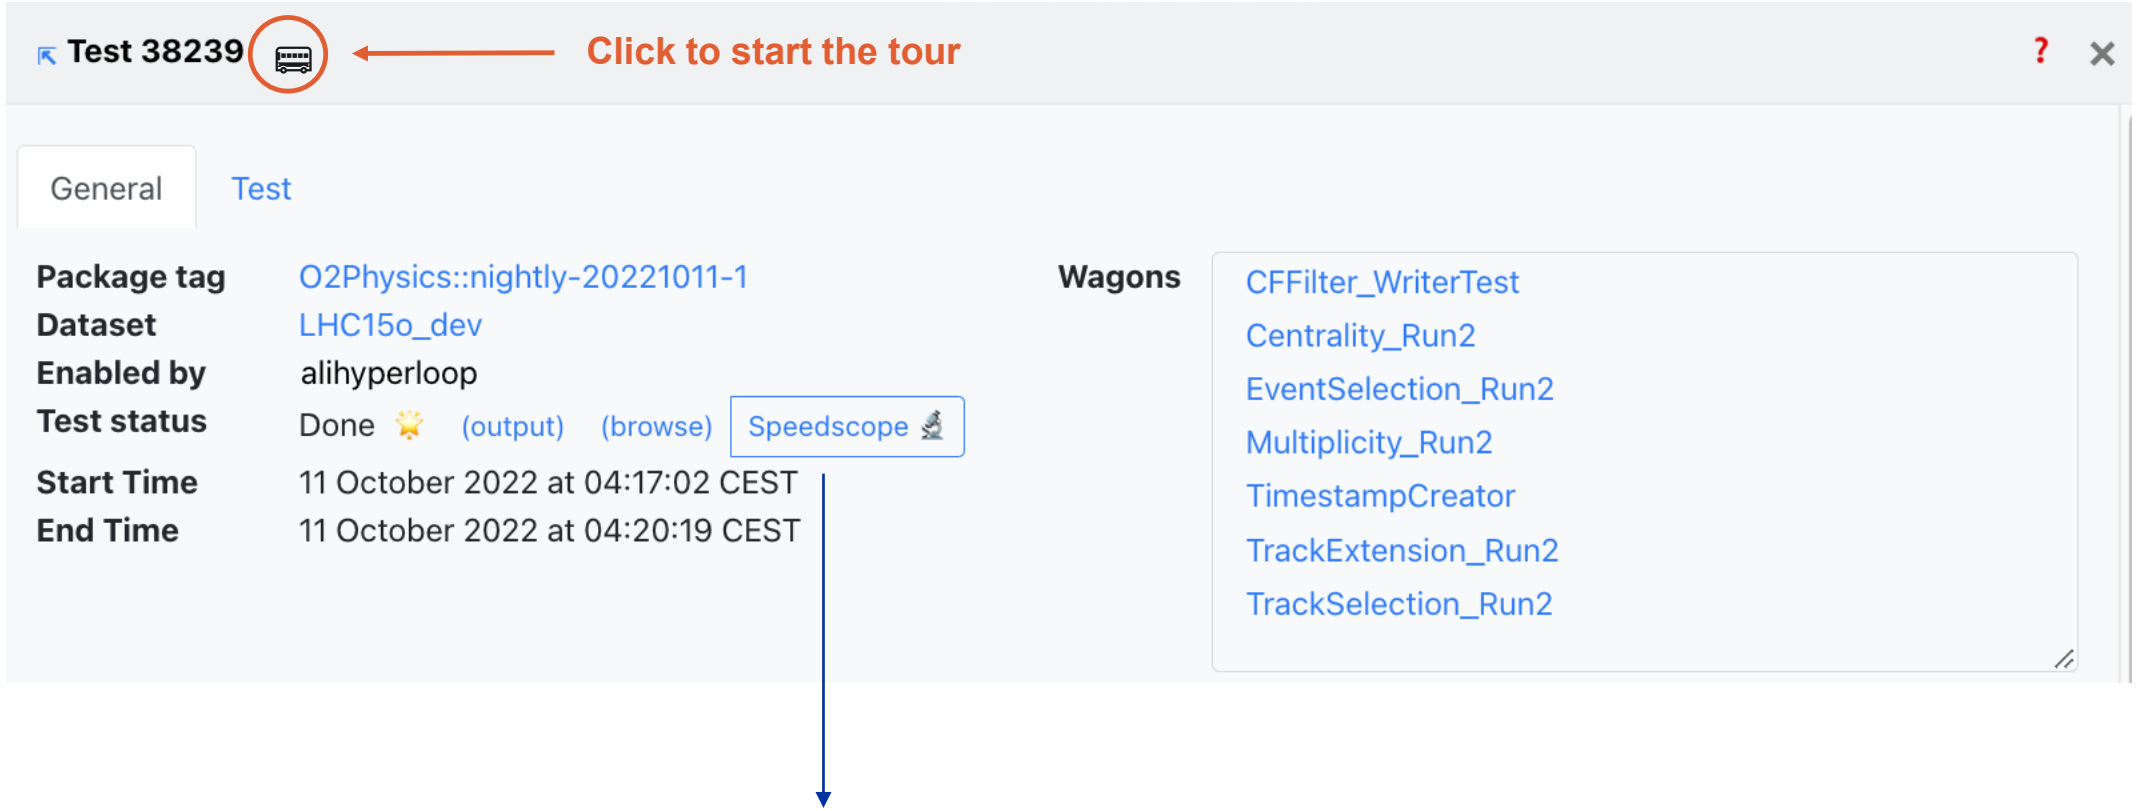

- Link to the Speedscope view of the profiling of your task (also available in the Train run result)
- Speedscope is a web tool to display the profiling determined using Linux perf

Wagon test – Derived data

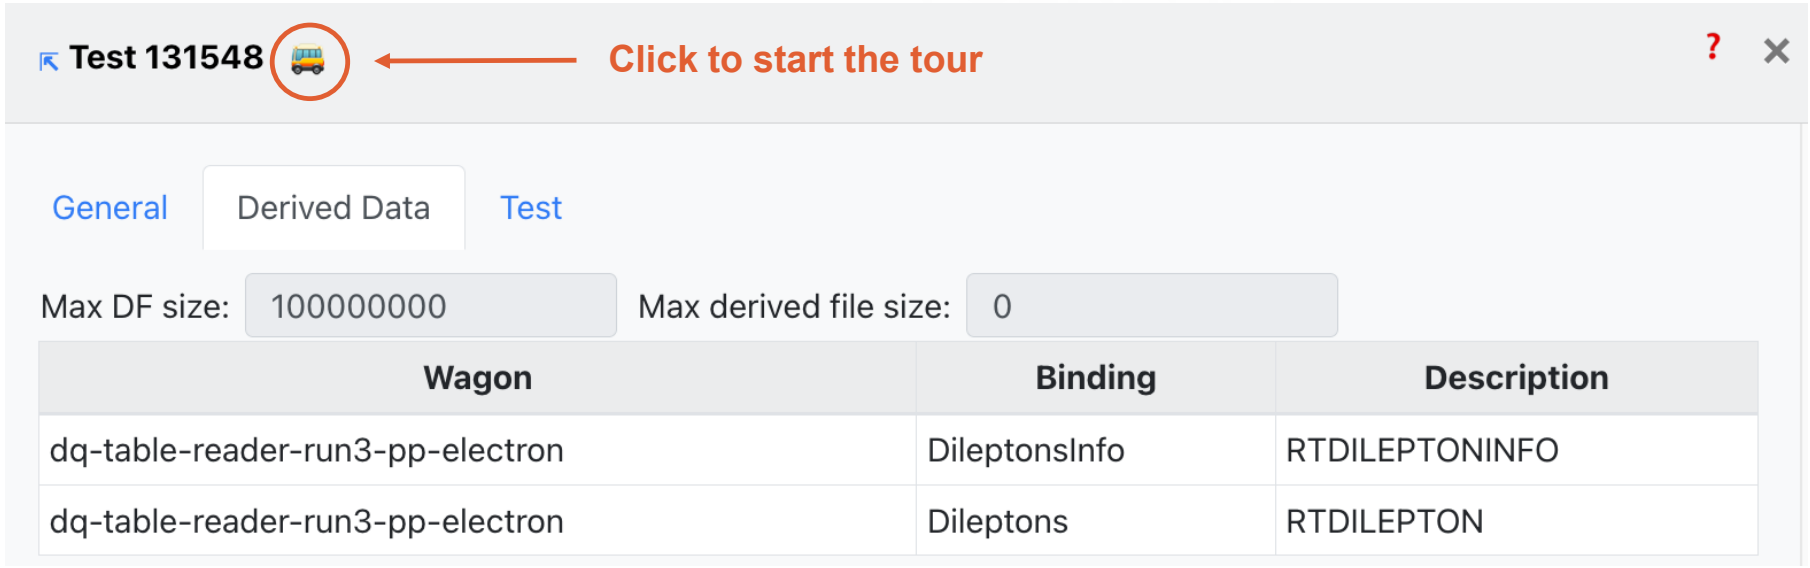

- Displays the tables which are produced by the task
- Max data frame size and max derived file size

Wagon test results

Full Test metrics

- Expected resources and Interactive graphs

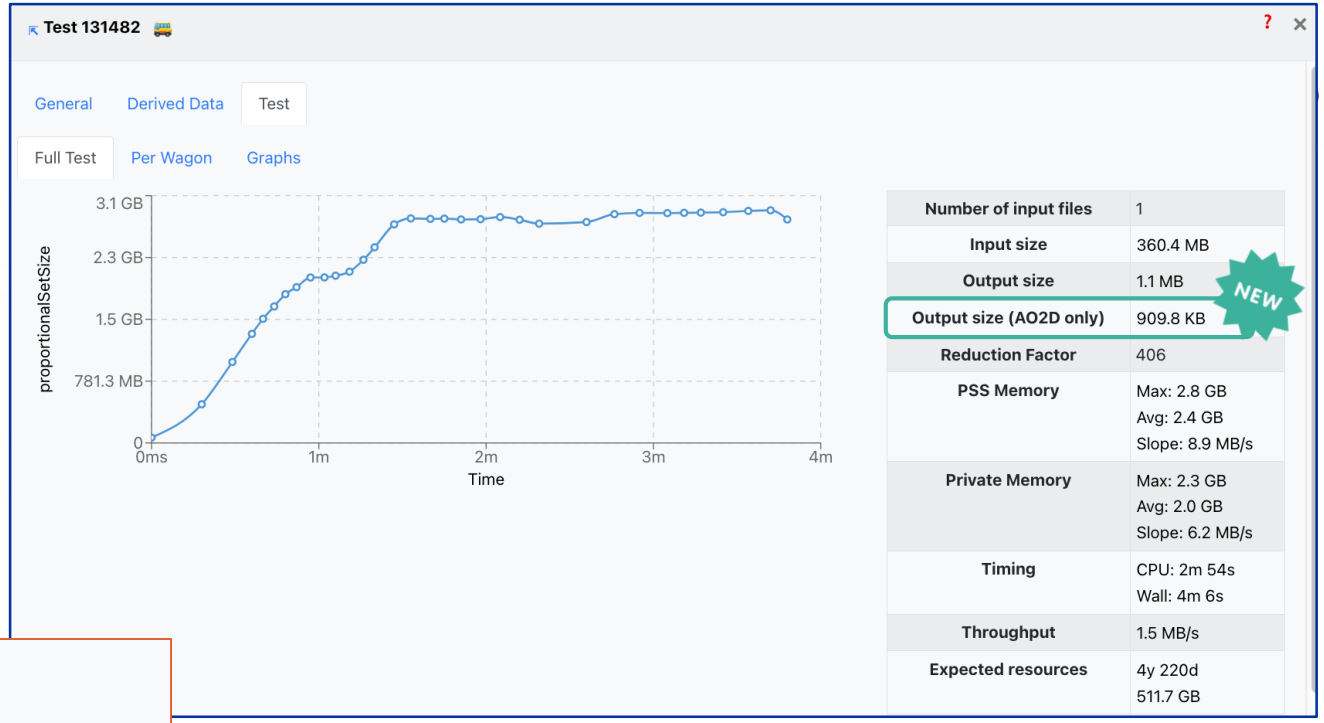

Number of input files	1
Input size	360.4 MB
Output size	1.1 MB
Output size (AO2D only)	909.8 KB
Reduction Factor	406
PSS Memory	Max: 2.8 GB Avg: 2.4 GB Slope: 8.9 MB/s
Private Memory	Max: 2.3 GB Avg: 2.0 GB Slope: 6.2 MB/s
Timing	CPU: 2m 54s Wall: 4m 6s
Throughput	1.5 MB/s
Expected resources	4y 220d 511.7 GB

Full Test | Per Wagon | Graphs

Clear all filters

Wagon		PSS Memory	Private Memory	CPU Time	CCDB C
Search 15 records...					
Reader	Max Avg Slope	527.5 MB 424.0 MB -166.1 KB/s	425.3 MB 299.1 MB -361.8 KB/s	21s (12%)	0
bcConverter	Max Avg Slope	22.1 MB 19.0 MB -12.3 KB/s	20.7 MB 17.2 MB -18.5 KB/s	2s (1%)	0
CorrectedFT0	Max Avg Slope	19.0 MB 18.6 MB 4.4 KB/s	18.0 MB 17.9 MB 1.6 KB/s	2s (1%)	0
DQ-DalitzSel	Max Avg Slope	316.9 MB 253.6 MB 1.0 MB/s	254.2 MB 218.3 MB 754.2 KB/s	13s (8%)	1
DQ-TPCPID-DalitzReducedMaker	Max Avg Slope	288.4 MB 225.5 MB 980.2 KB/s	208.1 MB 181.8 MB 550.5 KB/s	15s (9%)	1

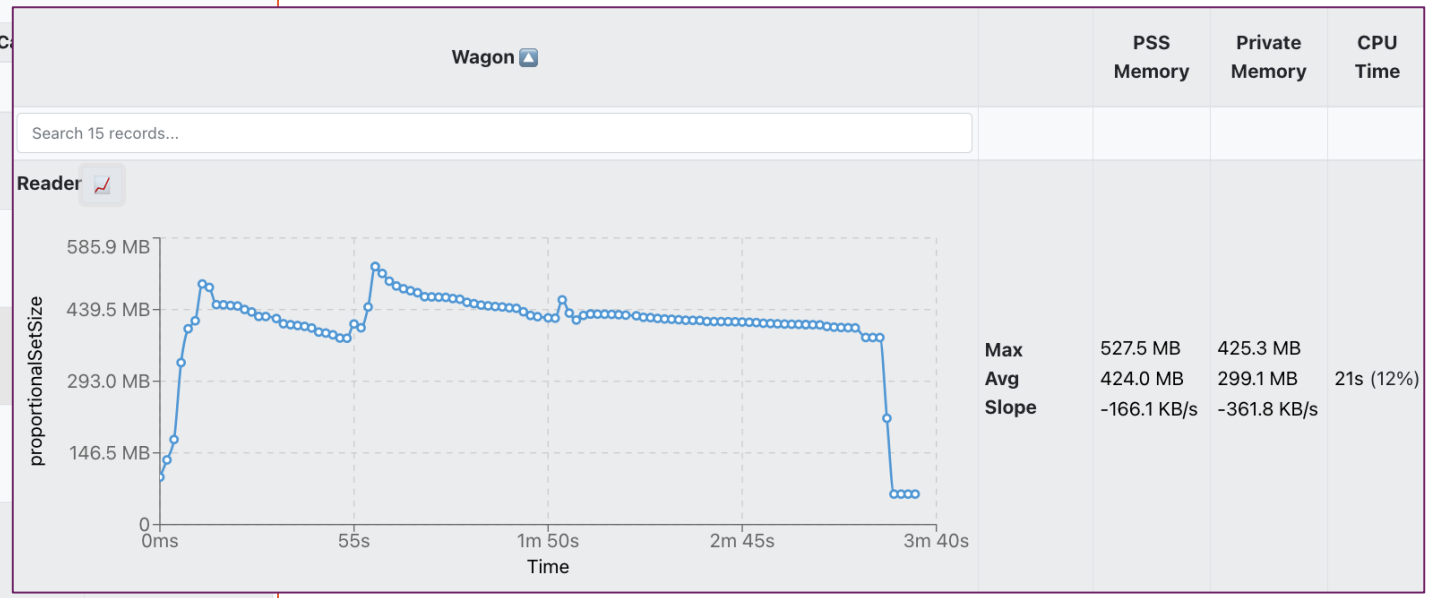

Performance graphs

Performance Graphs page
https://alimonitor.cern.ch/hyperloop/performance-graphs

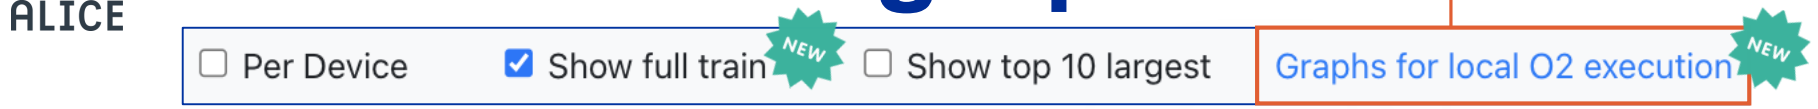

- **Interactive graphs**
- Check box to see performance per device
- Use the dropdown to select the different types of metrics
- **Show / hide full train**

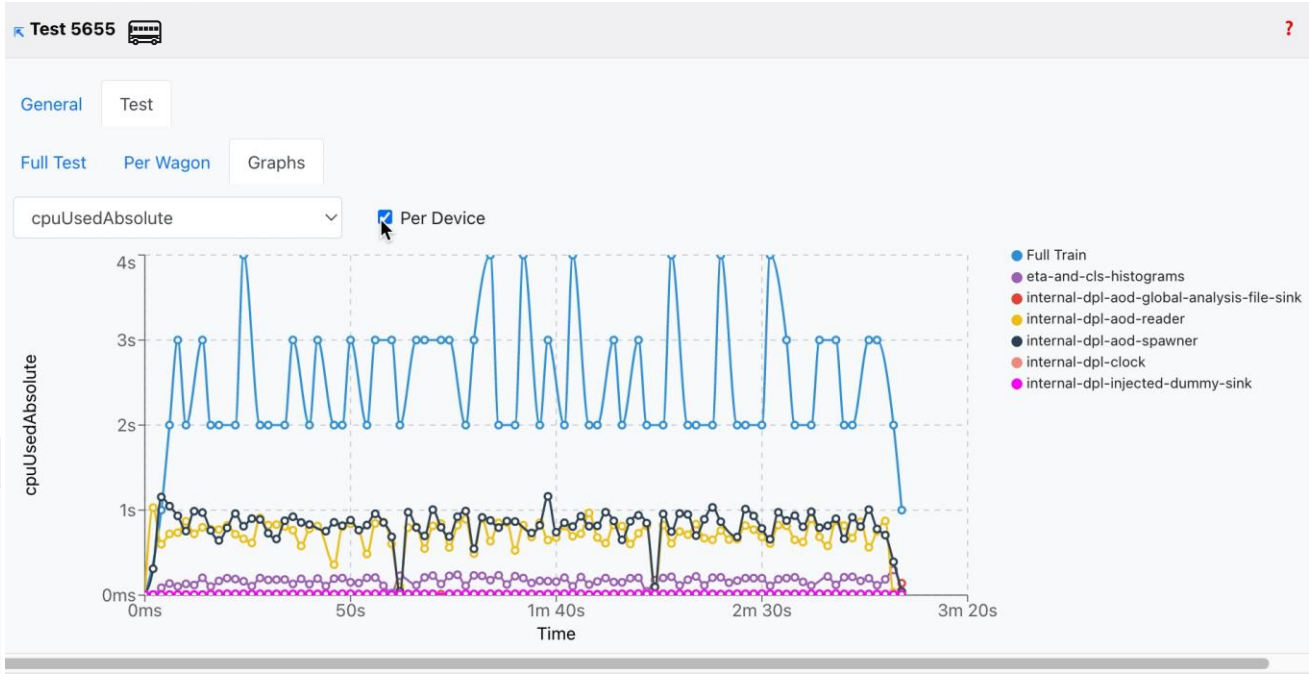

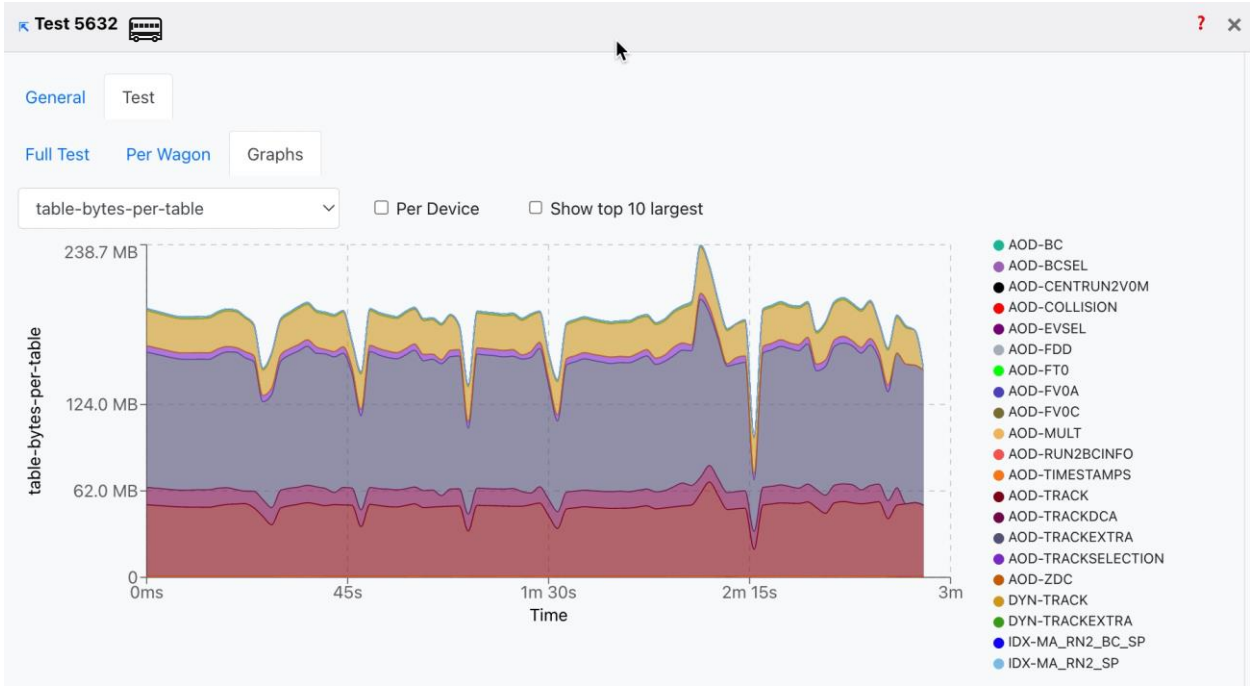

- **table-bytes-per-table**
  - Shows the size of the different tables which contain the processed data
- Zooming
- Show 10 largest

Performance graphs (dedicated page)

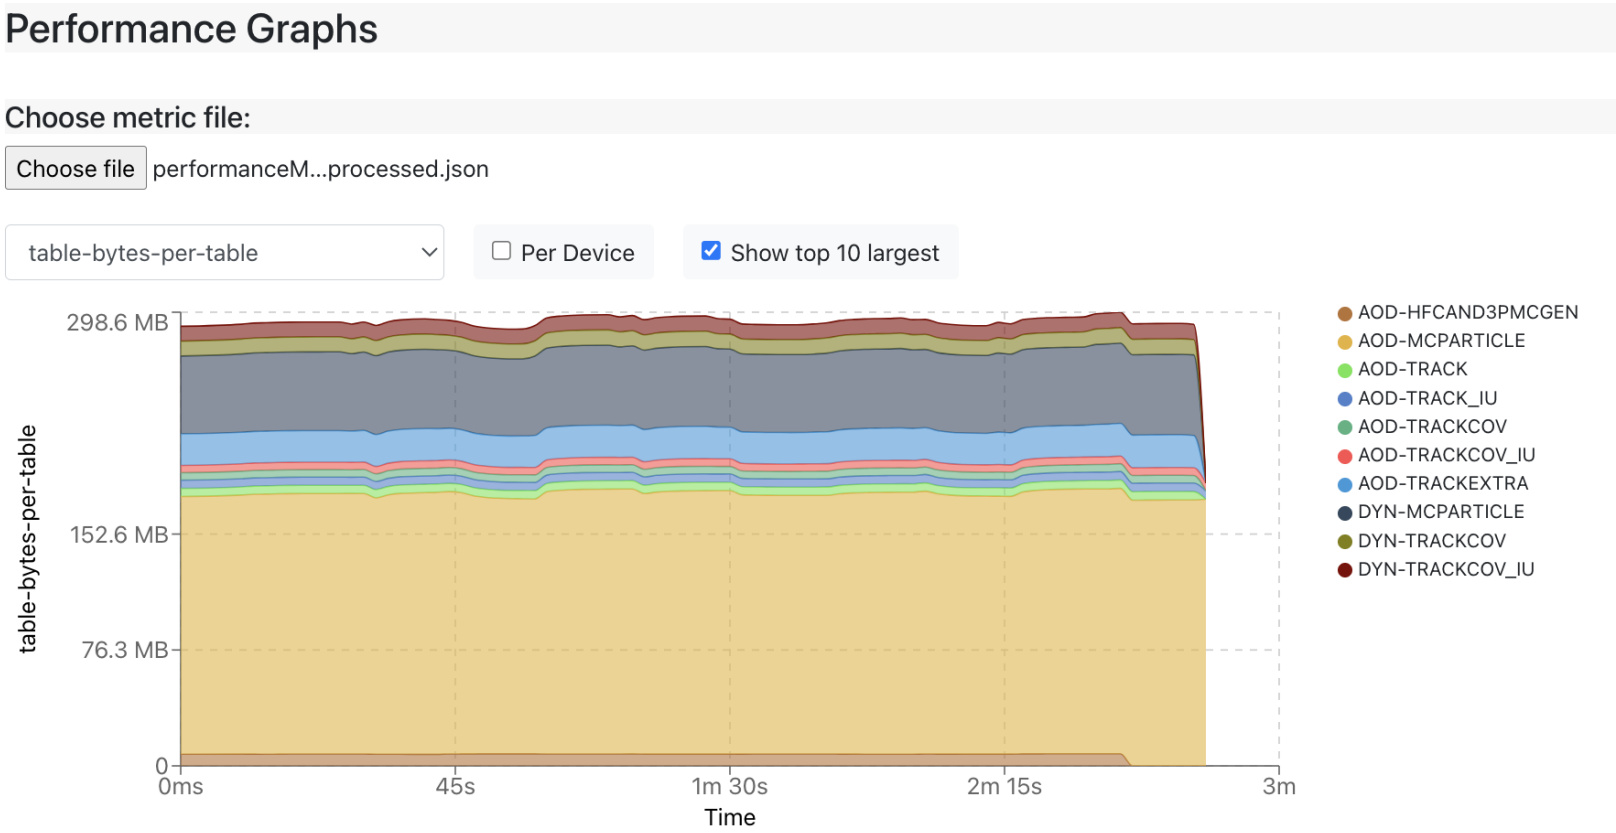

- You can upload your own local metrics file, and then generate the test graphs specific to this file
- These are the same type of graphs produced in the *Test Graphs* tab of the train run and wagon test

Train runs

- **Automatic train composition** can be scheduled per dataset – see the schedule in the dataset view

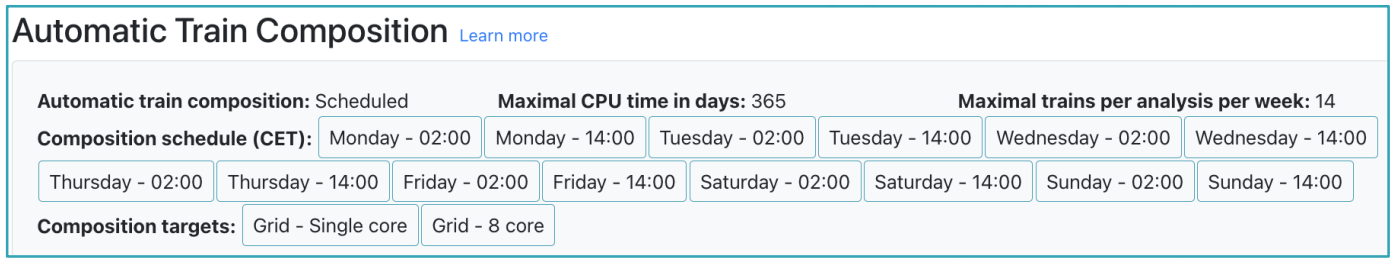

- Trains will be composed based on *target memory*, *wagon configuration* and *dependencies*
- Automatic **train submission** to **GRID** or **Analysis Facilities**
- **Wagons will not be considered for automatic composition if:**
  - Automatic submission is not enabled for the dataset
  - Derived data output is activated and train is not slim
  - Wagon test finished with warning

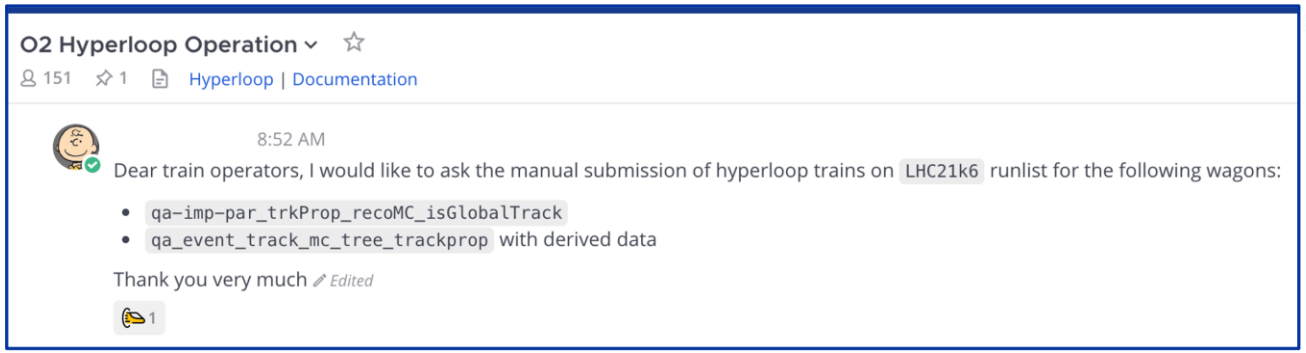

Train Support

- **O2 Hyperloop Operation Channel**
- 24/5 Operation (different timezones)
- Institutes: 1 in Americas, 2 in Europe,1 in Asia

Train runs

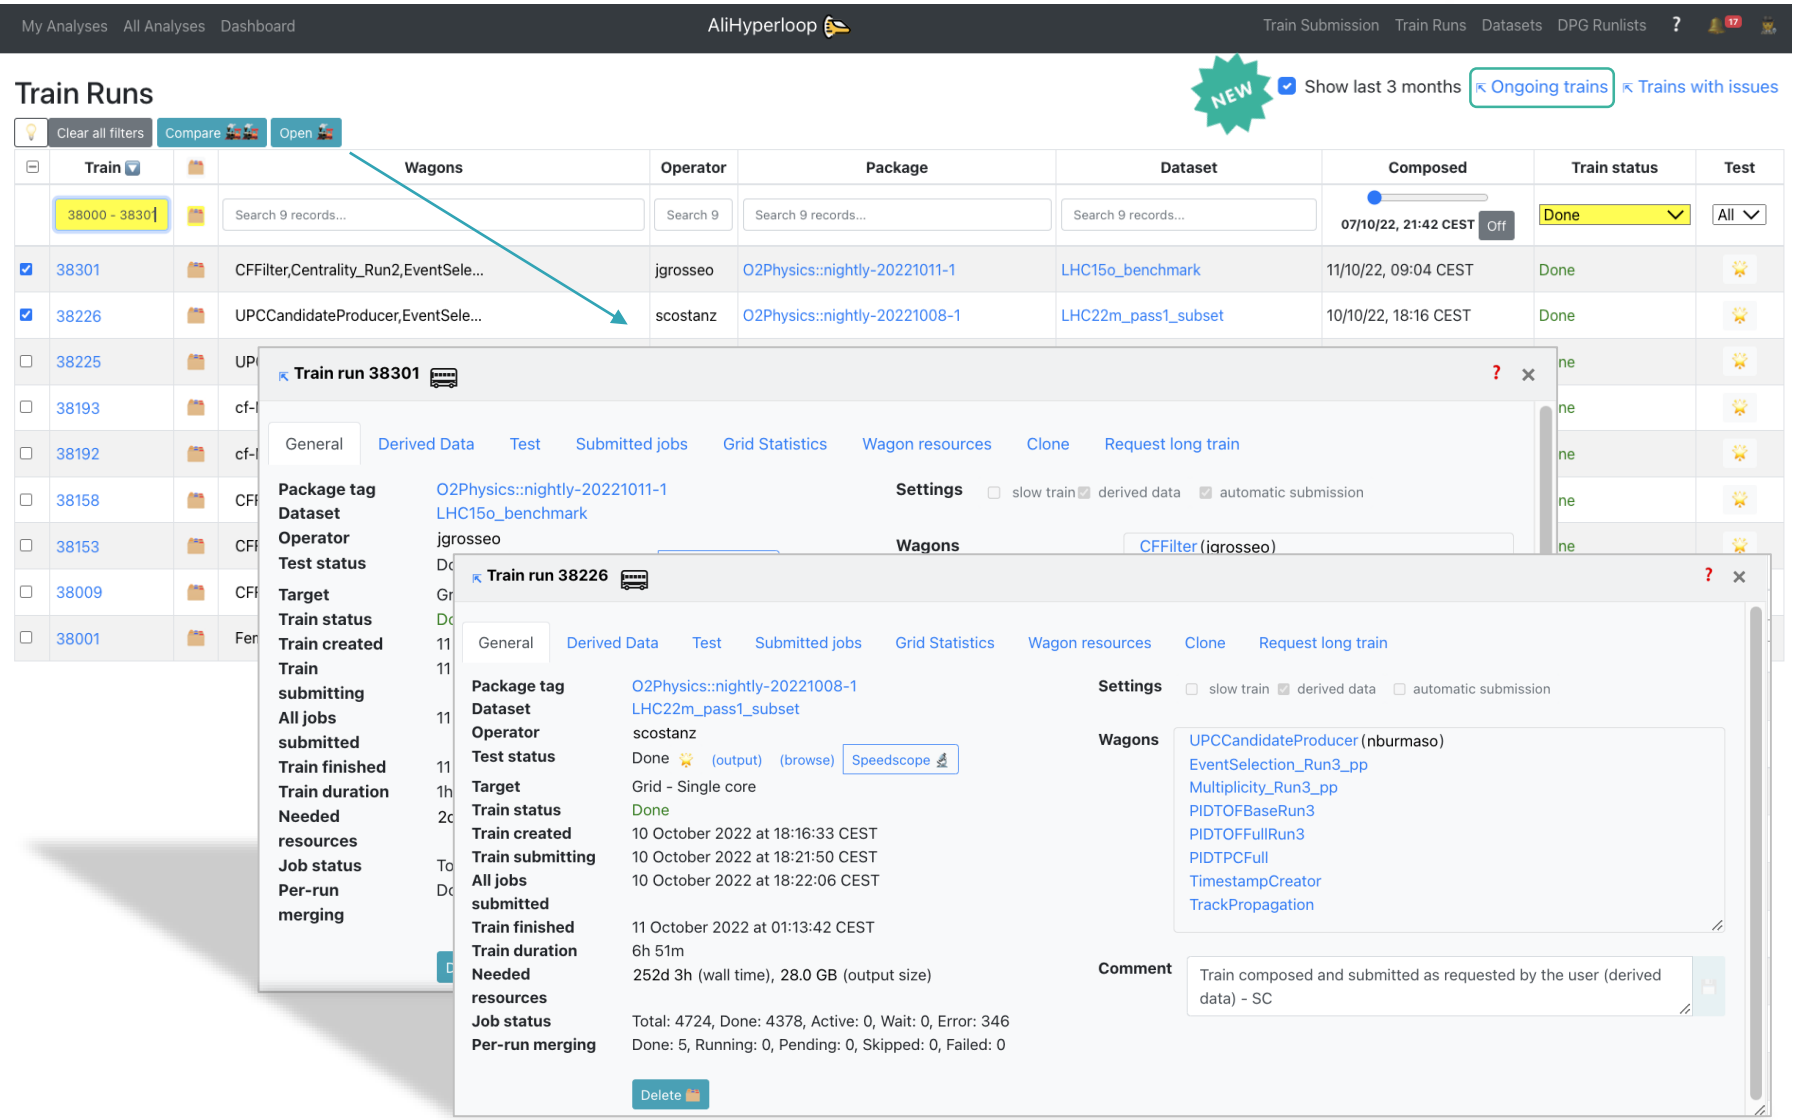

- Displays all the train runs **(by default, last 3 months of activity)**
- Use the searching and filtering features to find your trains
- Click on the train number to access the train run details
- To open several train runs views, select all the trains of interest and click Open
- To compare two trains, select them in the Compare column and click *Compare.* This will open a new tab displaying the differences between the two trains

Train runs

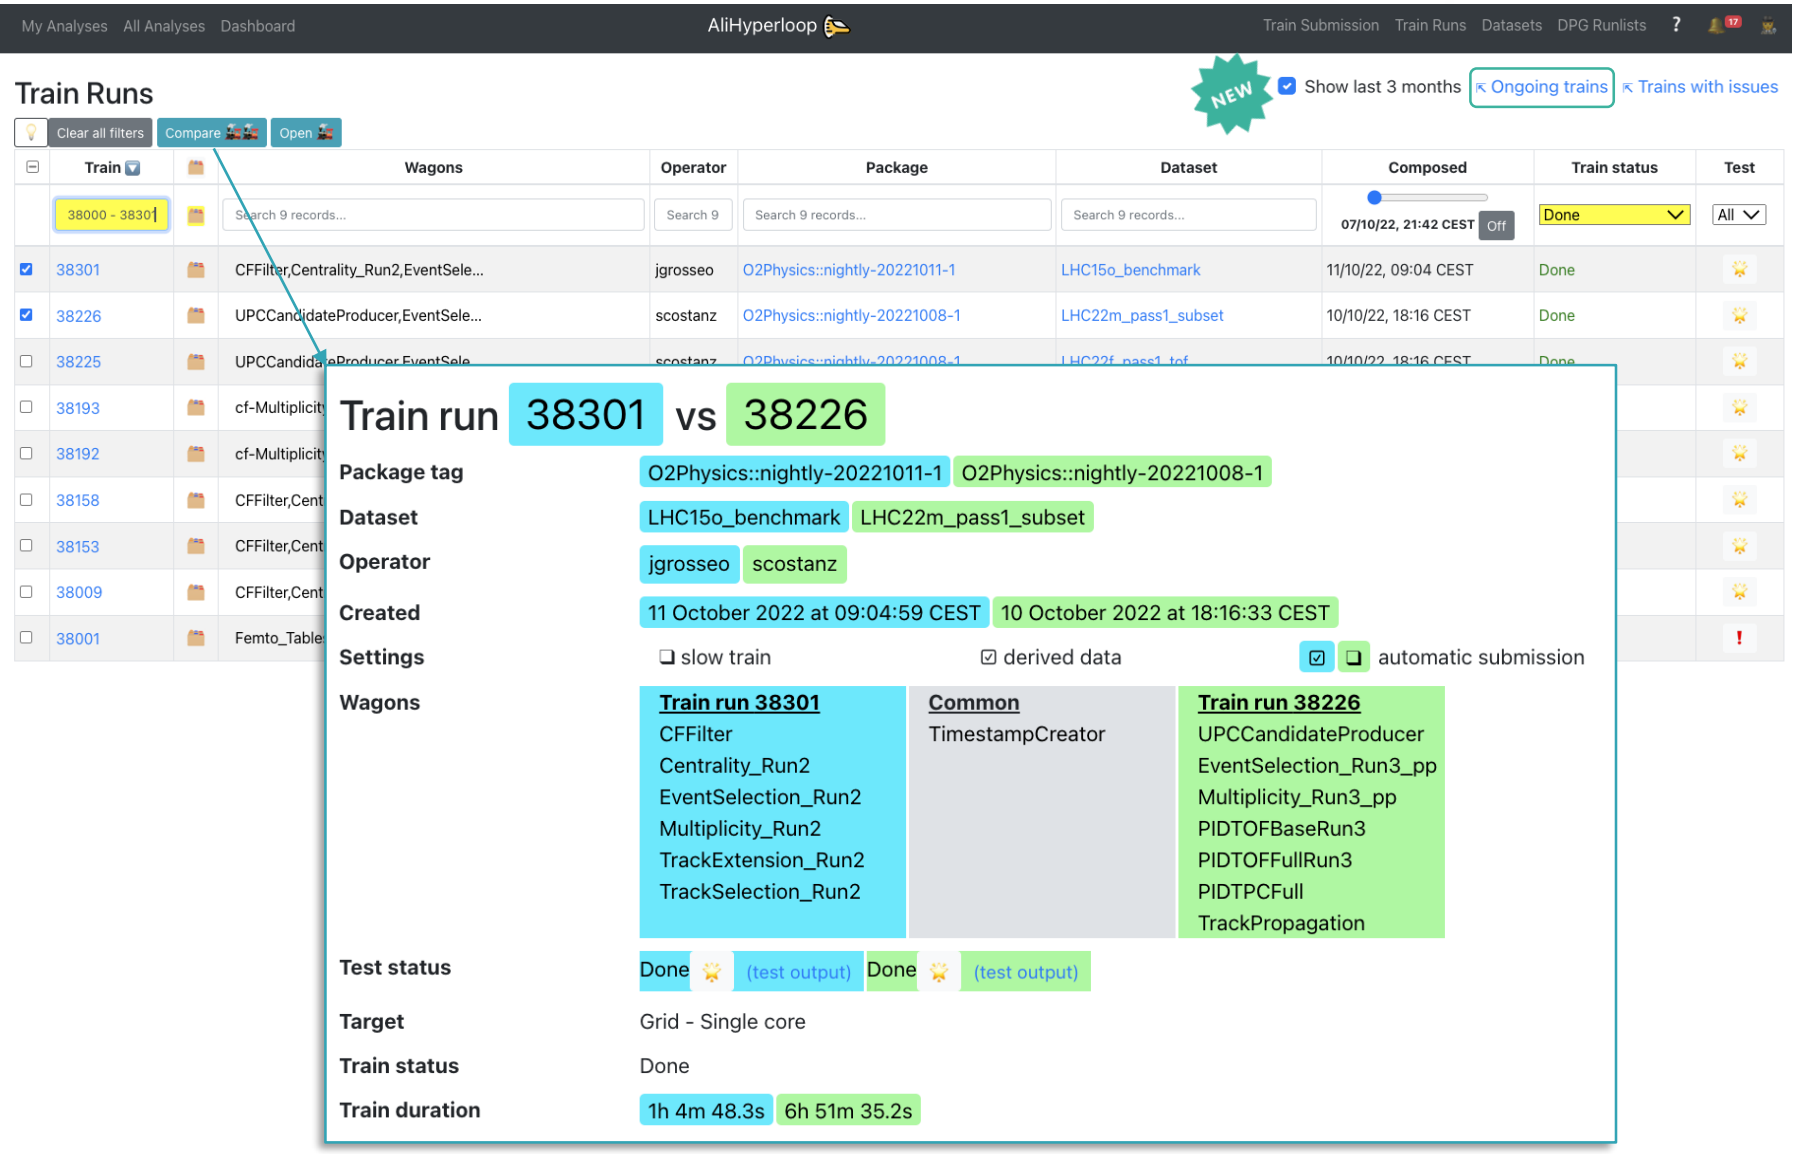

- Displays all the train runs
  (by default, last 3 months of activity)
- Use the searching and filtering features to find your trains
- Click on the train number to access the train run details
- To open several train runs views, select all the trains of interest and click Open
- To compare two trains, select them in the Compare column and click *Compare.* This will open a new tab displaying the differences between the two trains

Train run view – General

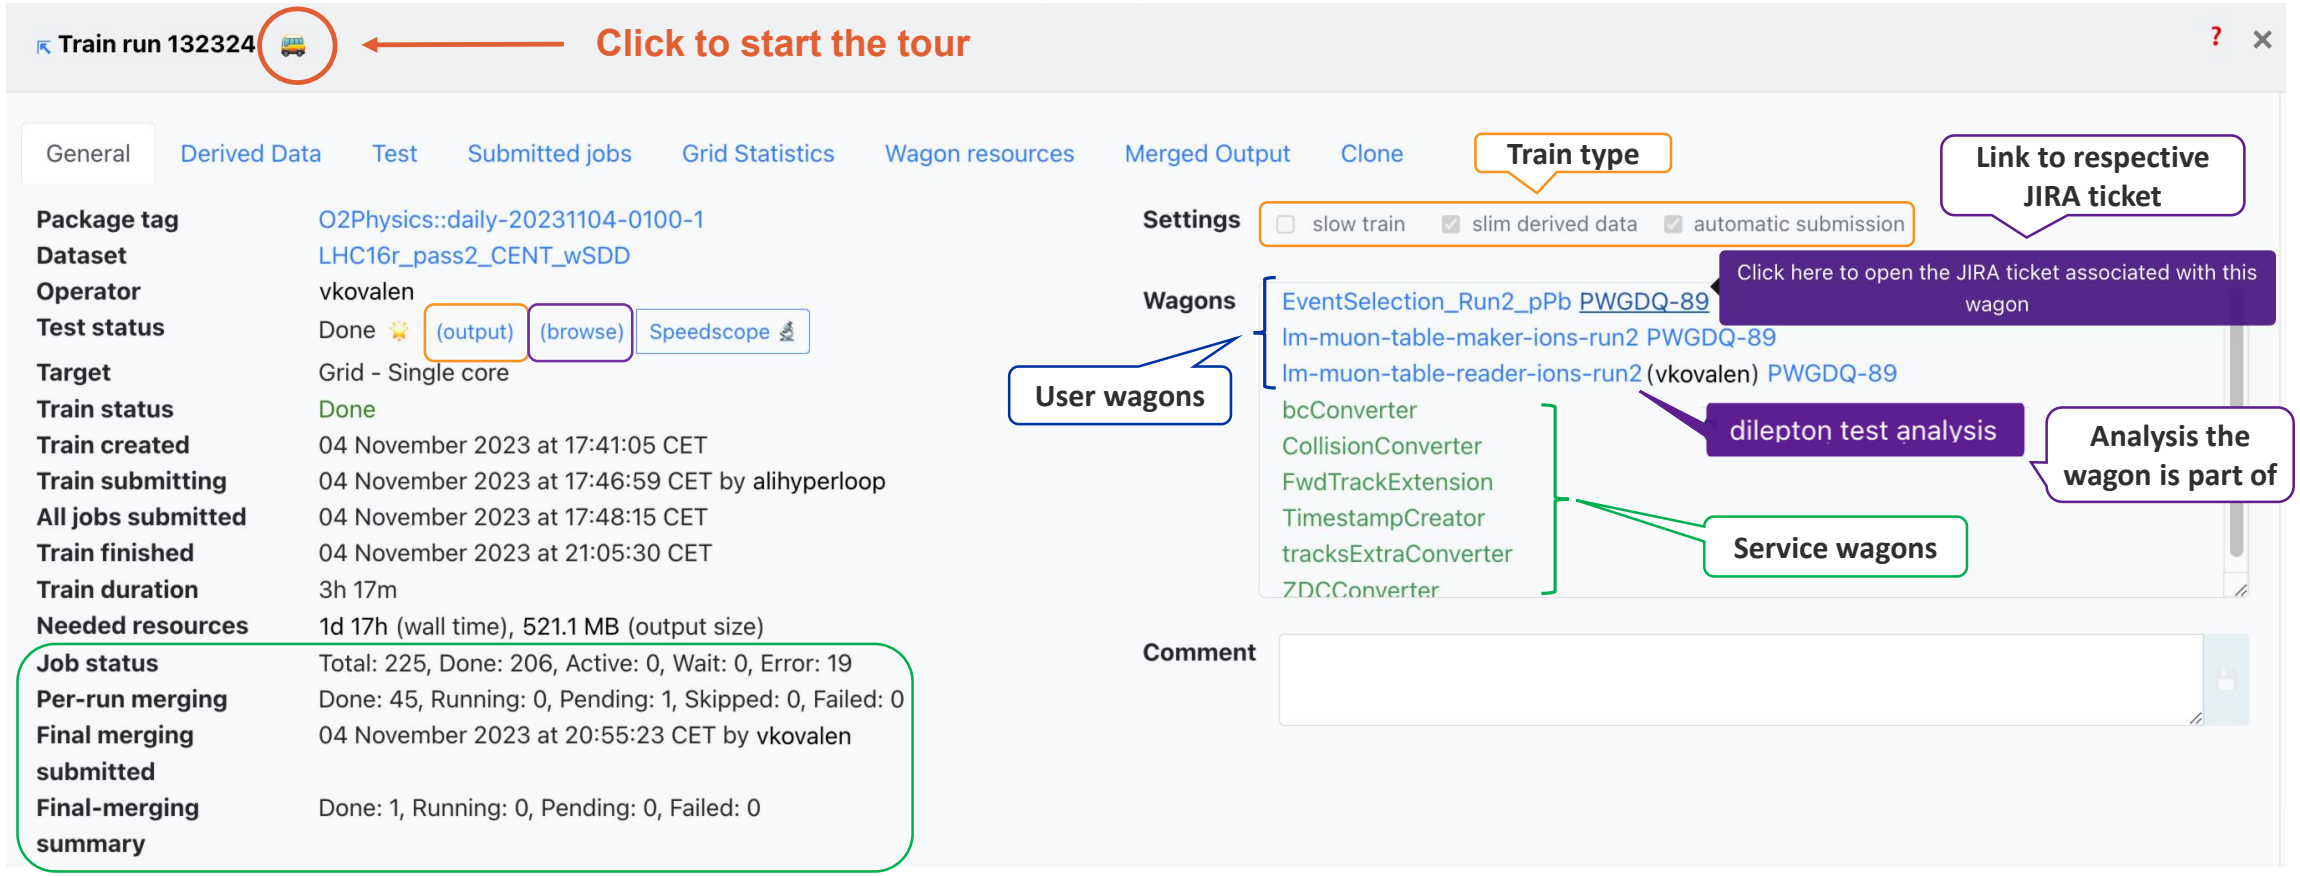

- Direct links to dataset, package tag and participating wagon configuration
- Train settings, progress timeline and current state
- Duration and needed resources

- Job status, per-run merging and final-merging summary
- Direct links to **test output**, **task profiling**, **AnalysisResults.root**, **JIRA ticket**

Train run view – Derived data

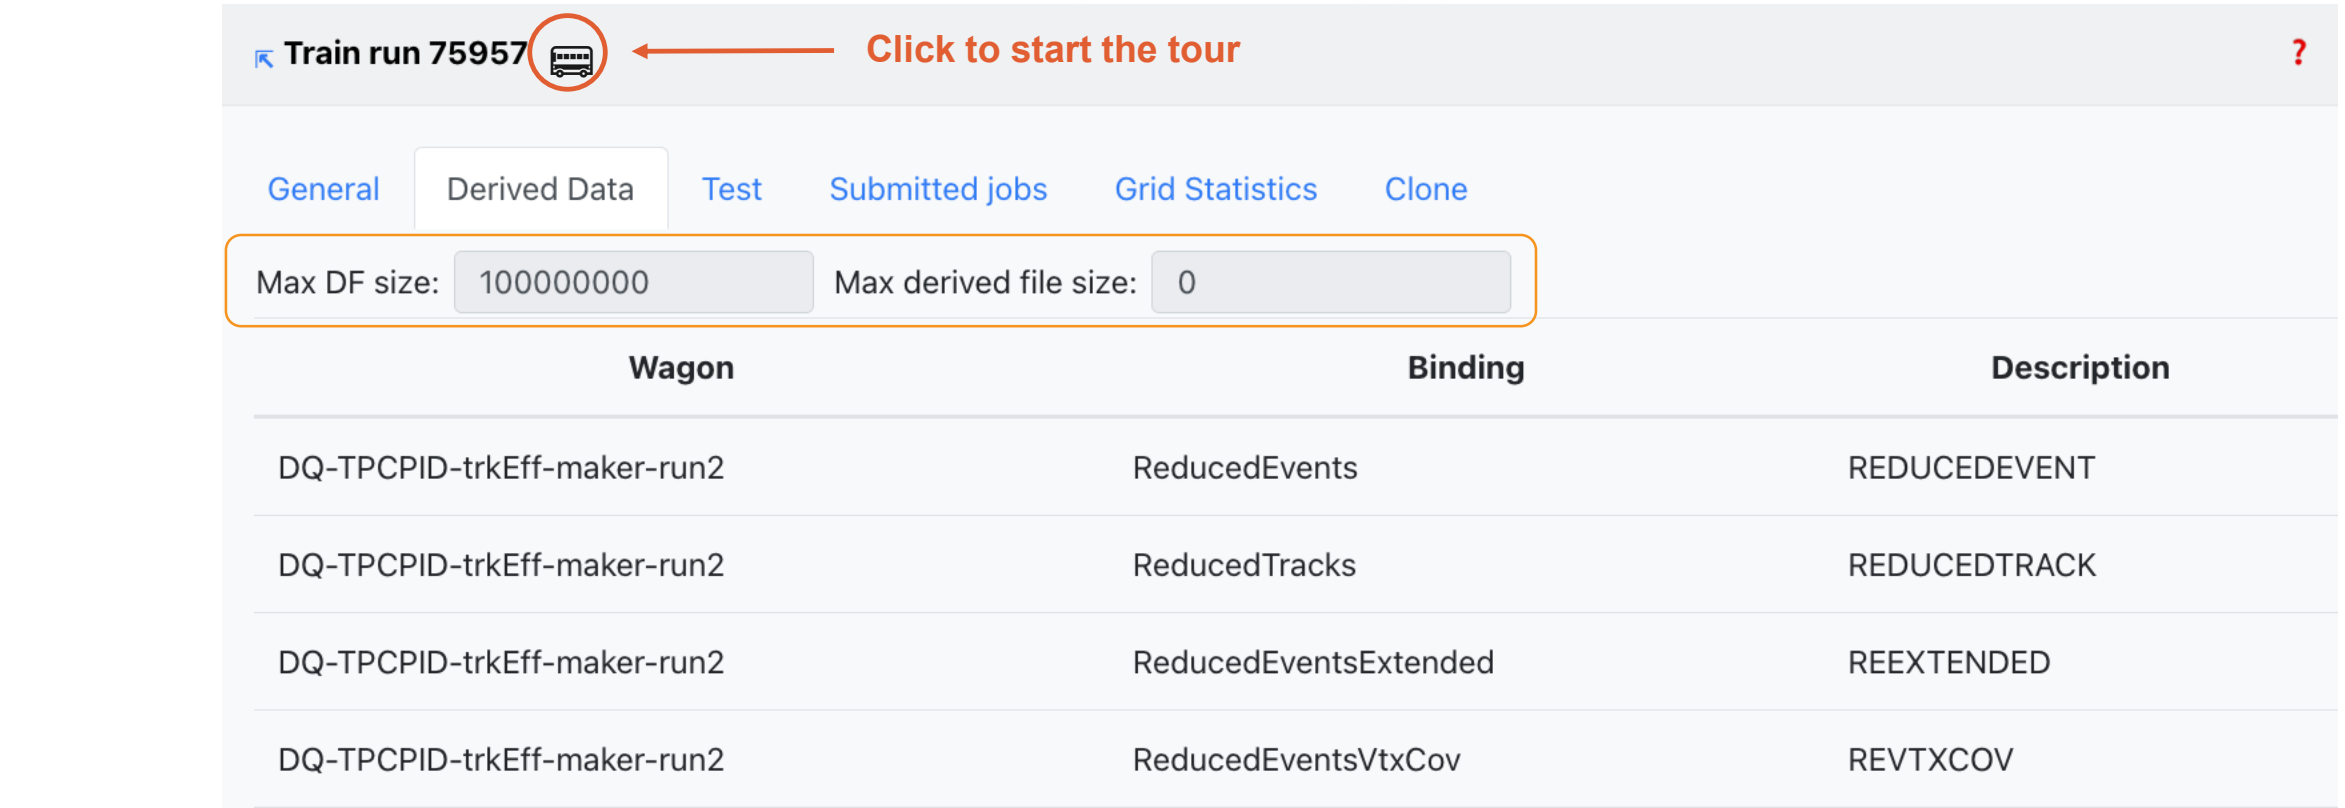

- Displays the activated derived data output
- **Max DF size and max derived file size**

Train run view – Submitted jobs

- **List of jobs**
- Use the **Search box** to search for *PIDs*, *Run number* and *Output directories*
- Download **job json configuration**
- **IO statistics**
- **Stack trace**

Copy the entire list of output directories (regardless of the active filters)

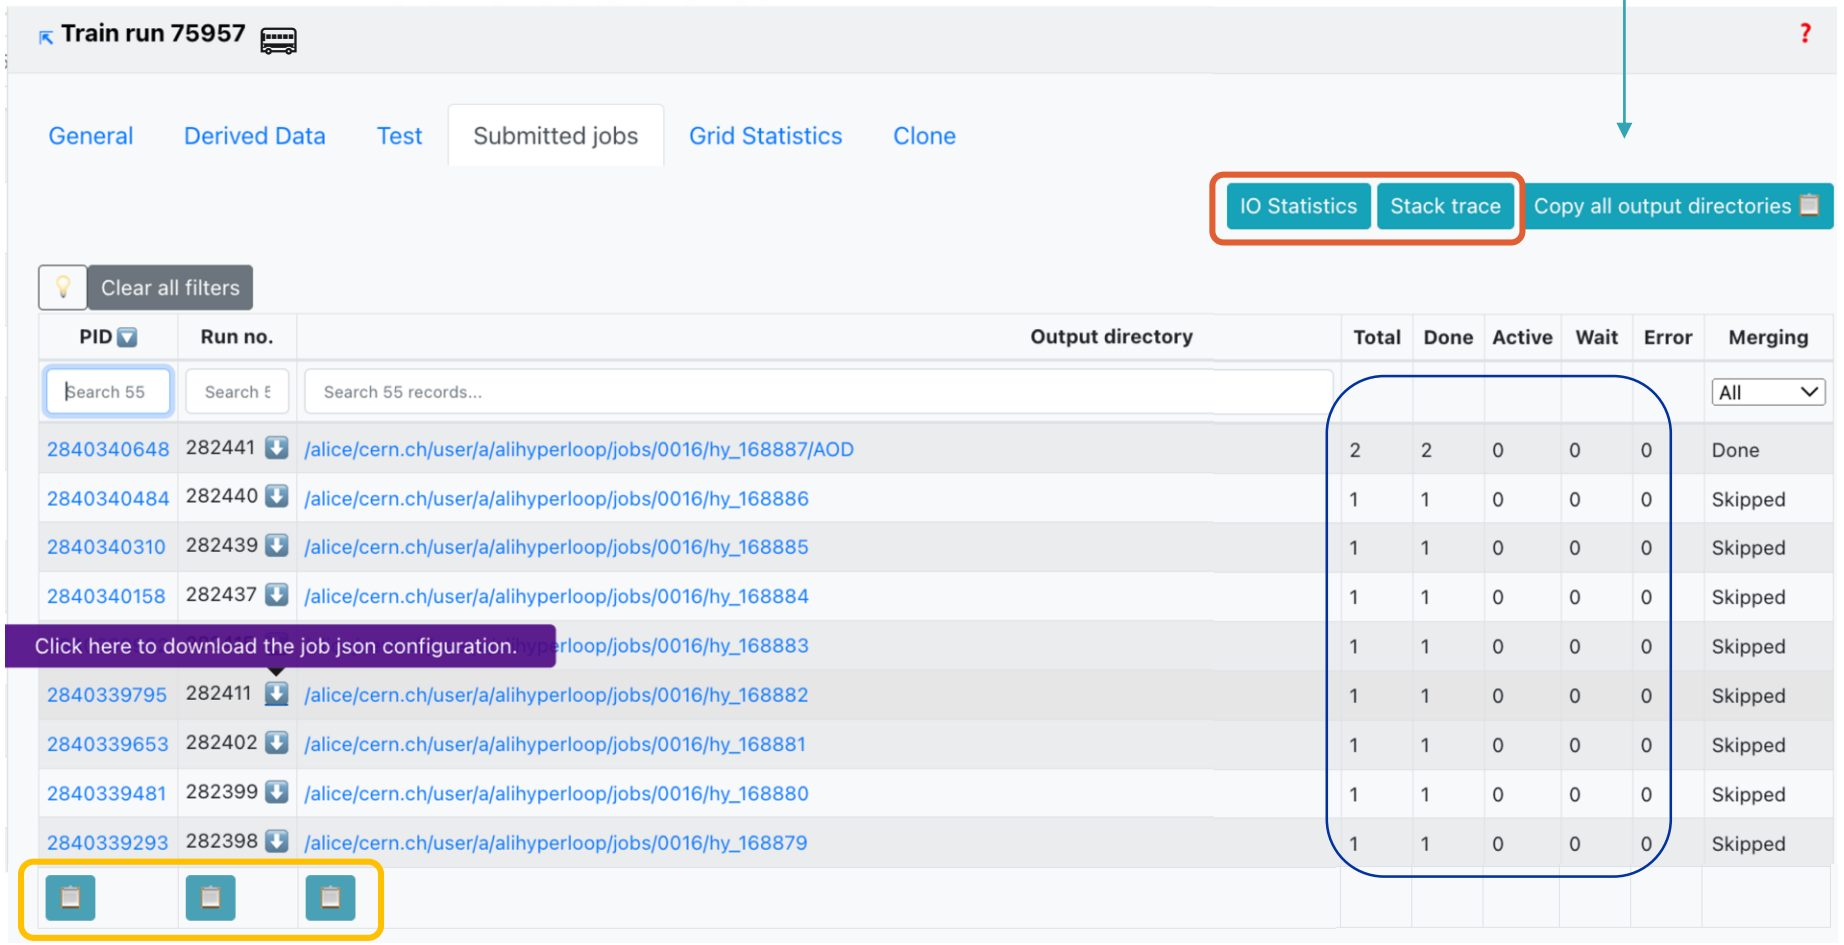

- Search and copy the filtered lists by clicking on the dedicated buttons

Train run view – Submitted jobs

- IO statistics

Train run 75967

General | Test | Submitted jobs | Grid Statistics | Wagon resources | Merged Output | Clone

IO Statistics | Stack trace | Copy all output directories

Clear all filters

PID	Run no.	Output directory	Total	Done	Active	Wait	Error	Merging
Search 12	Search 1:	Search 12 records...						All
2840407805	529043	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168943	190	187	0	0	3	Done
2840407804	529039	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168942	598	592	0	0	6	Done
2840407786	529038	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168941	240	224	0	0	16	Done
2840407762	529037	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168940	48	47	0	0	1	Done
		/0016/hy_168939	6	6	0	0	0	Done
		/0016/hy_168938	9	9	0	0	0	Done
		/0016/hy_168937	9	9	0	0	0	Done
		/0016/hy_168936	8	8	0	0	0	Done
		/0016/hy_168935	18	18	0	0	0	Done
		/0016/hy_168934	17	17	0	0	0	Done
		/0016/hy_168933	31	31	0	0	0	Done
		/0016/hy_168932	26	26	0	0	0	Done

IO stats for **12** master jobs executed between **20.04.2023 12:11** and **20.04.2023 13:43** (**1:31**)

Site activity					SE access		
Site	Job eff.	All files	Local files	Remote files	CERN EOS	CCIN2P3 SE	CNAF SE
CERN 1133 jobs (96.43%)	123.3%	2644 files 7.486 MB/s	2644 (100%) 7.486 MB/s		**2644 (100%) 7.486 MB/s**		
CCIN2P3 26 jobs (2.213%)	139.4%	61 files 19.3 MB/s	58 (95.08%) 23.83 MB/s	3 (4.918%) 3.244 MB/s	3 (4.918%) 3.244 MB/s	**58 (95.08%) 23.83 MB/s**	
CNAF 16 jobs (1.362%)	128.4%	38 files 17.02 MB/s	38 (100%) 17.02 MB/s				**38 (100%) 17.02 MB/s**
TOTAL 1175 jobs	123.7%	**2743 files 7.643 MB/s 3.547 TB**	**2740 (99.89%) 7.652 MB/s 3.544 TB**		**2644 (96.5%) 7.486 MB/s 3.424 TB**	**58 (2.117%) 23.83 MB/s 77.45 GB**	**38 (1.387%) 17.02 MB/s 45.62 GB**
				3 (0.109%) 3.244 MB/s 2.98 GB	**3 (100%) 3.244 MB/s 2.98 GB**		

Train run view – Submitted jobs

- **Stack trace**

Collects the errors from your failed jobs and groups them by error message.

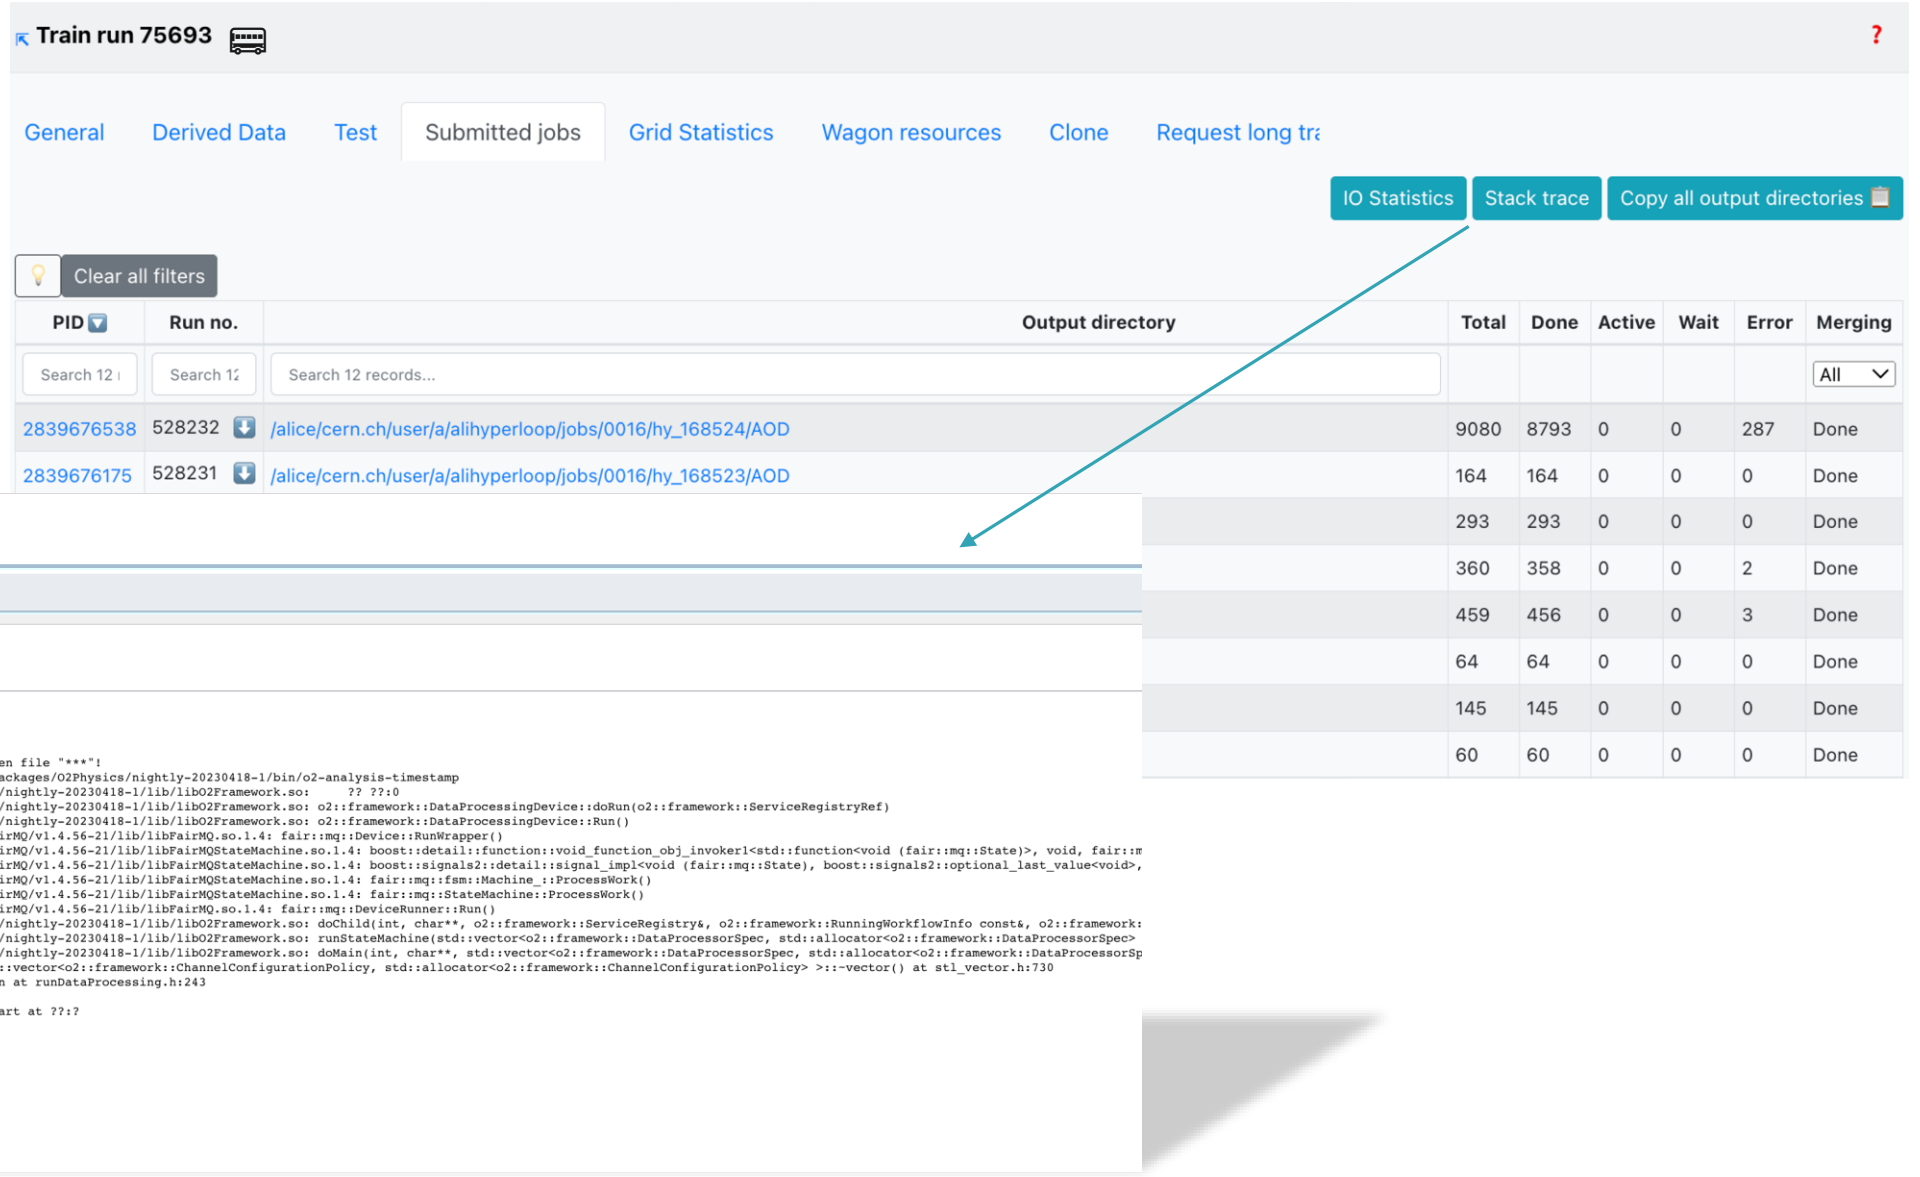

Train run view – Grid Statistics

- Summarizes the jobs performance
- Plots:
  - Files/Job
  - CPU time/Job
  - Wall time/Job

Merged Output

- Displays the jobs status after submitting the train. The mergelists are defined in the ***dataset settings***
- When the final merge is started manually by the operator, some of the runs may not be merged
- You can copy the list of **merged runs** or the total **list of runs** by clicking on the (red) number

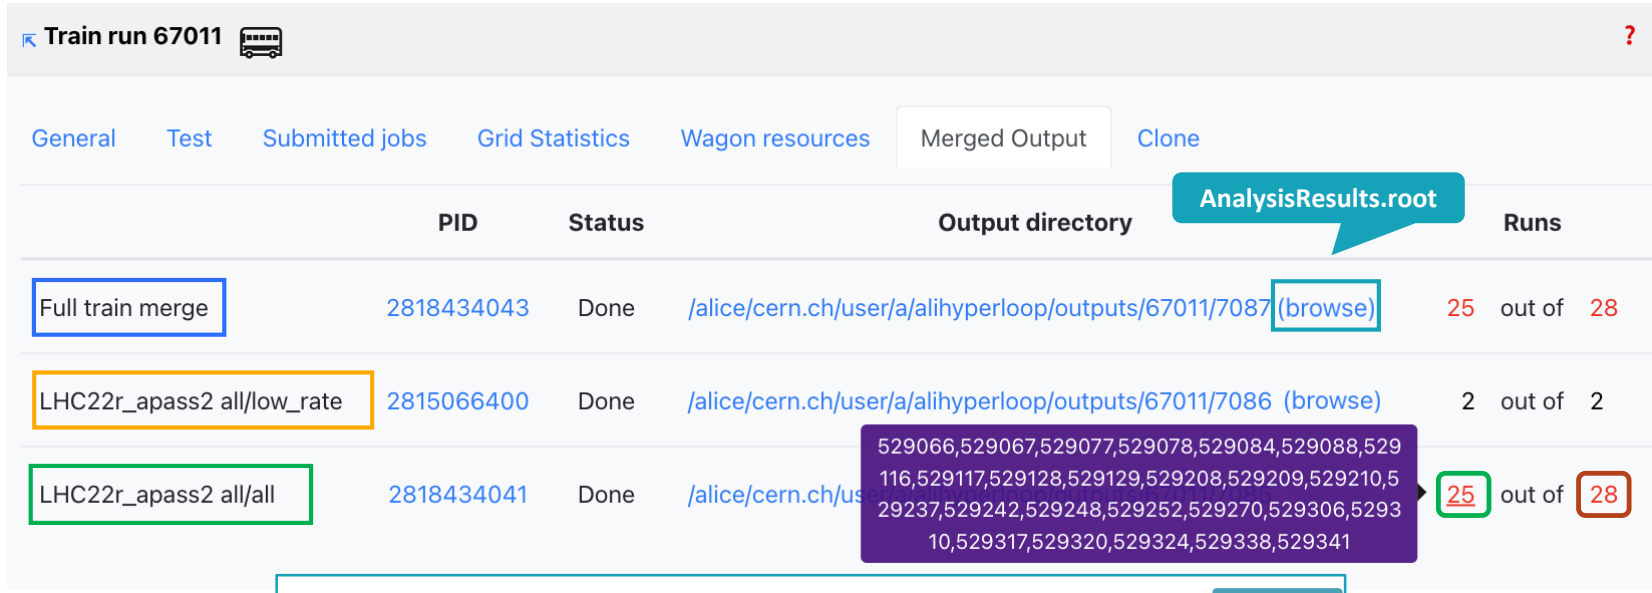

Staged submission – Long train request

- Linked datasets are subsets of a large dataset (set up in the Dataset settings)
- First, a train run needs to be **Done** on a smaller linked dataset before being run on a larger dataset

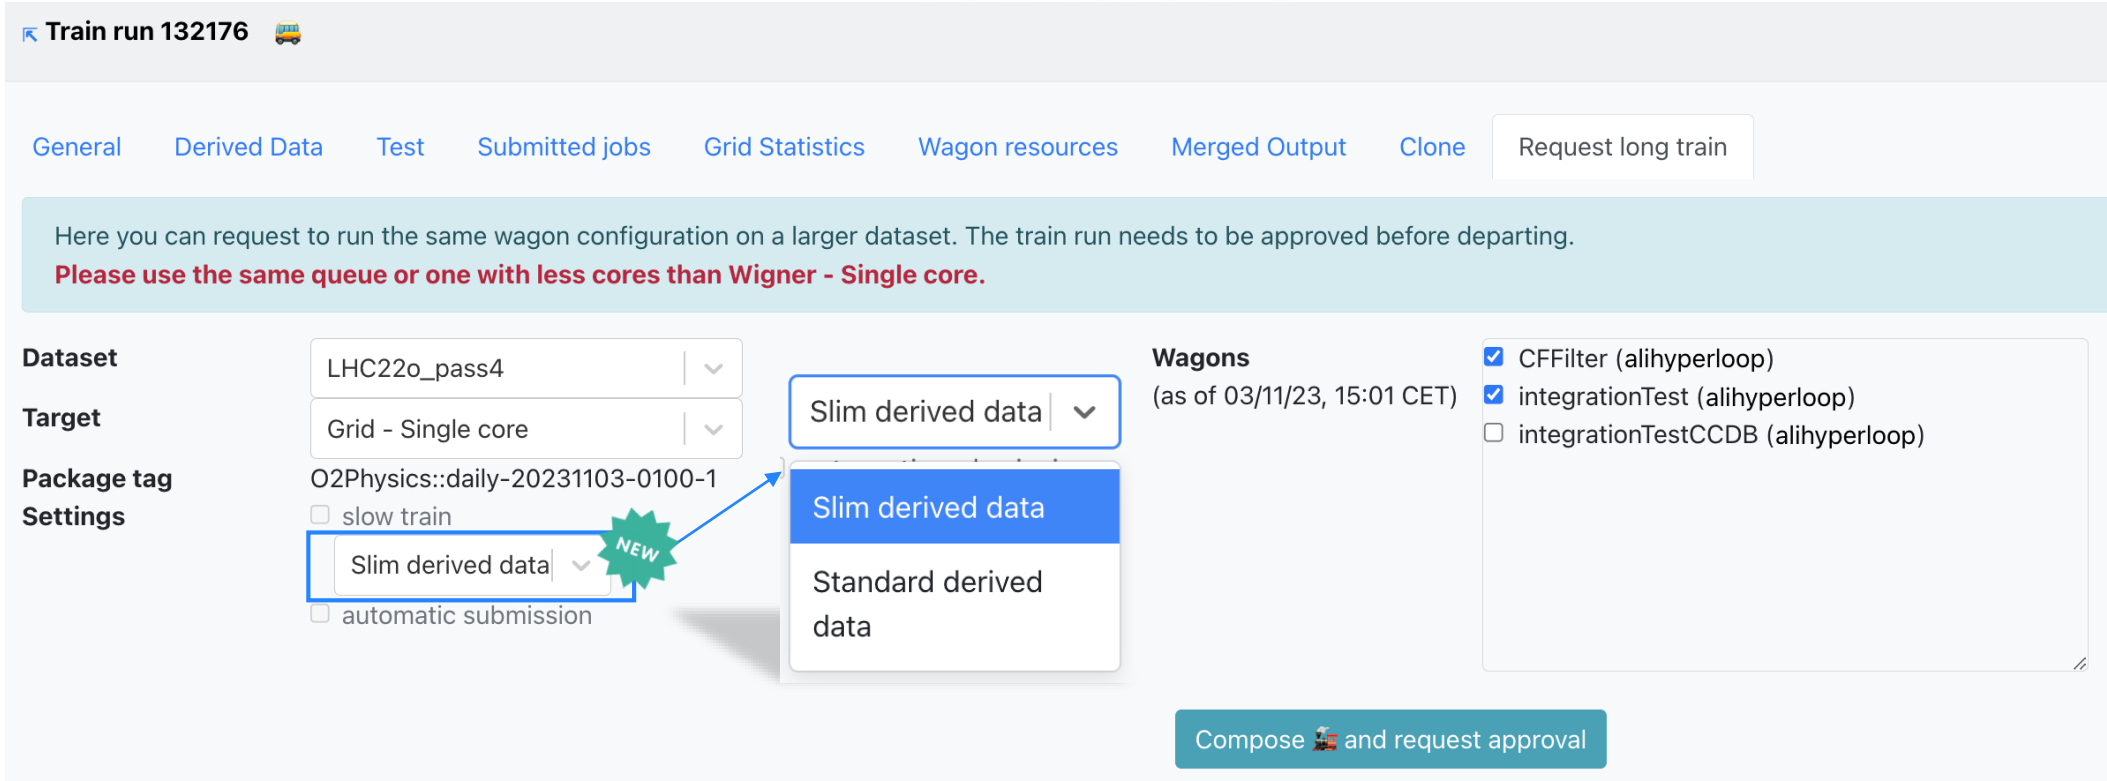

- Any user who is part of the analysis can request a long train
- You can change from slim to standard derived data (recommended if output is large)
- **Approval from the participating analyses PWGs conveners is required in order to submit a long train. Train operators or admins can also approve but usually to be done by the PWG**

Staged submission – Long train request

The user can **decompose** their own long train.

You may wish to do this when:

- The wrong wagons were selected to be part of the long train
- The selected train type is incorrect (e.g. slim instead of standard)
- The long train request was created by mistake

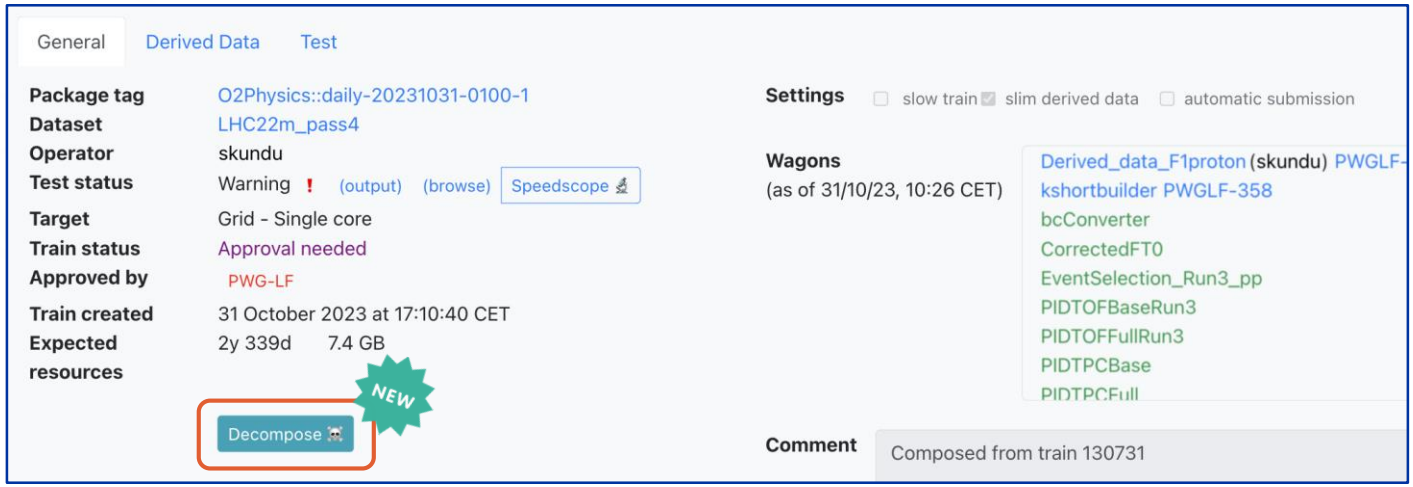

Once the long train is approved:

- If **Automatic Submission** is enabled and the train test finished *without memory warning* and *within the CPU time limit*, the train is submitted
- **Otherwise**, you may request the submission by the operators on the Operation channel

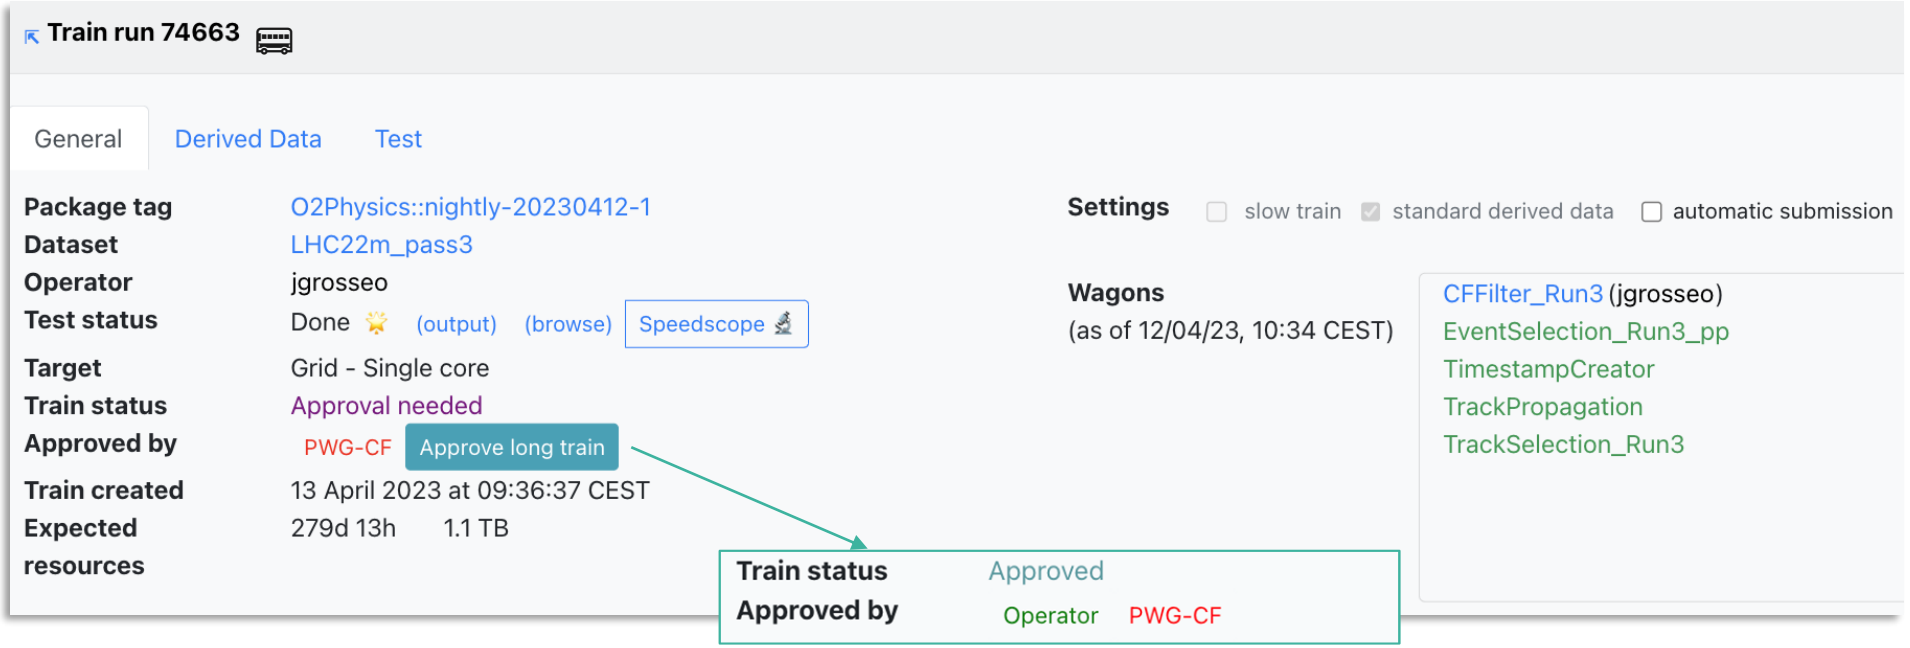

Dataset view

- **Options**
  - Allowed parent dataset access level (HY)
  - Parent datasets (HY) => size with parents NEW
  - **Dataset sampling** NEW
- **Linked datasets** (staged submission)
- **Staging info**, dataset size
- **Composition schedule**
  - CPU limit
  - Max trains/analysis/week
- **Productions**
  - Link to production details in MonALISA (MC and DATA)
  - Production - List of runs (x out of y selected)
  - Runs marked as used / not used
  - Mergelists marked as activated / deactivated
  - Total number of runs per mergelist

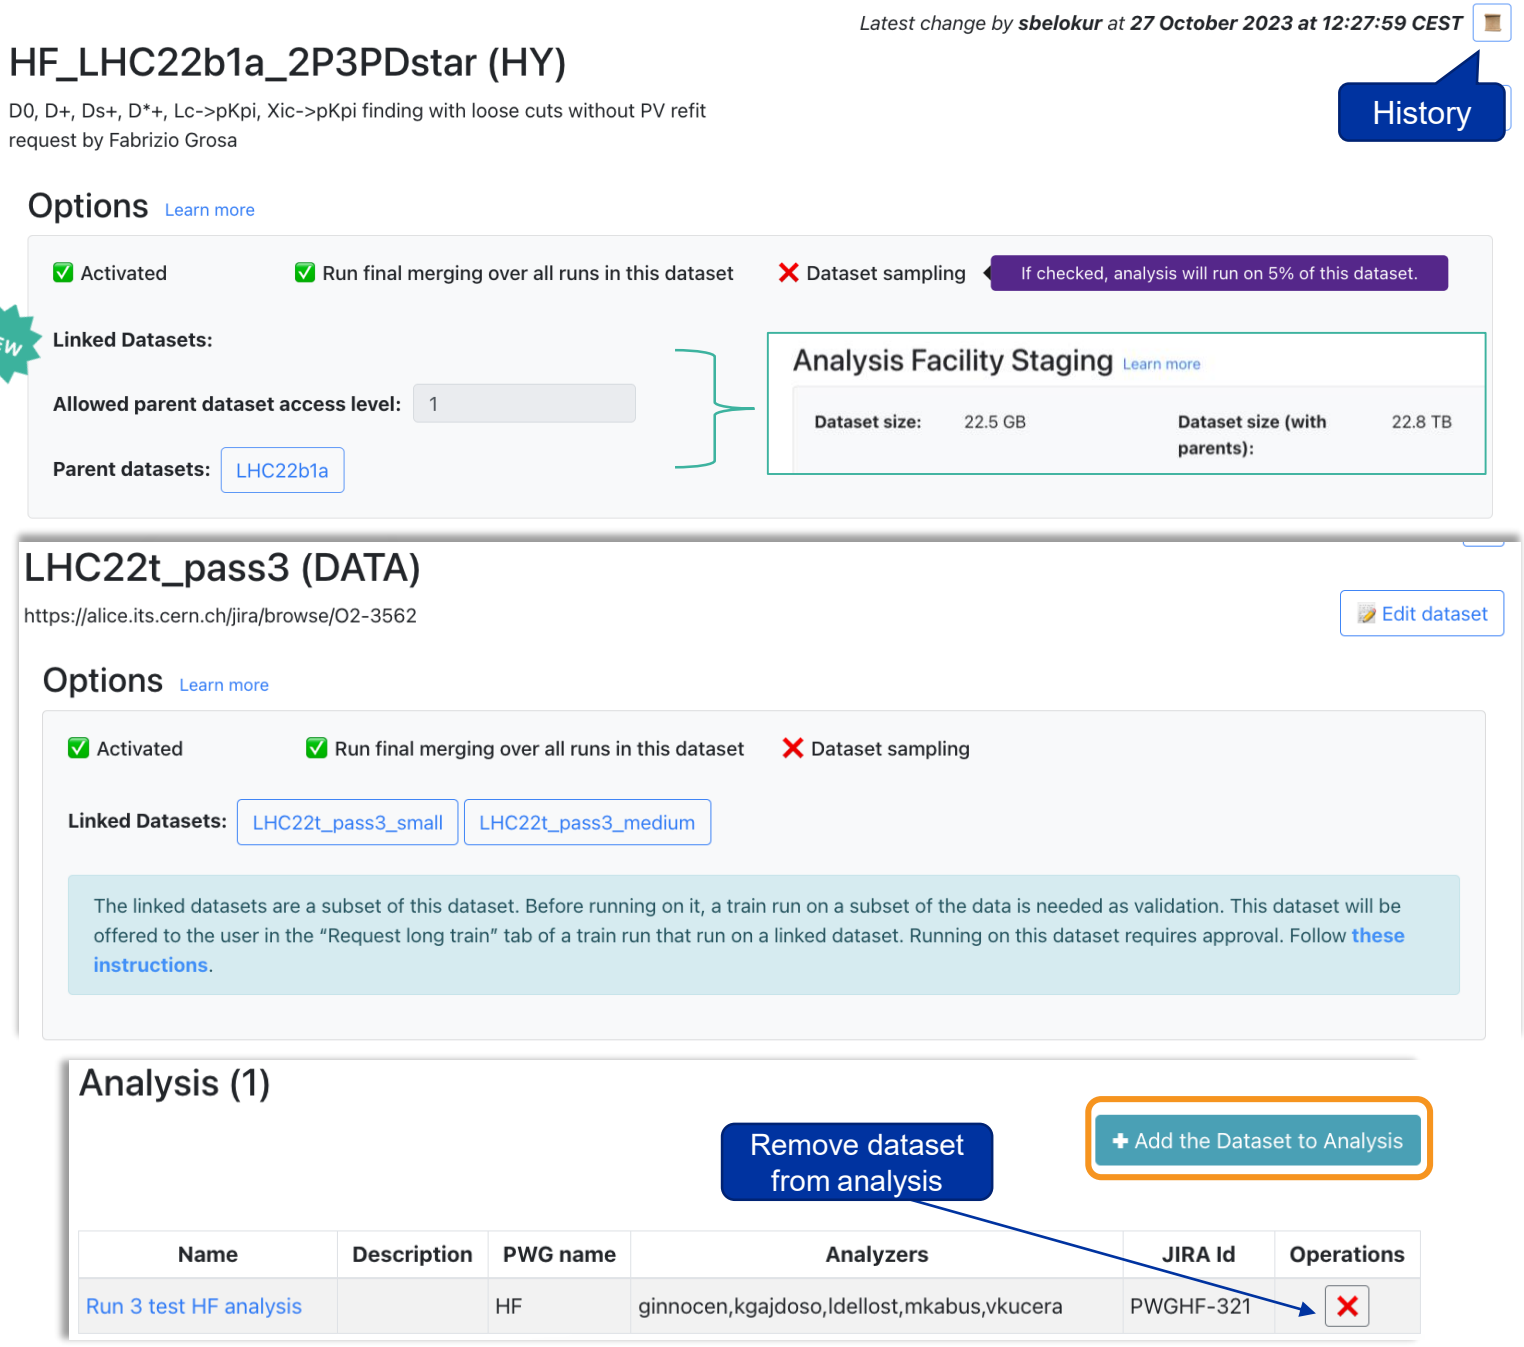

Dataset view

- **Options**
  - Allowed parent dataset access level (HY)
  - Parent datasets (HY) => size with parents NEW
  - **Dataset sampling** NEW
- **Linked datasets** (staged submission)
- **Staging info**, dataset size
- **Composition schedule**
  - CPU limit
  - Max trains/analysis/week
- **Productions**
  - Link to production details in MonALISA (MC and DATA)
  - Production - List of runs (x out of y selected)
  - Runs marked as used / not used
  - Mergelists marked as activated / deactivated
  - Total number of runs per mergelist

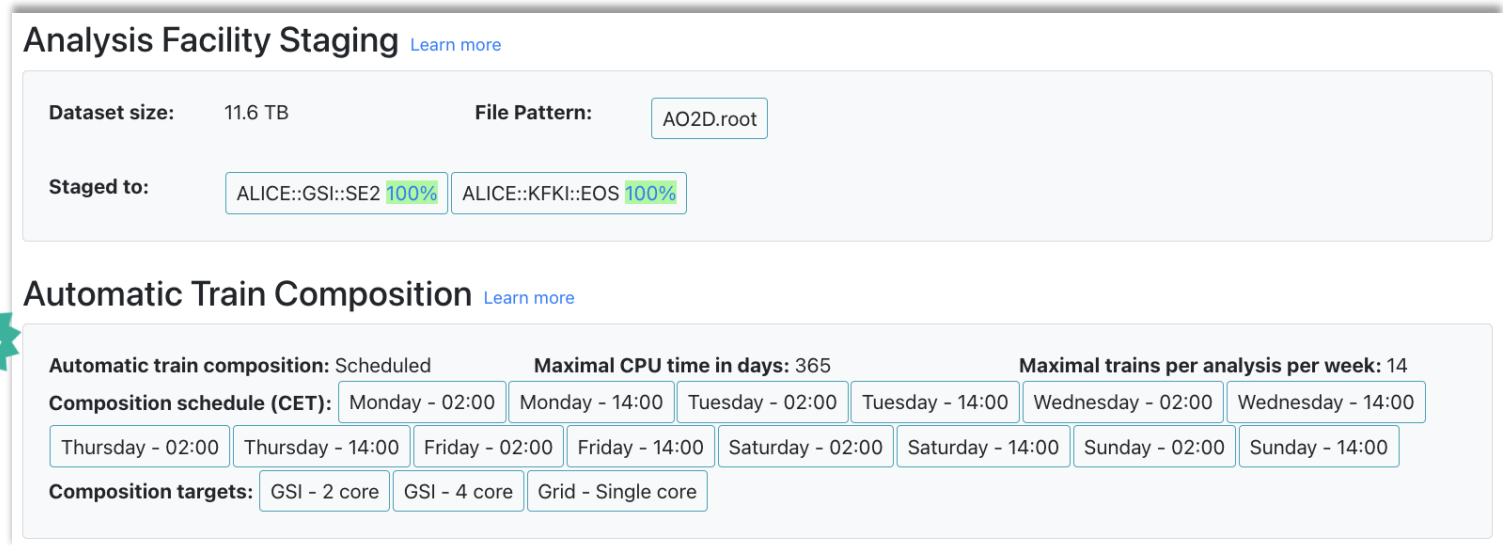

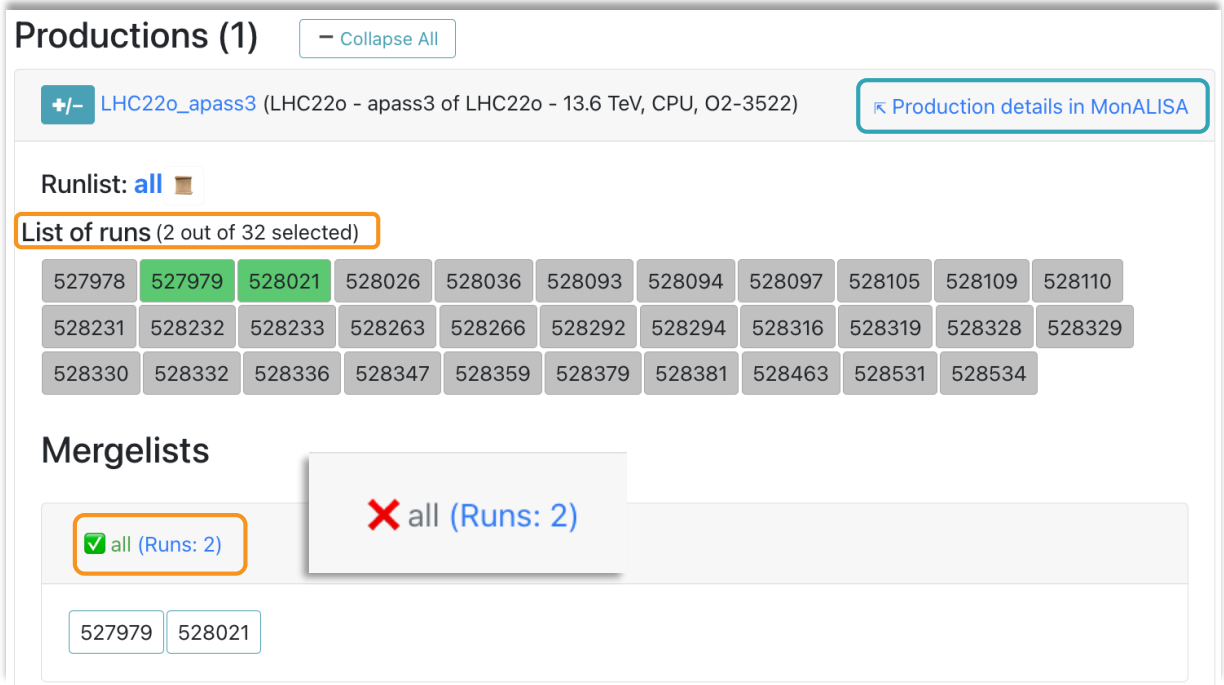

Message of the day

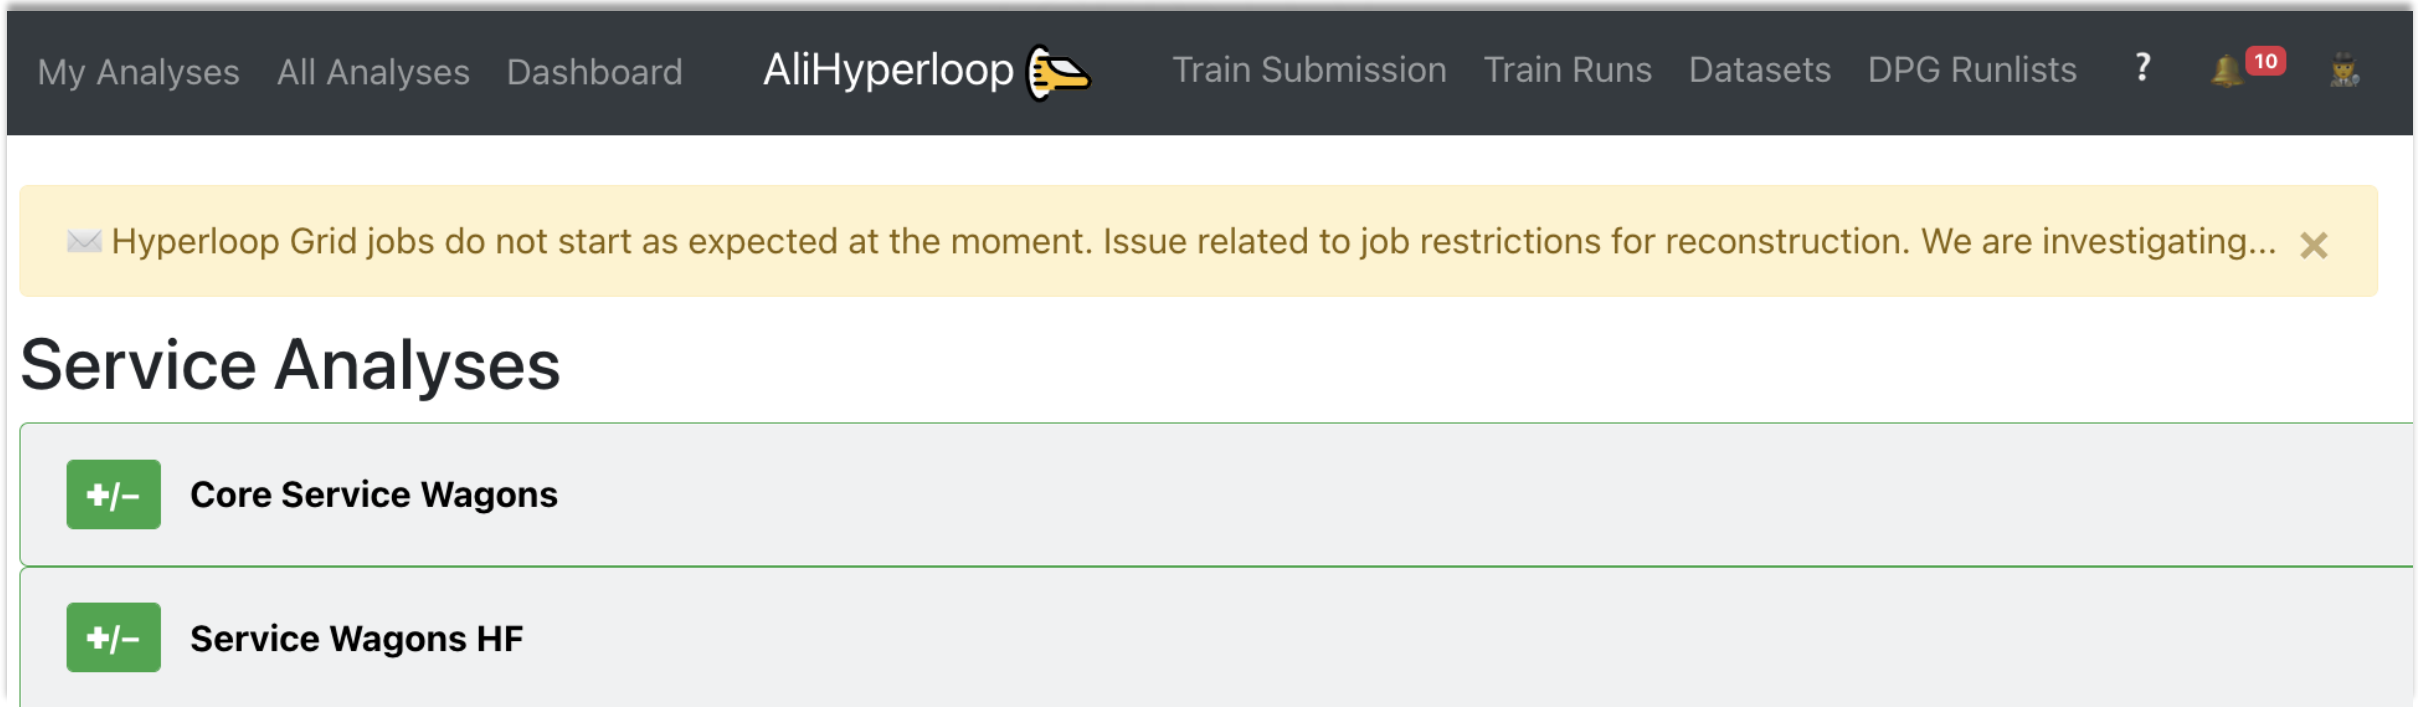

General news / issues concerning the system appear on top of the page and can be dismissed

Look out for new deployments

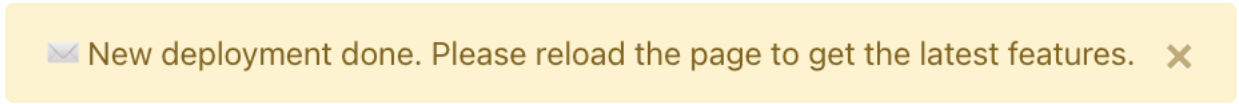

Notifications

- There are different types of notifications that the user can receive.
- Displayed in the *My Analyses* page and in the *Notifications* page, accessed by clicking in the menu bar.
- Not all notifications require an action, there can also be updates and warnings
- You can click the ✕ button to remove a notification.
- To remove all the notifications, go to the Notifications page, and click the ✕ Dismiss all button.

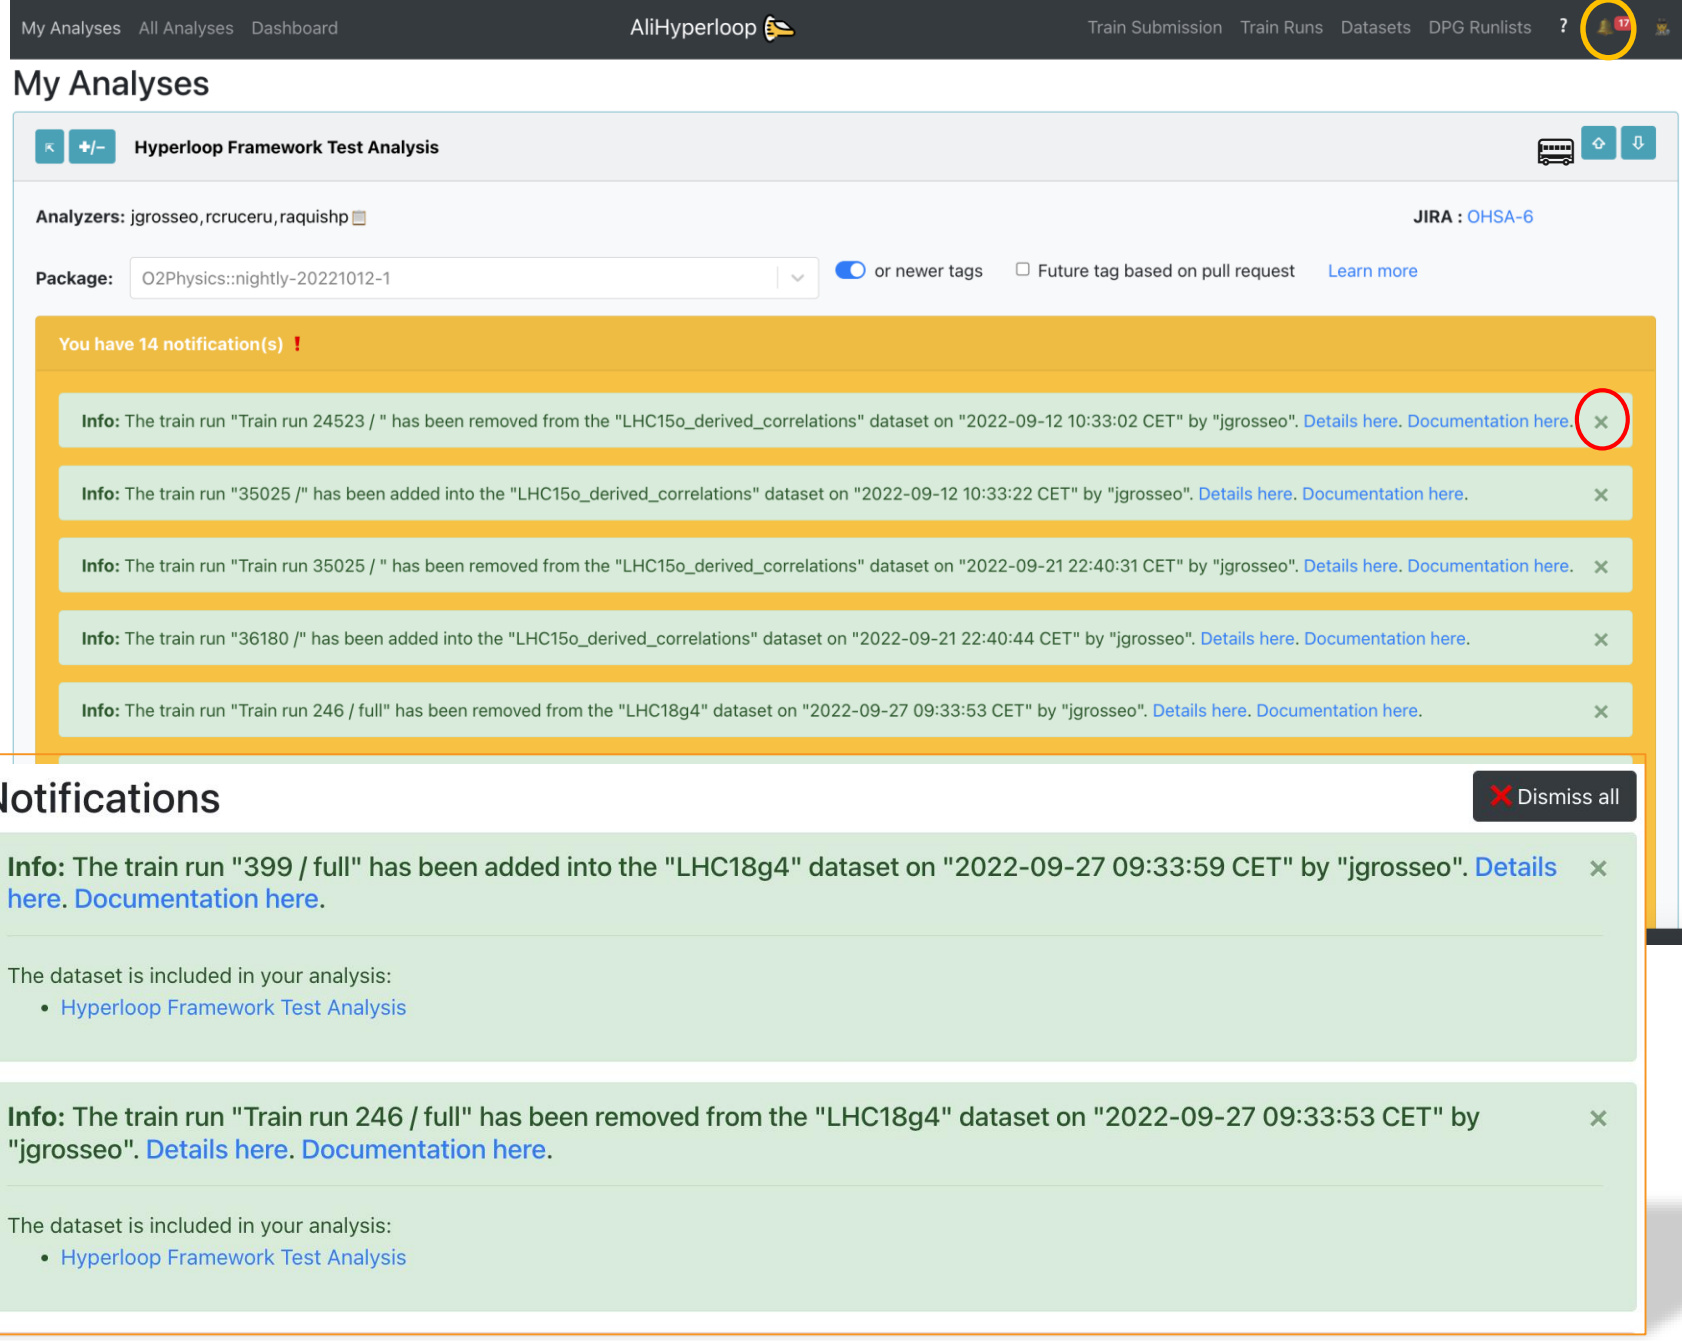

Notifications – Service wagons

- **When a service wagon is updated:**
  - Tests are reset: for all enabled wagons that depend on it (e.g. use it as dependency)
  - Users who own wagons that depend on the respective service wagon will receive the following notification, with direct links to:
    - Latest service wagon configuration
    - Dependent wagons' configurations and the analyses they are part of
    - Details of the change in question (wagon changelog)

Info: The configuration of the service wagon EventSelection was adjusted. This is a dependency of your wagon(s): correlations_depends_on_centrality_clone. The change was made on "2023-10-13 19:29:55 CET" by "rcruceru". Details here.

This is relevant to your analysis:

- HF O2 developments for ALICE3 pp Open HF 2.0 T

Documentation

- Accessed from the Navbar or everywhere across Hyperloop where you see the ? icon
- Can also be accessed from the notifications

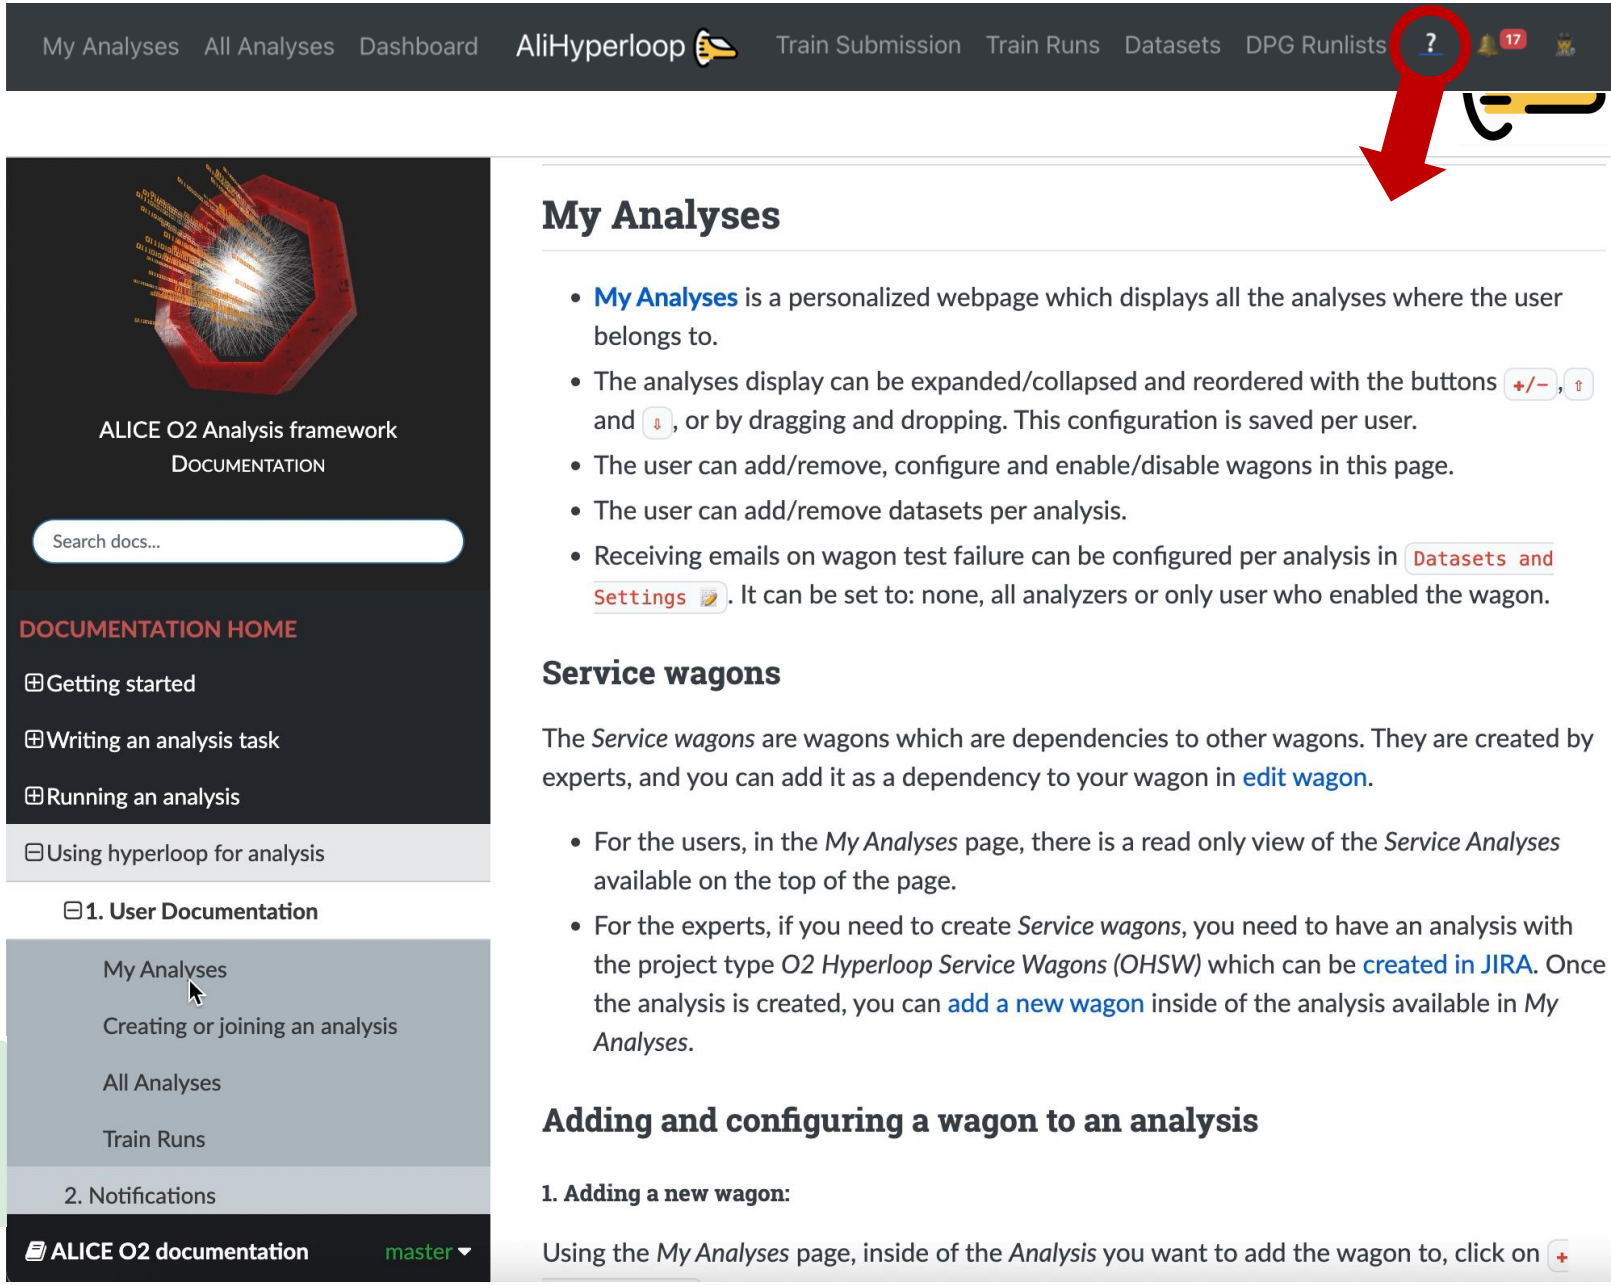

Summary

- Make use of the **My Analyses** page to see your analyses, enable *datasets*, and *choose package tag*
- Create / Edit / Delete / Clone / Compare your wagons
- **Enable** your wagons to start the **immediate testing**
- Wait for ***automatic composition*** (if enabled for the dataset, wagon test is done and no derived data , unless slim) **or**
- Request train composition in **O2 Hyperloop Operation**
- Follow up / compare trains in My Analyses or Train runs page
- Checkout the **notifications** (remember not all require action)
- Heads up to the message of the day
- Make use of the **?** button to check the documentation
- When in doubt, hop on the bus!

Back-up Slides

Ongoing trains

Ongoing Trains from last 2 weeks (ready, submitted)

Show done and killed | All trains | Trains with issues

Clear all filters | Compare | Open

Train	Job Status	Operator	Package	Dataset	Composed	Train status	Test
	Search 24 records...	Search 24	Search 24	Search 24 records...	07/03/23, 14:01 CET Off	All	All
70223	Total: 4145, Done: 164, Active: 2482, Wait: 1431, Error: 68 Done: 0, Running: 0, Pending: 9, Skipped: 0, Failed: 0 Done: 0, Running: 0, Pending: 1, Failed: 0	alihyperloop	nightly-20230316-1	LHC22p_pass2	17/03/23, 09:01 CET	Submitted	
70222	Total: 2810, Done: 139, Active: 97, Wait: 2573, Error: 1 Done: 0, Running: 0, Pending: 12, Skipped: 0, Failed: 0	rsingh	nightly-20230316-1	LHC22q_pass3	17/03/23, 07:16 CET	Submitted	
70221	Total: 1368, Done: 544, Active: 665, Wait: 18 Done: 0, Running: 0, Pending: 1, Skipped: 0,						
70220	Total: 85, Done: 76, Active: 8, Wait: 0, Error: Done: 5, Running: 1, Pending: 5, Skipped: 0, Done: 0, Running: 0, Pending: 1, Failed: 0						

Ongoing Trains from last 2 weeks (ready, submitted, done, killed)

Show done and killed | All trains | Trains with issues

Clear all filters | Compare | Open

Train	Job Status	Operator	Package	Dataset	Composed	Train status	Test
	Search 295 records...	Search 29	Search 29	Search 295 records...	28/02/23, 17:20 CET Off	Done	All
70211	Total: 63, Done: 60, Active: 0, Wait: 0, Error: 3 Done: 15, Running: 0, Pending: 0, Skipped: 0, Failed: 0 Done: 1, Running: 0, Pending: 0, Failed: 0	alihyperloop	nightly-20230316-1	LHC16i	17/03/23, 00:01 CET	Done	
70202	Total: 2498, Done: 2452, Active: 0, Wait: 0, Error: 46 Done: 1, Running: 0, Pending: 0, Skipped: 0, Failed: 0 Done: 1, Running: 0, Pending: 0, Failed: 0	nvalle	nightly-20230312-1	LHC22p_pass3_subset	16/03/23, 17:20 CET	Done	

- Accessed from Train Runs page
- Shows ongoing trains in the last 2 weeks
- Checkbox to add trains which are done and killed

Staging status

Accessed from Datasets page

Datasets

Derived data | Show staging status | Show removed datasets | + Add Dataset

Clear all filters

Name	Description	Type	Size	Staging status
pass3	Search 6 records...	DATA	Search 6 records...	
LHC22m_pass3_relval_cpu2	https://alice.its.cern.ch/jira/browse/O2-3540	DATA	13.0 TB	ALICE::GSI::SE2 99.63%
LHC22p_pass3_subset	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	20.8 TB	ALICE::KFKI::EOS 100%
LHC22q_pass3	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	46.1 TB	ALICE::GSI::SE2 100%
LHC22r_pass3_subset	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	29.7 TB	ALICE::GSI::SE2 98.33%
LHC22t_pass3_small	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	25.4 TB	ALICE::KFKI::EOS 100%
PilotBeam_pass3	async pass 3	DATA	80.1 GB	ALICE::GSI::SE2 100% ALICE::KFKI::EOS 99.99%

Completed in warning state(99.63% success)

Staging status

Accessed from Datasets page

Datasets

Derived data | Show staging status | Show removed datasets | + Add Dataset

Clear all filters

Name	Description	Type	Size	Staging status
pass3	Search 6 records...	DATA	Search 6 records...	
LHC22m_pass3_relval_cpu2	https://alice.its.cern.ch/jira/browse/O2-3540	DATA	13.0 TB	ALICE::GSI::SE2 99.63%
LHC22p_pass3_subset	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	20.8 TB	ALICE::KFKI::EOS 100%
LHC22q_pass3	https://alice.its.cern.ch/jira/browse/O2-3562	DATA	46.1 TB	ALICE::GSI::SE2 100%
				ALICE::GSI::SE2 98.33%
				ALICE::KFKI::EOS 100%
				ALICE::GSI::SE2 100%
				ALICE::KFKI::EOS 99.99%

Completed in warning state(99.63% success)

Transfer requests (add new request)

	HY masterjob 2803959817%pattern AO2D.root,2803960202%patter	ALICE::GSI::SE2	- any -					Filter
ID	Path	Target SE	Status	Progress	Files	Total size	Started	Ended
21322.	HY masterjob 2803960204 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Done		8	24.34 GB	20 Feb 2023 11:05	20 Feb 2023 19:53
21321.	HY masterjob 2803959812 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Error		286	980.8 GB	20 Feb 2023 11:05	20 Feb 2023 21:13
21320.	HY masterjob 2803960203 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Error		358	1.203 TB	20 Feb 2023 11:05	20 Feb 2023 21:08
21319.	HY masterjob 2803960201 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Error		407	1.354 TB	20 Feb 2023 11:05	20 Feb 2023 21:18
21318.	HY masterjob 2803960200 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Done		416	1.377 TB	20 Feb 2023 11:05	21 Feb 2023 00:35
21317.	HY masterjob 2803960202 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Done		415	1.384 TB	20 Feb 2023 11:05	20 Feb 2023 14:07
21316.	HY masterjob 2803960117 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Done		420	1.4 TB	20 Feb 2023 11:05	20 Feb 2023 13:34
21315.	HY masterjob 2803959817 on ALICE::GSI::SE2 filepattern AO2D.root	ALICE::GSI::SE2	Done		420	1.406 TB	20 Feb 2023 11:05	21 Feb 2023 00:35
	8 requests				**2730**	**9.106 TB**		

Dashboard

- Information updated every 10 min
- Added number of composed trains waiting to be submitted

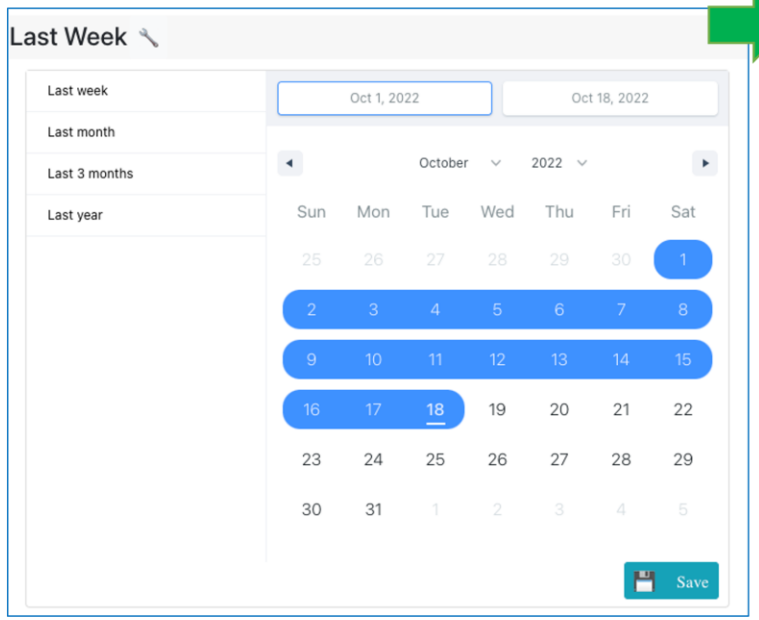

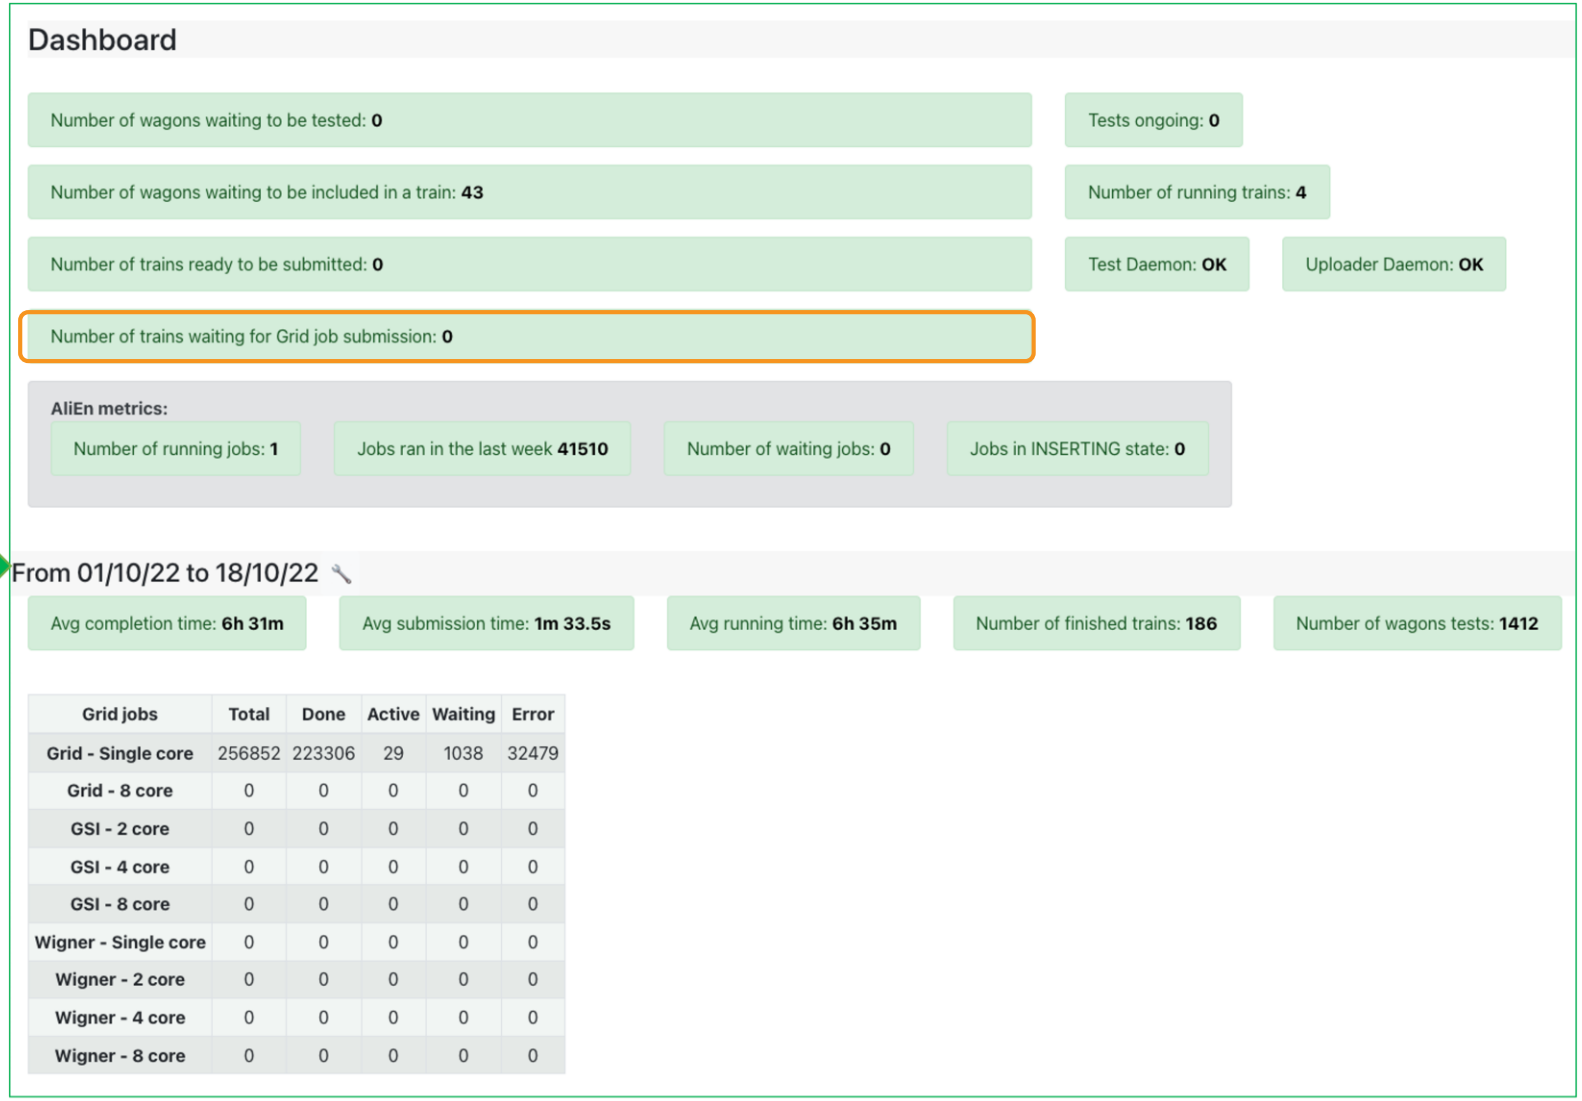

Grid jobs	Total	Done	Active	Waiting	Error
Grid - Single core	256852	223306	29	1038	32479
Grid - 8 core	0	0	0	0	0
GSI - 2 core	0	0	0	0	0
GSI - 4 core	0	0	0	0	0
GSI - 8 core	0	0	0	0	0
Wigner - Single core	0	0	0	0	0
Wigner - 2 core	0	0	0	0	0
Wigner - 4 core	0	0	0	0	0
Wigner - 8 core	0	0	0	0	0