Hyperloop - Accelerating Analysis

Nicolas Poffley (CERN)

10.11.25

- **Introduction & key concepts**
- ’My Analyses’, wagons & datasets
- Enabling a wagon to start a test
- Trains – automatic composition, debugging & long trains
- Key things to know

What is it?

- Analysis train system, allowing for cooperative large-scale analysis
- Concept: Individual workflows (‘wagons’) are combined into a ‘train’ to run over data together
- Analysis jobs processed using the **GRID** and **Analysis Facilities**
- Fully integrated with $\mathbf{O^2}$ → direct task configuration

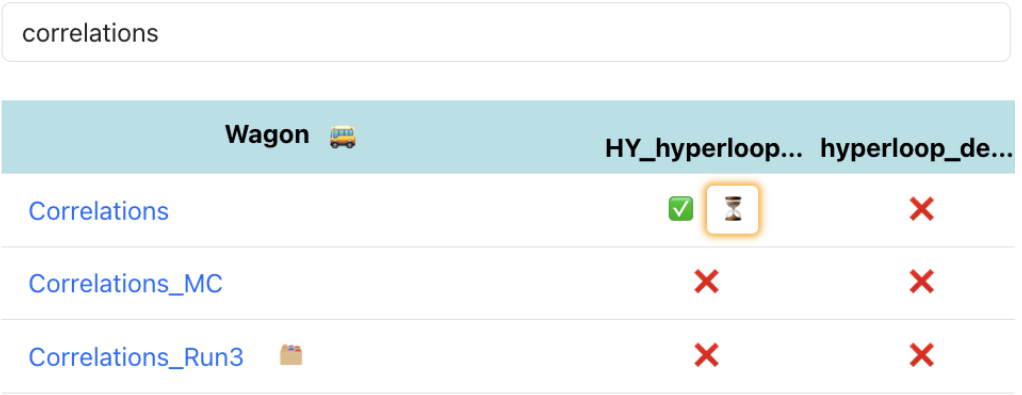

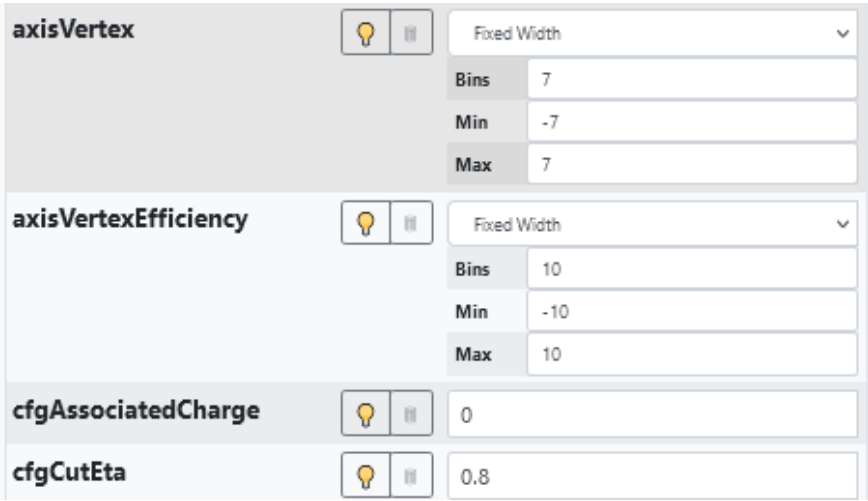

Train support

- Operators from different institutes available 24 hours per day Monday – Friday

Resources

- Detailed documentation
- Dedicated mattermost channel for requesting trains or support
- Integrated tours to guide you through on-the-go

click on: 🚌

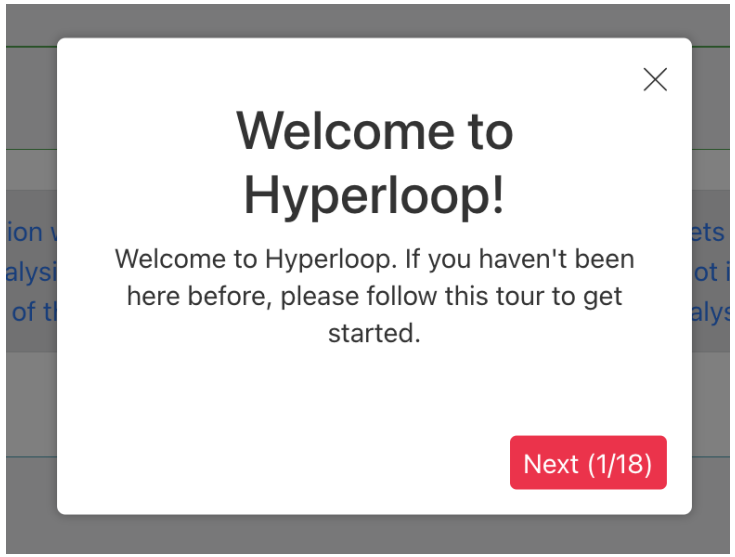

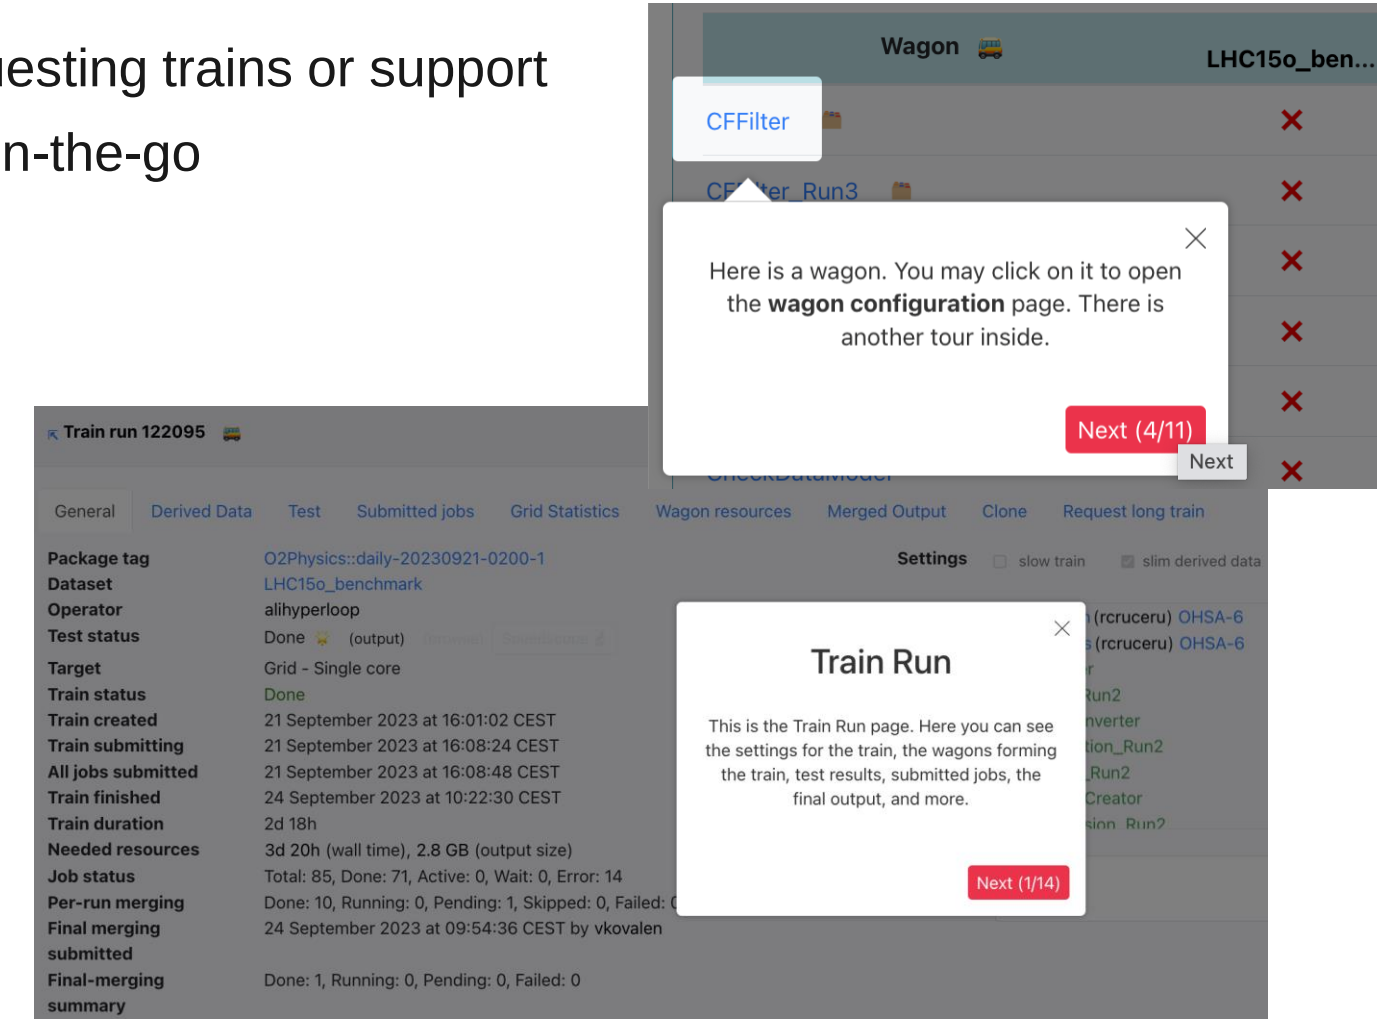

Key concept

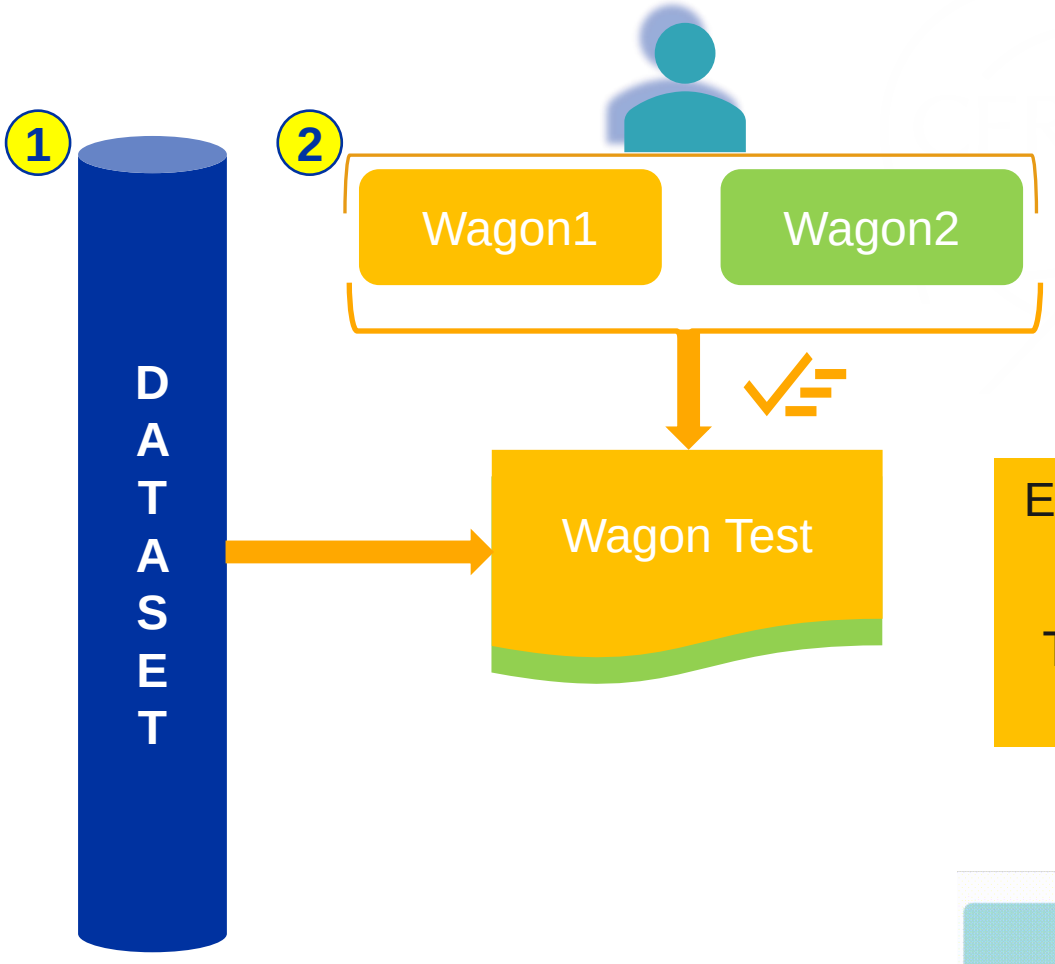

Enable your O^2 analysis workflows, known as wagons, on a dataset to trigger a **wagon test**.

Tests ensure that your workflows are valid and aren't too resource intensive.

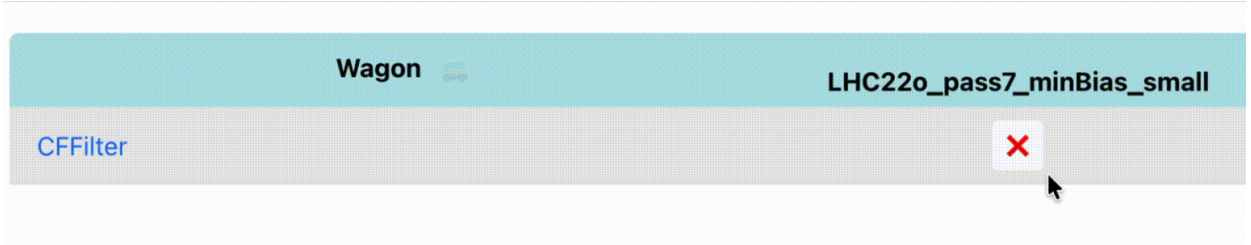

Key concept

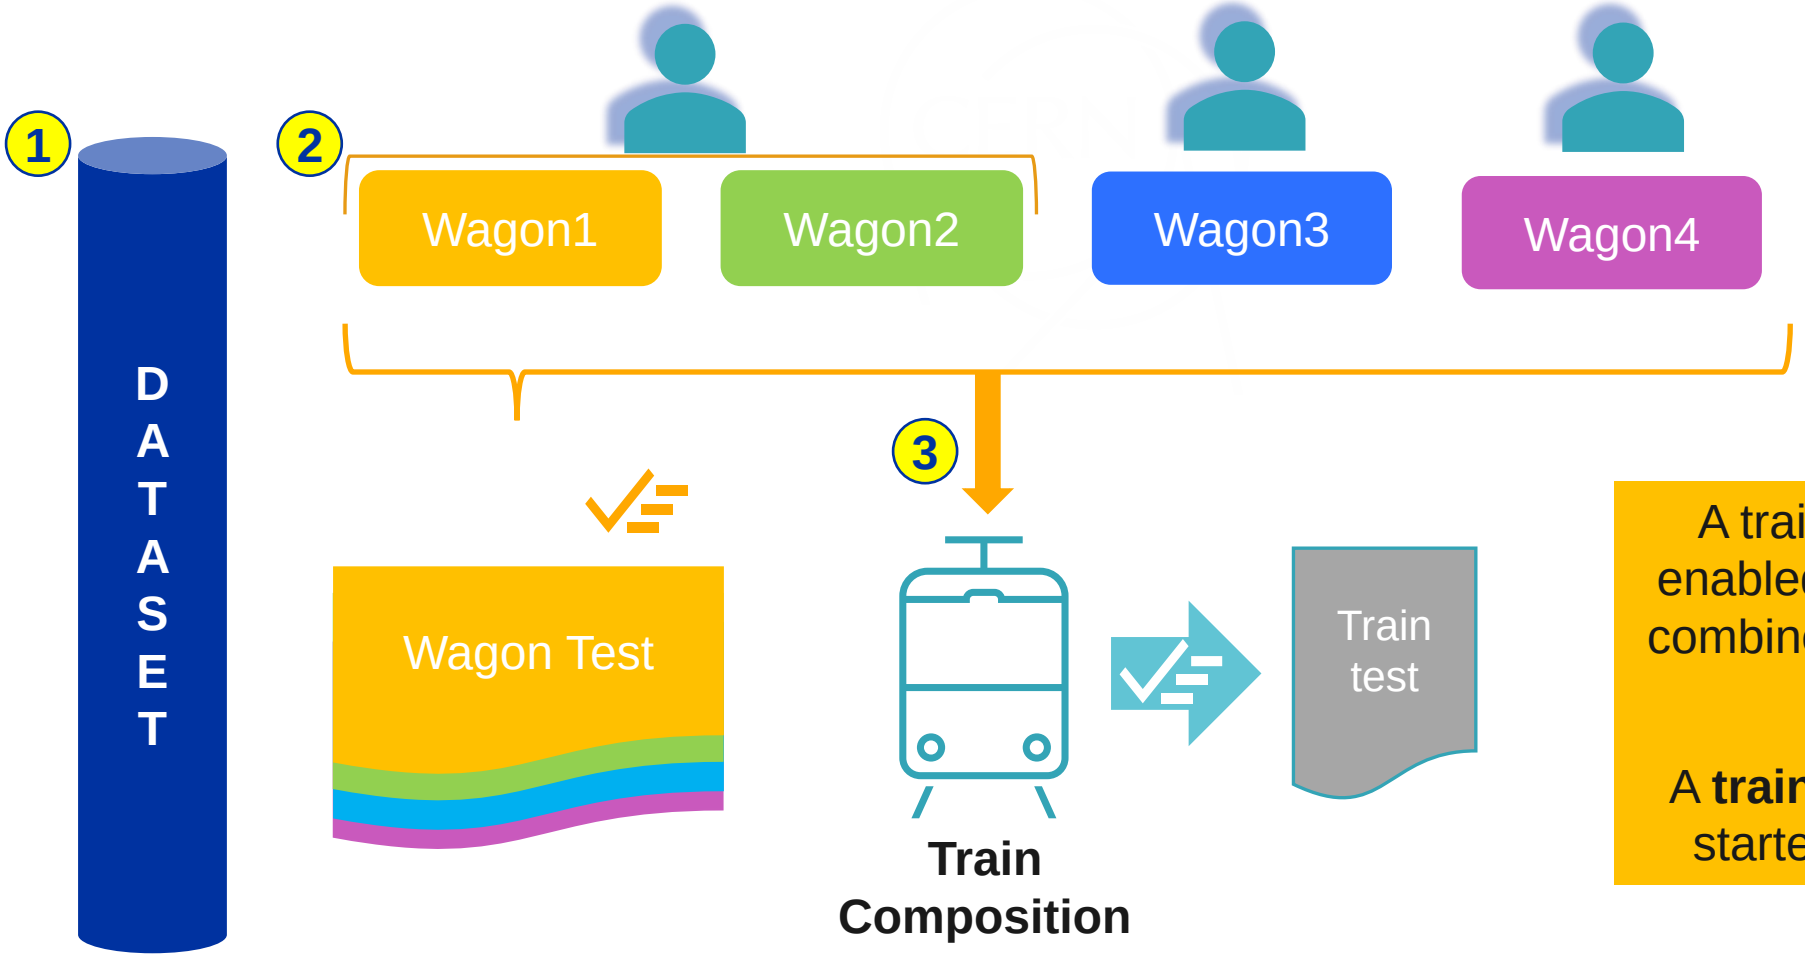

A train is composed from enabled wagons. Trains may combine multiple wagons into one train.

A **train test** is automatically started after composition.

Individual O^2 analysis workflows - known as wagons - are combined into trains to run over the data together

Key concept

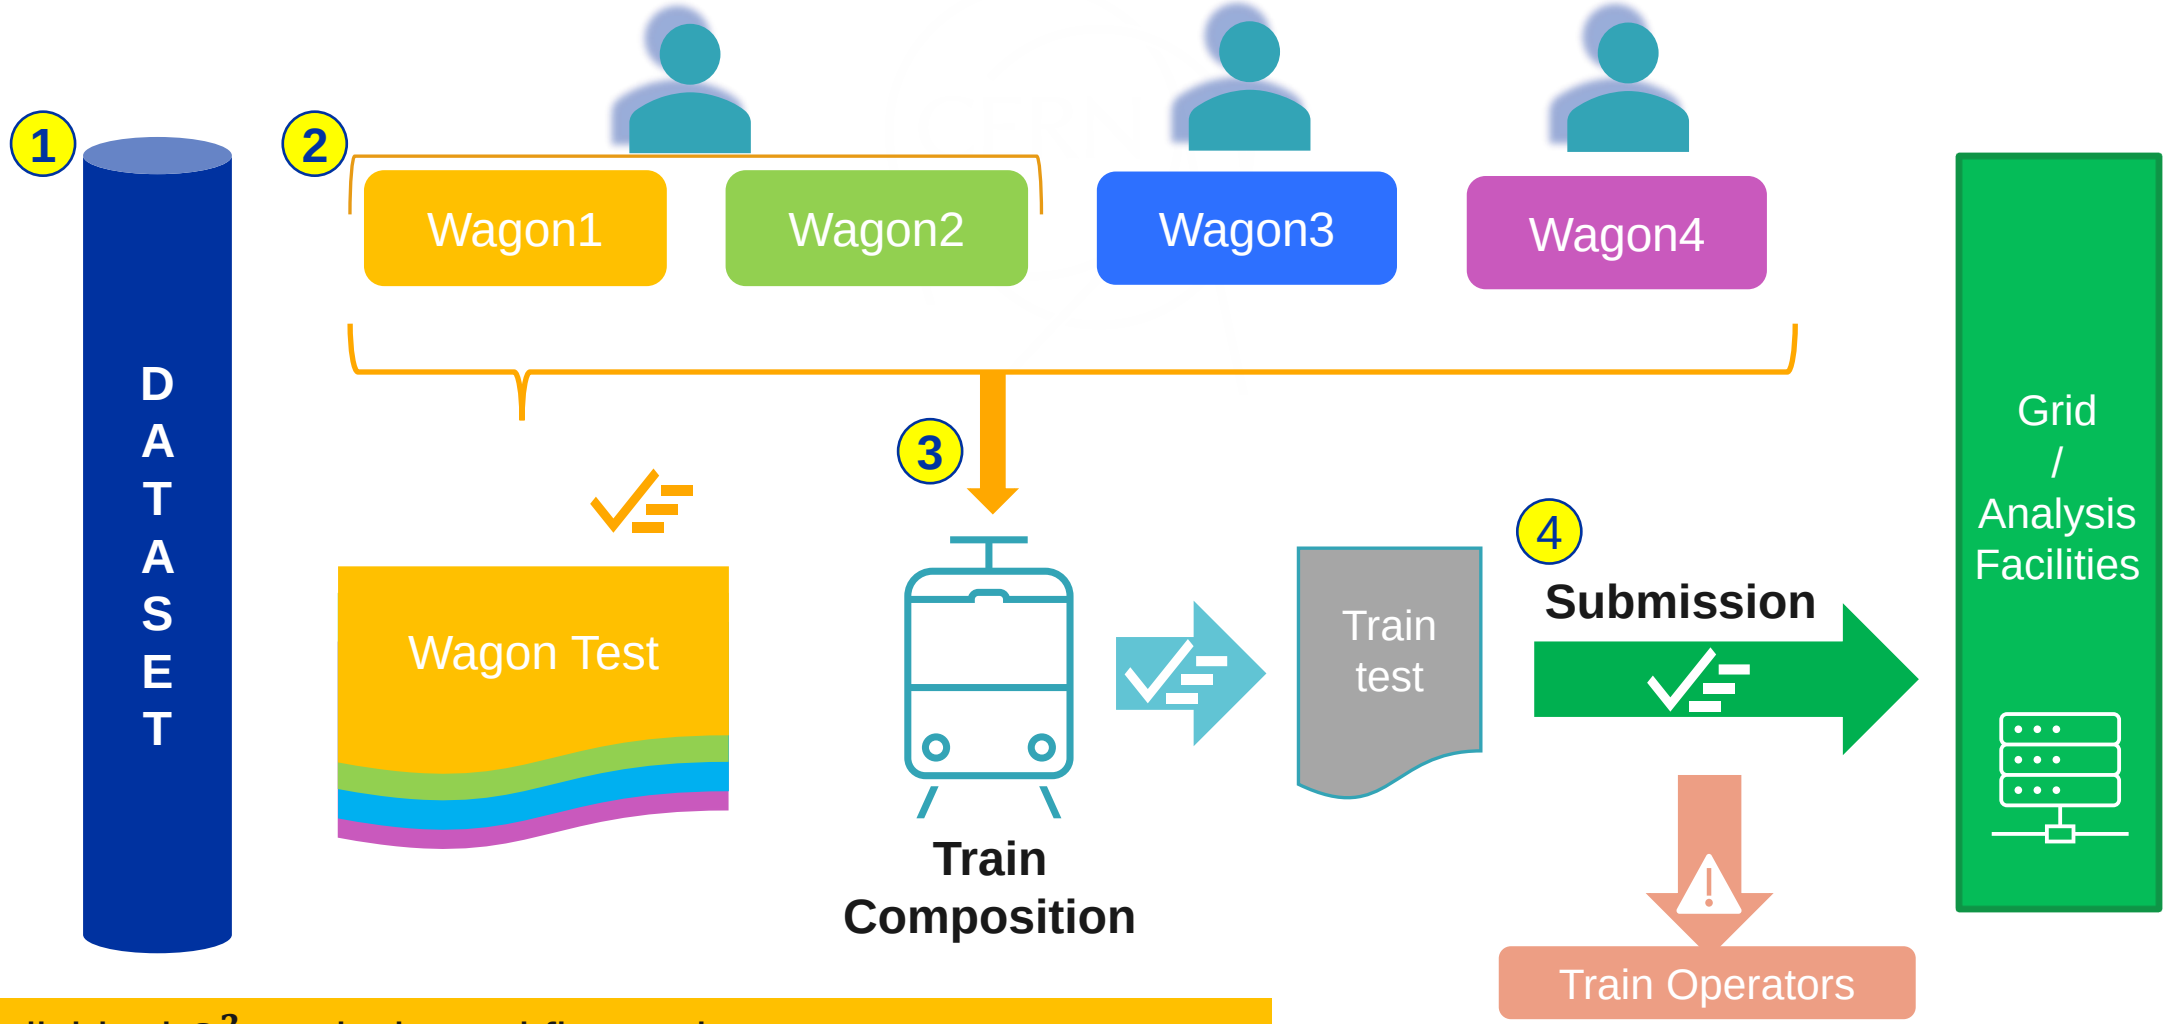

Individual O^2 analysis workflows - known as wagons - are combined into trains to run over the data together

Train runs

Four types of trains can run on Hyperloop

1. Analysis trains
2. Slim derived data trains
3. Standard derived data trains
4. Linked derived data trains

Let’s see the differences between these train types

Analysis Train

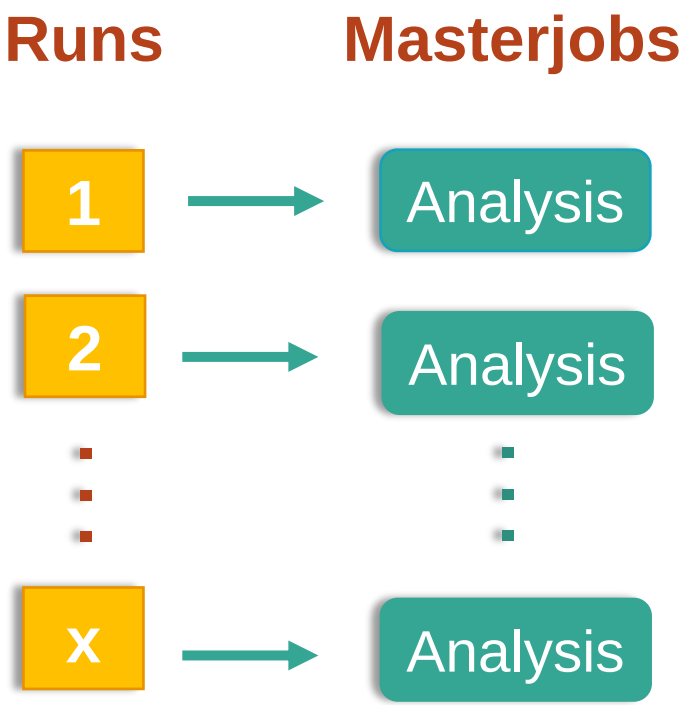

Each run from the dataset is given a **masterjob**

Analysis Train

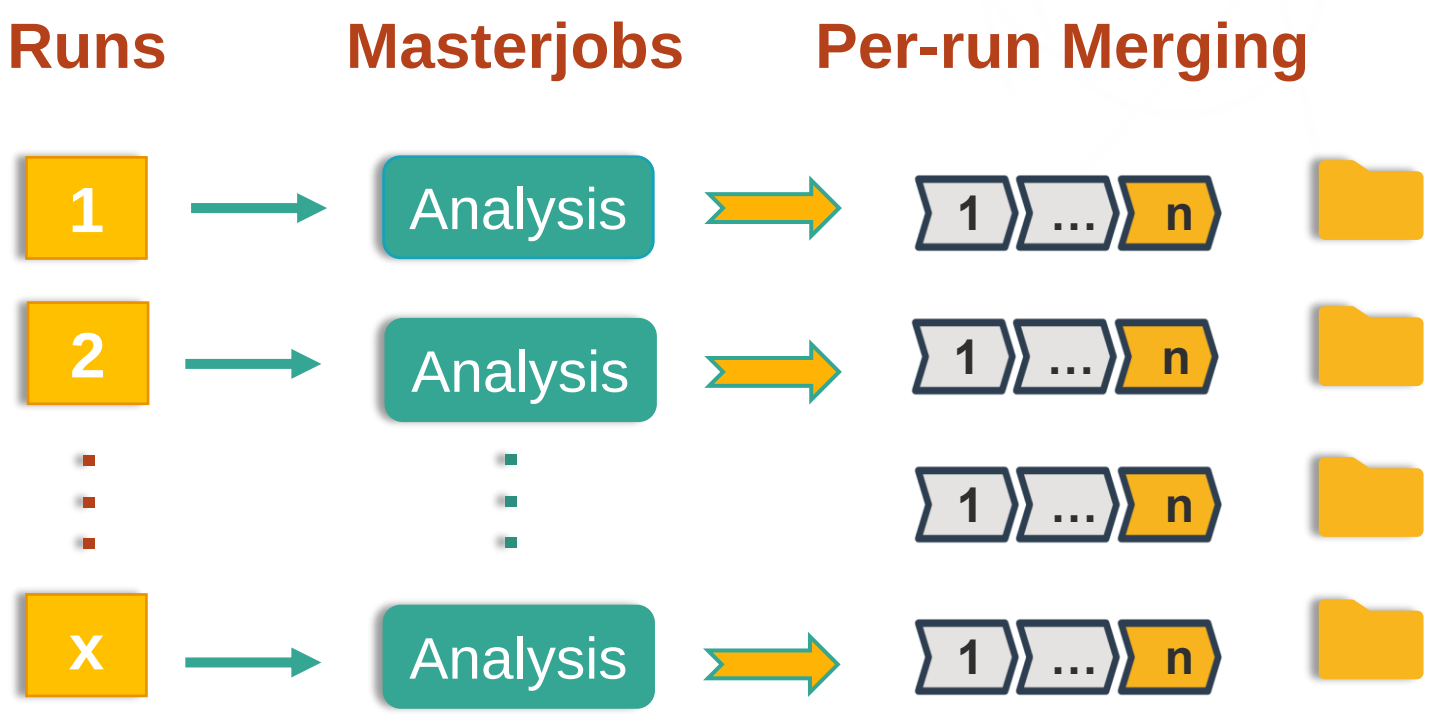

Each run from the dataset is given a **masterjob**

Masterjobs contain **subjobs**

Subjobs are merged in per-run merging

You may access the individual results of each run

Analysis Train

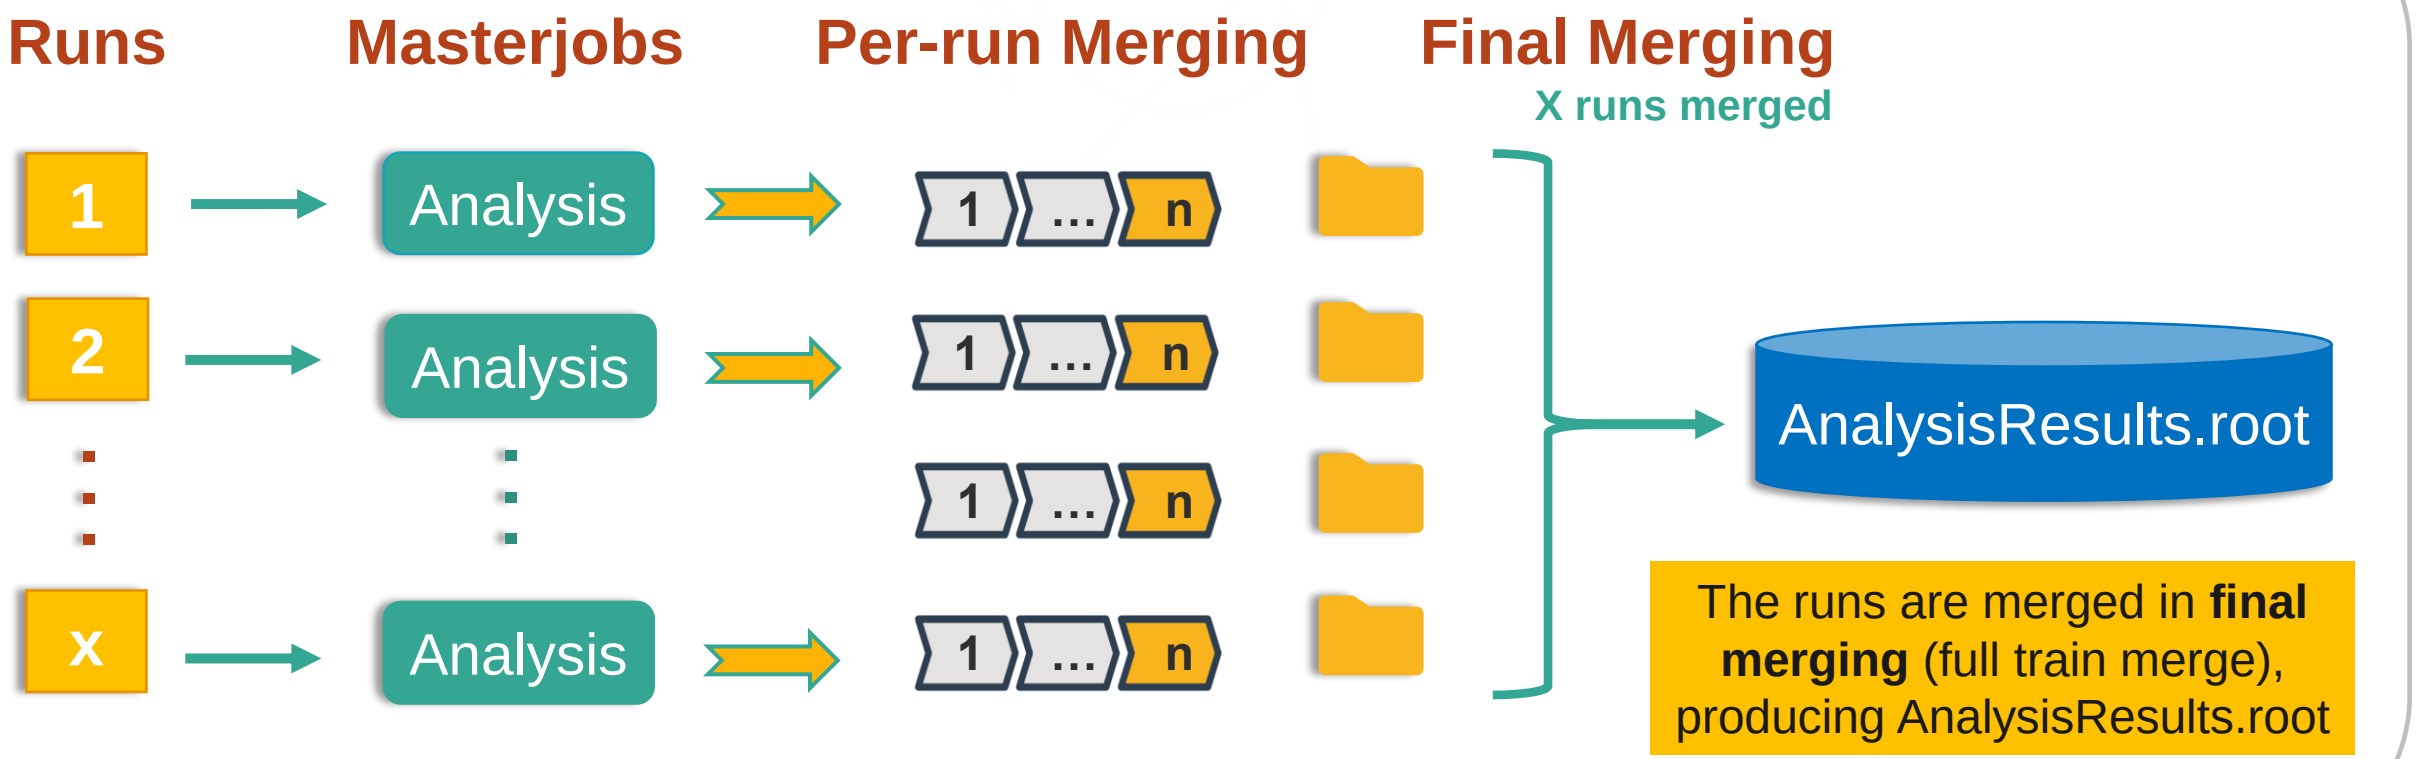

Analysis Train

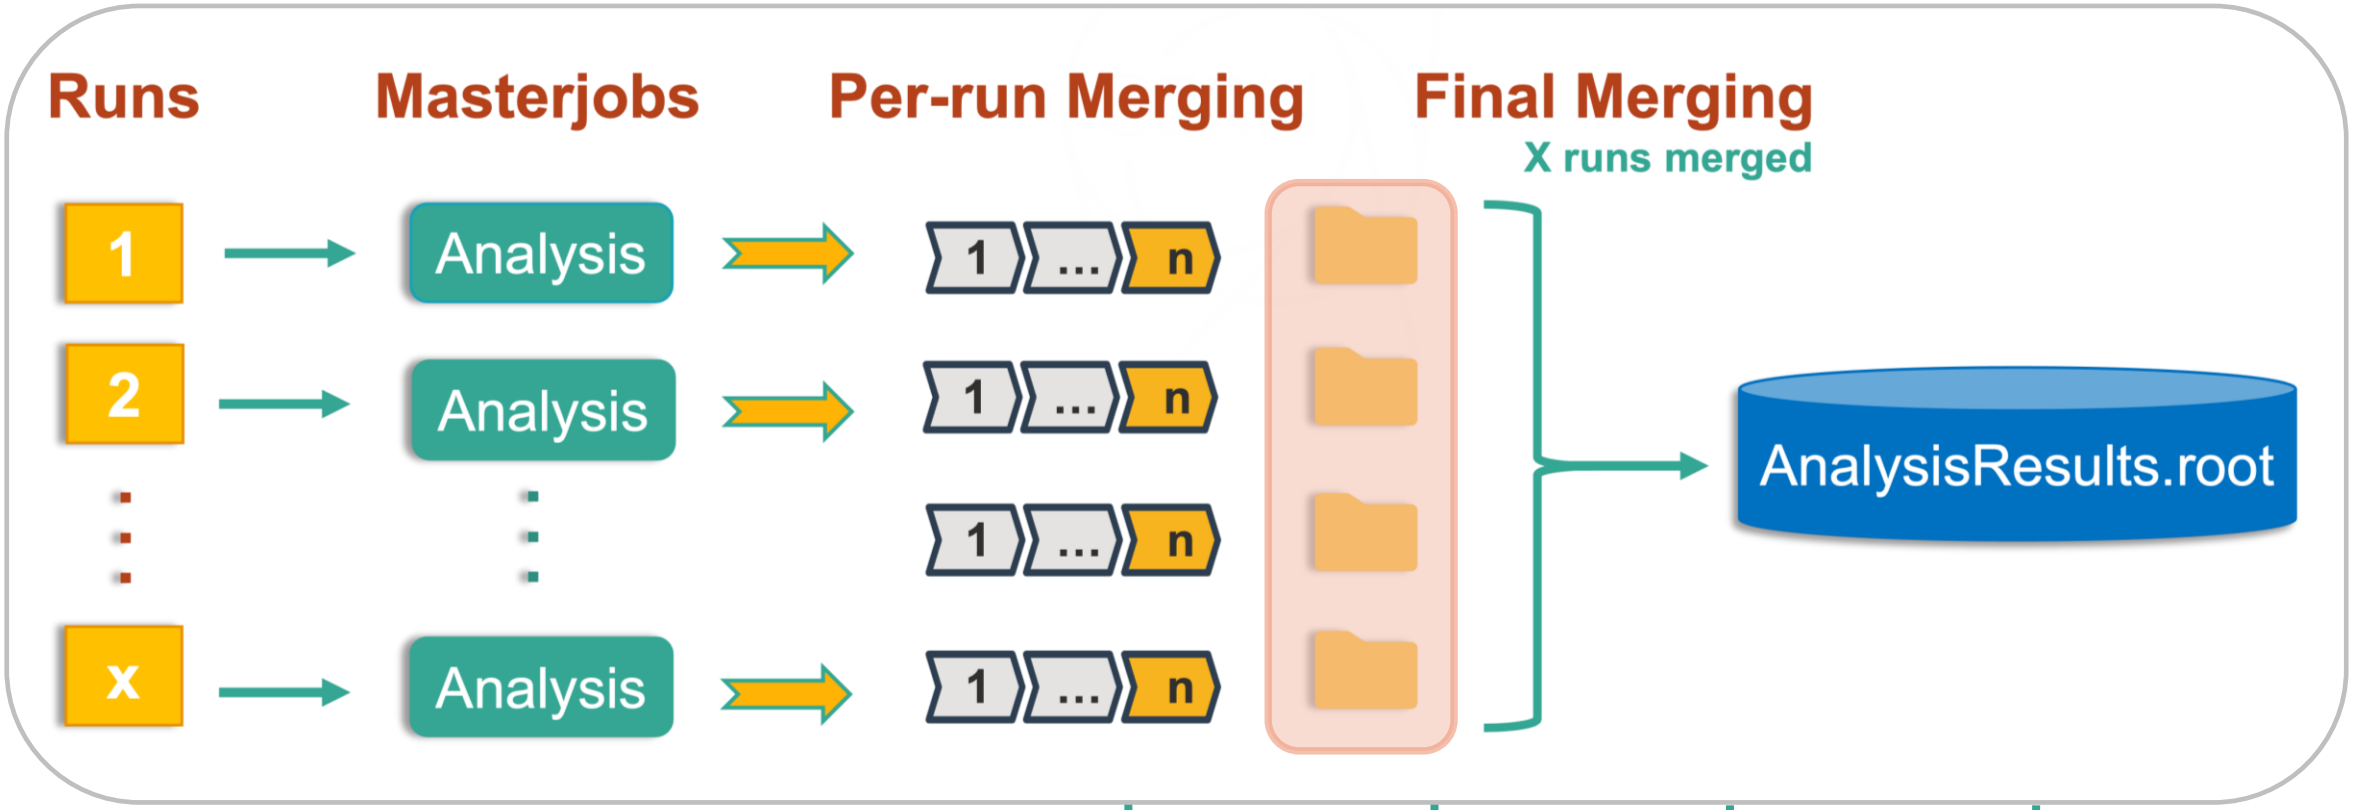

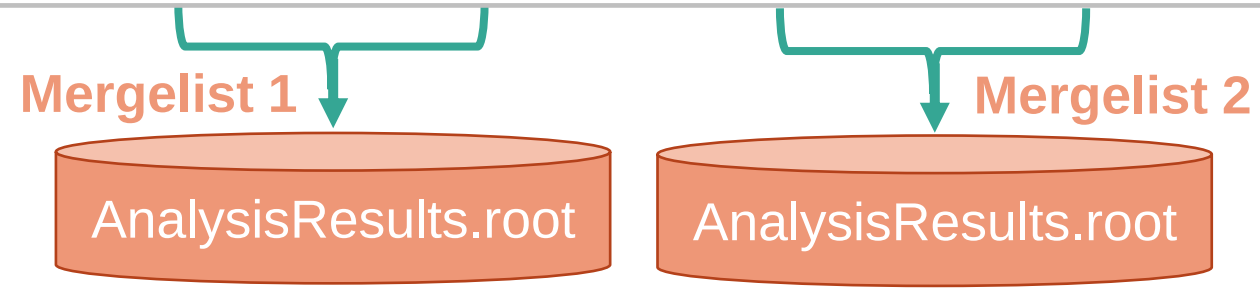

In datasets with multiple **mergelists** or **final mergelists**, a separate merged output is provided for each

Derived Data

- Reconstruction output is not meant for final analysis. (Too much input = the AO2D will be too large!)
- Most analyses need to produce derived data (“trees / tuples”)
  - Why? If analysis ran over AO2D, e.g. for Pb-Pb, we could have only 135 trains/month

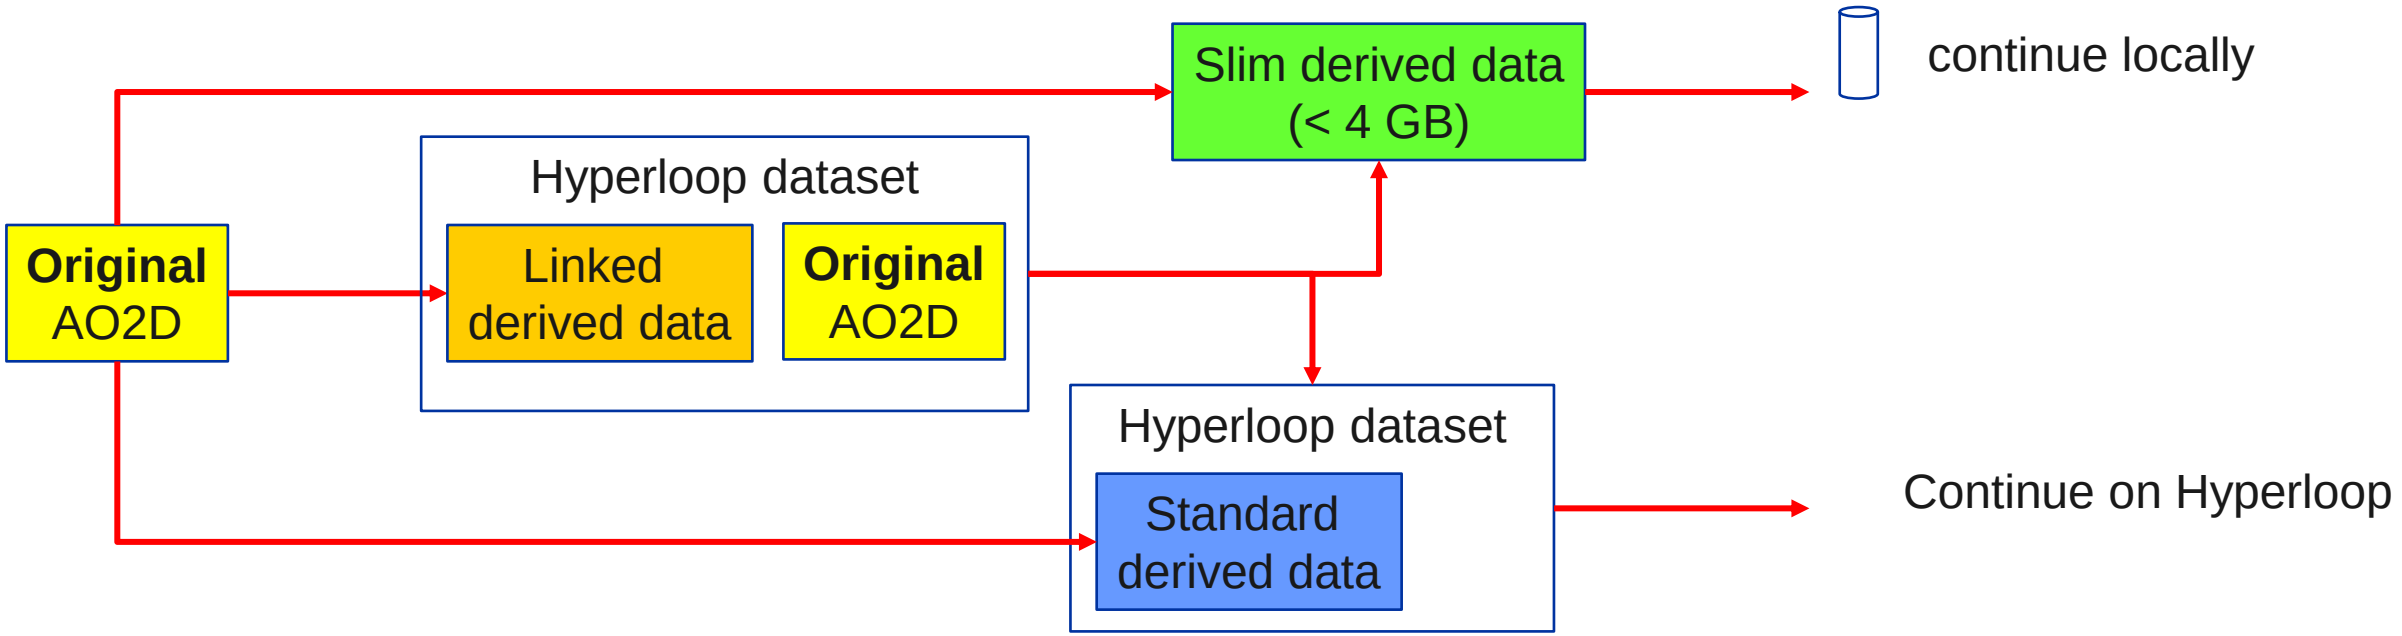

Slim Derived Data Train

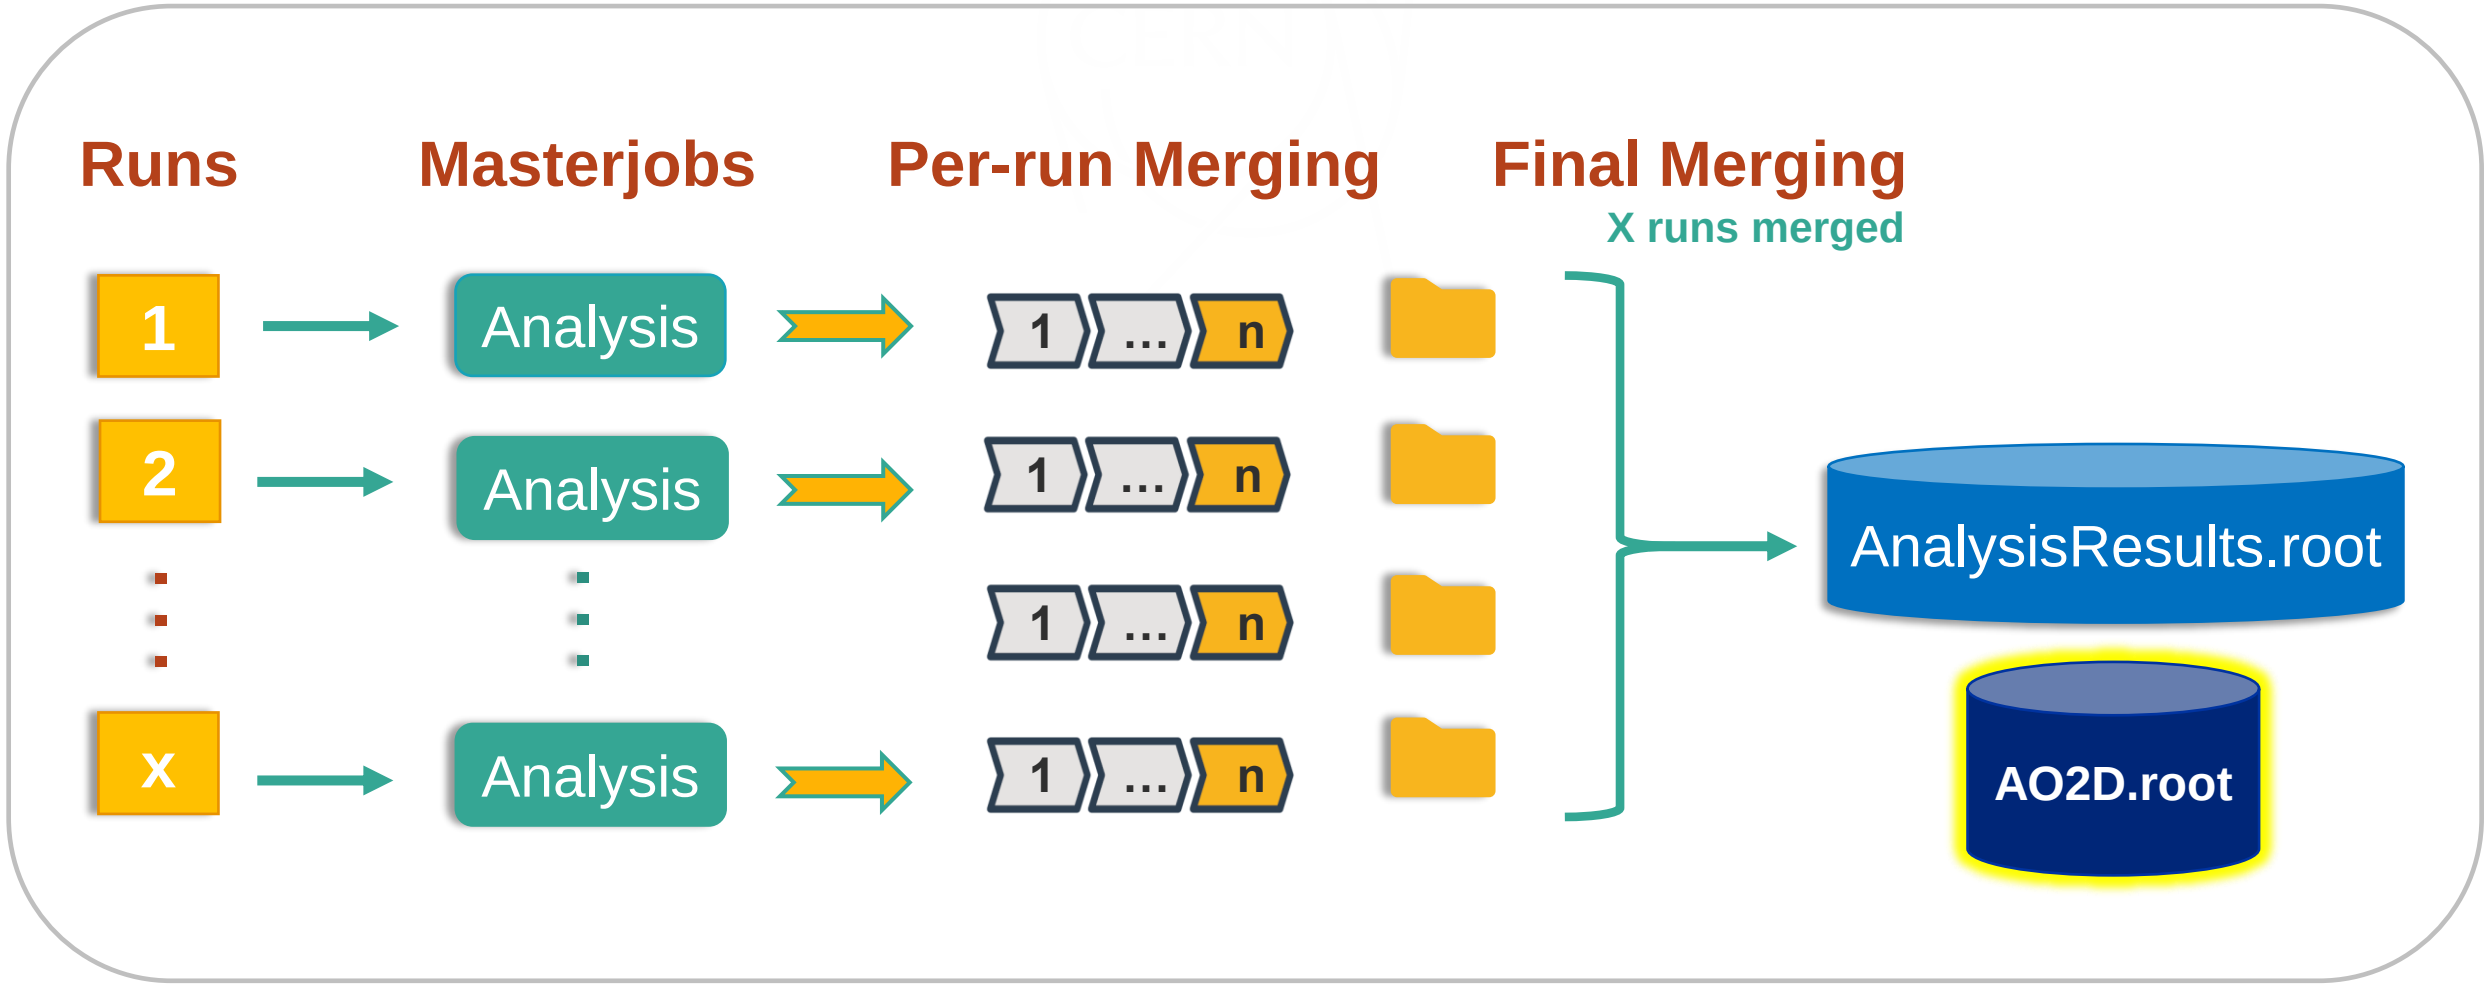

Slim derived trains provide an AO2D.root to be used locally. **Only possible when output < 4GB.**

Derived Data Train

Significant differences!

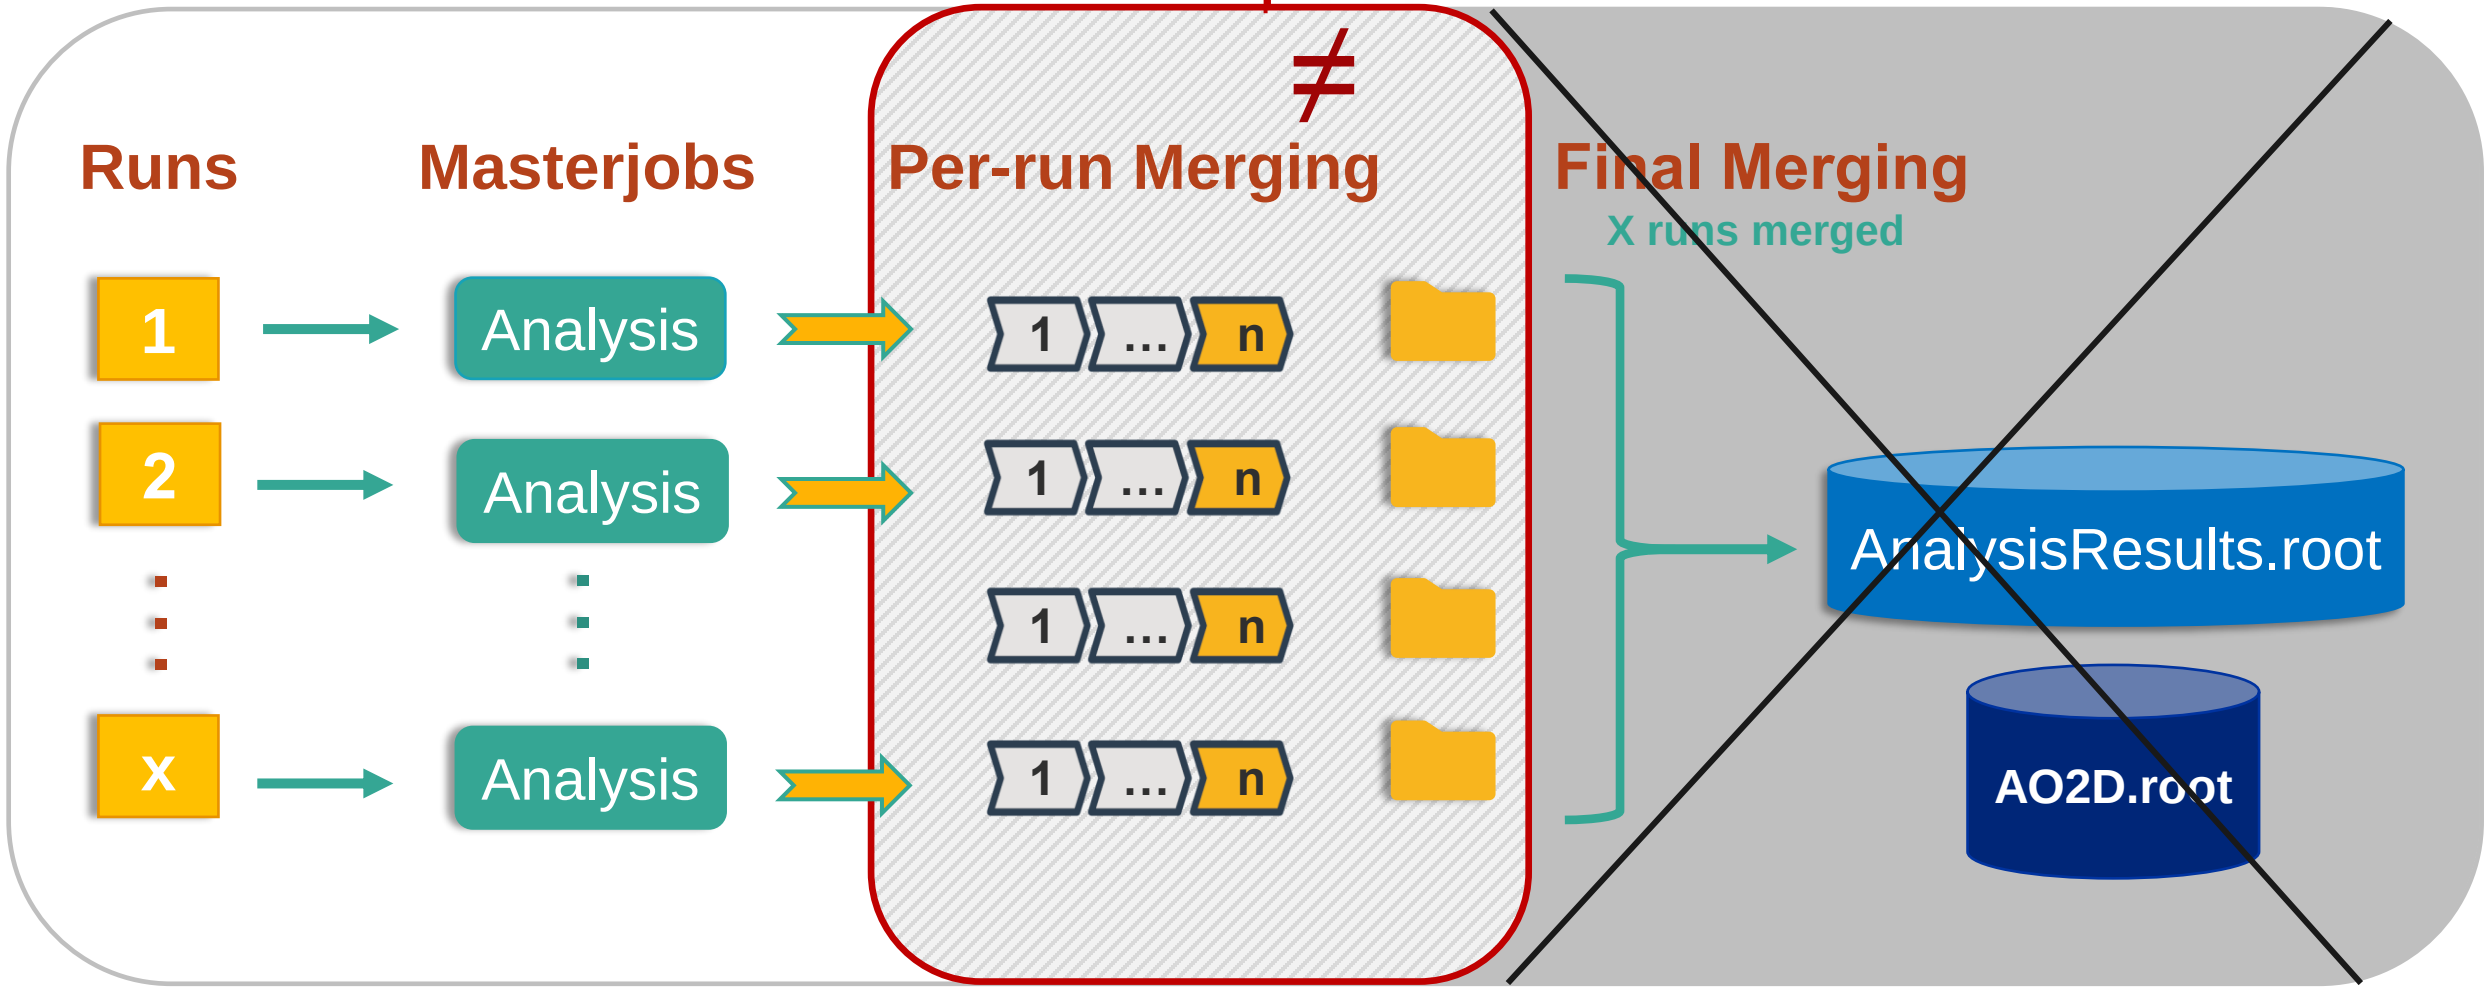

Derived Data Train

To be used as input in future train runs.

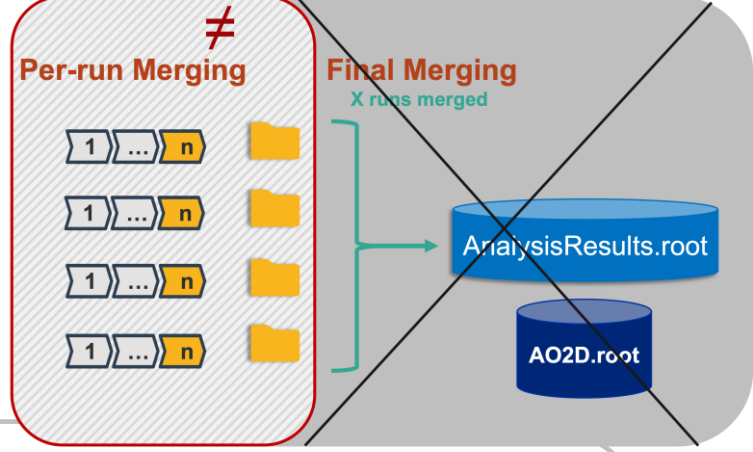

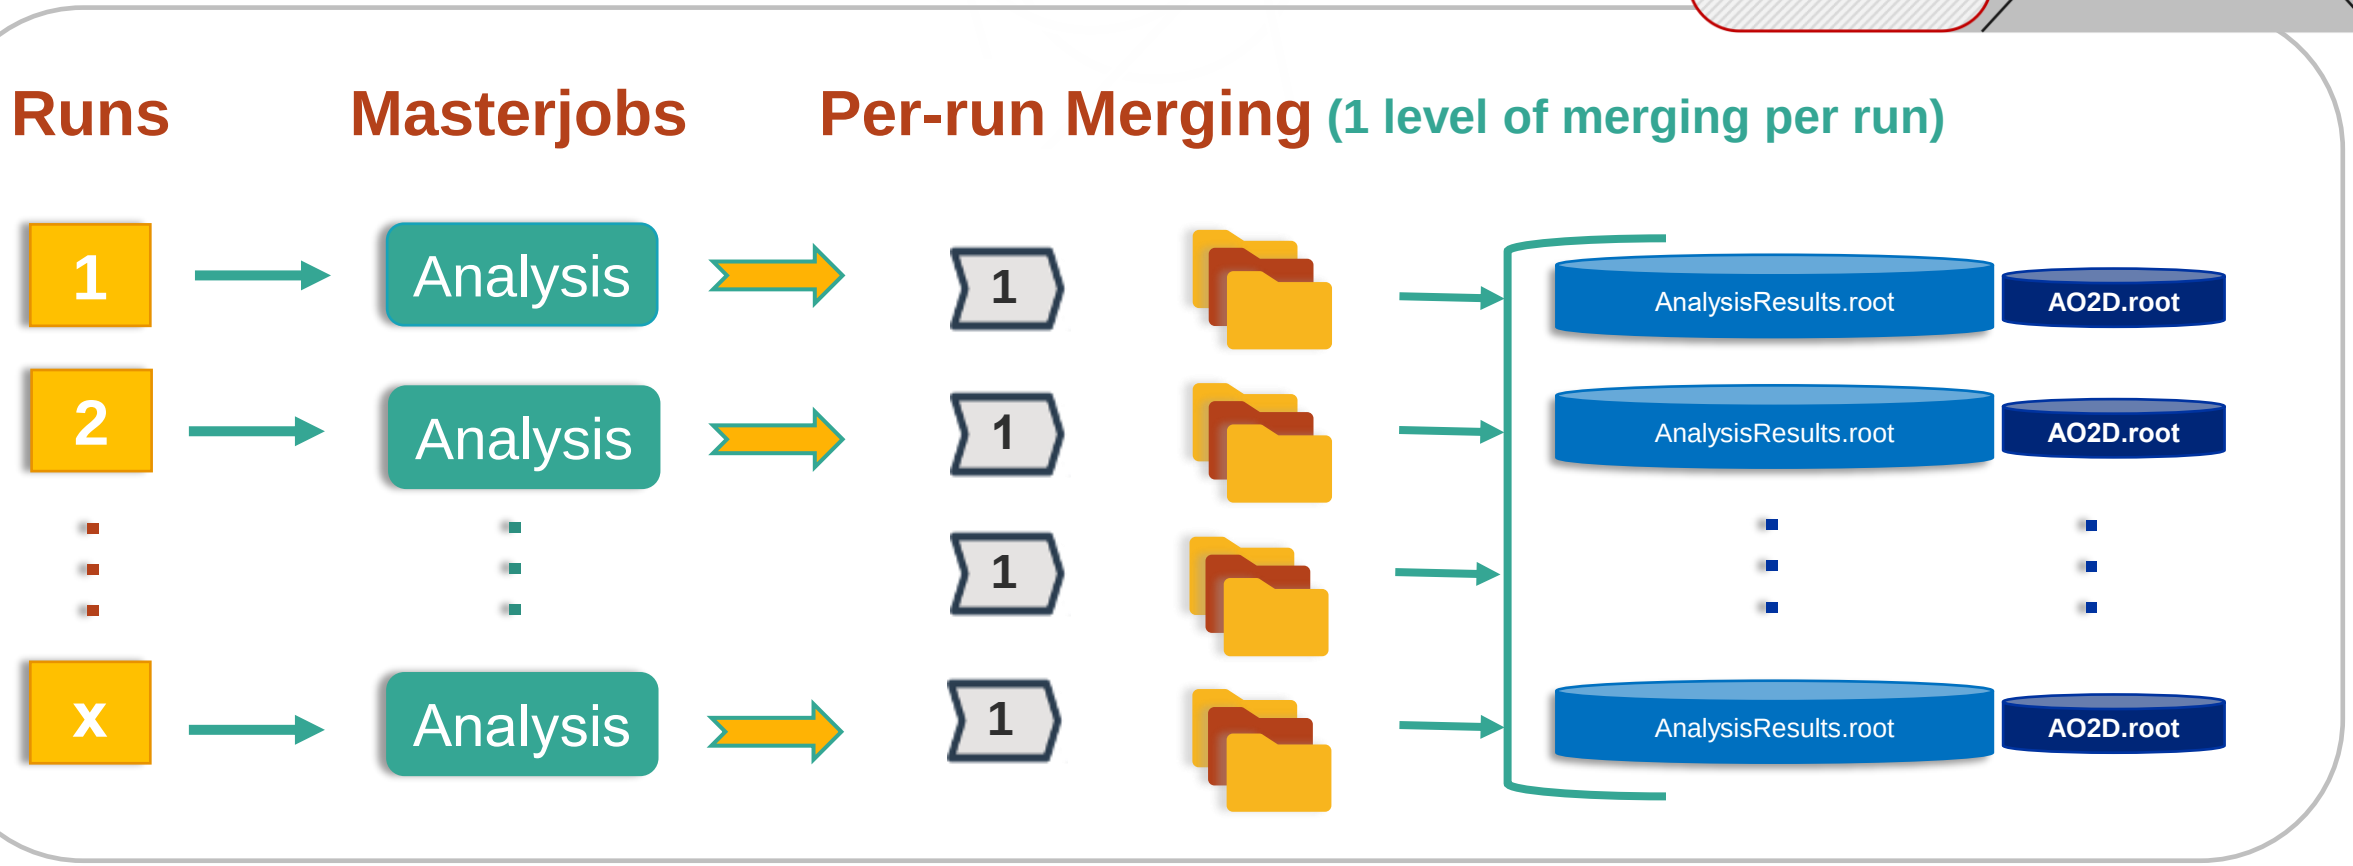

Linked Derived Data Train

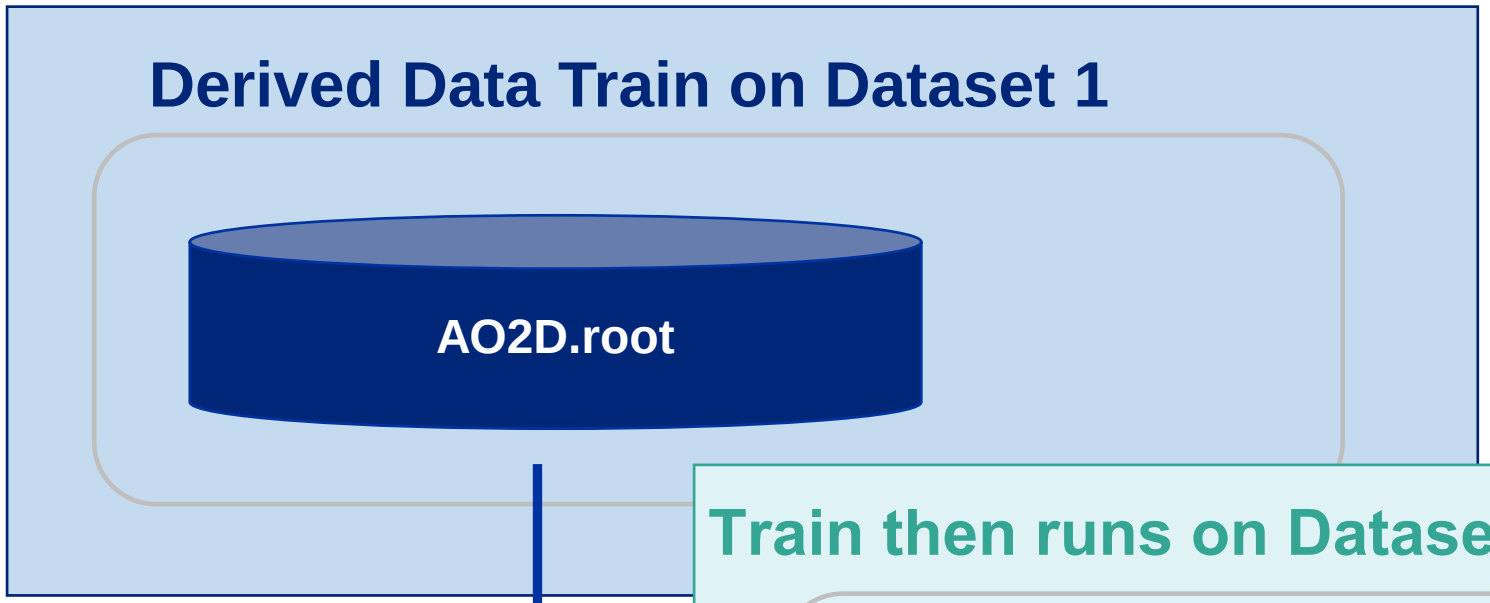

Linked derived data has access to the files from the parent.

So, how do we use Hyperloop?

- Introduction & key concepts
- **’My Analyses’, wagons & datasets**
- Enabling a wagon to start a test
- Trains – automatic composition, debugging & long trains
- Key things to know

My Analyses

- **Displays all the analyses you are part of**
- Analyses are defined in JIRA for the respective PWG, along with the users that will be part of them
- Several analyzers can share an analysis
- Datasets are enabled per analysis

Instantaneous wagon test when enabling a wagon

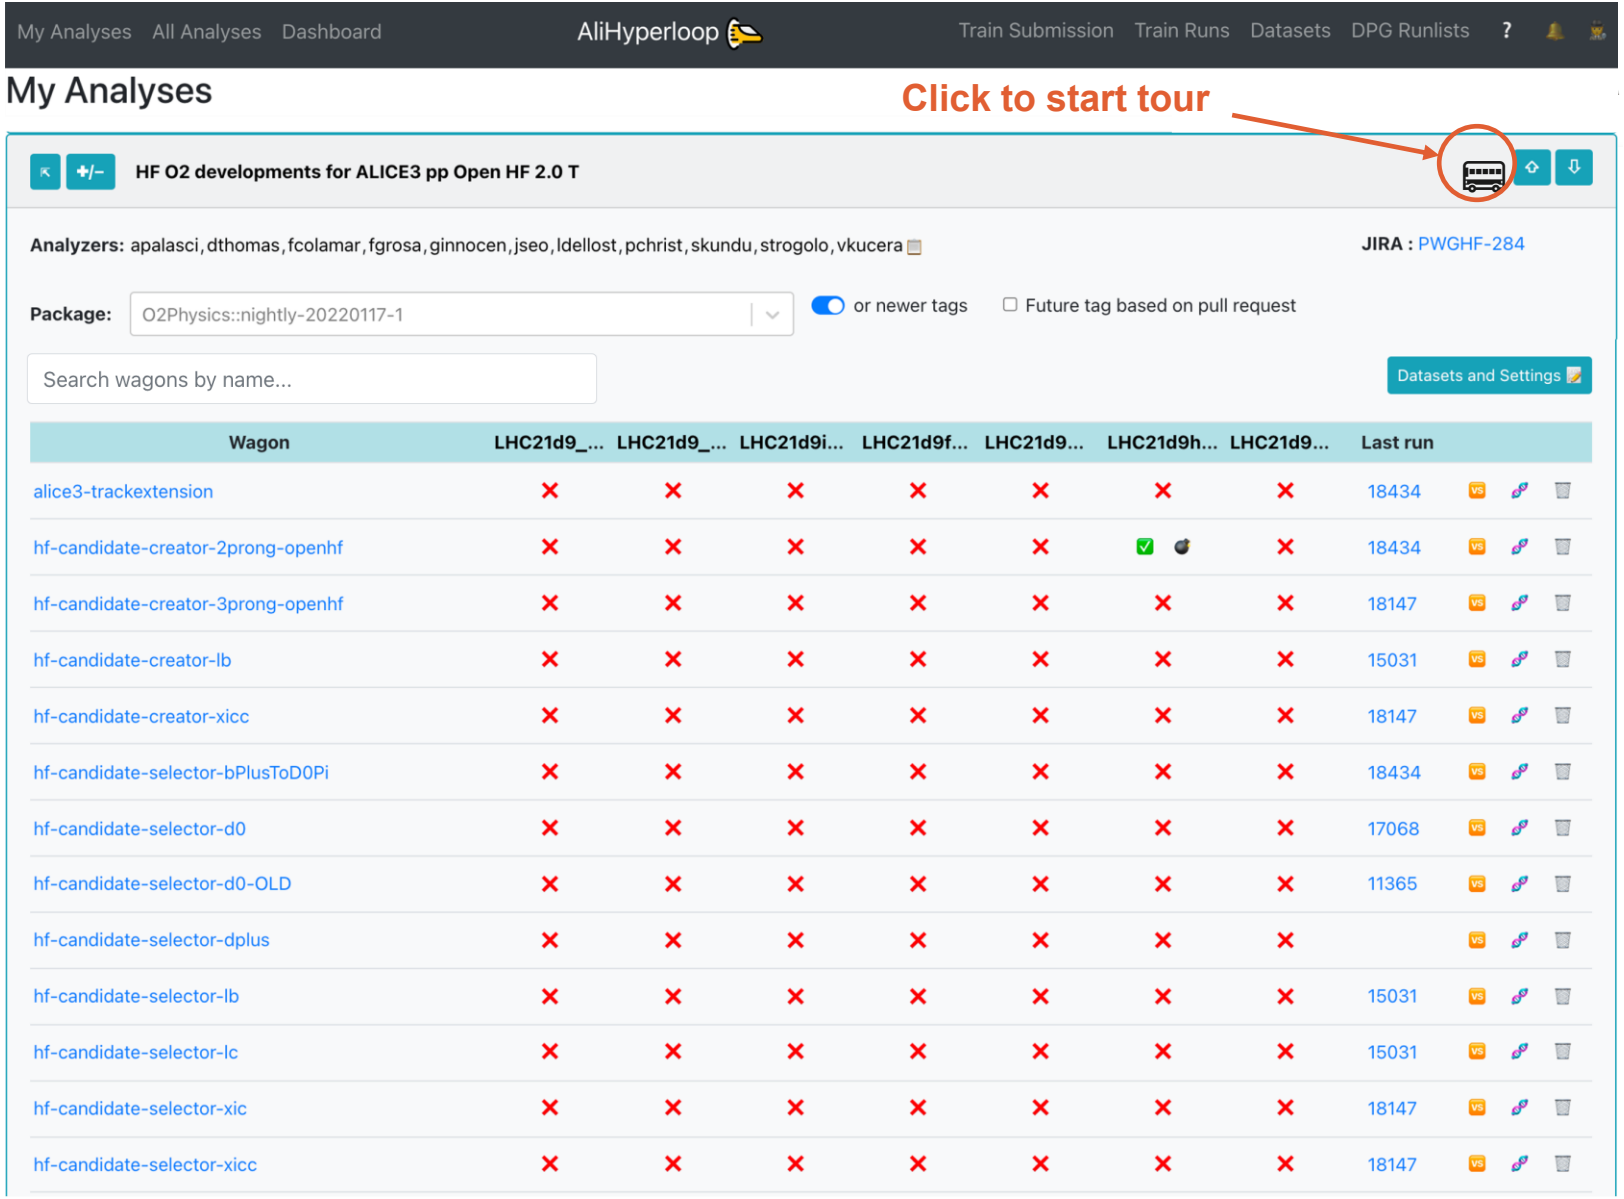

My Analyses

- Displays all the analyses you are part of
- **Analyses are defined in JIRA for the respective PWG, along with the users that will be part of them**
- **Several analyzers can share an analysis**
- Datasets are enabled per analysis

Instantaneous wagon test when enabling a wagon

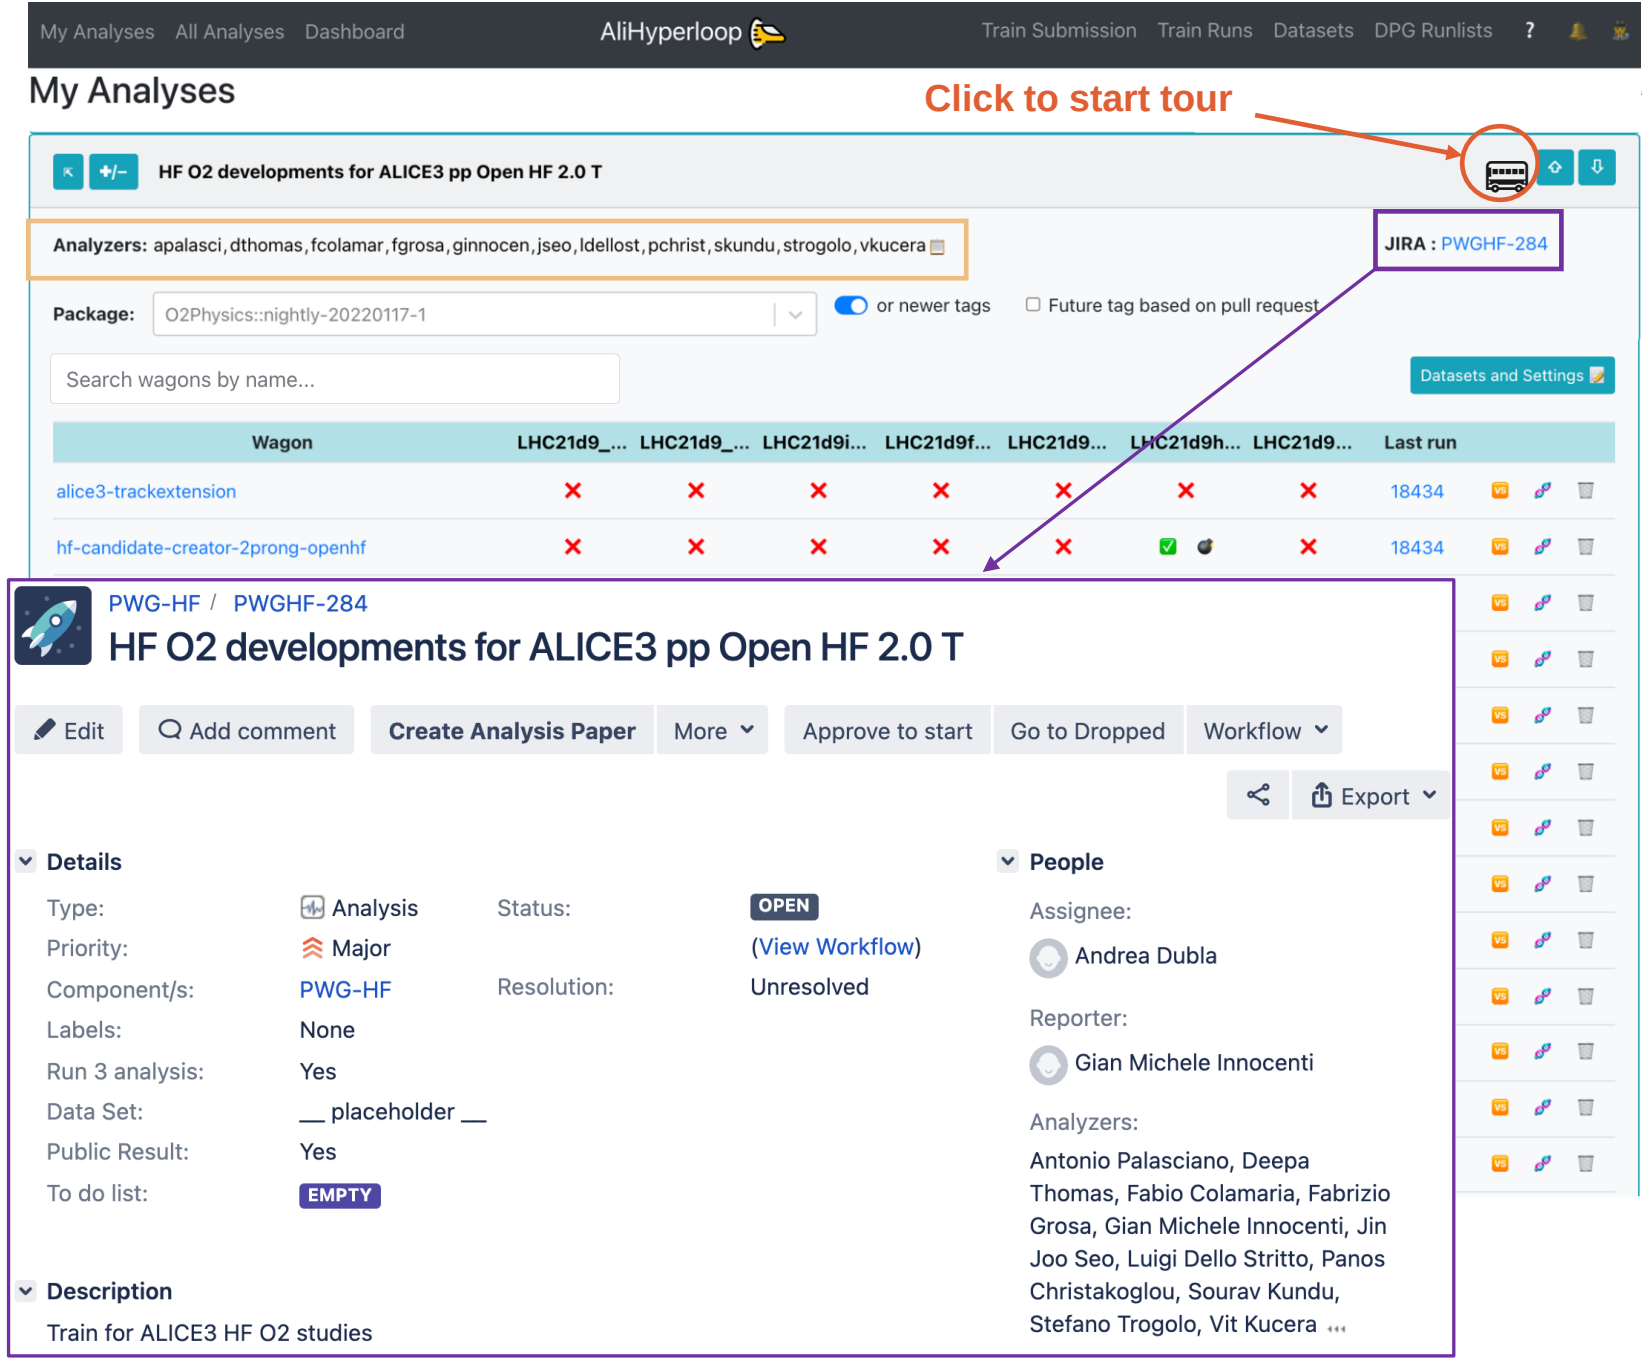

My Analyses

- Displays all the analyses you are part of
- Analyses are defined in JIRA for the respective PWG, along with the users that will be part of them
- Several analyzers can share an analysis
- **Datasets are enabled per analysis**

Instantaneous wagon test when enabling a wagon

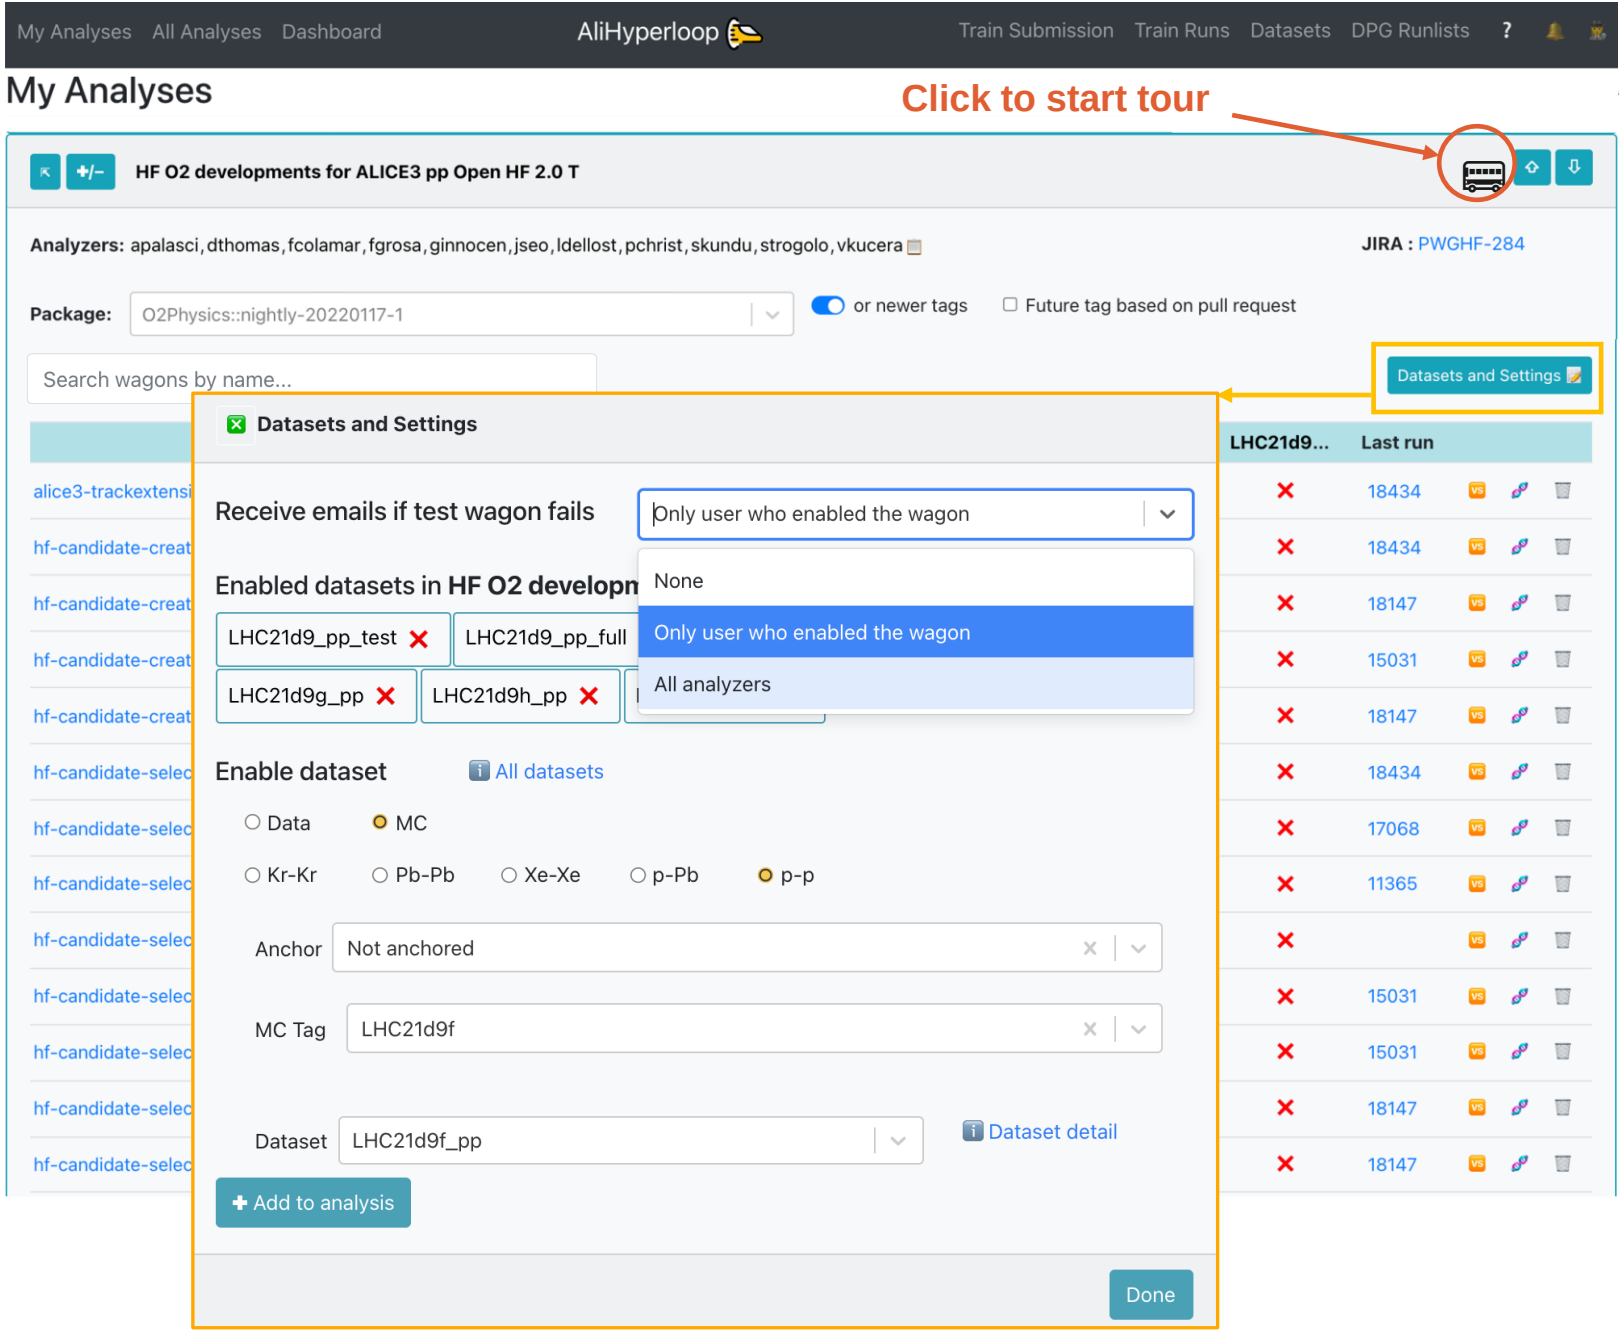

Wagons

Wagons are defined from O^2 analysis workflows

You may:

- Add/remove wagons
- Enable/disable wagon
- Clone/compare wagons

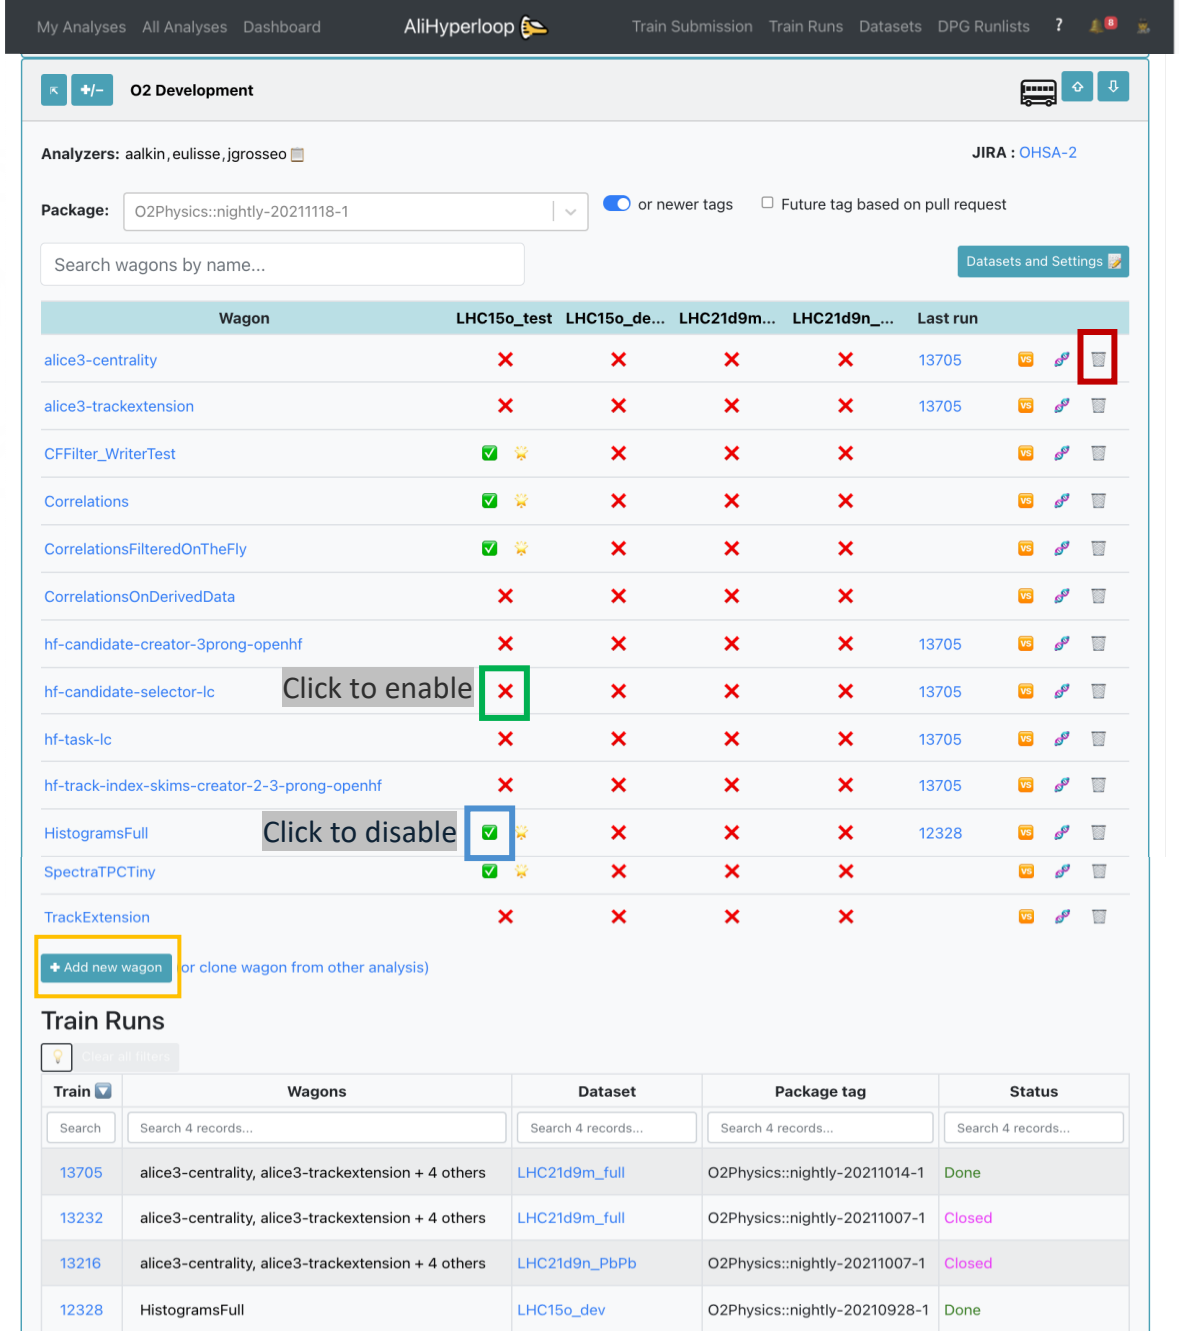

Wagons

Wagons defined from O^2 analysis workflows

You may:

- Add/remove wagons
- Enable/disable wagon
- Clone/compare wagons

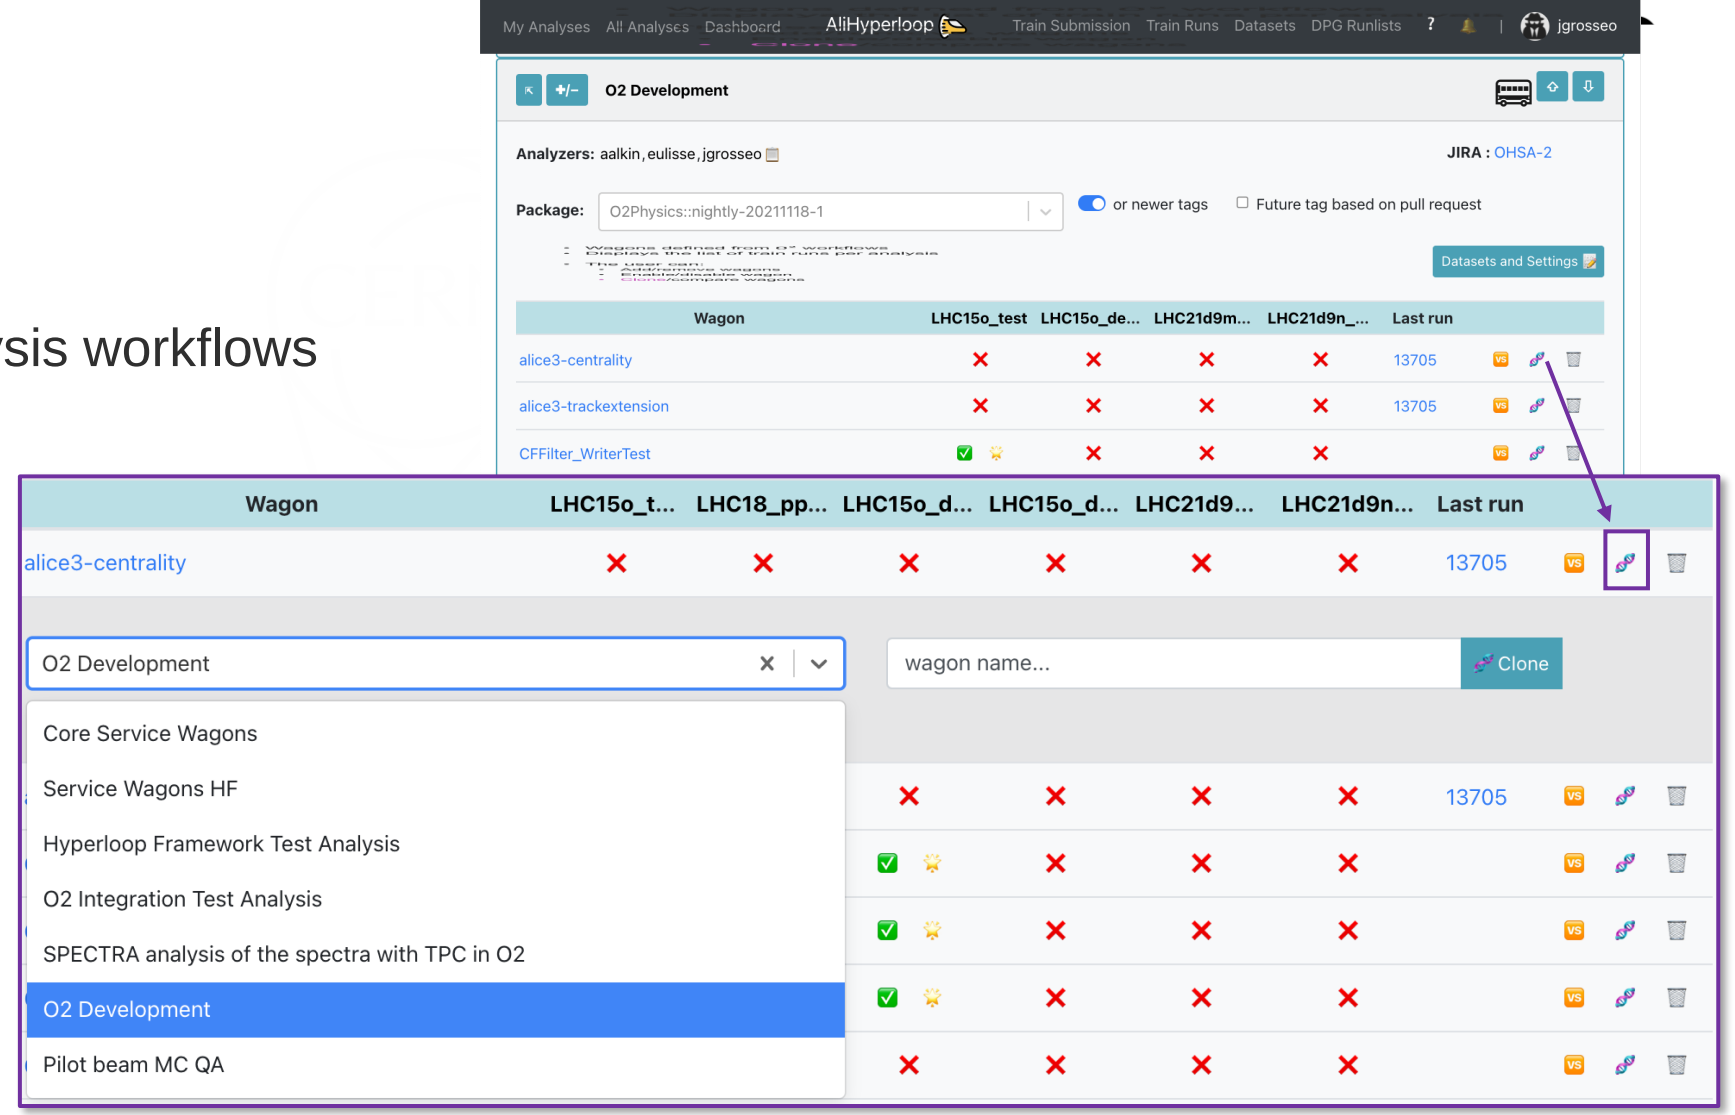

Wagons

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

- You may:
  - Add/remove wagons
  - Enable/disable wagon
  - Clone/compare wagons

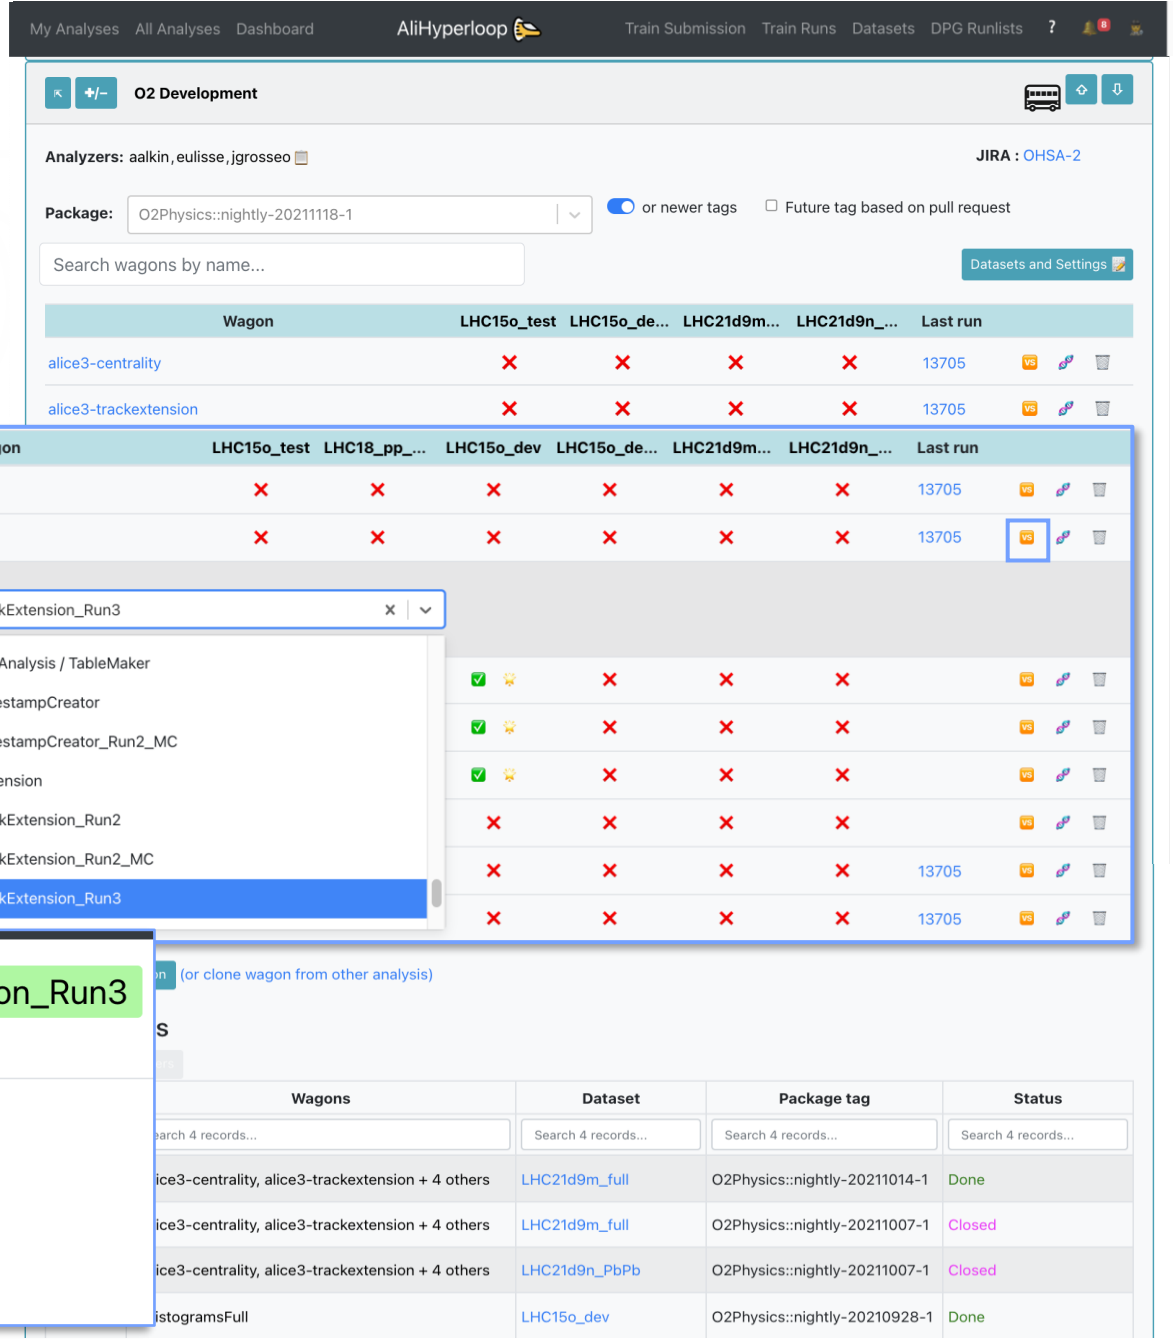

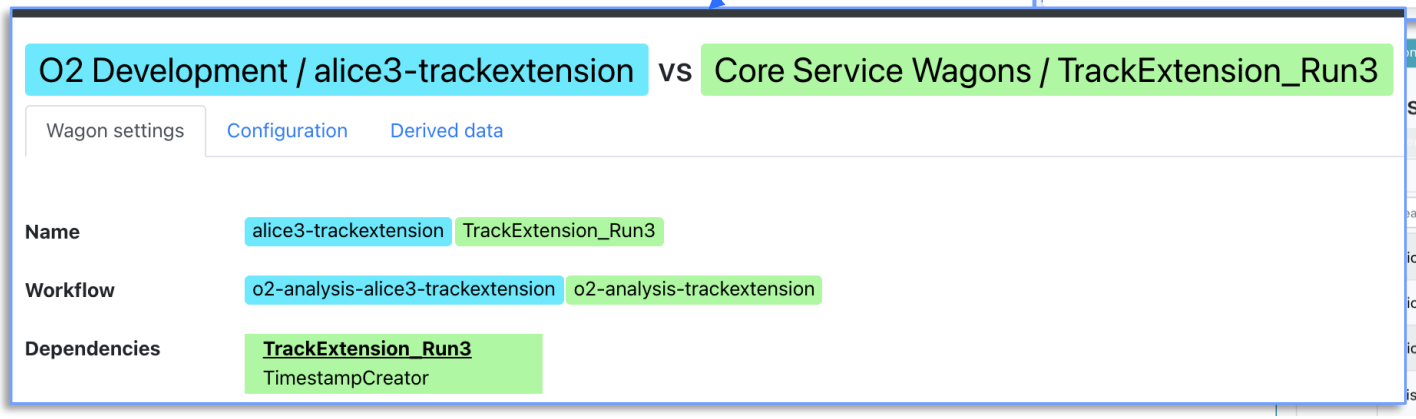

Create a wagon

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

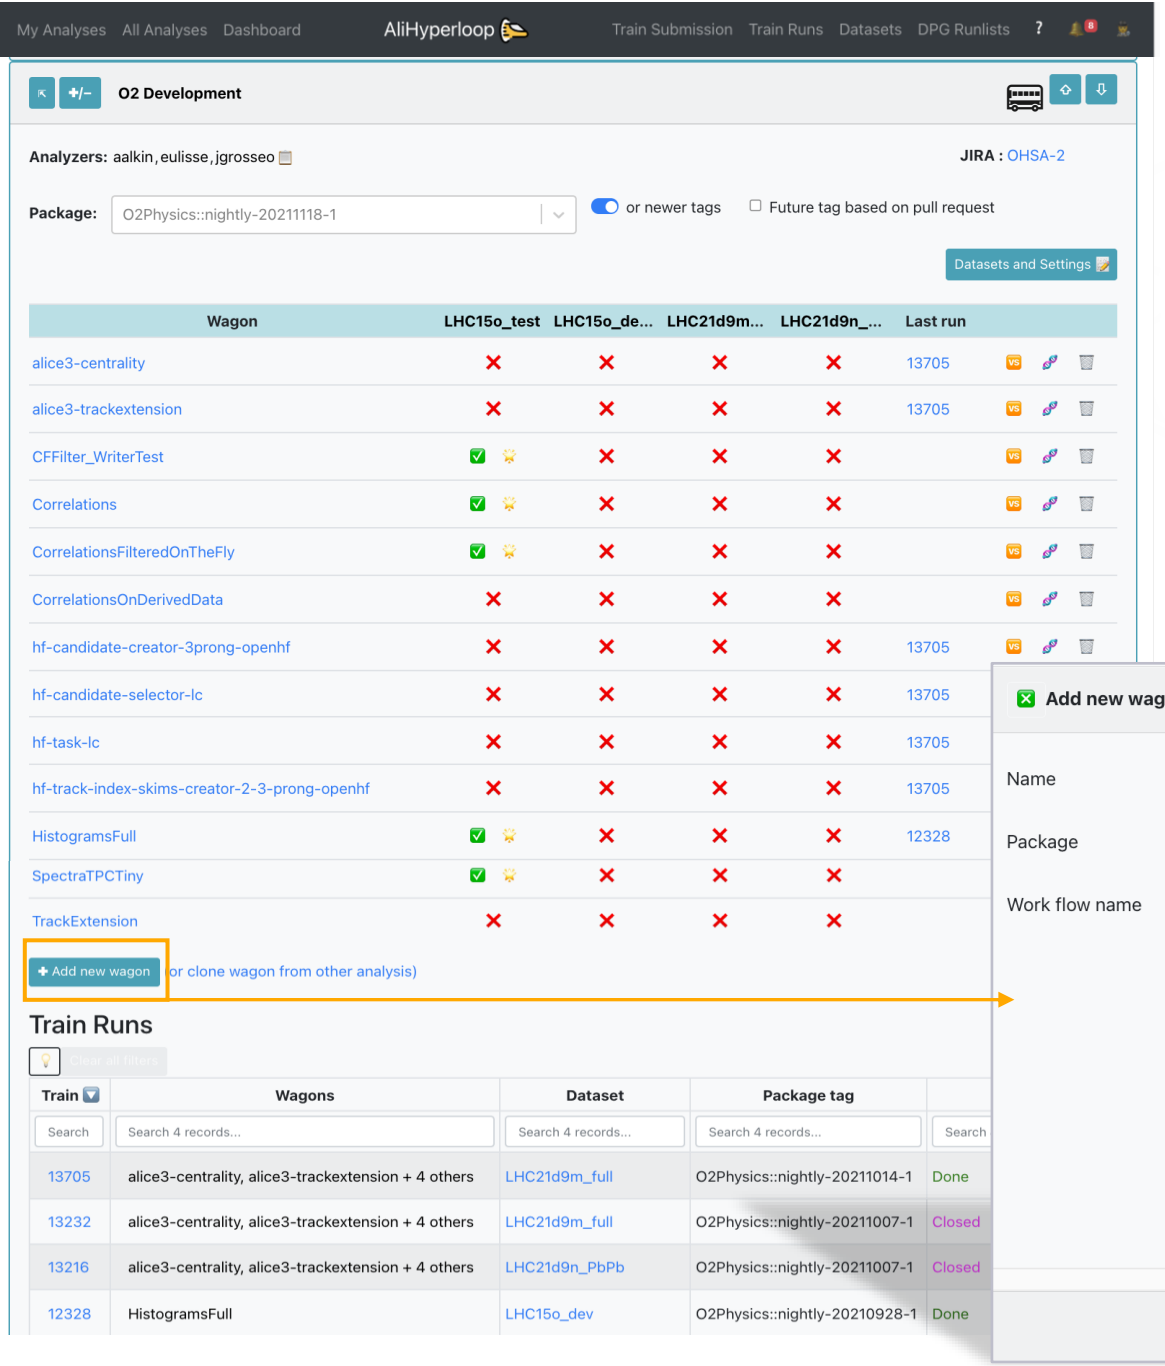

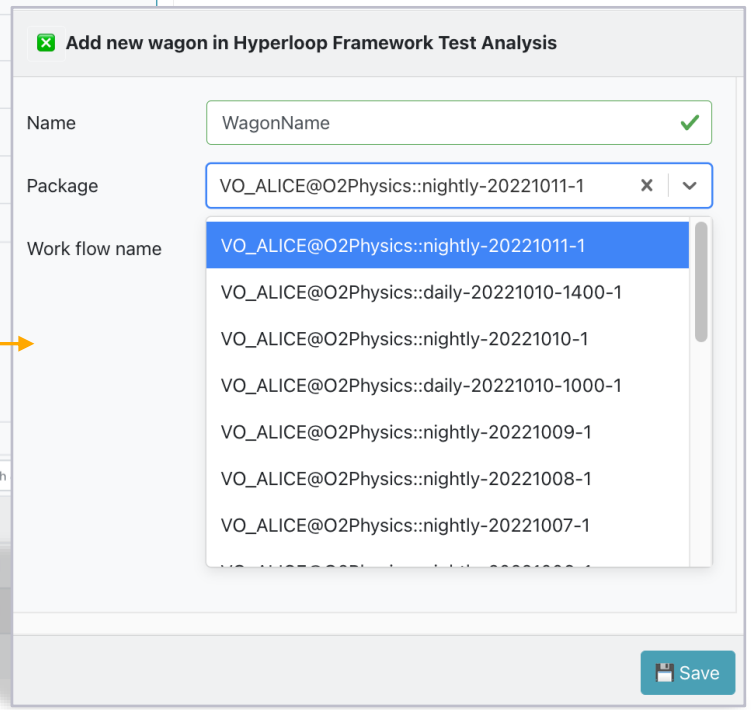

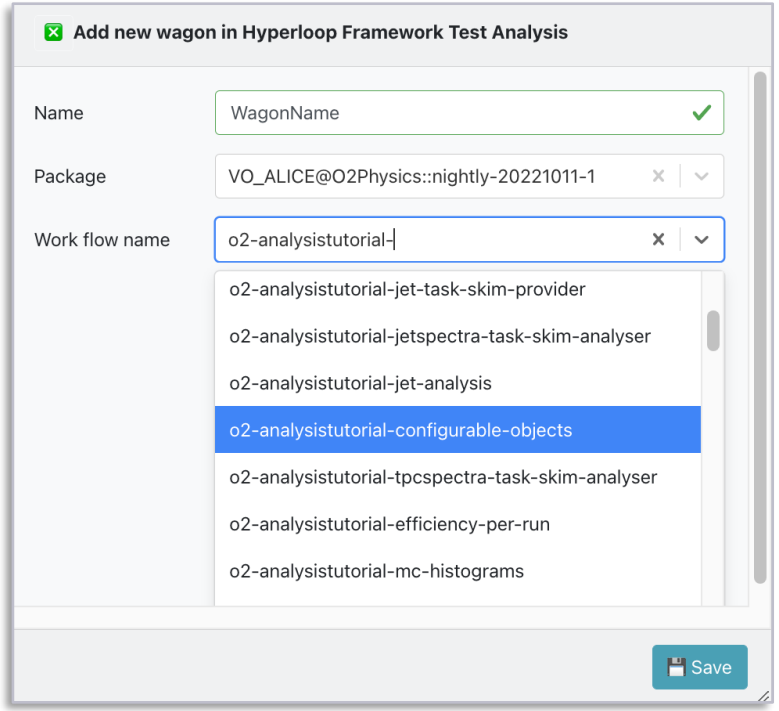

Edit wagon

Click to start the tour

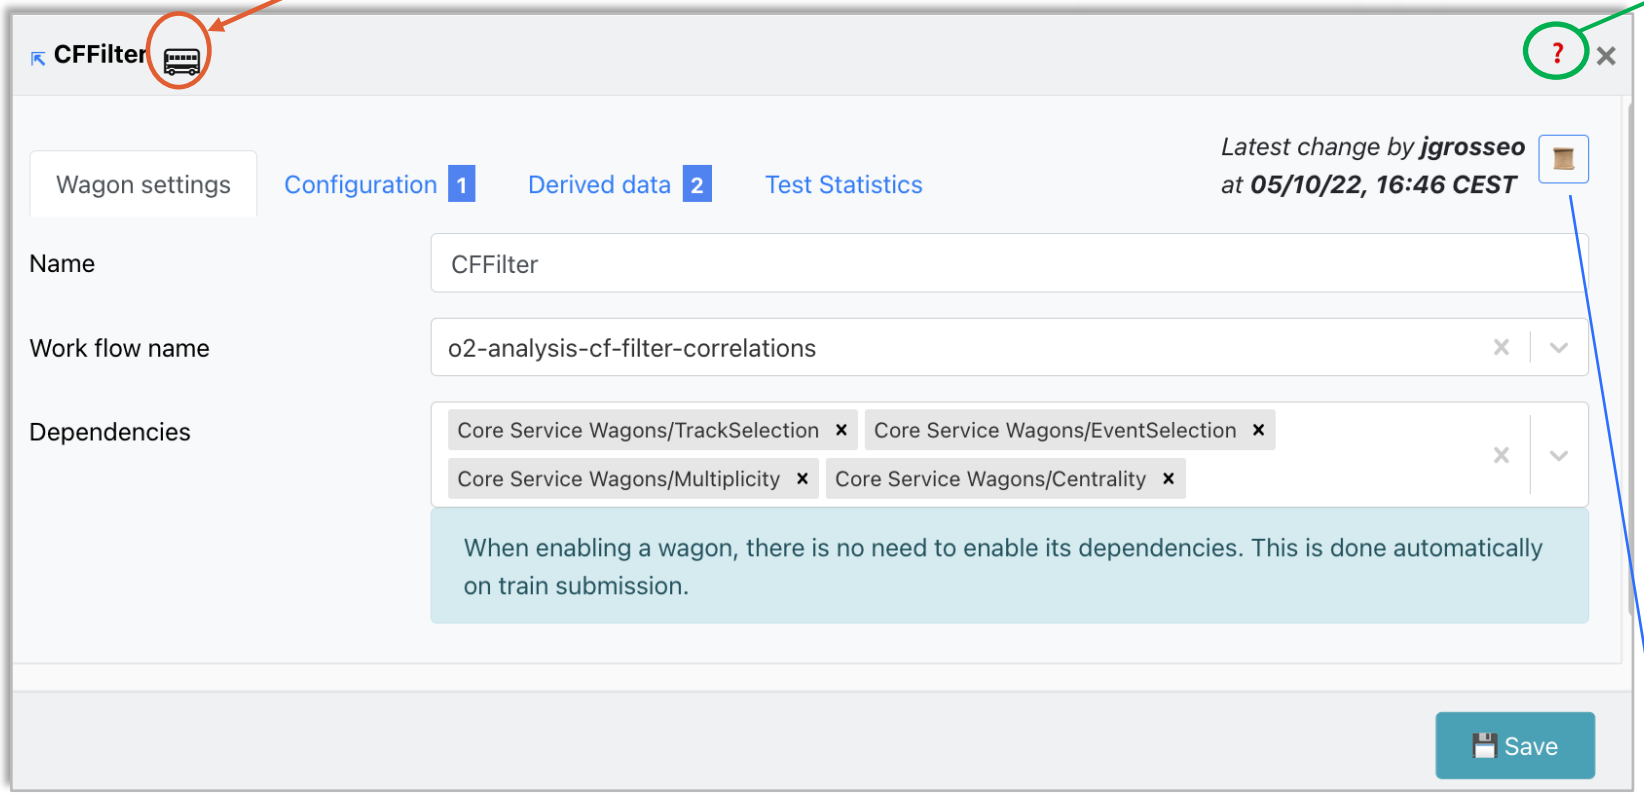

Documentation

You can edit a wagon by clicking on it 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
- 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
- Useful for sending to colleagues or to the mattermost channel for help

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

Edit wagon

- Task configuration fully integrated. Wagon configuration corresponding to the workflow is available in Base
- Supports a variety of parameter types defined in task as Configurable including arrays, matrices, labelled matrices and histogram binning
- If you need to run the same workflow, but modify one (or more) of the parameter's value, you can use a subwagon instead of creating a new wagon
- The configuration of a subwagon will be kept the same as the *Base* and will overwrite only the parameters that you provide.
- Sync to update the base and subwagon configuration with the latest version of the workflow

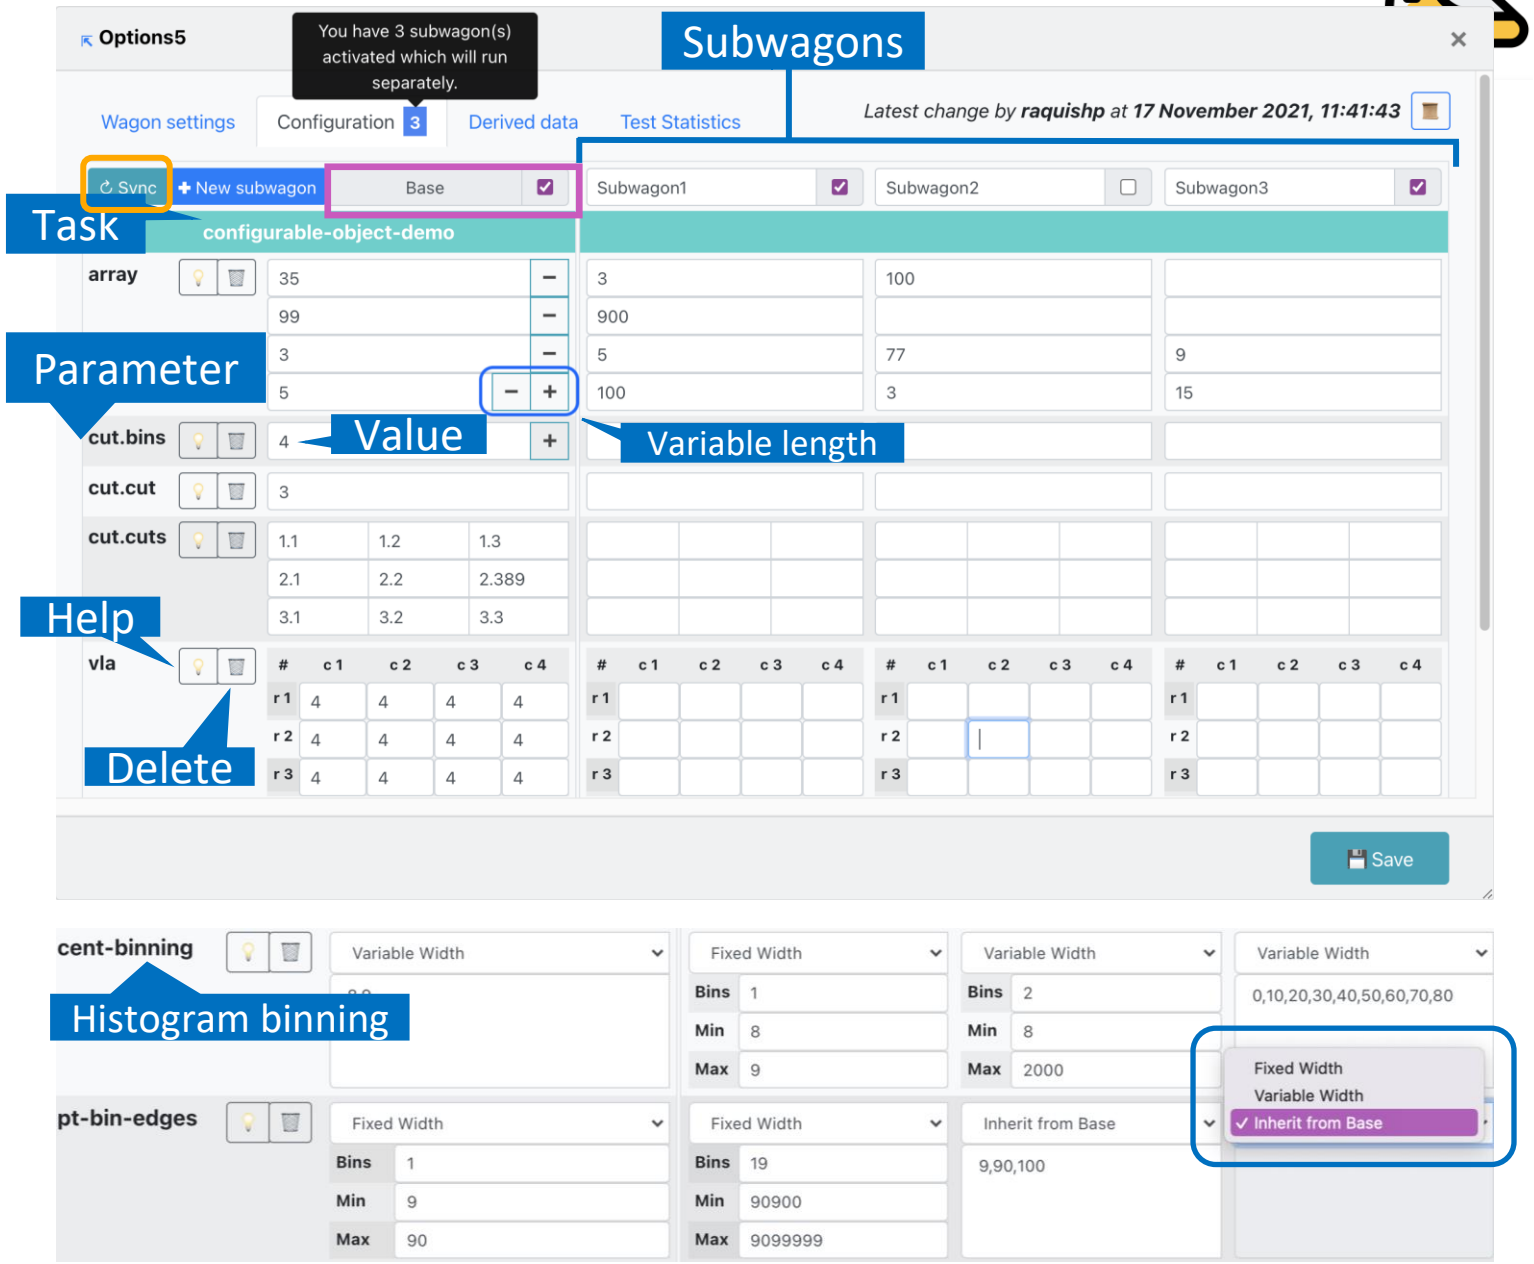

Datasets in Hyperloop

- From the ‘datasets’ page you can see all datasets in Hyperloop
  - Open one to see the configuration!

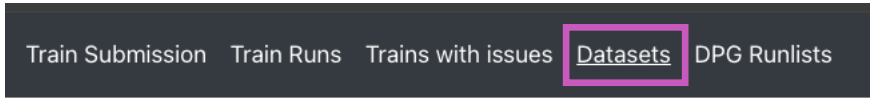

- Available datasets:
  - Run 2: Converted data and general-purpose MC
  - Run 3: DATA, MC, MCGEN (MC on the fly)
- Tiered structure of datasets with automatic composition schedules
  - Automatic composition: small & medium
  - large, requires run on small/medium first and PWG approval
  - PB (Physics Board) approval if too slow (see this slide)

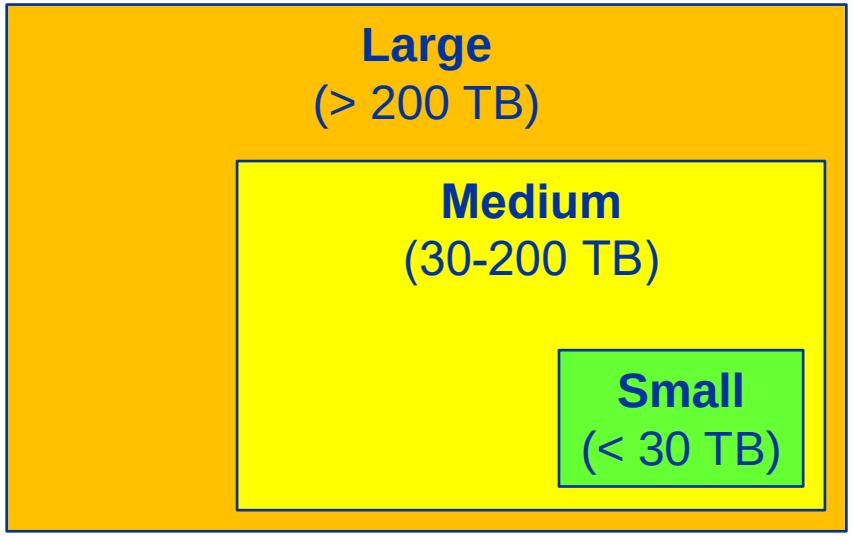

Dataset view – General

- **General**
  - JIRA ticket
  - Dataset sampling (for QC datasets only which are not for analysis)

- **Analysis Facility staging**
  - All datasets can run on the GRID
  - But datasets with heavy usage are staged to the analysis facilities, to run there instead

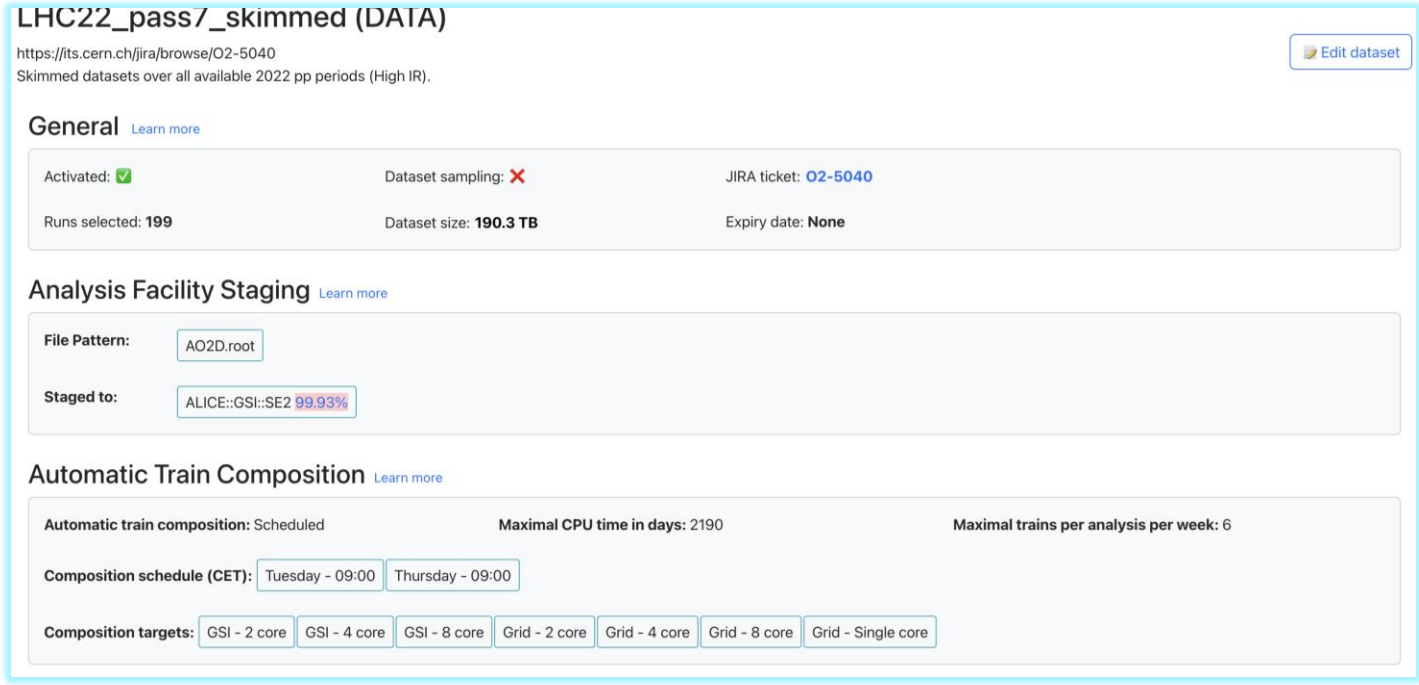

‘Short’ datasets

- Large datasets (> 200 TB) cannot be directly used. They are linked to ‘short’ datasets (subsets) which must first be run on as validation.

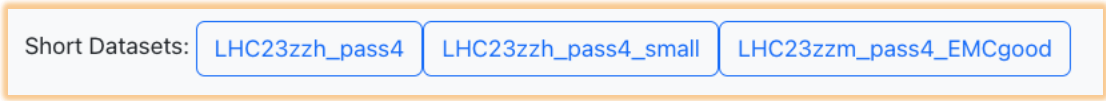

Dataset view – Automatic composition

- **Automatic composition** schedules are found in each dataset

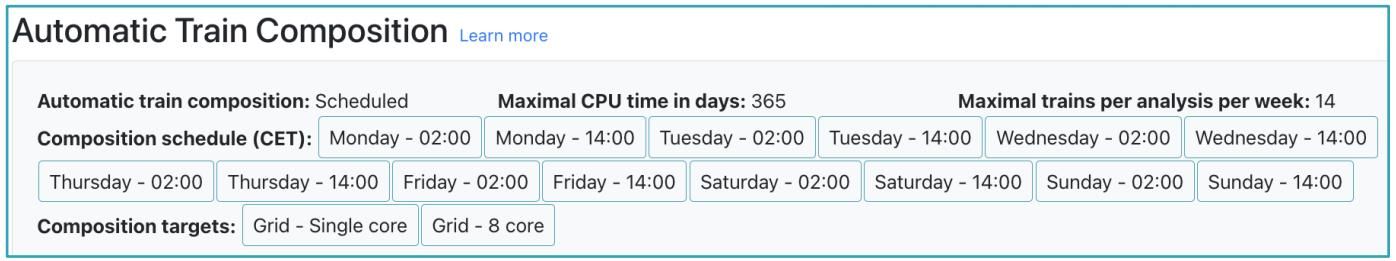

- At the scheduled times, Hyperloop gathers all eligible enabled wagons and compose them into a train.
- **Wagons will not be considered for automatic composition if:**
  - They have **derived data** output which is not slim
  - The wagon test finished with a warning

Derived data wagon composition can be requested as such:

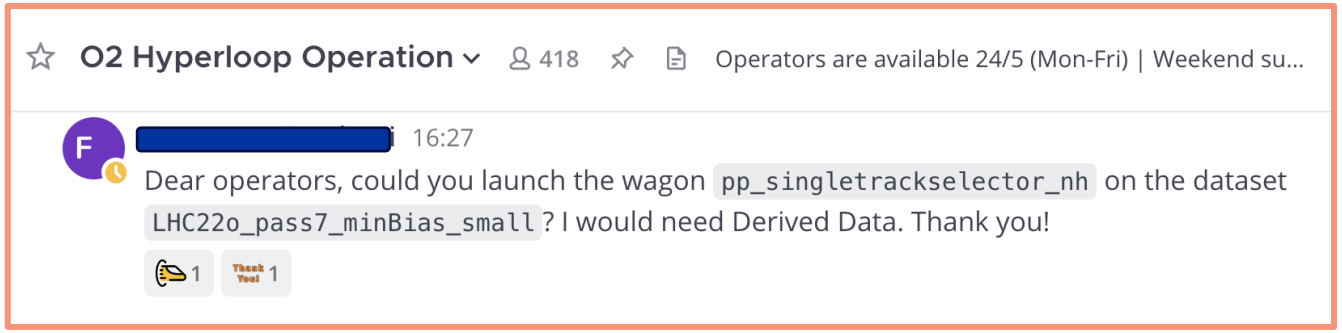

Add a dataset to your Analysis

- Click on '**Datasets and Settings**' in **My Analyses**
- Use the filtering options to find the dataset
- Click '**Add to analysis**'

My Analyses

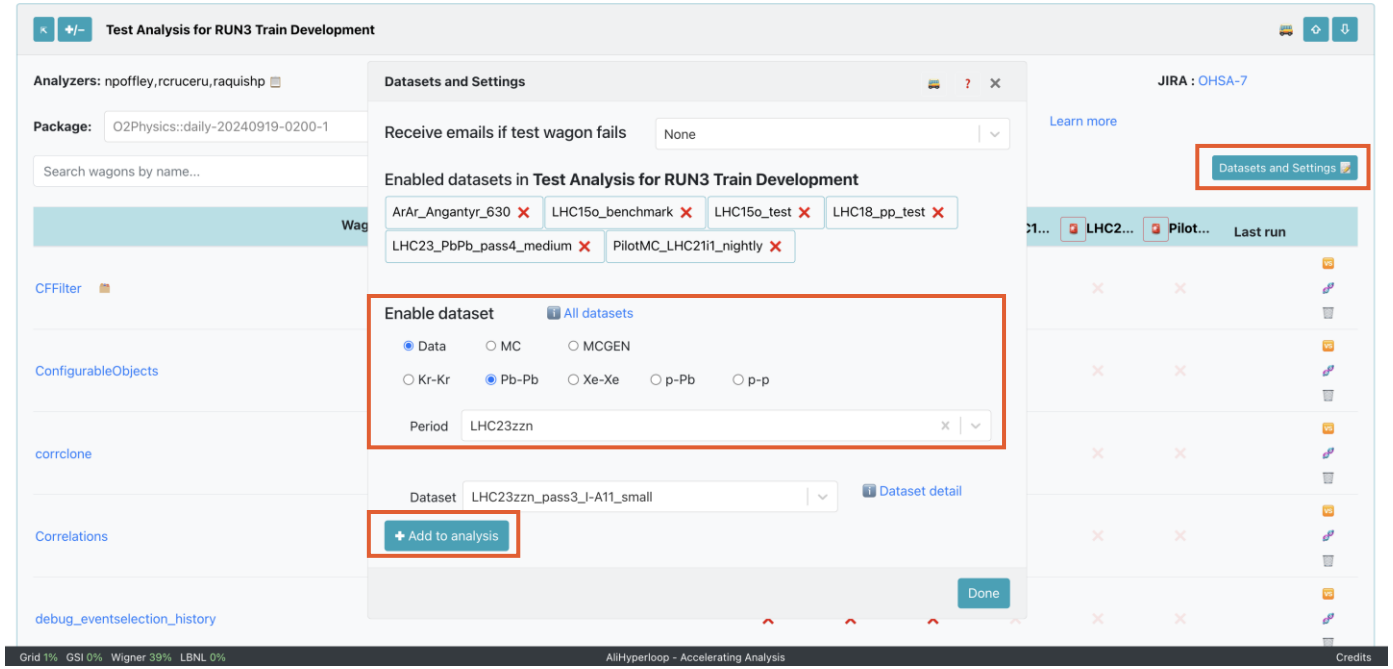

- Or open a dataset and click here:

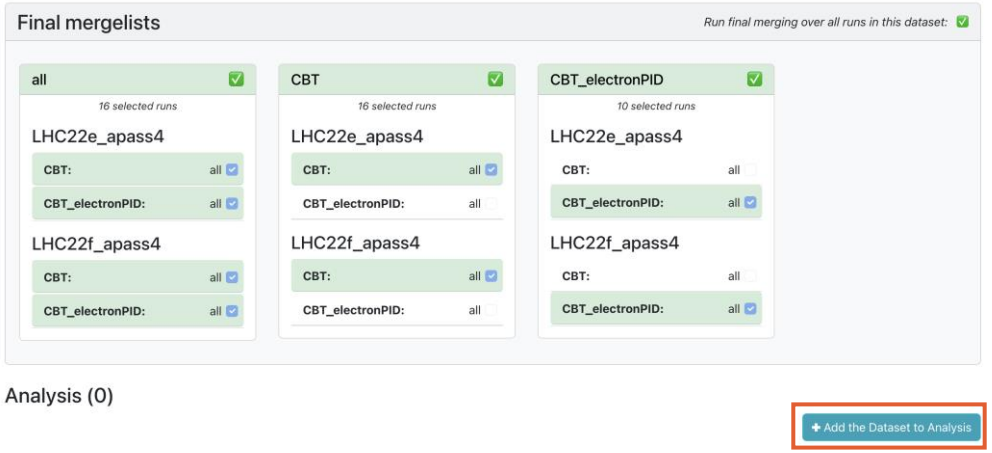

- Introduction & key concepts
- ’My Analyses’, wagons & datasets
- **Enabling a wagon to start a test**
- Trains – automatic composition, debugging & long trains
- Key things to know

Choose a package tag

- Choose a tag from the ***dropdown*** and the wagon will be tested only with that specific tag
- Enabling the '***or newer tags'*** switch will allow the wagon to be included in a train run using the selected tag **or a newer tag** *(Group more wagons together in a train -> resource efficient)*

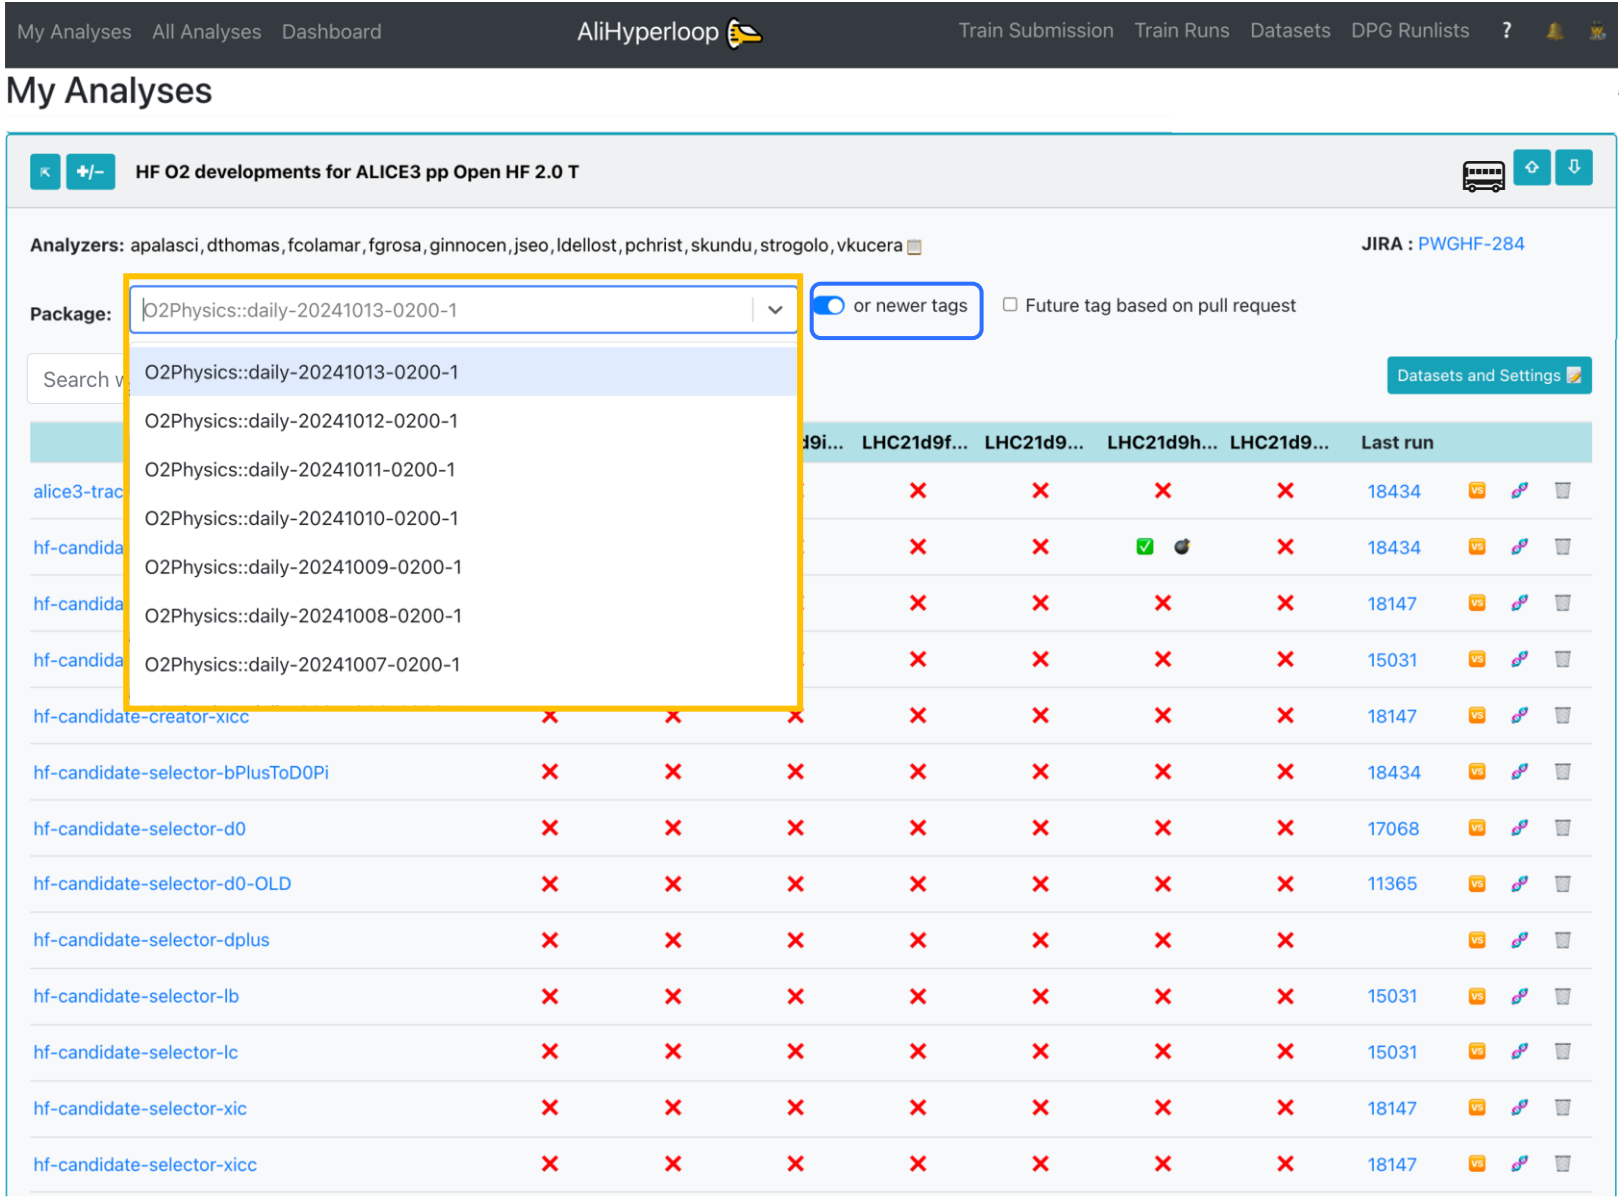

Choose a package tag

- Choose a tag from the ***dropdown*** and the wagon will be tested only with that specific tag
- Enabling the '***or newer tags'*** switch will allow the wagon to be included in a train run using the selected tag **or a newer tag** *(Group more wagons together in a train -> resource efficient)*
- ***Future tag based on PR*** **:**
  - The wagon is queued for a test to start once the *pull request is included in a tag*
  - Once the PR is part of the tag, the wagon can use that tag or a future tag

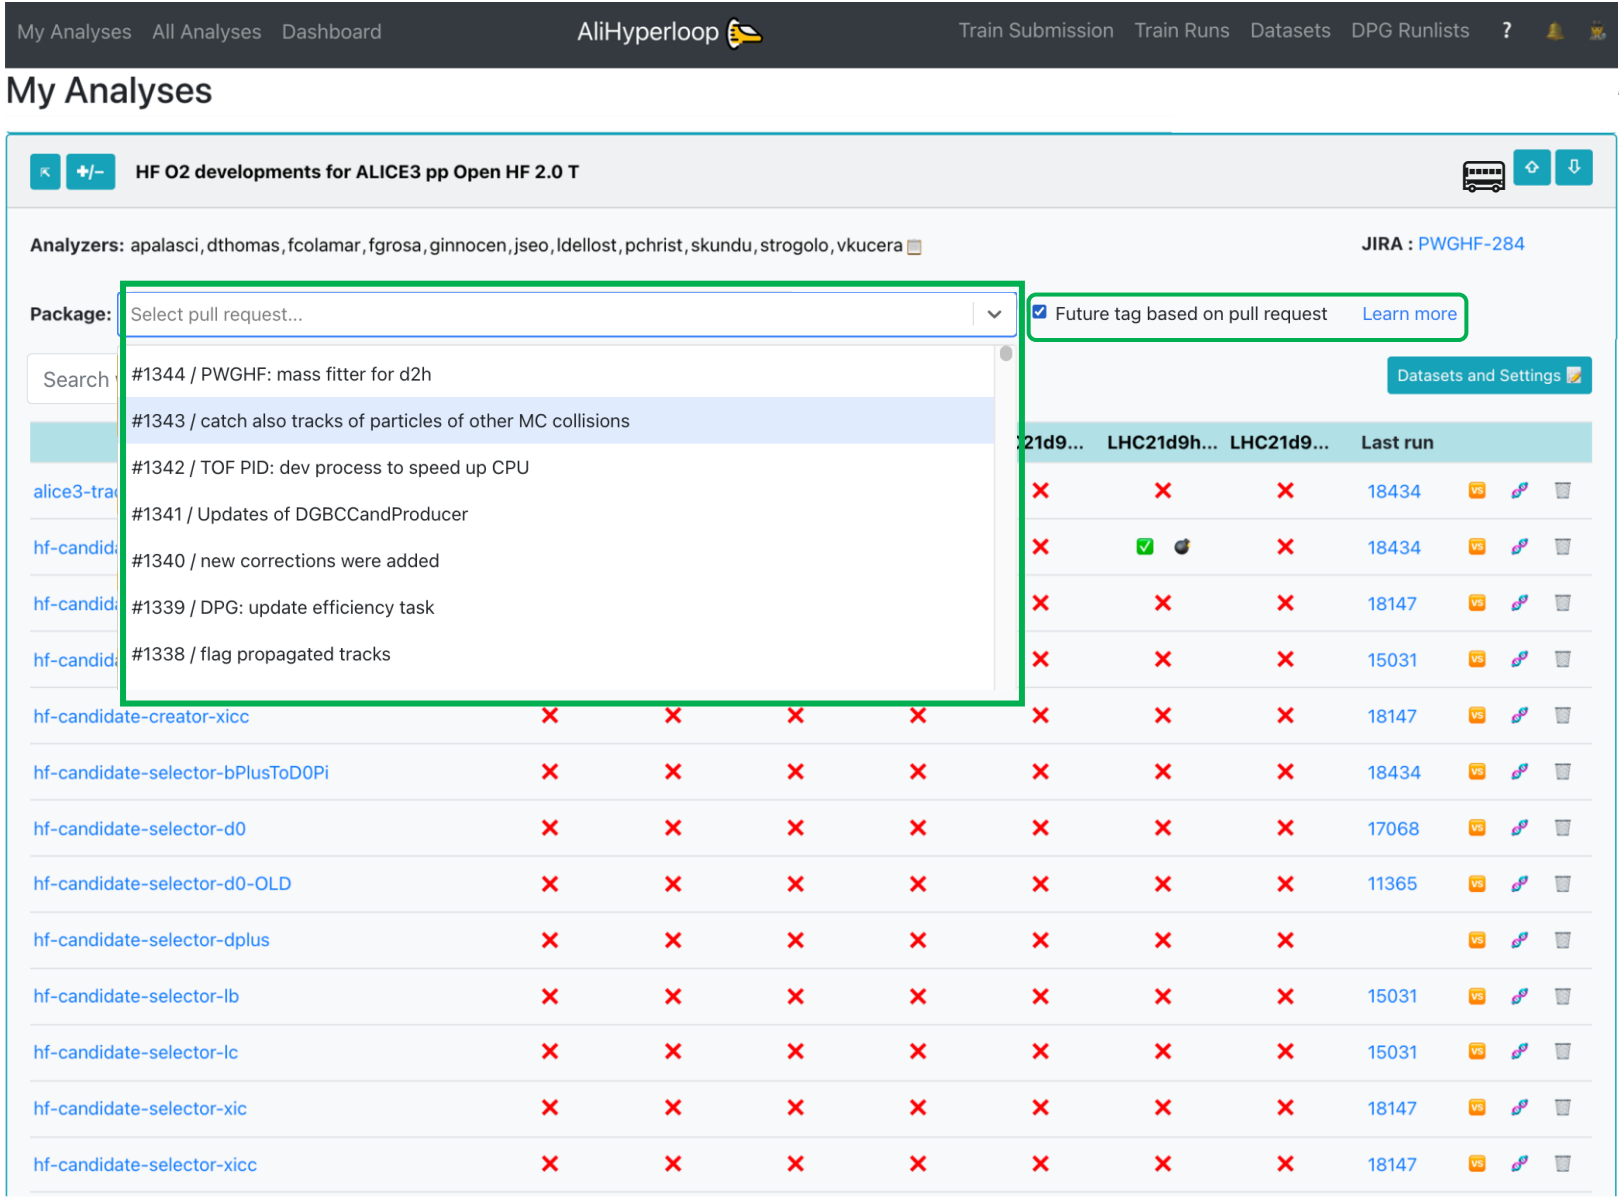

Enable the wagon to start a wagon test

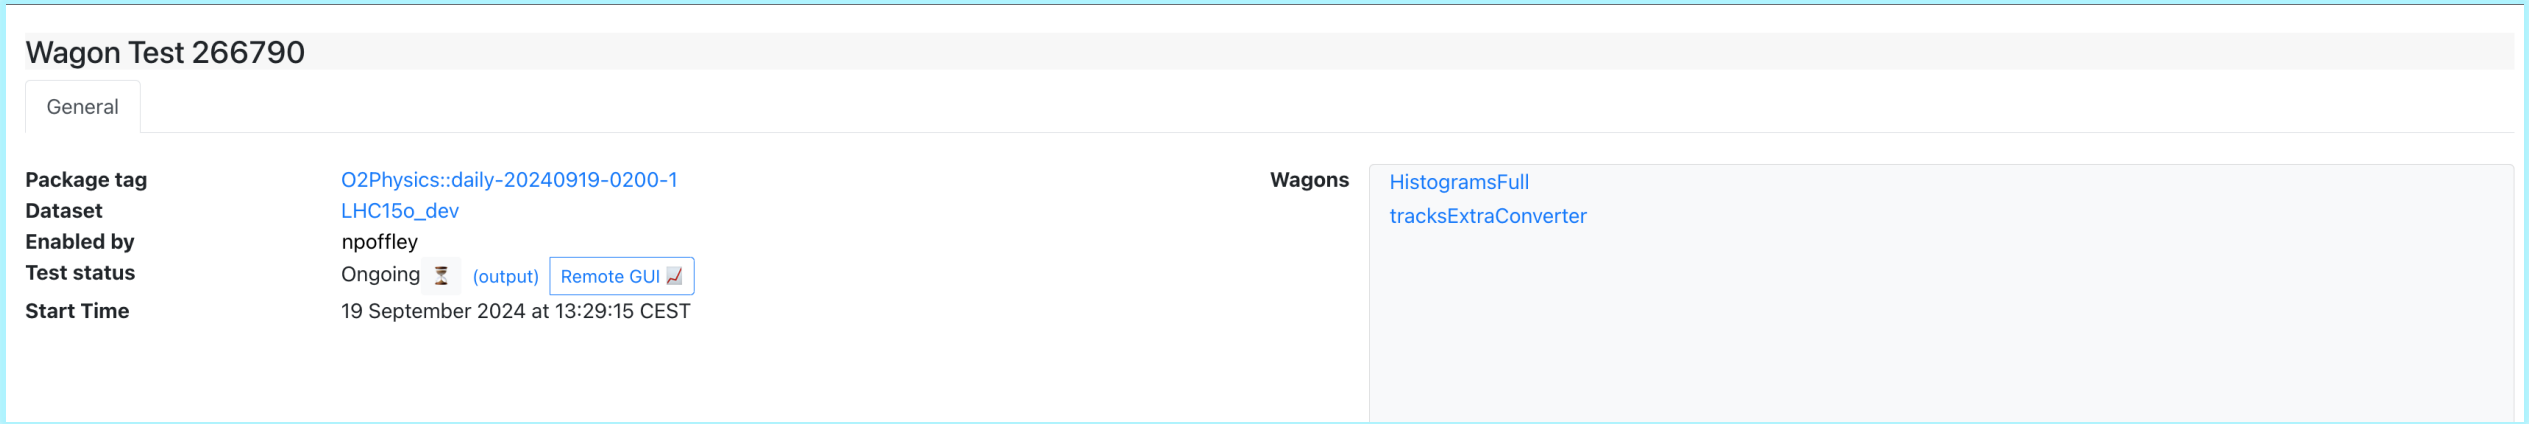

- Enable the wagon on a dataset by clicking on ✕ in **My Analyses**
- When the wagon test starts, click on the ⌛ to open the wagon test page
- View the test output by clicking on (output)

Wagon tests

- Wagons may finish successfully 🌟, finish with warnings ❗, or fail 💣.

- Successful 🌟 wagons are automatically queued for submission if the dataset has an automatic composition schedule.

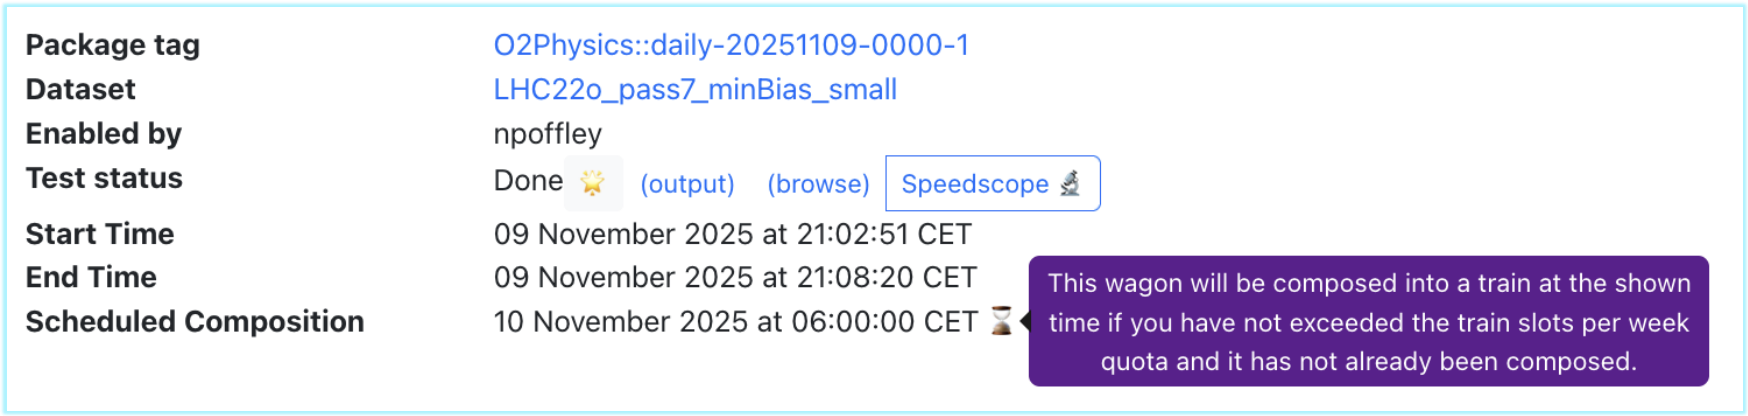

- **Derived data** wagons as well as wagons running on **large datasets** are not automatically submitted.

Wagon tests - warnings

- Wagons with warnings will not be automatically submitted.
- Some wagons with warnings can be submitted with operator intervention.
- Others cannot be submitted at all.
- If unsure, ask the operators for assistance!

Wagon tests – failed wagons

If the wagon test failed, click (output) -> stdout.log to see what went wrong.

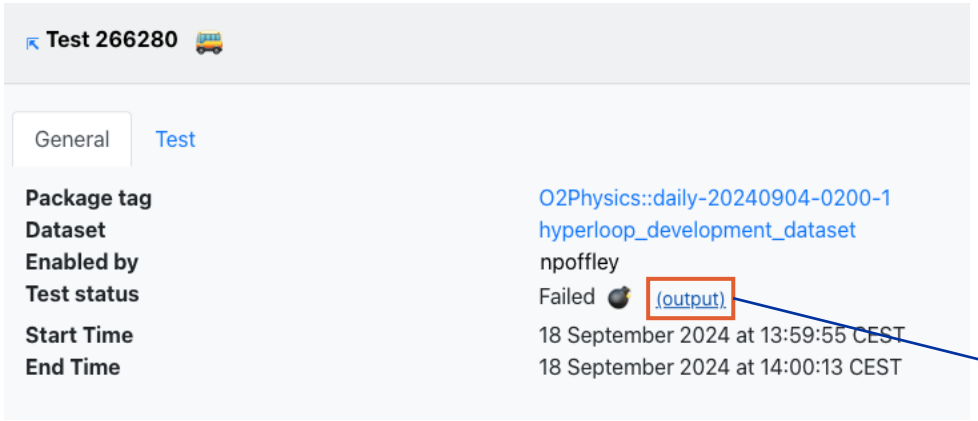


```
[331854:internal-dpl-aod-reader]: [14:00:10][ERROR] Exception while running: Couldn't get TTree "DF_2235406567711662/O2track" from "/home/alitrain/train-
workdir/testdata/LFN/alice/data/2022/LHC22f/520471/apass4/1450/o2_ctf_run00520471_orbit0644189358_tf0000000301_epn069/002/AO2D.root". Please check https://aliceo2group.github.io/analysis-framework/docs/troubleshooting/#tree-not-found for more
information.. Rethrowing.
[331854:internal-dpl-aod-reader]: [14:00:10][STATE] Exiting FairMQ state machine
[331854:internal-dpl-aod-reader]: [14:00:10][FATAL] Unhandled o2::framework::runtime_error reached the top of main of o2-analysistutorial-histograms, device shutting down. Reason: Couldn't get TTree "DF_2235406567711662/O2track" from
"/home/alitrain/train-workdir/testdata/LFN/alice/data/2022/LHC22f/520471/apass4/1450/o2_ctf_run00520471_orbit0644189358_tf0000000301_epn069/002/AO2D.root". Please check https://aliceo2group.github.io/analysis-framework/docs/troubleshooting/#tree-
not-found for more information.
```

- Introduction & key concepts
- ’My Analyses’, wagons & datasets
- Enabling a wagon to start a test
- **Trains – automatic composition, debugging & long trains**
- Key things to know

Train composition flow chart

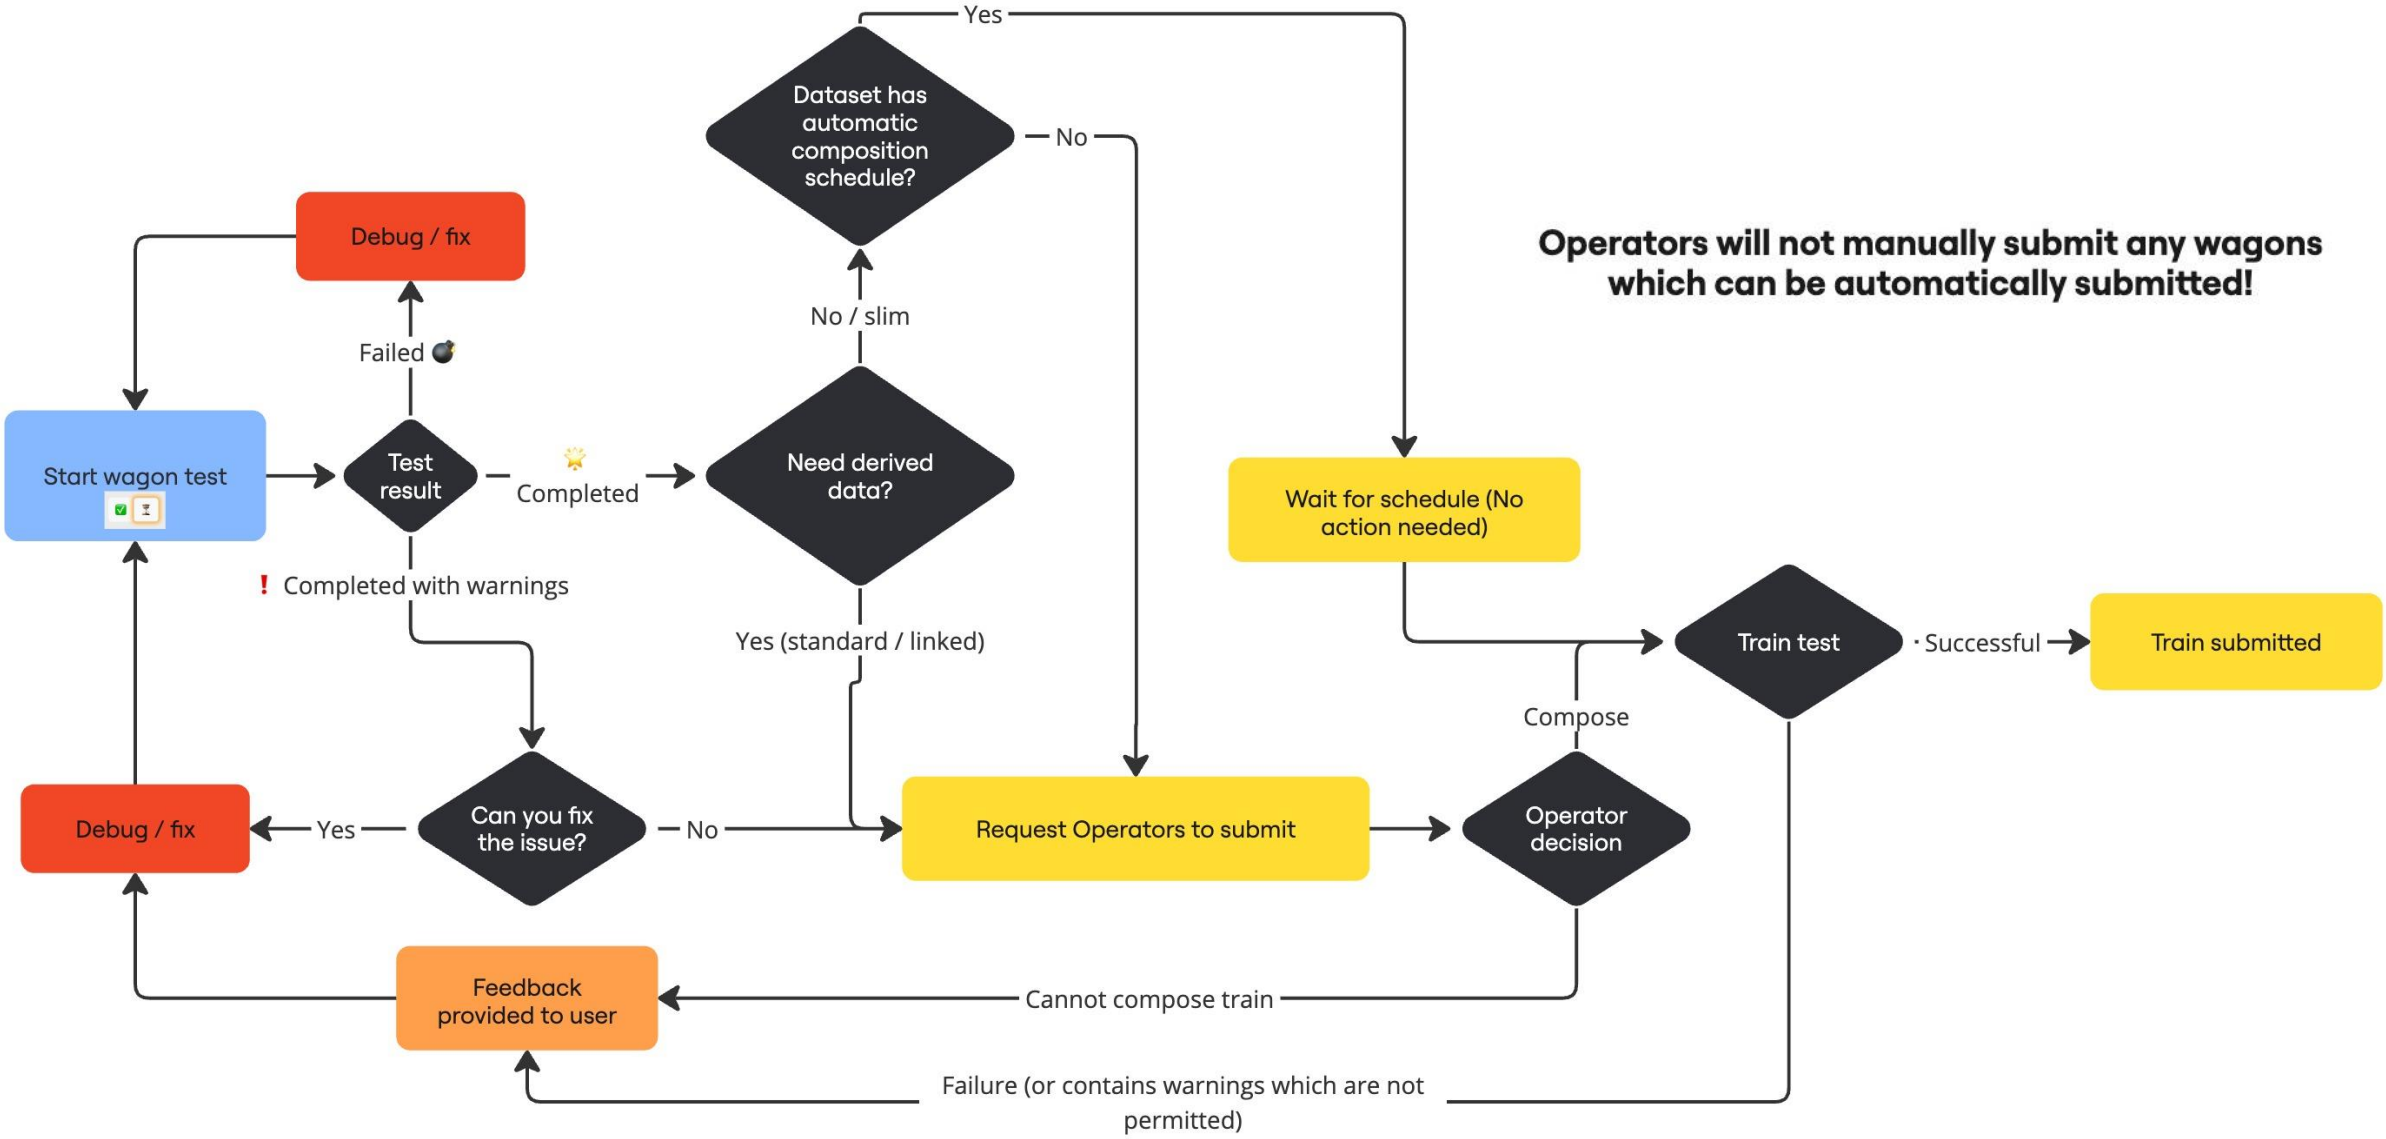

Train Runs page

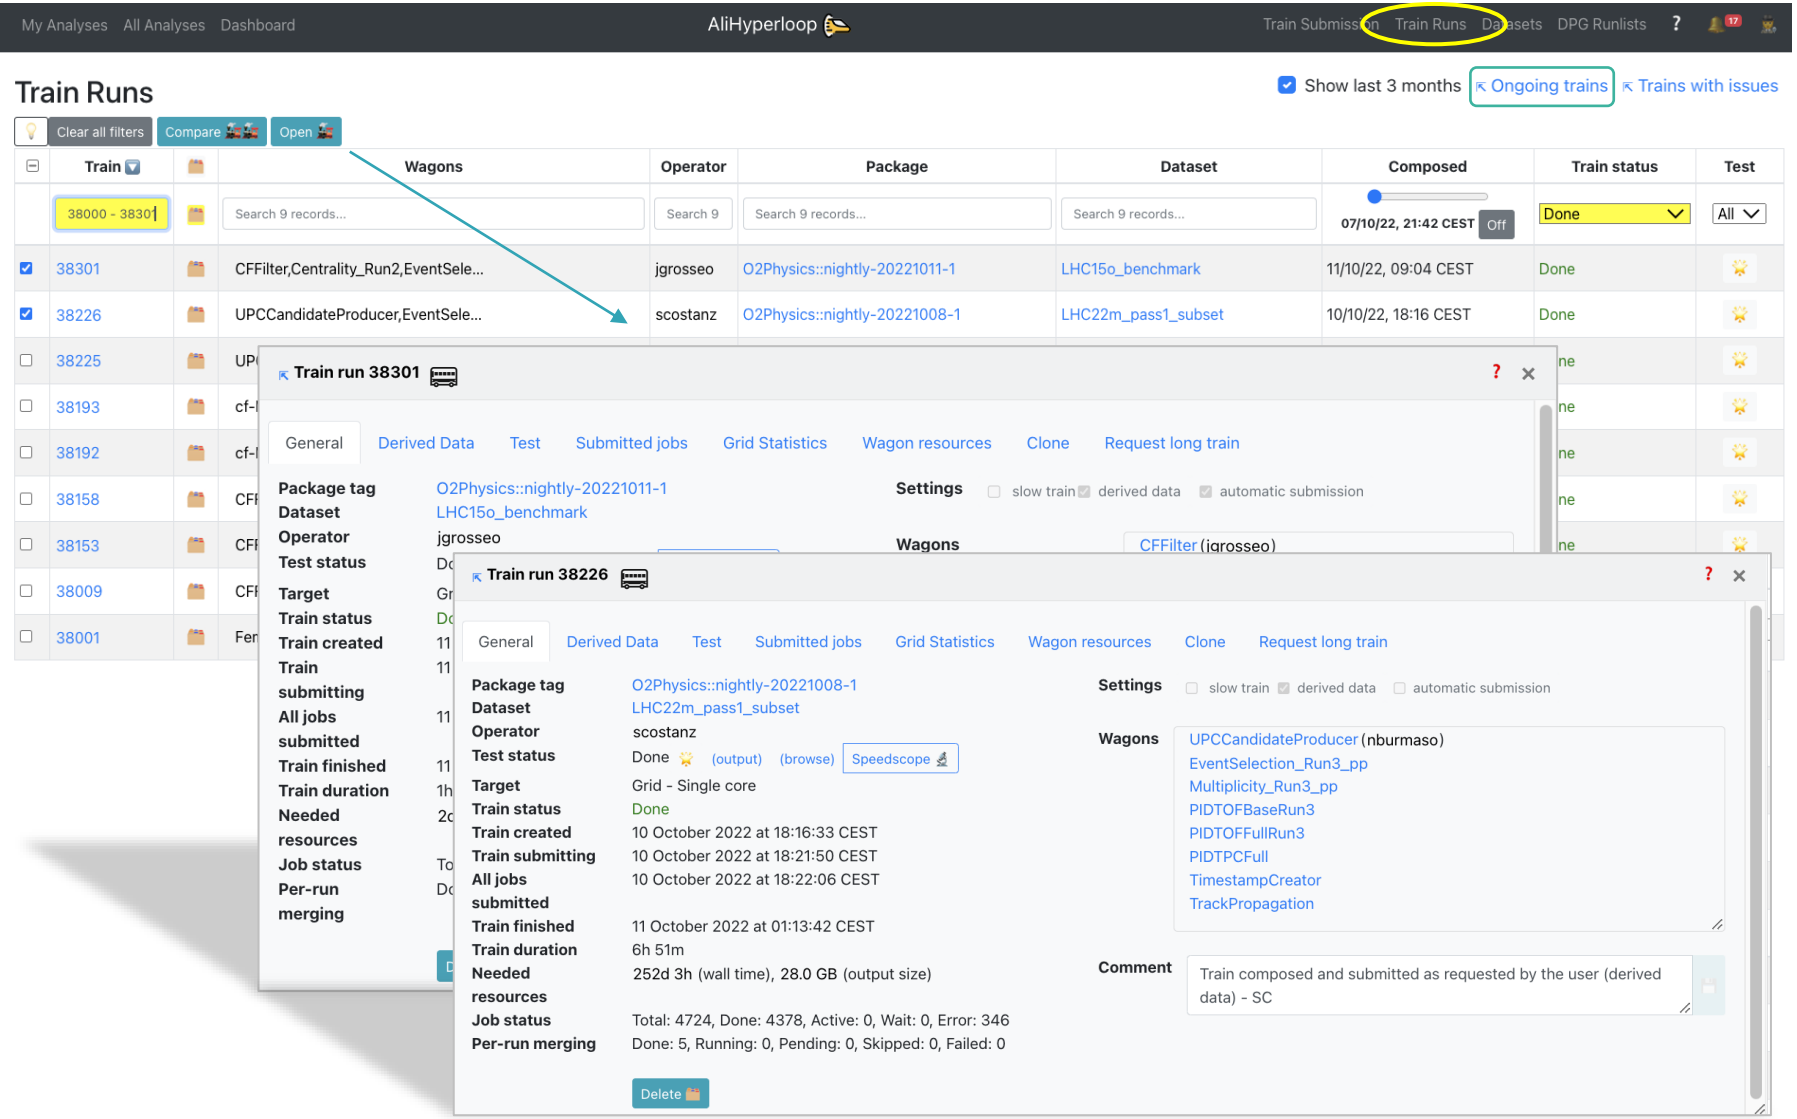

- Displays all the trains
- To compare two trains, select them in the Compare column and click *Compare.* This will open a new tab displaying the differences between the two trains

Train run view – General

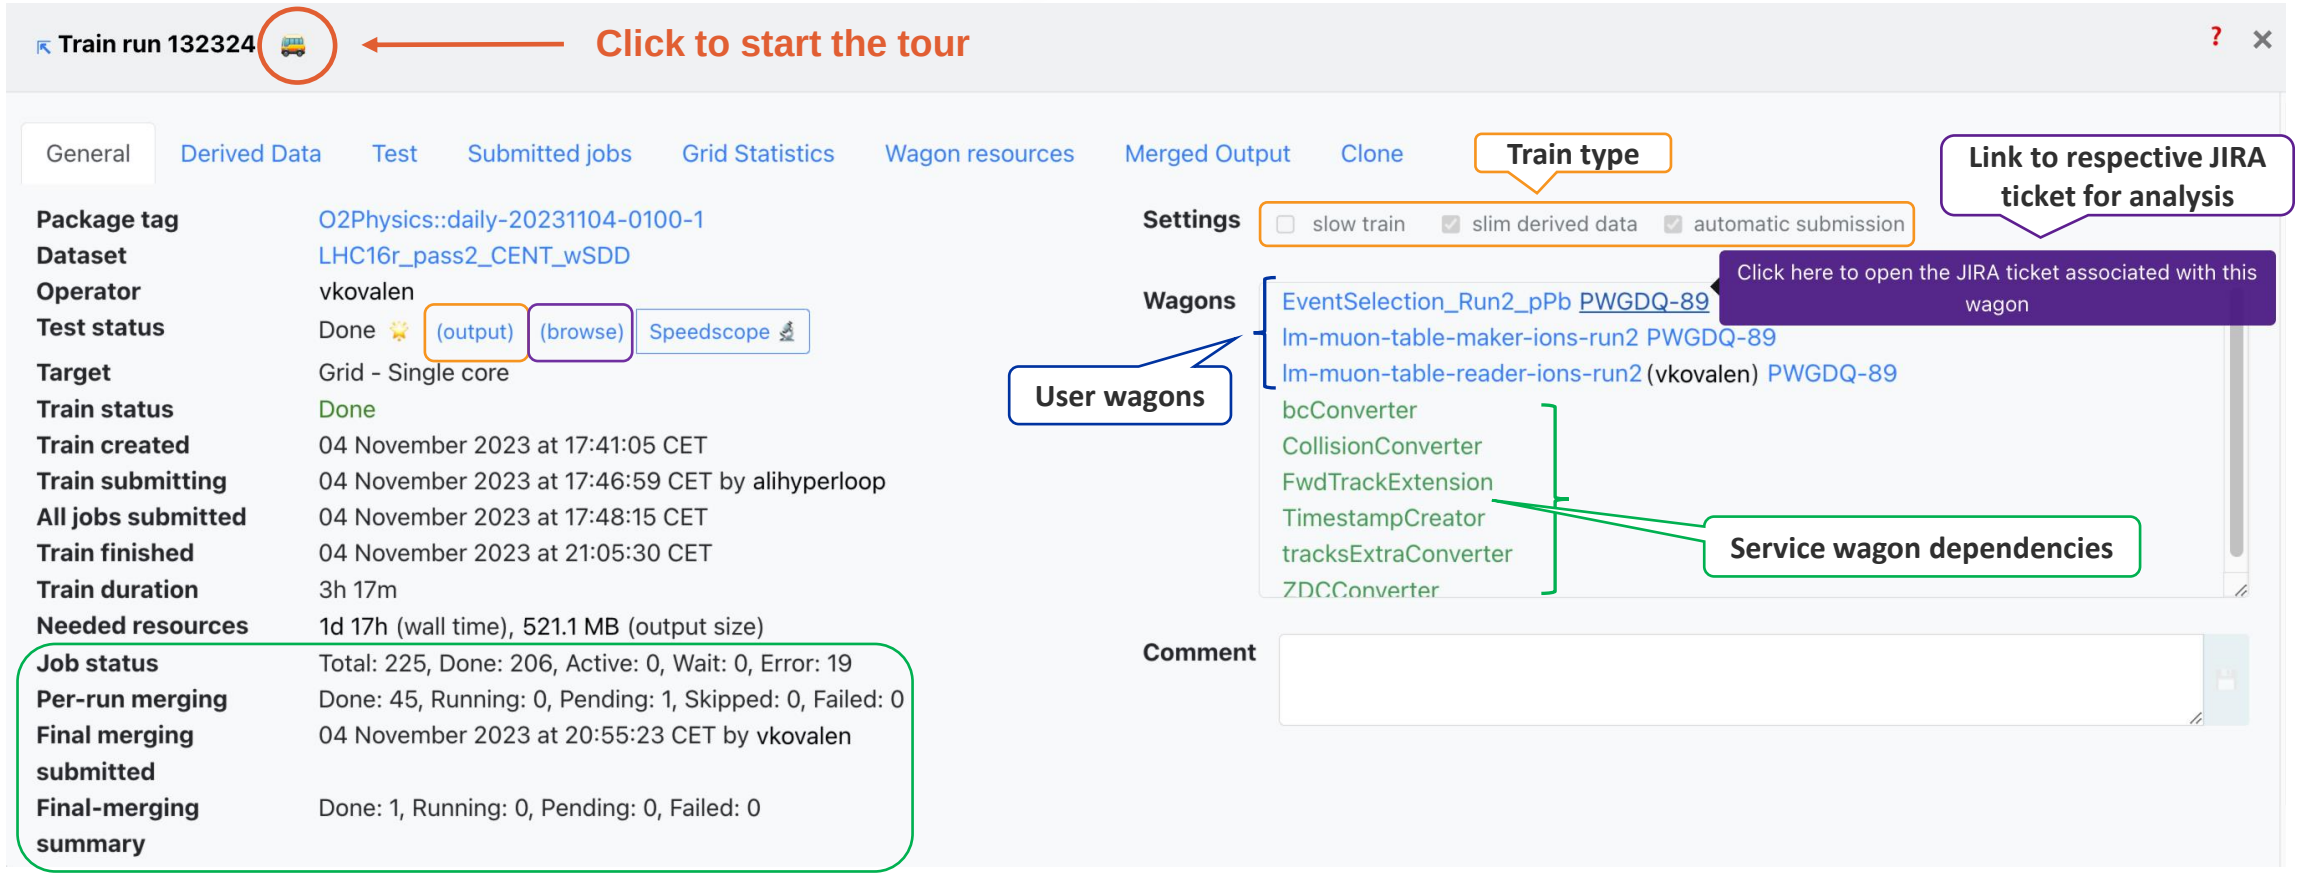

- Direct links to dataset and wagons **at the timestamp of the train (preserved state – changes after train run ignored)**
- Duration and needed resources

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

Train run view – Train test

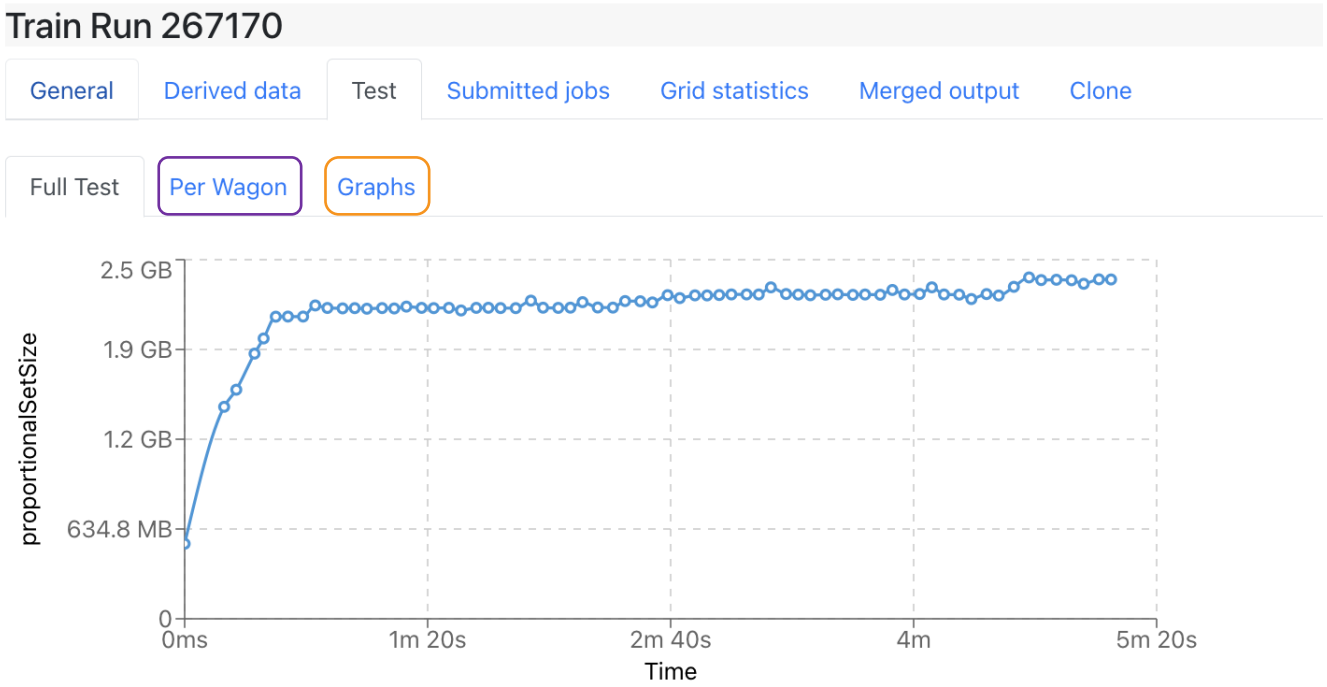

Number of input files	3
Input size	8.3 GB
Output size	63.0 KB
Output size (AO2D only)	15.4 KB
Reduction factor	562572
PSS Memory	Max: 2.4 GB Avg: 2.2 GB Slope: 1.1 MB/s
Private Memory	Max: 1.9 GB Avg: 1.8 GB Slope: -38.0 KB/s
Timing	CPU: 3m 59s Wall: 5m 10s
Throughput	27.3 MB/s
Expected resources	11d 14h 48.7 MB
Estimated emissions	2 kgCO2eq
Weekly slots used	*Non-prompt cascades with strangeness tracking:* 1

- Shows statistics of the train test, similar to the wagon test
- View graphs and per wagon statistics
- Expected resources of the full train run
- Estimated kgCO2eq emissions of the full train run
- Weekly slots used by the train from the dataset quota

Train run view – Train test

View statistics per wagon

View graphs on different statistics per wagon

Train Run 267170

General | Derived data | Test | Submitted jobs | Grid statistics | Merged output | Clone

Full Test | Per Wagon | Graphs

Wagon		PSS Memory	Private Memory	CPU Time
Reader	Max Avg Slope	700.0 MB 610.8 MB 136.4 KB/s	558.9 MB 446.9 MB -234.3 KB/s	1m 14s (30%)
CorrectedFT0	Max Avg Slope	23.7 MB 21.8 MB 14.4 KB/s	18.9 MB 18.8 MB 635.7 B/s	2s (1%)
EventSelection_Run3_pp	Max Avg Slope	318.7 MB 273.6 MB 319.8 KB/s	241.3 MB 232.9 MB 167.0 KB/s	19s (8%)
mcCollisionConverter	Max Avg Slope	19.9 MB 18.5 MB 9.0 KB/s	16.6 MB 16.5 MB 389.9 B/s	1s (0%)
Multiplicity_Run3_pp	Max Avg Slope	181.4 MB 157.0 MB 162.1 KB/s	134.4 MB 130.1 MB 66.1 KB/s	7s (3%)
new-non-prompt-cascade-mc	Max Avg Slope	227.7 MB 203.5 MB 142.7 KB/s	178.8 MB 176.5 MB 45.8 KB/s	6s (2%)

Train run view – Submitted jobs

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

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

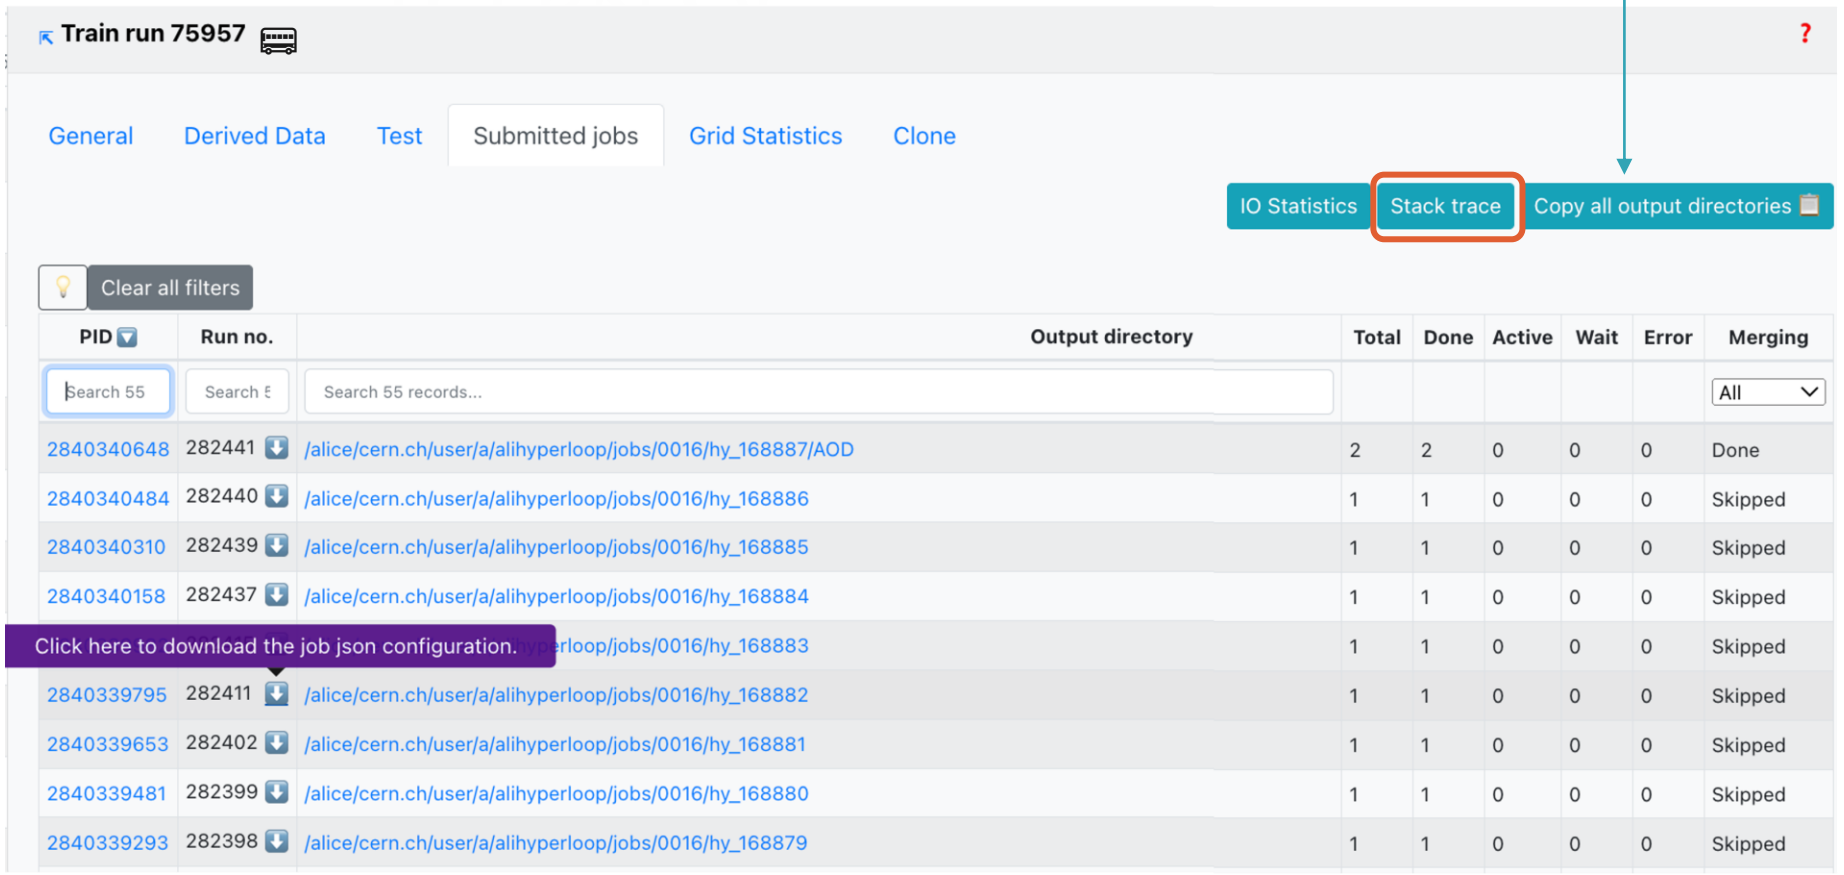

PID	Run no.	Output directory	Total	Done	Active	Wait	Error	Merging
2840340648	282441	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168887/AOD	2	2	0	0	0	Done
2840340484	282440	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168886	1	1	0	0	0	Skipped
2840340310	282439	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168885	1	1	0	0	0	Skipped
2840340158	282437	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168884	1	1	0	0	0	Skipped
		/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168883	1	1	0	0	0	Skipped
2840339795	282411	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168882	1	1	0	0	0	Skipped
2840339653	282402	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168881	1	1	0	0	0	Skipped
2840339481	282399	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168880	1	1	0	0	0	Skipped
2840339293	282398	/alice/cern.ch/user/a/alihyperloop/jobs/0016/hy_168879	1	1	0	0	0	Skipped

Train run view – Submitted jobs

- **Stack trace**

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

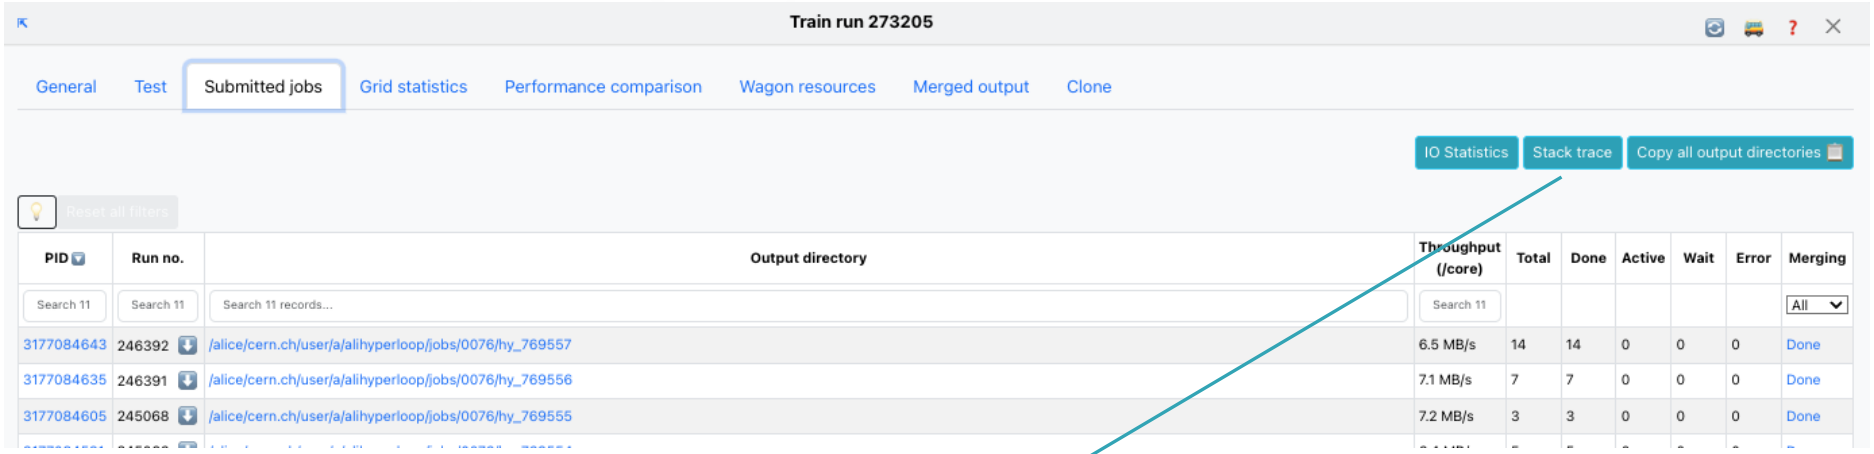

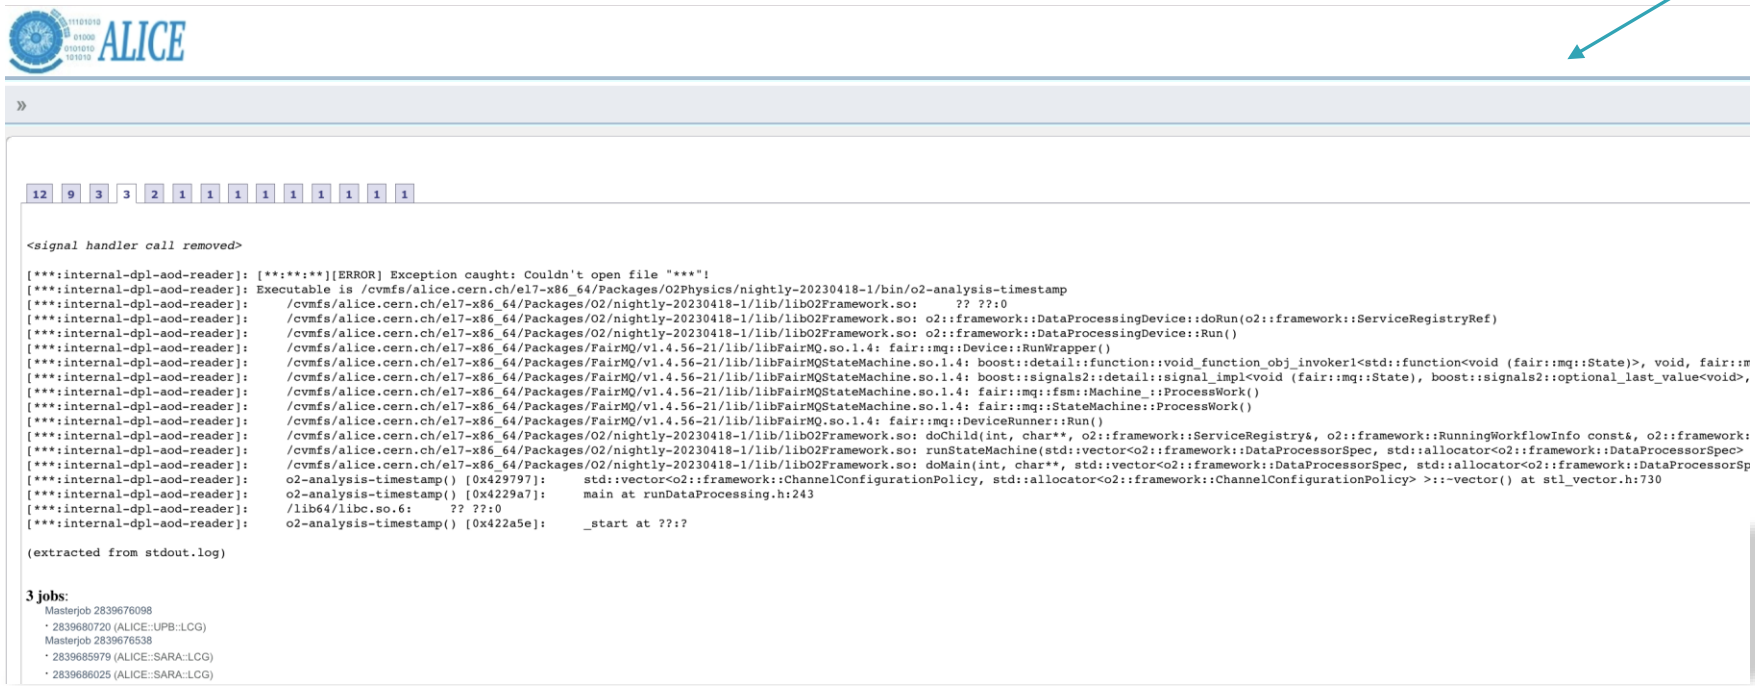

- **Per-run merge**
  - This shows the status of every per-run merging job.
  - Merging only after all jobs are finished
  - *Only occurs if > 50% of jobs succeeded*

Train run view – Derived data

Train run 75957

General | Derived Data | Test | Submitted jobs | Grid Statistics | Clone

Max DF size: 100000000 Max derived file size: 0

Wagon	Binding	Description
DQ-TPCPID-trkEff-maker-run2	ReducedEvents	REDUCEDEVENT
DQ-TPCPID-trkEff-maker-run2	ReducedTracks	REDUCEDTRACK
DQ-TPCPID-trkEff-maker-run2	ReducedEventsExtended	REEXTENDED
DQ-TPCPID-trkEff-maker-run2	ReducedEventsVtxCov	REVTXCOV

- Displays the activated derived data output (all the output enabled in the wagons of the train)

Train run view – Grid Statistics

- Summarizes job performance
- See a summary of how many jobs failed,

  and what error codes they failed with

- Plots:
  - Files/Job
  - CPU time/Job
  - Wall time/Job

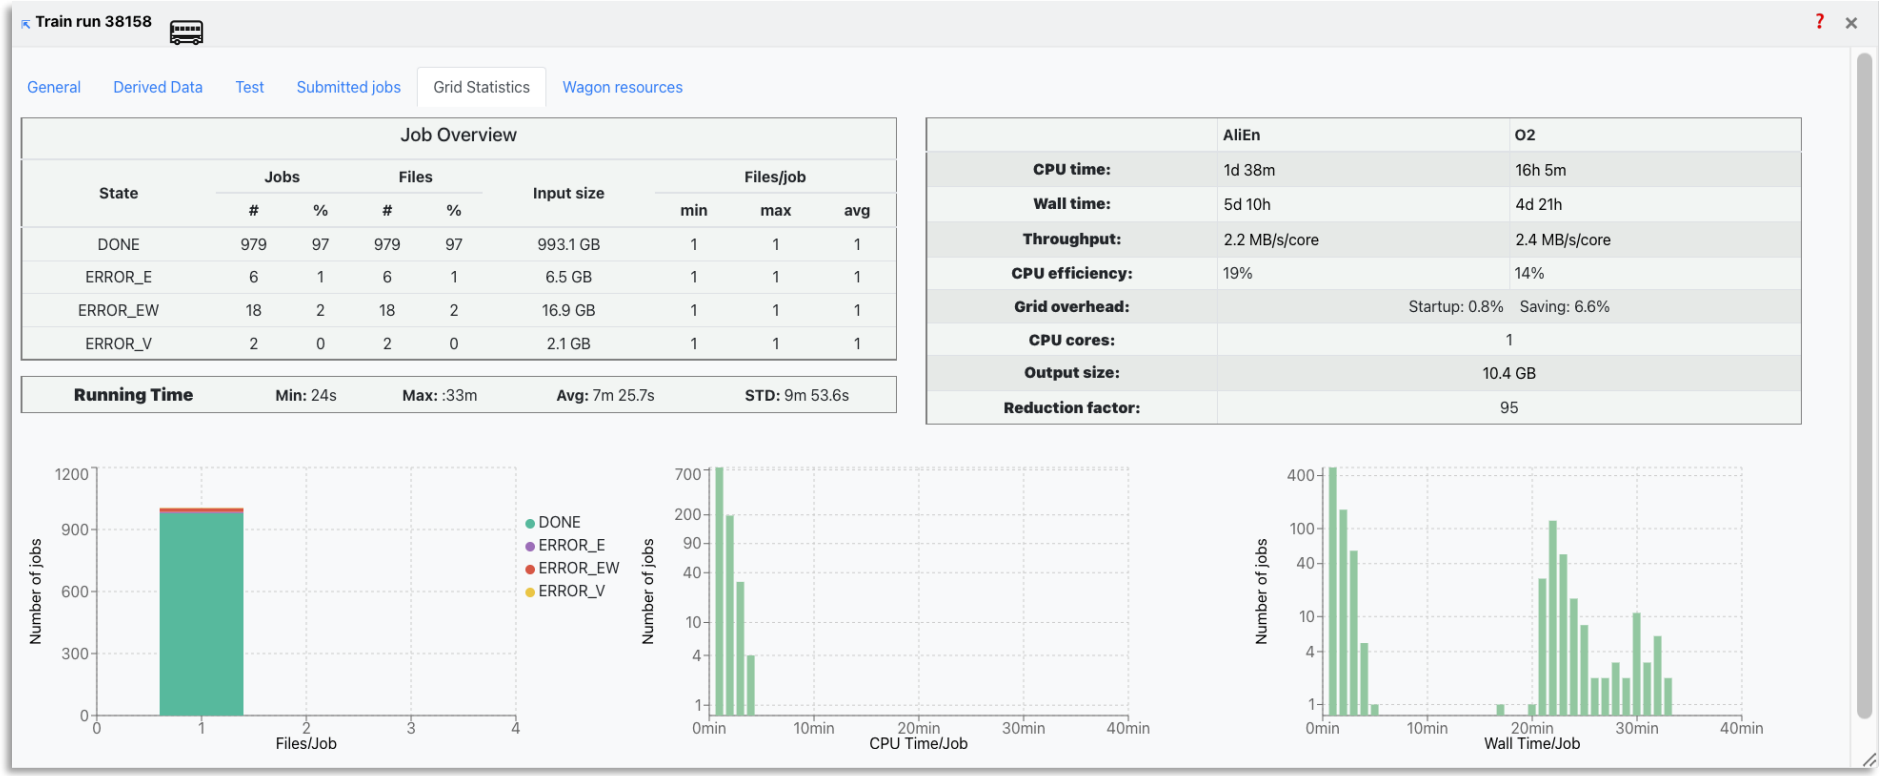

Train run view - Merged Output

- **For analysis and slim derived data trains only**
- Merged outputs generated for the full train, plus every individual mergelist and final mergelist. The mergelists are defined in the ***dataset***
- Runs may be excluded from the final merge if the per-run merges did not complete.

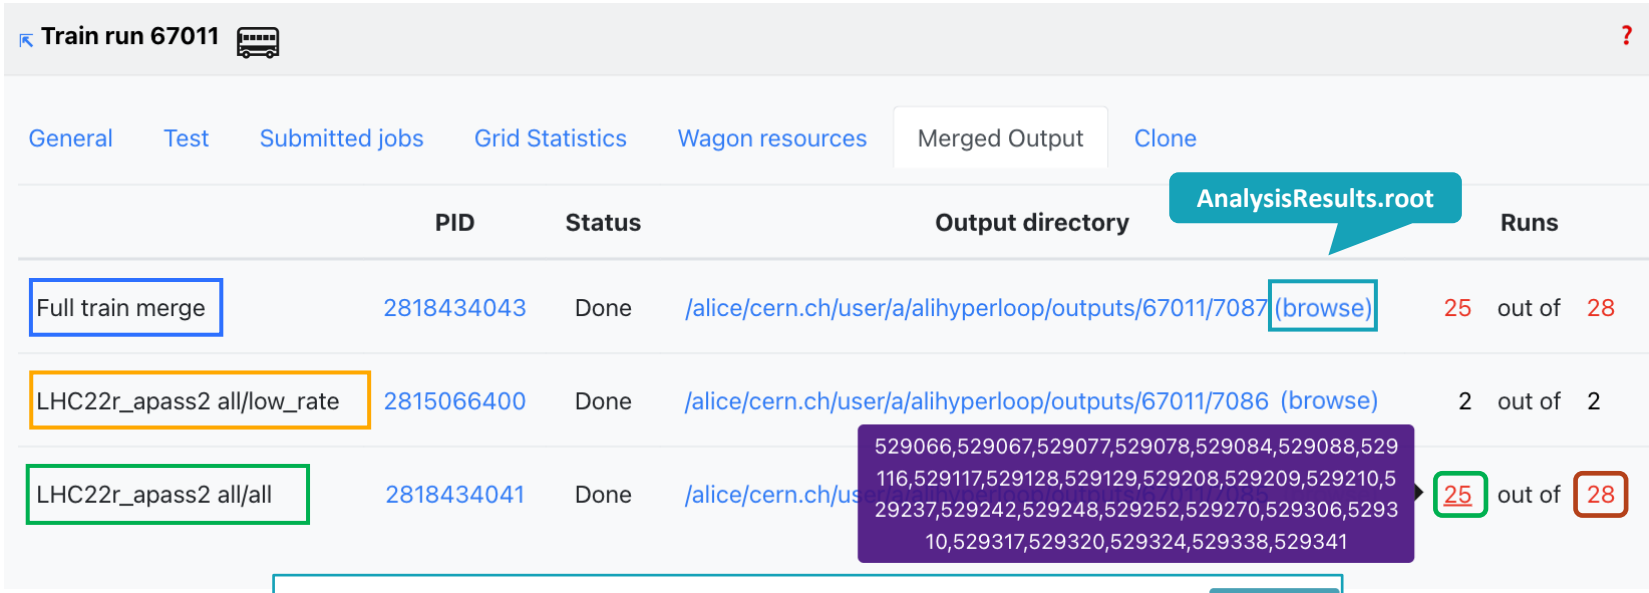

Long train request

- From the completed train on the **short** dataset, click on the ‘Request long train’ tab

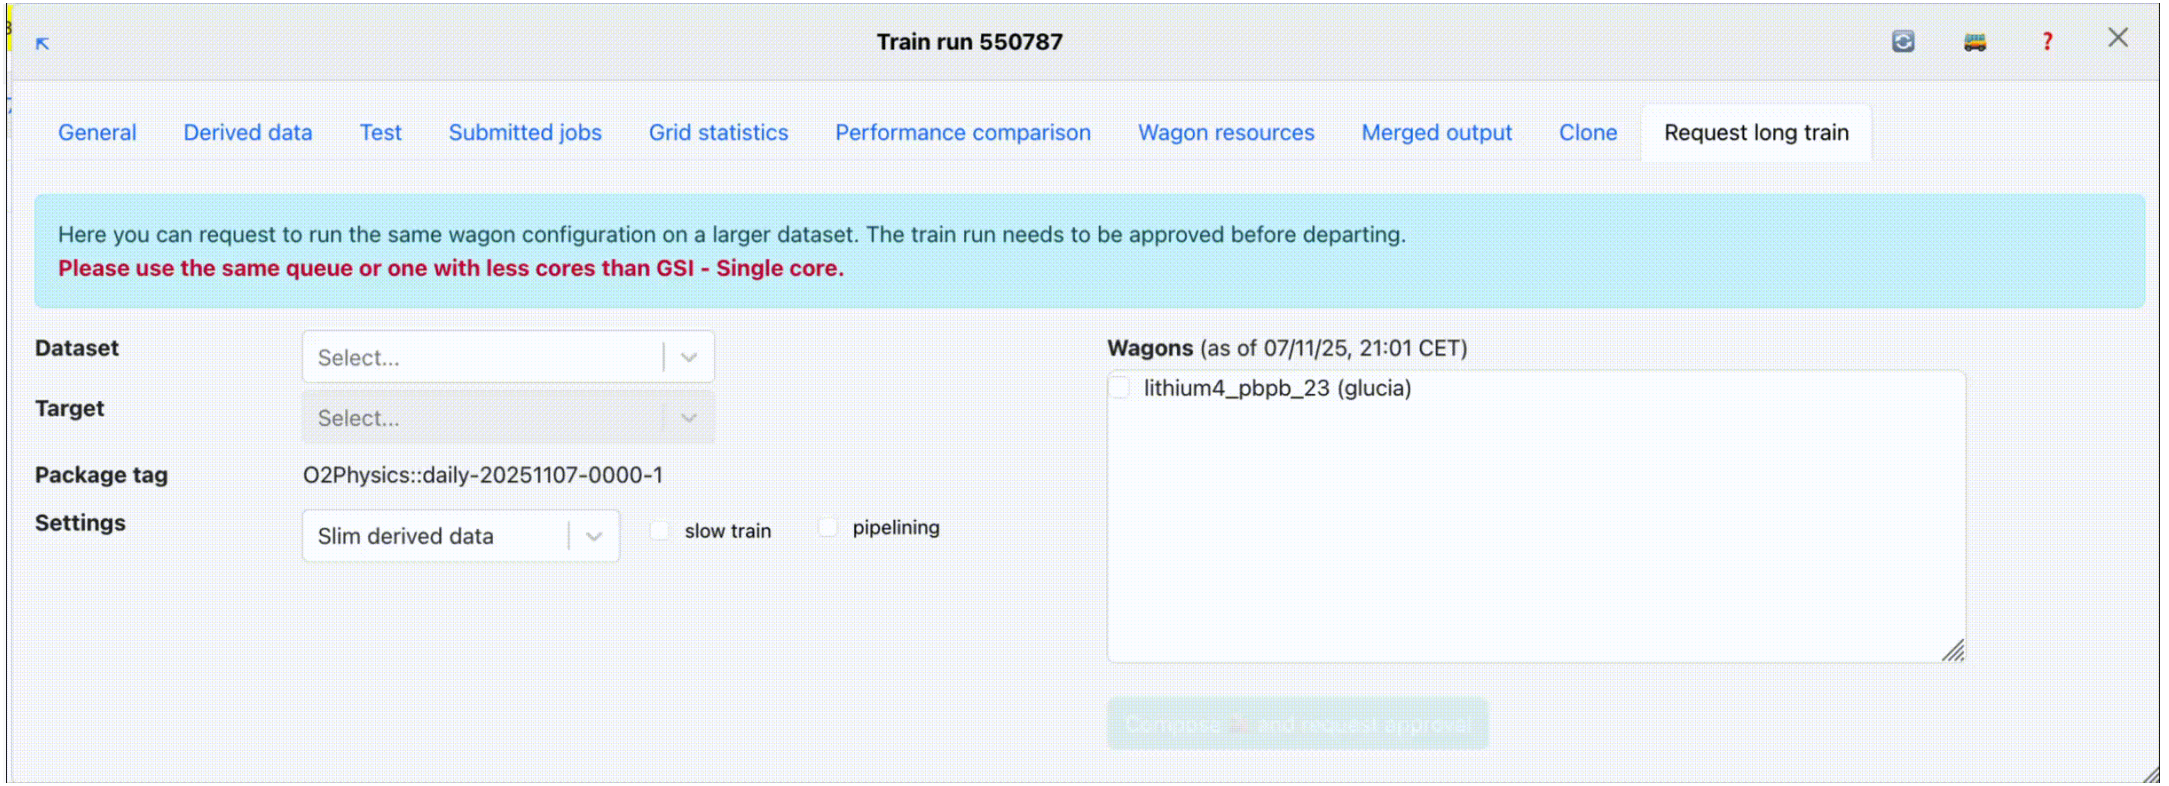

- Any user who is part of the analysis can request a long train
- You can change from slim to standard derived data (recommended if output is large)
- **Approval from the participating analyses PWGs conveners is required in order to submit a long train.**

Long train request

You may **decompose** your own long train if needed, if a mistake was made.

Once the long train is approved:

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

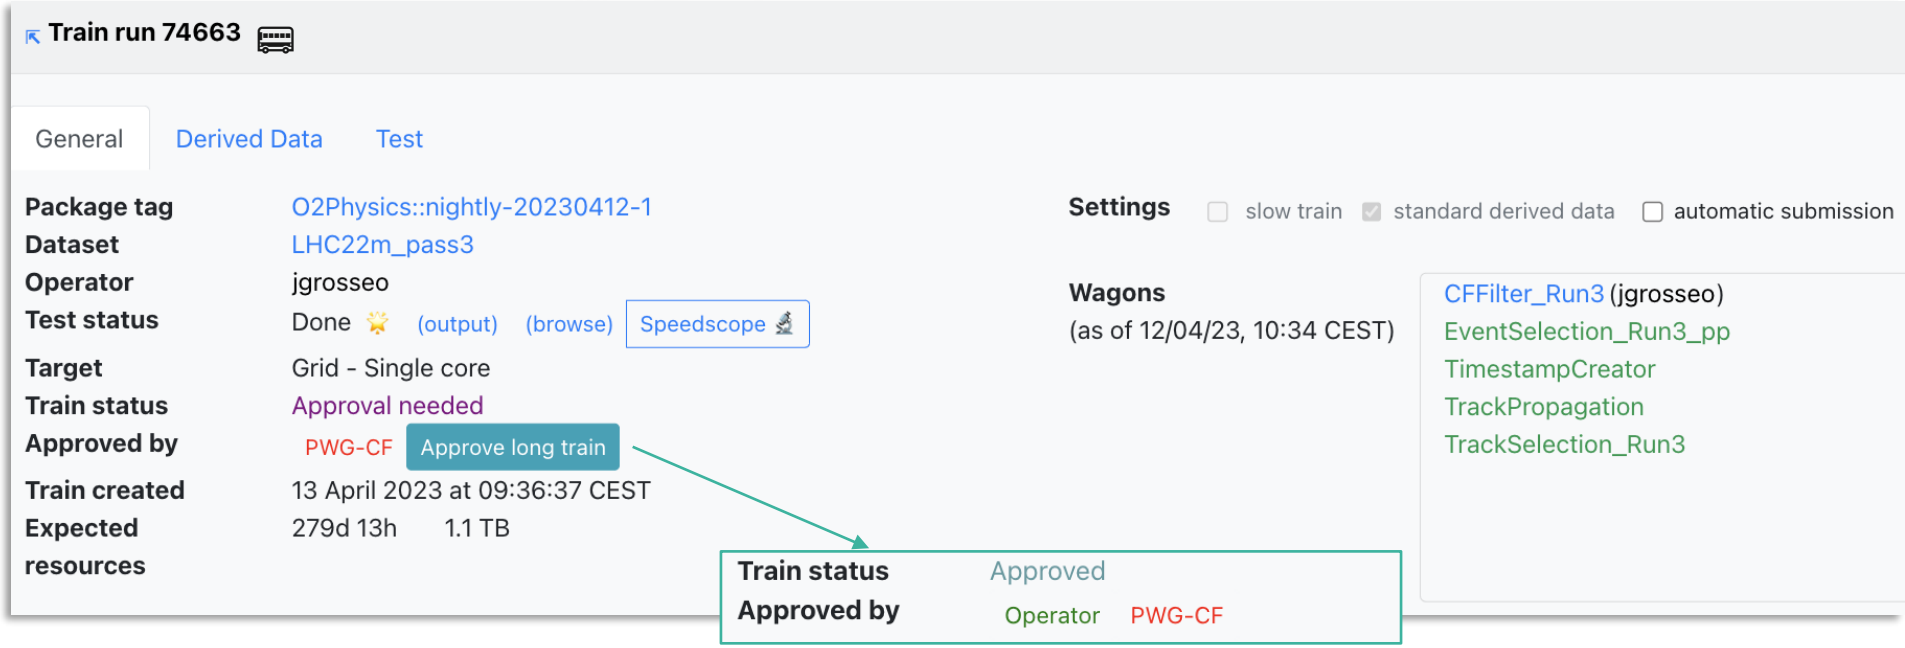

Train run - debugging

If your train has many failed jobs…

Check Grid statistics for an overview of the error codes

In this example, 34% of jobs failed with ERROR_V...

Train Run 245367

General | Test | Submitted jobs | Grid statistics | Performance comparison | Wagon resources | Merged output | Clone | Request long train

Job Overview

State	Jobs		Files		Input size	Files/job		
	#	%	#	%		min	max	avg
DONE	6970	63	679048	63	1.9 PB	1	100	97
ERROR_A	6	0	596	0	1.7 TB	99	100	99
ERROR_E	1	0	93	0	242.6 GB	93	93	93
ERROR_E (idle)	330	3	31823	3	90.4 TB	3	100	96
ERROR_E (TTL)	4	0	382	0	595.9 GB	92	99	96
ERROR_EW	6	0	280	0	784.4 GB	1	100	47
ERROR_V	3665	33	363541	34	0.9 PB	1	100	99
EXPIRED	1	0	99	0	271.0 GB	99	99	99
Running Time	**Min:** 45.9s		**Max:** :11h 48m		**Avg:** 1h 32m 26.1s		**STD:** 51m 17.1s	

Train run - debugging

- Check ‘Submitted jobs’ tab to find masterjobs with many subjob failures
- Click on a masterjob, then choose a subjob trace and log to see what went wrong.
- For some errors, you’ll find the cause in the **trace**. For jobs failing with ERROR_V, check the **log files** for stdout and the **stack trace** to see the errored jobs grouped by error type in the stdout.

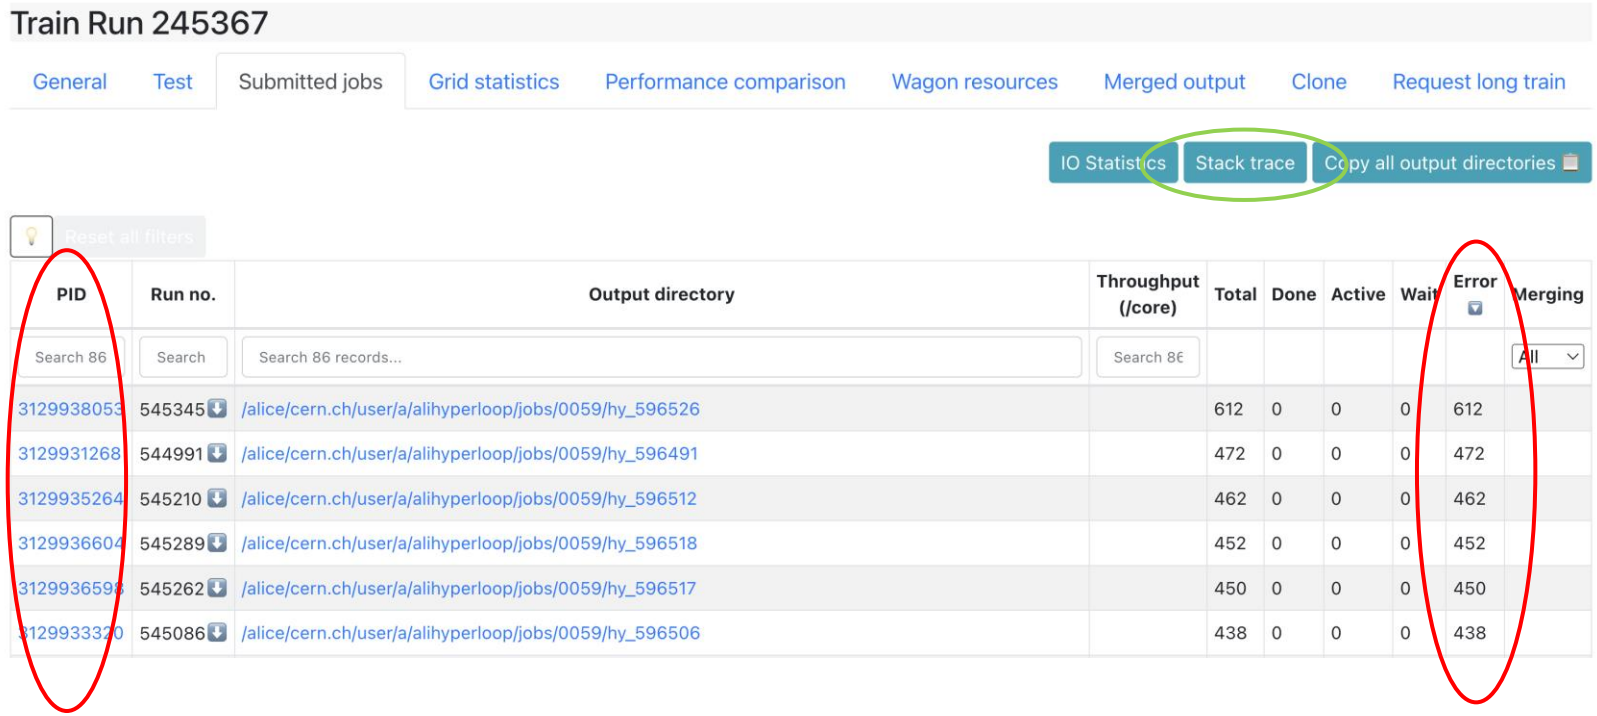

Train run - debugging

If you don’t understand your error..

- Check https://aliceo2group.github.io/analysis-framework/docs/troubleshooting/
- Ask in your working group
- Ask on mattermost

Train Run 245367

General | Test | Submitted jobs | Grid statistics | Performance comparison | Wagon resources | Merged output | Clone | Request long train

IO Statistics | Stack trace | Copy all output directories

Reset all filters

PID	Run no.	Output directory	Throughput (/core)	Total	Done	Active	Wait	Error	Merging
Search 86	Search	Search 86 records...	Search 8€						All
3129938053	545345	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596526		612	0	0	0	612	
3129931268	544991	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596491		472	0	0	0	472	
3129935264	545210	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596512		462	0	0	0	462	
3129936604	545289	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596518		452	0	0	0	452	
3129936598	545262	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596517		450	0	0	0	450	
3129933370	545086	/alice/cern.ch/user/a/alihyperloop/jobs/0059/hy_596506		438	0	0	0	438	

Masterjob 3157034637 of alihyperloop, status : **DONE** (refresh | JDL | trace | more details | show input data | Gantt)
Hyperloop analysis, /alice/cern.ch/user/a/alihyperloop/jobs/0072/hy_723006, O2Physics::daily-20240918-0200-1
Subjobs: 430

ERROR_V (32 jobs, 7.442%) : resubmit all
*Exit code: **1** (Snapshot ?)*: 9 jobs
3157034644 : trace | log files | resubmit
3157034689 : trace | **log files** | resubmit
3157034712 : trace | **log files** | resubmit
3157034719 : trace | log files | resubmit
3157034790 : trace | **log files** | resubmit
3157034812 : trace | log files | resubmit
3157034813 : trace | log files | resubmit
3157034820 : trace | log files | resubmit
3157034880 : trace | **log files** | resubmit
*Exit code: **21** (O2 exit code != 0 ?)*: 23 jobs
3157034661 : trace | **log files** | resubmit
3157034666 : trace | log files | resubmit
3157034681 : trace | **log files** | resubmit
3157034685 : trace | **log files** | resubmit
3157034694 : trace | **log files** | resubmit
3157034780 : trace | **log files** | resubmit
3157034783 : trace | **log files** | resubmit
3157034802 : trace | **log files** | resubmit
3157034831 : trace | **log files** | resubmit
3157034832 : trace | log files | resubmit
3157034855 : trace | **log files** | resubmit
3157034866 : trace | **log files** | resubmit
3157034868 : trace | log files | resubmit
3157034869 : trace | log files | resubmit
3157034916 : trace | **log files** | resubmit, 3.352 GB PSS, No Swap PSS
3157034917 : trace | **log files** | resubmit, 3.316 GB PSS, No Swap PSS
3157034932 : trace | **log files** | resubmit, 3.345 GB PSS, No Swap PSS
3157034942 : trace | **log files** | resubmit, 3.475 GB PSS, No Swap PSS
3157034949 : trace | **log files** | resubmit, 208.8 MB PSS, No Swap PSS
3157034957 : trace | **log files** | resubmit, 3.275 GB PSS, No Swap PSS
3157034985 : trace | **log files** | resubmit, 3.478 GB PSS, No Swap PSS
3157034986 : trace | **log files** | resubmit, 3.563 GB PSS, No Swap PSS
3157034998 : trace | **log files** | resubmit

- Introduction & key concepts
- ’My Analyses’, wagons & datasets
- Enabling a wagon to start a test
- Trains – automatic composition, debugging & long trains
- **Key things to know**

Train slots per week

Each dataset has a maximum number of train slots per week for your analysis.

- Currently (subject to change):
  - **14** for small datasets (< 30TB)
  - **6** for medium datasets (30 TB < x < 200 TB)
  - **2** for large datasets (> 200TB)

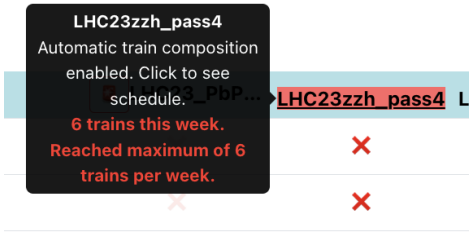

The quota works on a rolling basis. In each train, in the test tab you can see when it stops counting towards the quota:

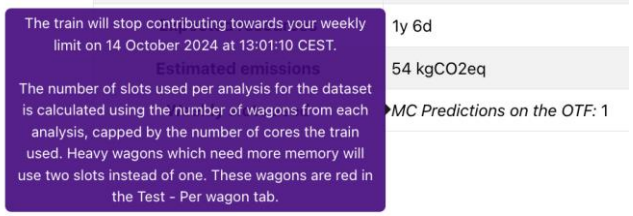

- Trains with many wagons may use > 1 slot. The slots used is never more than the number of cores the train ran with
- Slots are calculated per-analysis. Heavy wagons (which themselves need more memory than available in a single core queue) will use two weekly slots
- Shown in red in the Test – per wagon tab of the train.

Number of input files	1
Input size	2.1 GB
Output size	1.8 MB
PSS Memory	Max: 3.9 GB Avg: 3.7 GB Slope: 879.4 KB/s
Private Memory	Max: 2.8 GB Avg: 2.6 GB Slope: 103.8 KB/s
Timing	CPU: 7m 6s Wall: 10m 42s
Throughput	3.3 MB/s/core
Expected resources	125d 7h
Estimated emissions	18 kgCO2eq
Weekly slots used	*b-jet meausremt using heavy flavour tagging in pp collision at 13.6TeV:* 2

JetFinder_Charged_04_MCP	Max Avg Slope	240.1 MB 217.6 MB 302.4 KB/s	113.3 MB 107.2 MB 76.1 KB/s	8s (2%)	0
JetTaggerHf_Charged_04_MC_SV	Max Avg Slope	307.0 MB 277.5 MB 336.2 KB/s	177.7 MB 170.2 MB 14.2 KB/s	4m 15s (58%)	0
JetTaggerHfQA_Charged_04_MC_SV_weighted				5s (1%)	0
mcCollisionConverter	Max Avg Slope	19.5 MB 19.2 MB 2.7 KB/s	18.4 MB 18.3 MB 397.4 B/s	1s (0%)	0

This wagon needs more memory than available on a single core queue, therefore it counts as 2 train slots from your weekly running quota for LHC24g4.

Extensive comparison tools

Between different wagons

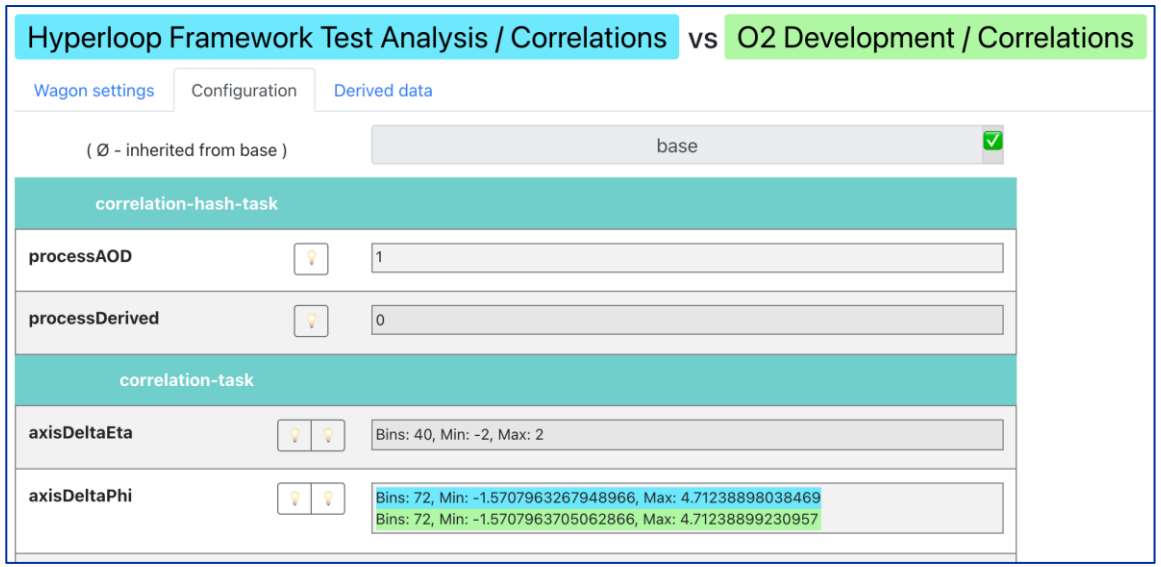

Between different trains

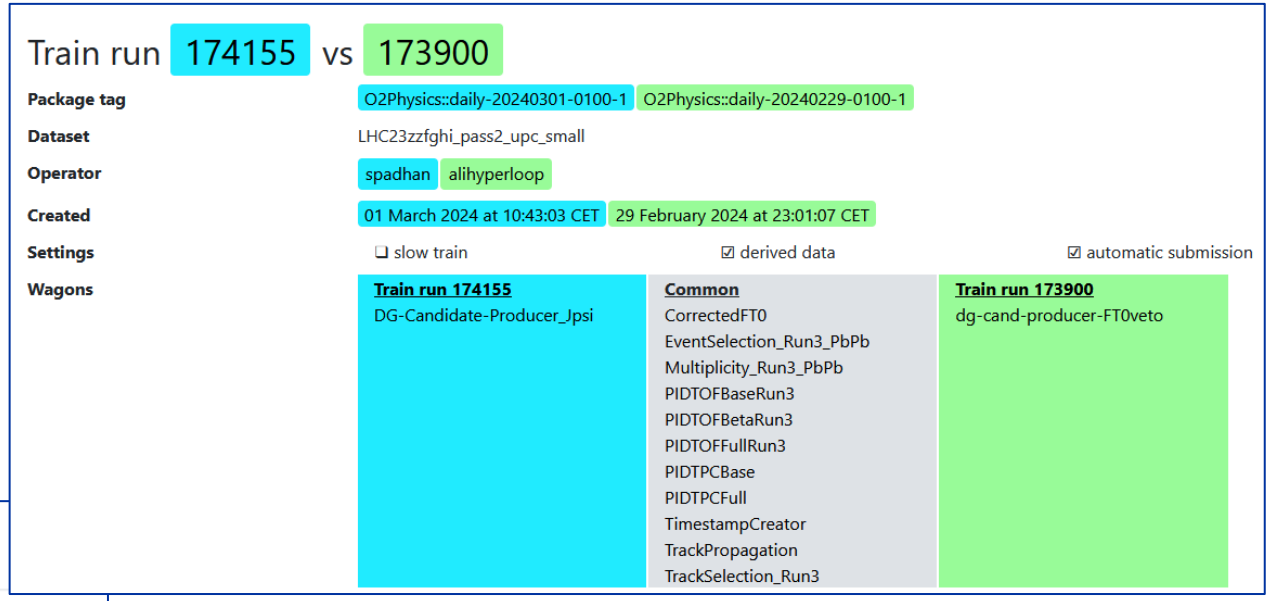

Between different moments in time

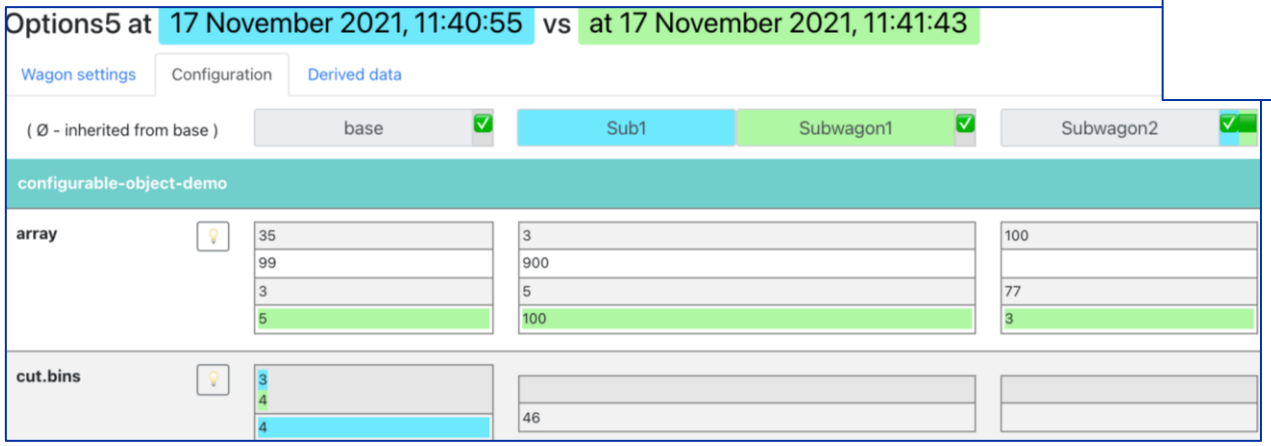

Full bookkeeping for analysis preservation:
Any wagon and train configuration kept forever

Message of the day

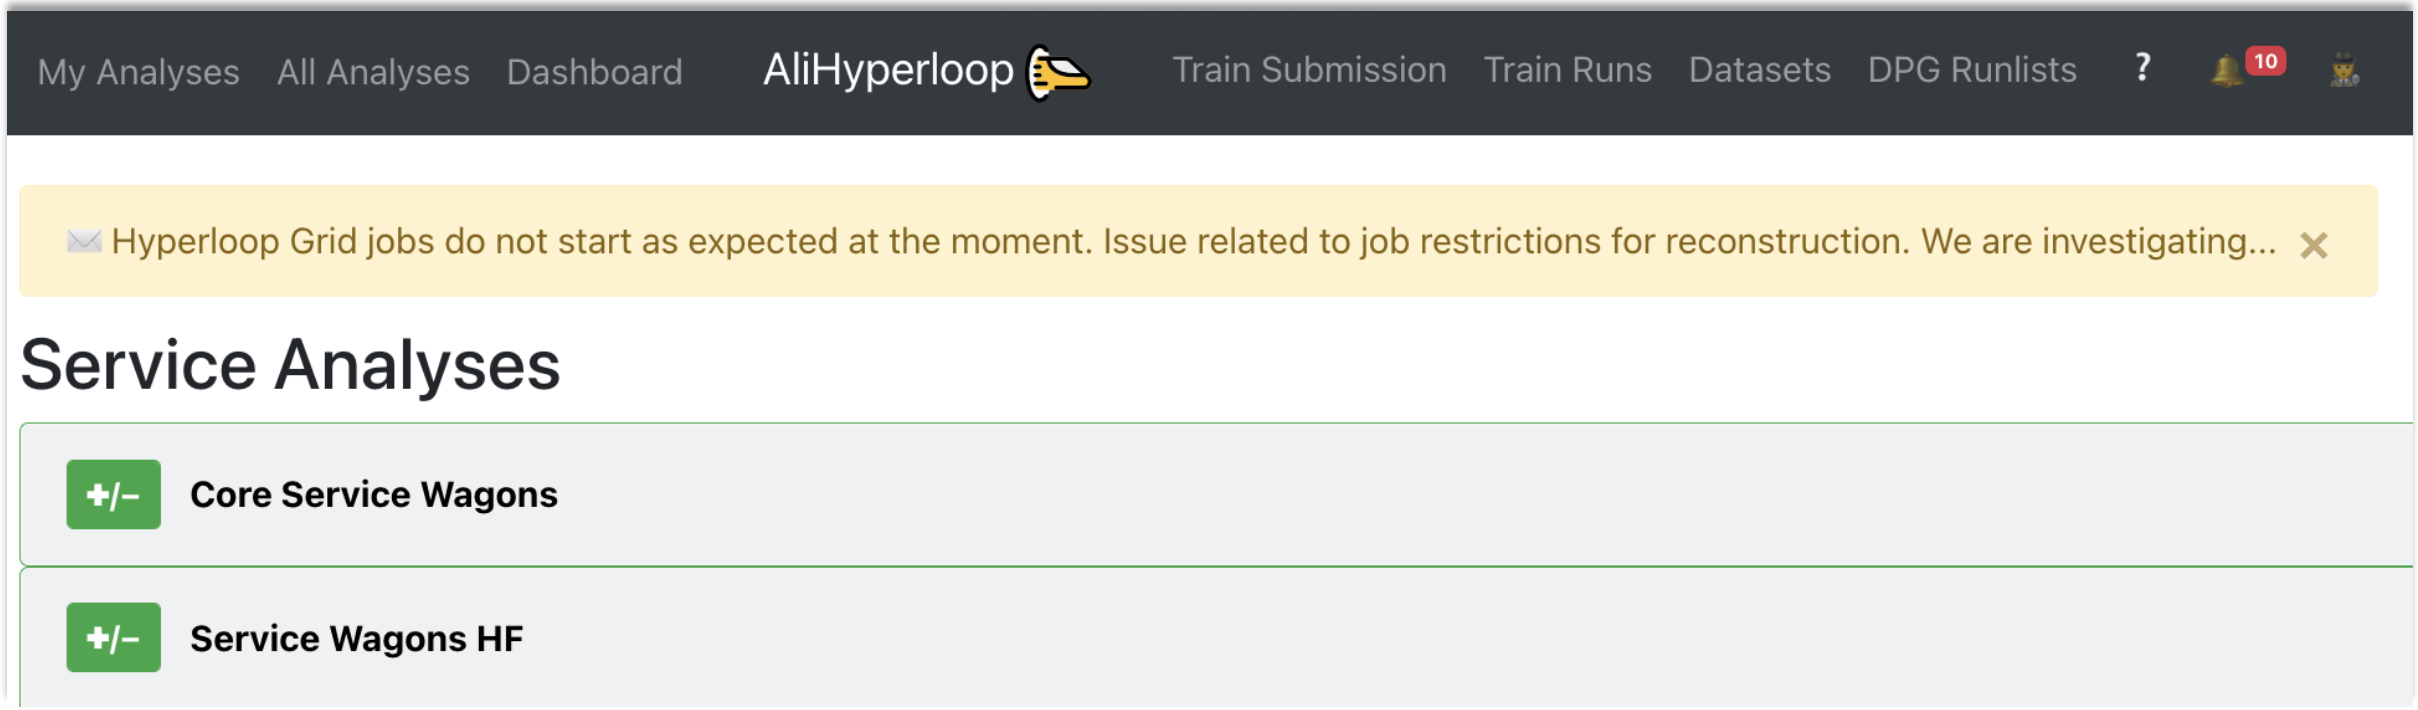

General news / issues concerning the system appear on top of the page and can be dismissed

Look out for new deployments

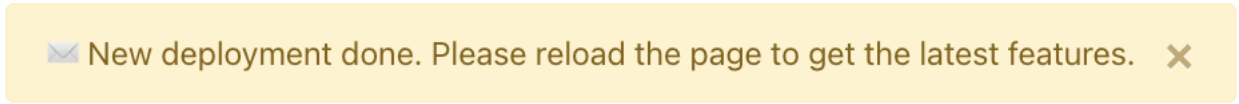

Notifications

- Displayed in the *My Analyses* page and in the *Notifications* page (accessed by clicking in the menu bar).
- Not all notifications require action, there can also be updates and warnings
- You can click the ✕ button to remove a notification.
- Click on ✕ Dismiss to remove all notifications **in the given analysis**
- To remove all the notifications **for all your analyses**, go to the Notifications page, and click the ✕ Dismiss all button.

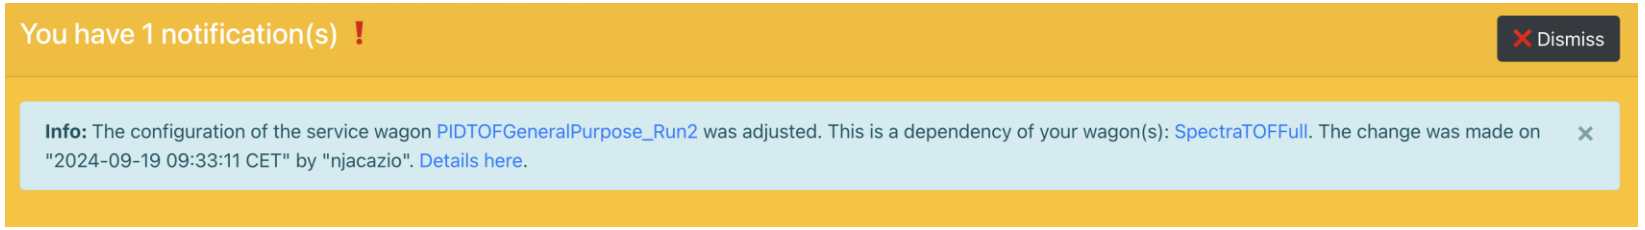

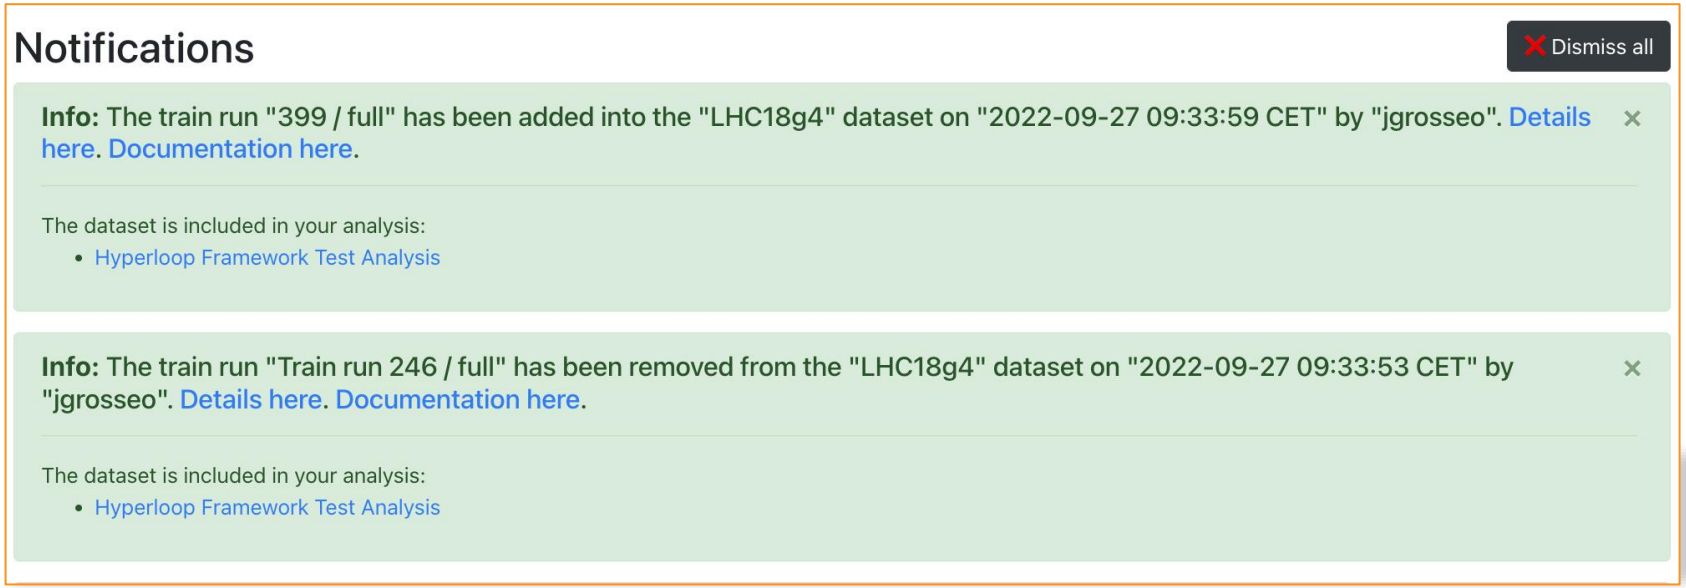

Notifications – Service wagons

- **When a service wagon is updated:**
  - Wagon tests are reset for enabled wagons which depend on it
  - Users who own wagons which 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)

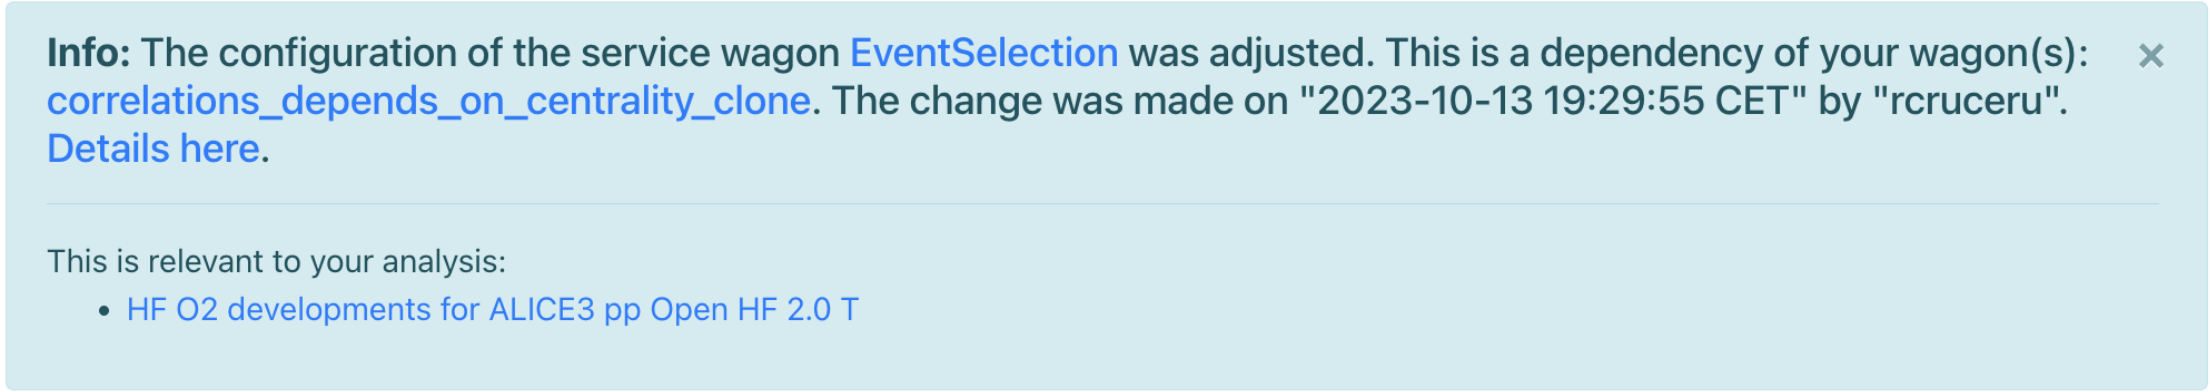

Documentation

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

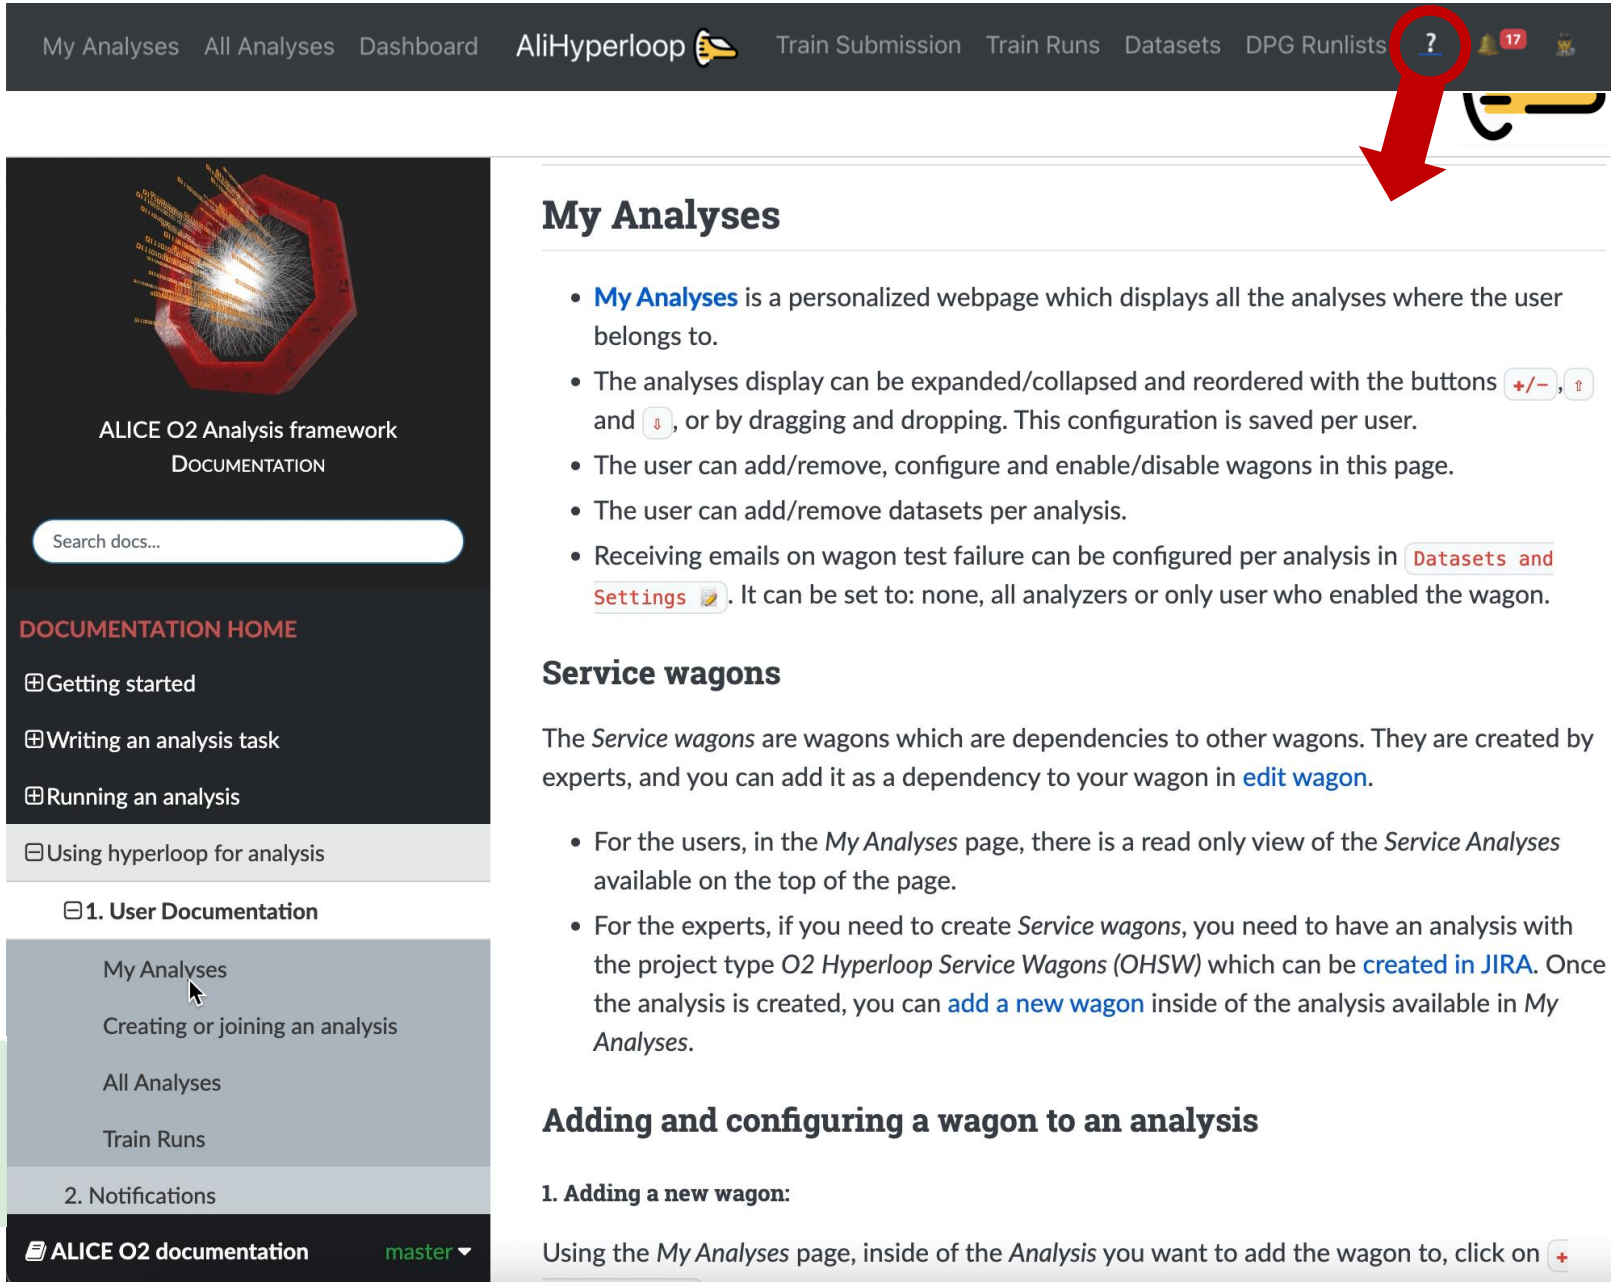

Fair usage policy

- Operators follow guidelines prepared by analysis coordination and approved by physics board (current policy documented here)
  - **Operators cannot grant exceptions, even if justified**
- Aim of guidelines
  - allow efficient analysis by everyone
  - share resources fairly
  - avoid excessive use; identify room for optimization

Dataset size	CPU limit	Trains / week	Automatic schedule
Small datasets			
< 30 TB	1.5 CPU year (550)	14	twice per day
Medium datasets			
< 100 TB	3 CPU years (1095)	6	once a day
< 200 TB	6 CPU years (2190)		twice per week
Large datasets			
< 300 TB	6 CPU years (2190)	2	none (PWG / PB approval)
< 400 TB	6 CPU years (2190)		

100 CPU years = more than a lifetime on your laptop!

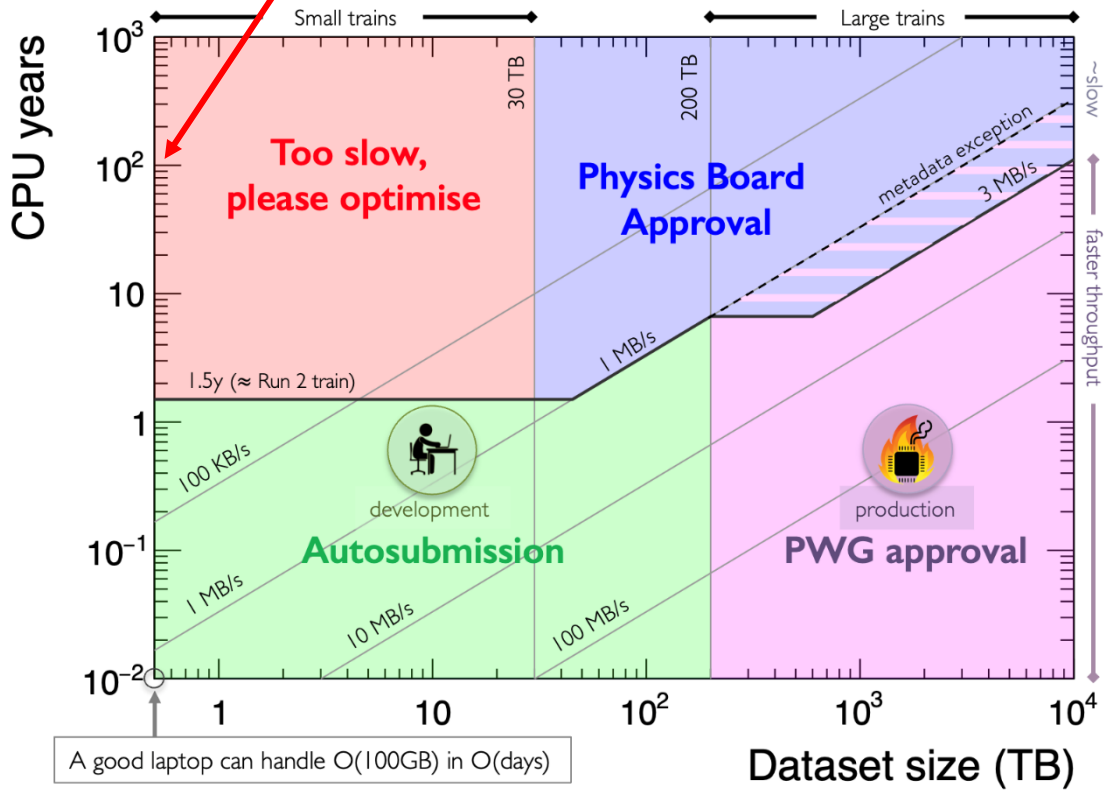

Summary

Summary

Thank you for your attention!

- Use the **My Analyses** page to see your analyses, enable datasets, and choose a package tag
- Create / Edit / Delete / Clone / Compare your wagons
- **Enable** your wagons to start **immediate testing**
- Wait for **automatic train composition** (if enabled for the dataset, wagon test is done and you don't want standard derived data or linked derived data)
- **Or** request train composition in **O2 Hyperloop Operation**
- Follow up / compare trains in My Analyses or Train runs page
- Checkout your **notifications** (remember not all require action)
- Make use of the **?** button to check the documentation
- When in doubt, hop on the bus!

USER	OPERATOR
MyAnalyses	Train Submission
AllAnalyses	Train Runs
Dashboard	Datasets
	Derived Data
Train Support	DPG runlists
24/5 Operation	Trains with issues

Resources

- User and operator documentation
- Public note with technical description
- Current fair usage policy

Backup

Dataset view - Productions

Productions

- Lists every production in the dataset
- Link to production details in MonALISA (MC and DATA)
- Productions contain mergelists
  - Mergelists are lists of runs from the production which are merged together in trains
  - Final mergelists contain multiple mergelists in order to merge runs across productions
- Runs marked as used / not used
- Mergelists marked as activated / deactivated
- Final Mergelists contain multiple mergelists, to merge runs across productions

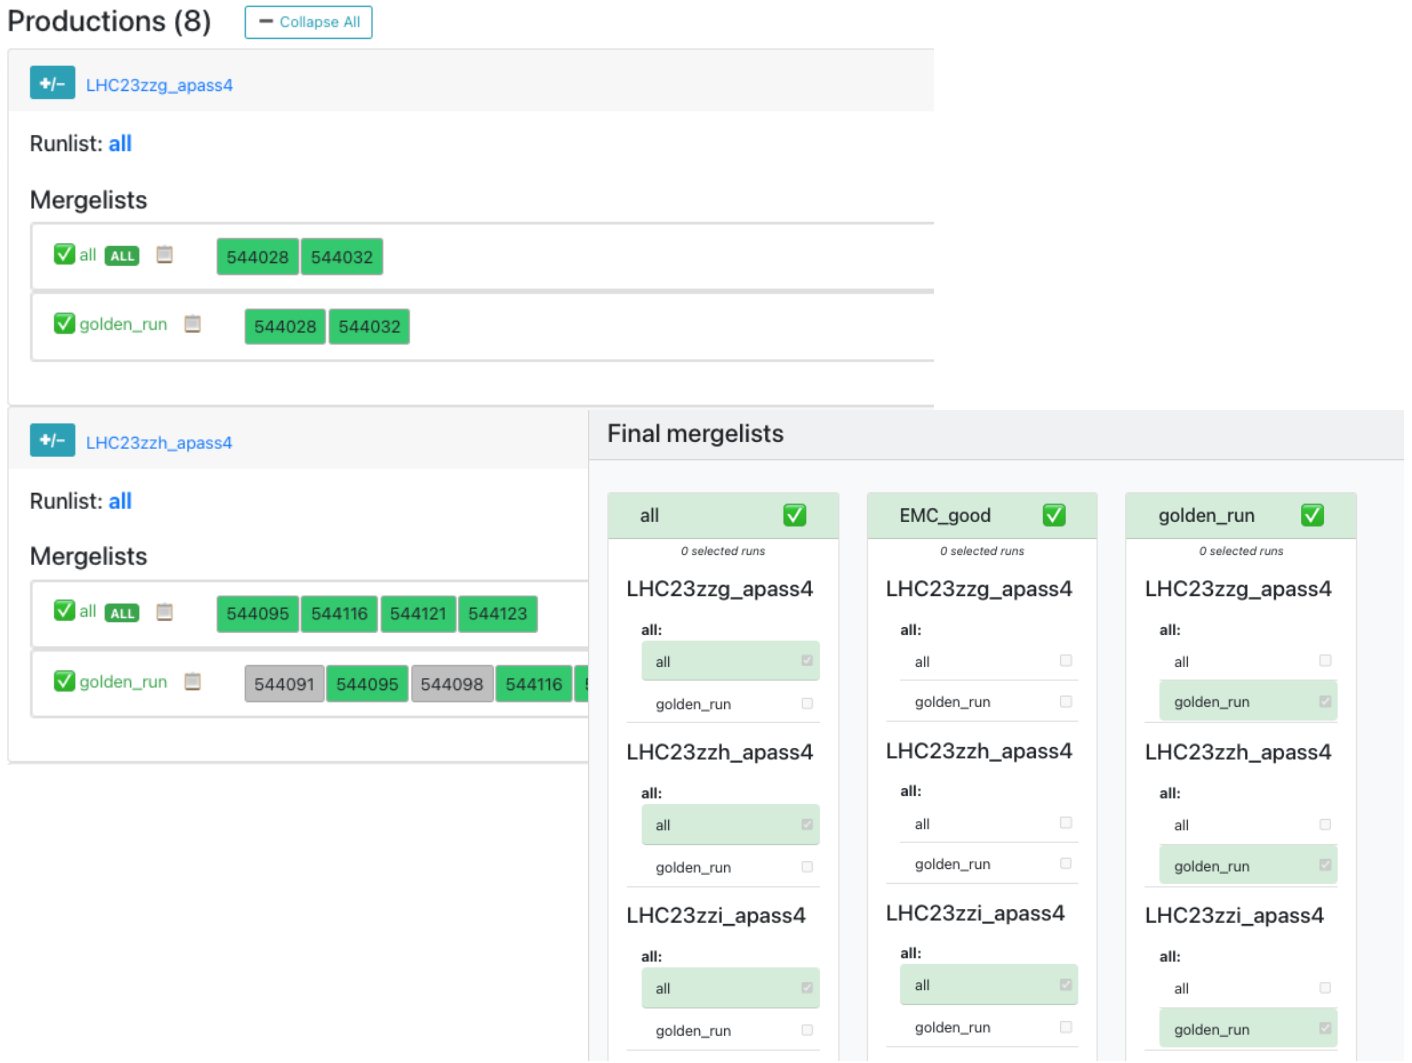

Train Runs page

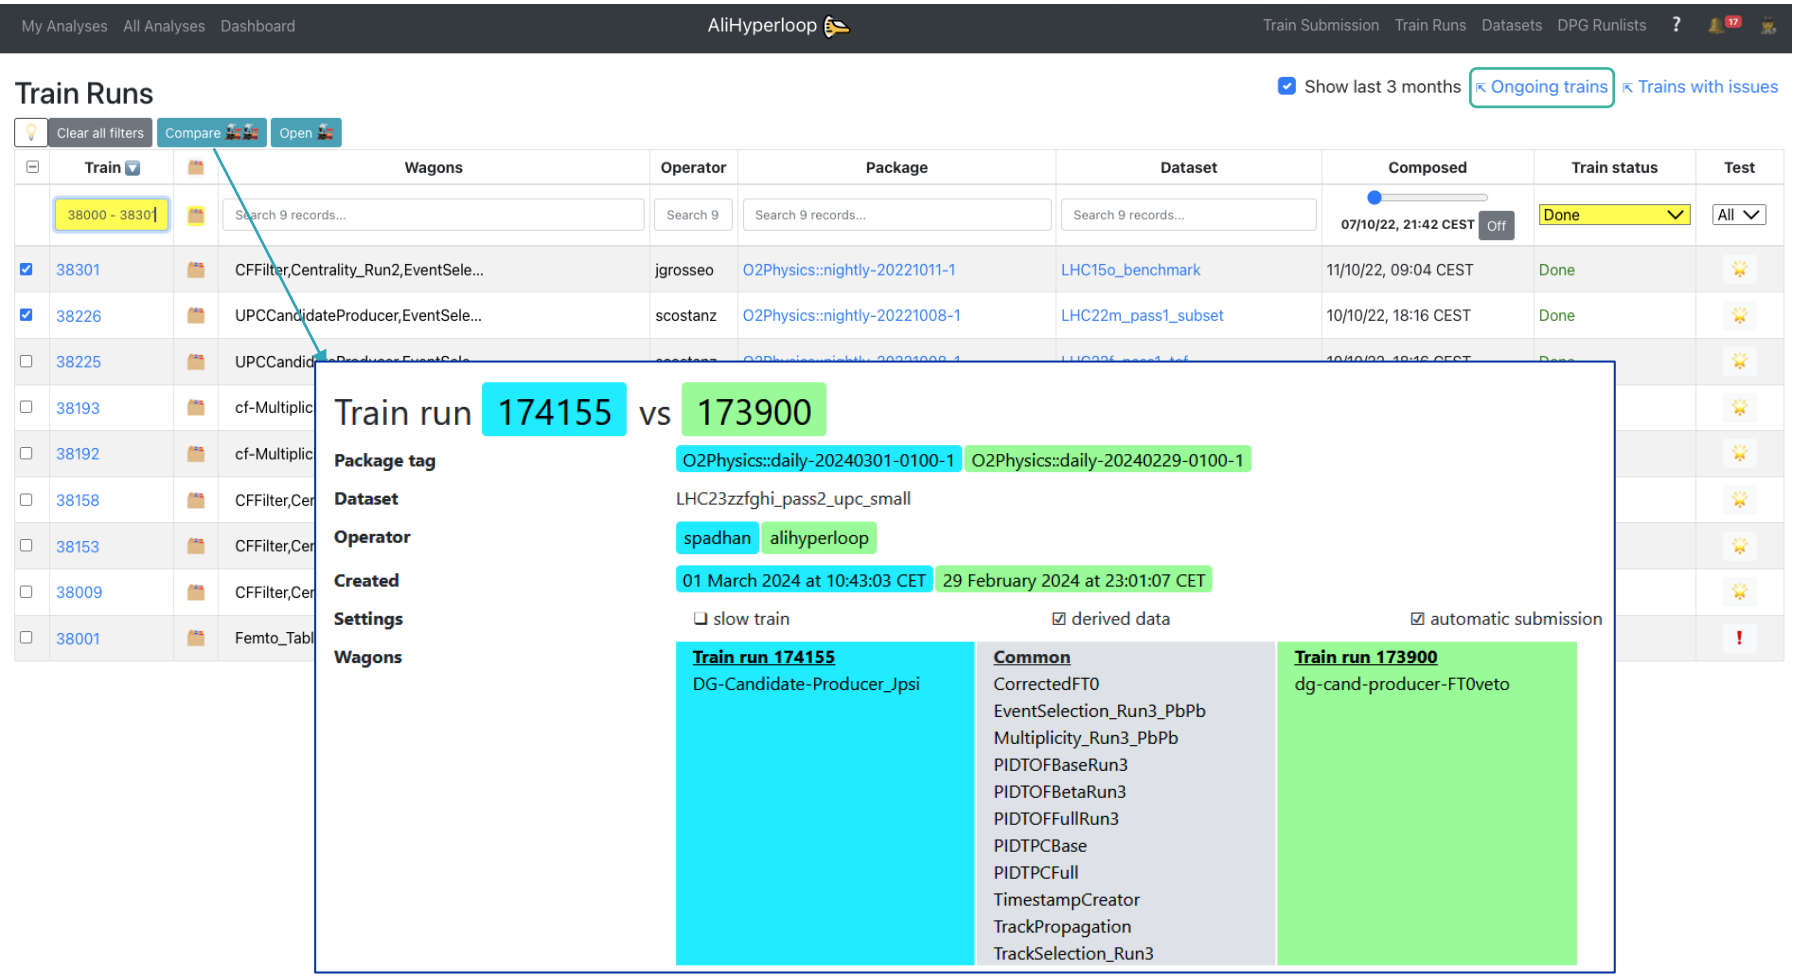

- Displays all the train runs
  (by default, last 3 months of activity)
- Use the searching and filtering features to find your trains
- Click on the train number to access the train run details
- To open several train runs views, select all the trains of interest and click Open
- To compare two trains,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