



New Event Index

ATLAScope 01.02.00+ [24/Apr/2021 at 10:12:16 CEST by atleivind for CERN] [Reset](#)

Search: ATLAS
Execute: g.V().has('lbl', 'canonical')

dataset: DAOD_HIGG2D1

Actions:

Graph Image Plot

Customize the interactions with the graph.

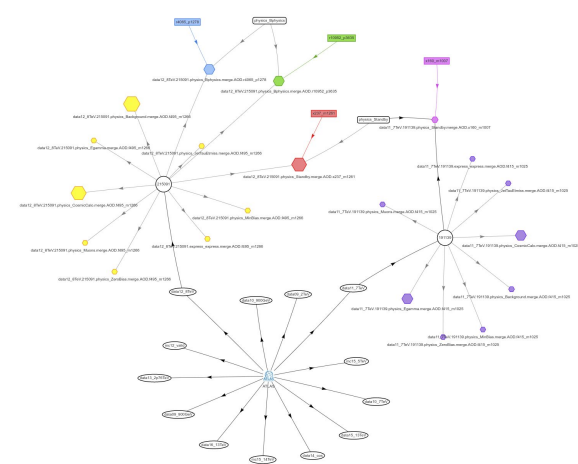
Cluster by group type Cluster by group size Expand all clusters Show all edges

clusterize zoom cluster stabilize get children get parents remove old

filter: Apply select: limit(10)

```

canonical:AOD
10221559 events
version: f832_m1812
runno: 326878
project: data17_13TeV
streamname: physics_Main
prodstep: merge
datatype: AOD
dspsid: 34930176
dstypeid: 8192
smk: 2573
events: ruco:10221559
ruco_at: Thu Nov 16 12:34:18 CET 2017
files: 742
updated_at: Tue Mar 30 09:29:07 CEST 2021
is_open: false
is_deleted: false
status: IMPORTED
has_raw: true
has_trigger: true
prov_seen: 2048
lbl: canonical
phenix: true
fulfill: true
{"id": "16560", "value": "10221559", "label": "data17_13TeV_3:10221559 events", "group": "f832_m1812", "actions": "", ""}
    
```



- Event Index
- Graph Databases
- Atlascope

*Julius Hrivnac, IJCLab
ATLAS, 6/Dec/2021*

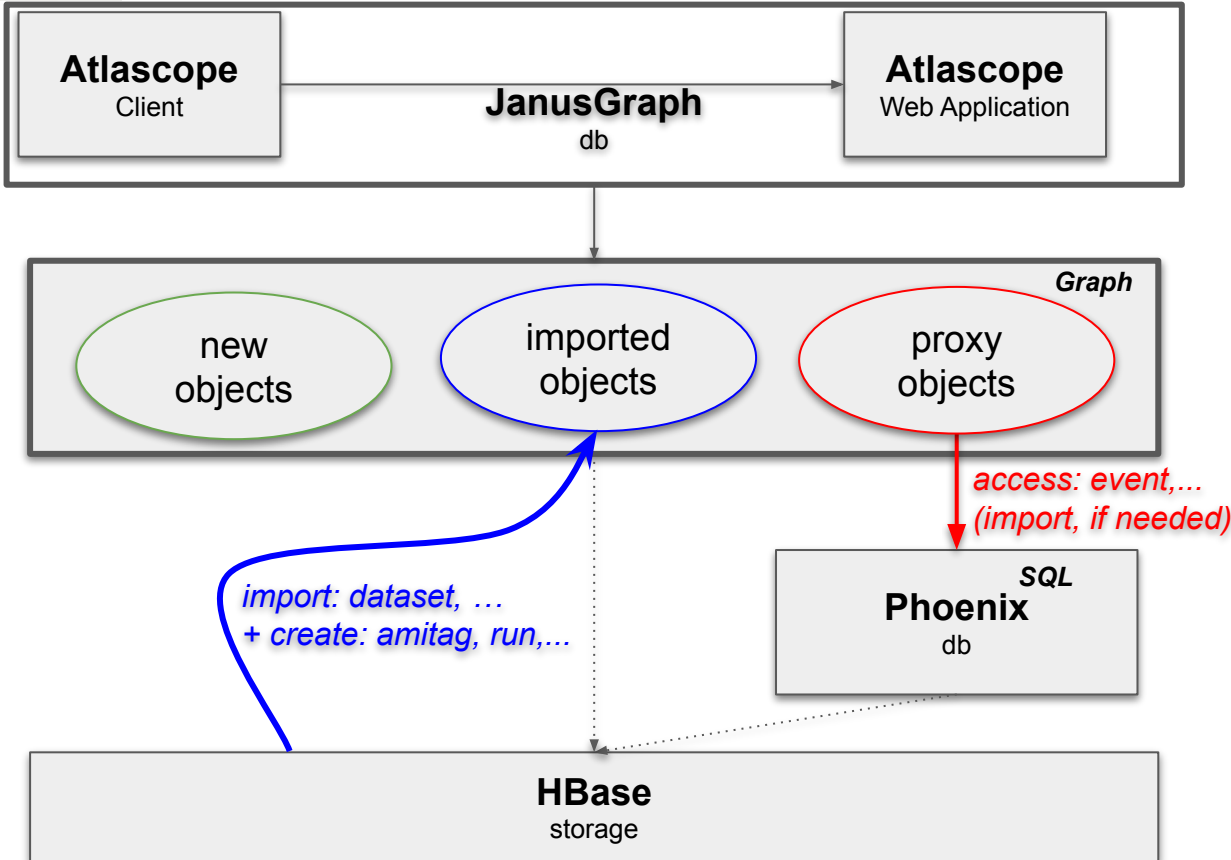


Event Index

- Catalog of all ATLAS events (data and MC) in all their versions (RAW -> ... -> DAOD)
 - With basic characteristics, triggers and references to physical datasets
 - $> 360 * 10^9$ entries
- Main Use cases:
 - Event Lookup: fast mapping of (set of) eventNumber+runNumber -> physical dataset
 - Grigori in Jan'22
 - Searching for events satisfying certain triggers
 - Overlaps between datasets (AOD, DAODs)
 - Overlaps between triggers within a dataset, triggers statistics
 - Finding errors in data processing or simulation (multiple copies of events,...)
- Incarnations:
 - Tag Database (retired): Included also volatile (physical) information, based on Oracle SQL
 - Current EI (in production): Based on Hadoop HBase (NoSQL database)
 - Run 3 EI (under development): Based on Phoenix SQL database + JanusGraph Graph database over HBase NoSQL database
- Event Index + AMI + Rucio form global metadata service allowing to locate all ATLAS data




Architecture



- Phoenix SQL API provides interface compatible with other ATLAS SQL databases (mostly Oracle)
- Graph API provides flexible and intuitive view of data
- **Some Phoenix objects are imported (replicated) into Graph, others are proxied (and can be replicated on request), and others can be created native to Graph**




Schema



- In Phoenix, mirrored to **Graph**:
 - **Dataset**: runno, project, streamname, prodstep, datatype, version, tid, dspid, dstypeid, smk, events_rucio, rucio_at, files, events, events_uniq, events_dup, files_dup, state, updated_at, dups_at, trigger_at, is_open, has_raw, has_trigger, prov_seen, sr_cnt, sr_clid, sr_tech
- In Phoenix, proxied to **Graph**:
 - **Event**: dspid, dstypeid, eventno, seq, tid, sr, mcc, mcw, pv, lb, bcid, lpsk, etime, phid, tpb, tap, tav, lb1, bcid1, hpsk, lph, ph
- In **Graph**:
 - **Project** (data18_13TeV,...)
 - **Stream** (physics_Main,...)
 - **Run** (348894,...)
 - **Amitag** (n0002_r13084_p4397,...)

*More General Remarks
about
Storage & Graphs*

HEP Storage



- Traditional data structures in HEP:
 - tuples (tables)
 - trees
 - nested tuples (trees of tuples)
 - relational (SQL-like)
- Schema-based or schema-less
- **But many of HEP data are graph-like & schema-less**
 - **Entities with relations**
- Not handled by standard tree-ntuple storage
 - Relations should be added and interpreted outside storage
- Not well covered by relational (SQL) databases
 - We need to add new relations, not covered by schema
- Difficult to manage by Object Oriented (OO) databases or serialisation
 - Problem to distinguish essential relations from volatile ones

Graph Databases

- Storing Graphs in a database
- **Graph = (Vertexes, Edges), $G = (V, E)$**
- Vertexes and Edges have properties



Graph databases have existed for a long time

- Matured only recently thanks to Big Data & AI (Graph NN)
- Very good implementations & (de-facto) standards available
- Rapid evolution

Moving essential structure from code to data

- Together with migration from imperative to declarative semantics
- Things don't **happen**, but **exist**
- Structured data with relations facilitates **Declarative Analyses**

Data elements appear in a **Context**

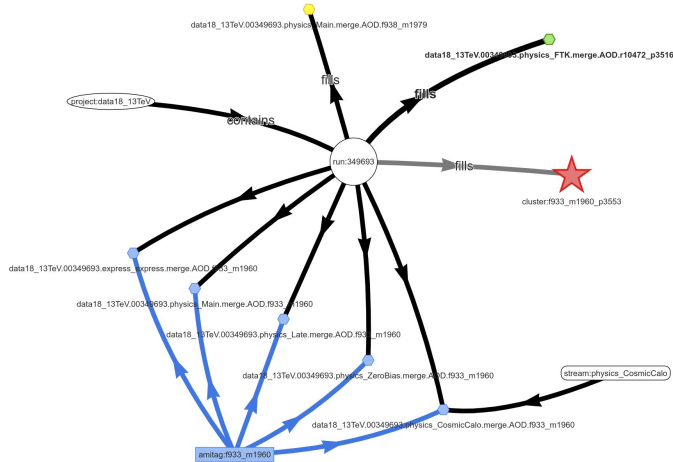

- Which simplifies understanding, analyses and processing

The difference between SQL and Graph database is similar as between Fortran and C++/Java


- On one side, a rigid system, which can be very optimized
- On the other side, a flexible dynamical system, which allows expressing of complex structures

Graph database is a synthesis of OO and SQL databases

- Expressing web of objects without fragility of OO world
- Capturing only essential relations, not an object dump



Graph DB: Languages



➤ Direct manipulation of Vertices and Edges

- Always available from all languages
- Doesn't use full graph expression power

➤ Cypher (and GQL)

- Pure declarative
- Inspired by SQL and OQL
 - But applied to schema-less database
- Available to all languages via JDBC-like API
 - Semantic mismatch, passed as String
 - There is a wall between coder and database, with a thin tunnel, only Strings can pass
- Coming from Neo4J
 - Accepted as a standard
 - Neo4J can be also used with Gremlin

```
MATCH (a:run)-[:has]->(b:dataset)
WHERE a.rnumber = 98765
RETURN b.name
```

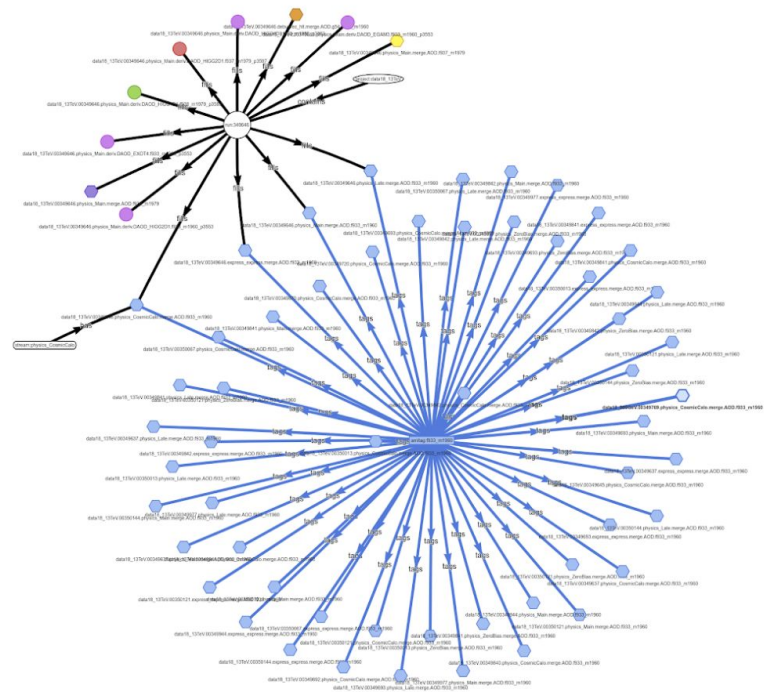

➤ Gremlin

- Functional syntax
- Originated from *Groovy*, but available to all languages supporting functional programming
- Integrated in the language

```
g.V().has('run', 'rnumber', 98765)
.out('has')
.values('name')
```


Graph Databases for HEP

- A lot of ongoing HEP effort to make execution more structured and parallel
 - Parallel programming
 - Functional programming
- Less effort (so far) to structure the data
 - More structured data => simpler and faster access
- Graphical Database advantages
 - More transparent code
 - Stable data structure is handled in the storage layer
 - Suitable for **Functional Style** and **Parallelism**
 - Suitable for **Deep Learning**
 - Suitable for **Declarative Analyses**
 - Can help with **Analysis Preservation**
 - Language & Framework neutral
- How to proceed
 - Store (all) data in a real Graph database
 - Build a Graph layer on top of the existing storage
 - Close to DB layer
 - In the application layer





Part of Event Index is stored in a Graph Database



Schema


Back to Event Index



- In **Phoenix**, mirrored to **Graph**:
 - **Dataset**: runno, project, streamname, prodstep, datatype, version, tid, dspid, dstypeid, smk, events_rucio, rucio_at, files, events, events_uniq, events_dup, files_dup, state, updated_at, dups_at, trigger_at, is_open, has_raw, has_trigger, prov_seen, sr_cnt, sr_clid, sr_tech
- In **Phoenix**, proxied to **Graph**:
 - **Event**: dspid, dstypeid, eventno, seq, tid, sr, mcc, mcw, pv, lb, bcid, lpsk, etime, phid, tpb, tap, tav, lb1, bcid1, hpsk, lph, ph
- In **Graph**:
 - **Project** (data18_13TeV,...)
 - **Stream** (physics_Main,...)
 - **Run** (348894,...)
 - **Amitag** (n0002_r13084_p4397,...)



Gremlin Client

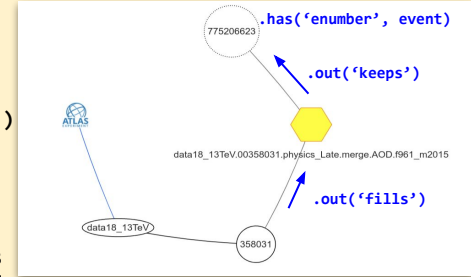


- Functional syntax
 - On top of other languages
- **Functional & navigational semantics**
- **Very intuitive, no special syntax needed** (using existing functional syntax), easy integration.
- Database just accessed as objects with structure and relations.
 - Nested collections with links.
- Can use functional API (streams) and Lambda.
- **No semantic mismatch.**
 - Using one language.
- Came from **Groovy**
 - (Almost) identical for other supported languages (Python, Scala, Go,...).
- Both **search** and **traversal** steps.
- Search steps can be boosted by indexes.
- Functions can be loaded on server for faster execution.

```
# add a vertex 'experiment' with the name 'ATLAS'
g.addV('experiment').property('ename', 'ATLAS')
# add edges 'owns' from all vertices 'project' to vertex 'experiment' 'ATLAS'
g.V().hasLabel('project')
  .addE('owns')
  .from(g.V()
  .hasLabel('experiment')
  .has('ename', 'ATLAS'))
# show datasets with more events or number of events in an interval
g.V().has("run", "number", 358031)
  .out()
  .has("nevents", gt(7180136))
  .values("name", "nevents")
g.V().has("run", "number", 358031)
  .out()
  .has("nevents", inside(7180136, 90026772))
  .values("name", "nevents")

# Event-Lookup function (server side UDF)
def el(run, event, g) {
  e = g.V().hasLabel('run')           # all runs
  .has('rnumber', run)                # selected run
  .out('fills')                       # all datasets filling that run
  .out('keeps')                       # all events kept in that dataset
  .has('enumber', event)              # selected event
  .values('guid')                     # its guid
}

# CLI command
curl -XPOST -d '{"gremlin": "el(run, event)"}' http://ei-gremlin-server.cern.ch:8182
# or using standard gremlin client
gremlin << EOF
:remote connect tinkpop.server $janusgraph_home/conf/remote.yaml
el(run, event)
EOF
```





Gremlin Client Examples

// Event Lookup using Graphs

```
g.V().has('lbl', 'dataset'). // all datasets
  has('runno', 801122). // datasets with runno == 801122
  out('has'). // all 'has' children (= events)
  has('eventno', 15379) // all events with eventno == 15379
```

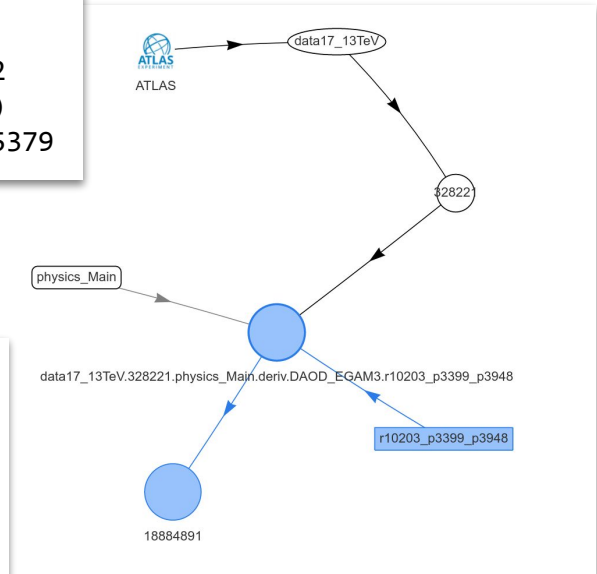

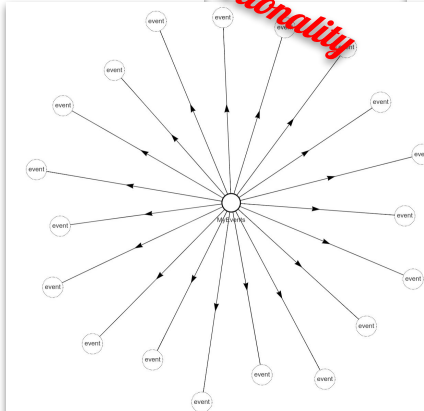
// Create new collection of events

```
eventsCollection = g.addV('ecollection')
  .property('name', 'MyEvents');
```

// Find all events satisfying certain conditions // and connect them to the event collection

```
g.V().has('lbl', 'event')
  .has(...some selection...)
  .collect {
    eventsCollection.addEdge('contains', it)
  };
```

Whiteboard functionality





(C) Python Client


```
#pip install gremlinpython
```

```
from gremlin_python import statics
from gremlin_python.process.anonymous_traversal import traversal
from gremlin_python.process.graph_traversal import __
from gremlin_python.process.strategies import *
from gremlin_python.driver.driver_remote_connection import DriverRemoteConnection
from gremlin_python.process.traversal import T
from gremlin_python.process.traversal import Order
from gremlin_python.process.traversal import Cardinality
from gremlin_python.process.traversal import Column
from gremlin_python.process.traversal import Direction
from gremlin_python.process.traversal import Operator
from gremlin_python.process.traversal import P
from gremlin_python.process.traversal import Pop
from gremlin_python.process.traversal import Scope
from gremlin_python.process.traversal import Barrier
from gremlin_python.process.traversal import Bindings
from gremlin_python.process.traversal import WithOptions
```

```
statics.load_statics(globals())
```

```
g = traversal().withRemote(DriverRemoteConnection('ws://aiatlas073.cern.ch:8182/gremlin','g'))
```

```
x = g.V().has('lbl', 'dataset').has(...).valueMap().next()
```




*Easy integration in
Atlas Framework*

Clients exist in most languages



Web Client



Atlascope 01.02.00+ [24/Apr/2021 at 10:12:16 CEST by atlevind for CERN] Reset

http://localhost:8182 ✓add

Search: ATLAS

Execute: g.V().has('lbl', 'canonical')

dataset: DAOD_HIGG2D1

Actions:

Graph Image Plot

Customize the interactions with the graph.

Cluster by group type Cluster by group size Expand all clusters Show all edges

clusterize zoom cluster stabilize get children get parents remove old

filter: select: limit(10) Apply

canonical:AOD 10221559 events

Physics_Main

data17_13TeV.326870.physics_Main.merge.AOD.832_m1812

data17_13TeV.327342.physics_Main.deriv.DAOD_HIGG2D1.r10203_p3399_p3917

data17_13TeV.326870.physics_Main.merge.AOD.832_m1812

data17_13TeV.327342.physics_Main.deriv.DAOD_HIGG2D1.r10203_p3399_p3917

326870

data17_13TeV

327342

data16_13TeV

data15_13TeV

ATLAS

data16_13TeV

data18_13TeV

```
canonical:AOD
10221559 events


version:f832_m1812
runno:326870
project:data17_13TeV
streamname:physics_Main
prodstep:merge
datatype:AOD
dspid:34930176
dstypeid:8192
smk:2573
events:10221559
rucio_at:Thu Nov 16 12:34:18 CET 2017
files:742
events:10221559
updated_at:Tue Mar 30 09:29:07 CEST 2021
is_open:false
is_deleted:false
status:IMPORTED
has_raw:true
has_trigger:true
prov_seen:2048
lbl:canonical
phoenix:true
fullfill:true

{"id":16560,"value":10221559,"label":"data17_13TeV.3:10221559 events","group":"f832_m1812","actions":",",}
```

- Built on top of Gremlin Client
- Generic Graph Browser
 - Customised for ATLAS Event Index
 - Another customisation exists for LSST



Web Client - Initial Top Panels



Database to use
(*'Proxy' tunnels requests to the database through the Web Service to by-pass firewalls*)

Initial Gremlin request

Options for interactive Graph manipulation



Web Client - Graph with Table View

A Table operation will show all visible elements of the same type in a tabular form

Table data can be plotted in various ways




Table View Customisation (visible columns, searches,...)

| lbl | sizeIn | sizeOut | intersection |
|----------|---------|---------|--------------|
| overlaps | 221269 | 709863 | 100939 |
| overlaps | 154827 | 592284 | 70168 |
| overlaps | 128979 | 106537 | 98308 |
| overlaps | 1282413 | 1427638 | 1122960 |
| overlaps | 821549 | 123581 | 83027 |
| overlaps | 502766 | 3101439 | 314622 |
| overlaps | 493169 | 410177 | 406733 |
| overlaps | 583587 | 2410109 | 130074 |
| overlaps | 42071 | 3403375 | 2695 |


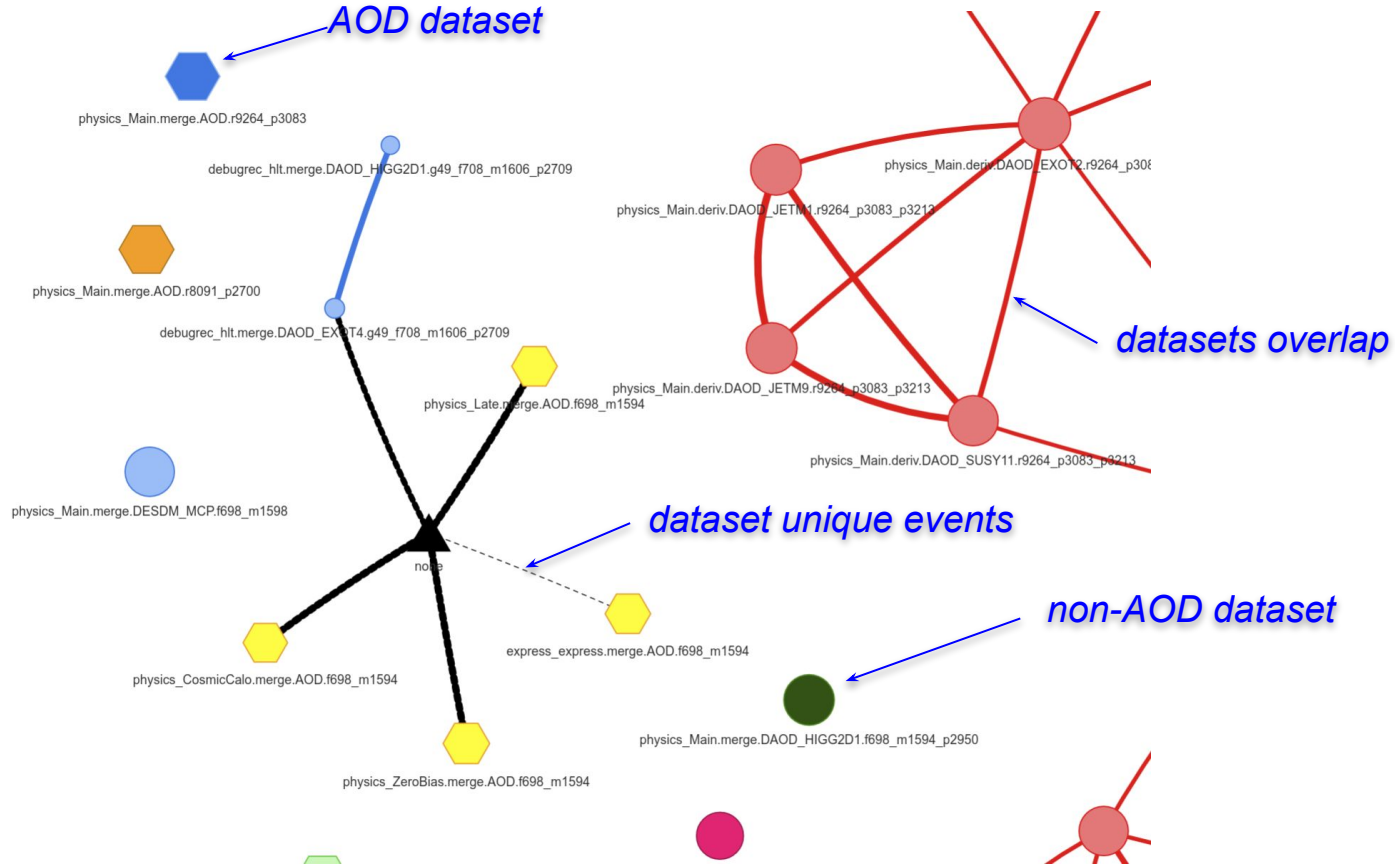
Executed actions (so they can be re-used from the command line)

```
Showing 1 to 9 of 9 rows  
Sending Gremlin request to //atlas-event-index.cern.ch/AtlasScopeProxy.jsp?server=http://atlas-event-index.cern.ch:8183: g.E("3yd-36w-6uj9-3a0")id(next()).toString().replaceFirst("","Edges.jsp?id=")  
Sending Gremlin request to //atlas-event-index.cern.ch/AtlasScopeProxy.jsp?server=http://atlas-event-index.cern.ch:8183: g.E("3yd-36w-6uj9-3a0")id(next()).toString().replaceFirst("","Edges.jsp?id=")
```

A Click on an element (Vertex or Edge) will offer a set of available operations (internal or external applications)

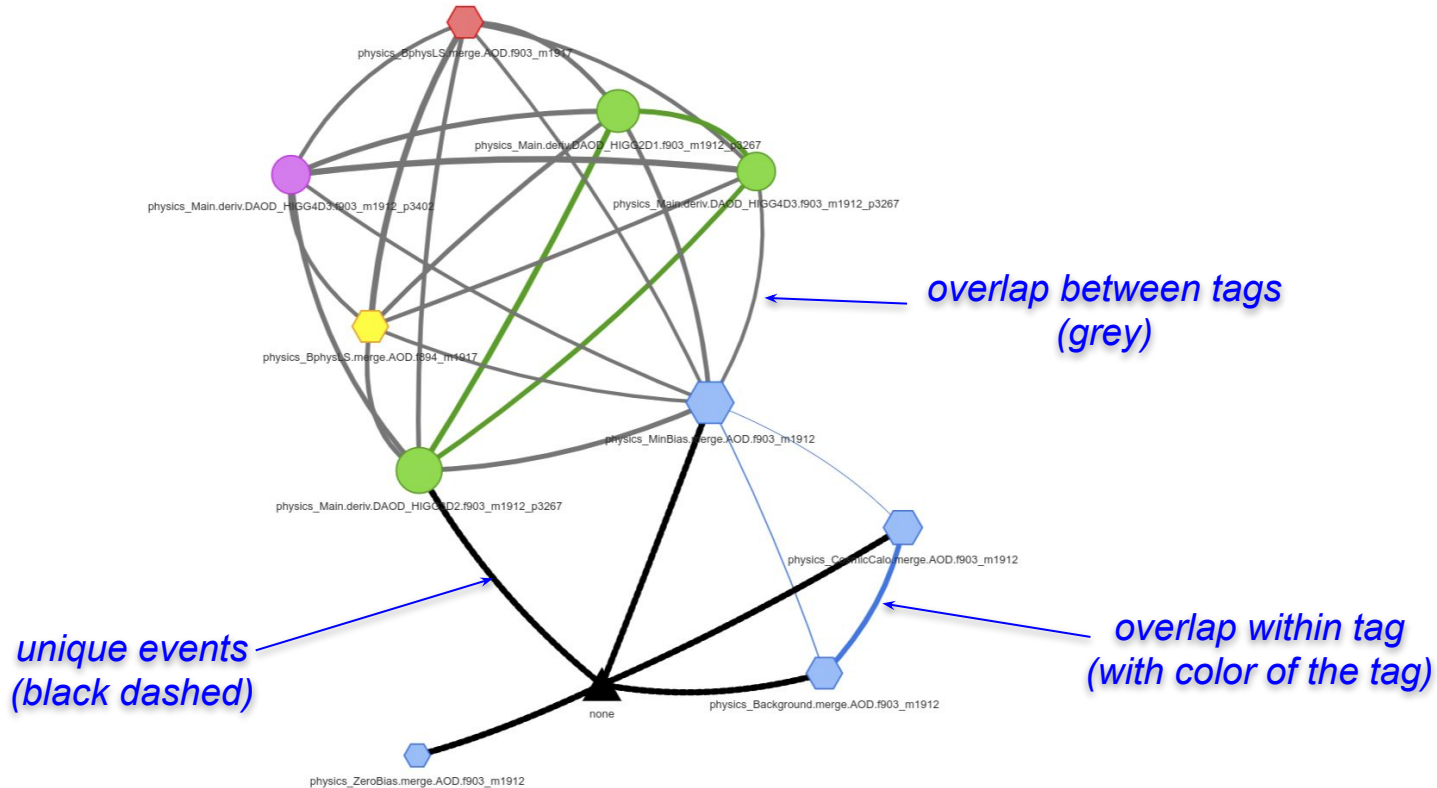



Datasets Graphical View





Overlaps between Tags, Unique Event





Dataset Details

[click here](#)

f903_m1912_p3336

data17_13TeV.00341419.physics_Background.merge.AOD.f903_m1912:
n = 282952

physics_CosmicCalo.merge.AOD.f903_m1912

physics_Background.merge.AOD.f903_m1912


f903_m1912_p3402

physics_MinBias.merge.AOD.f903_m1912

physics_ZeroBias.merge.AOD.f903_m1912

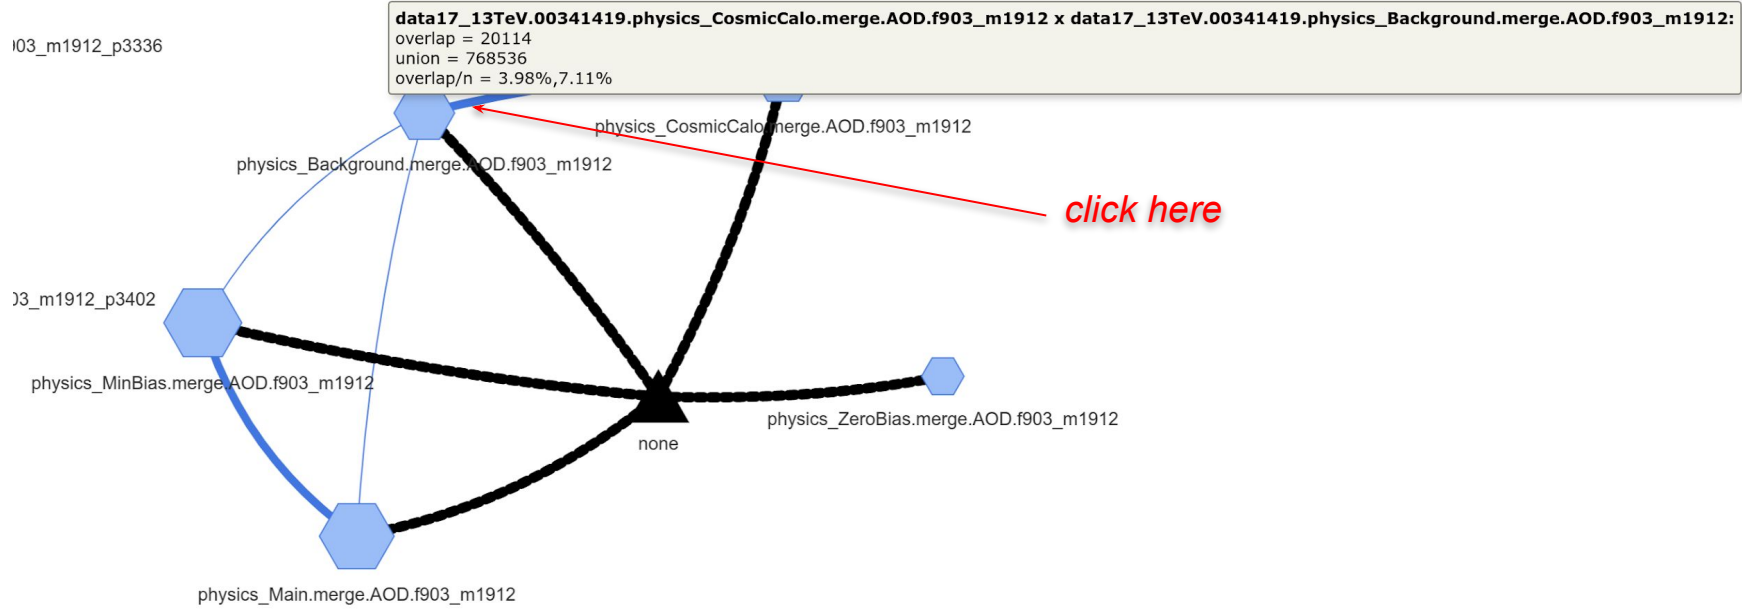

none

physics_Main.merge.AOD.f903_m1912






Overlap Details

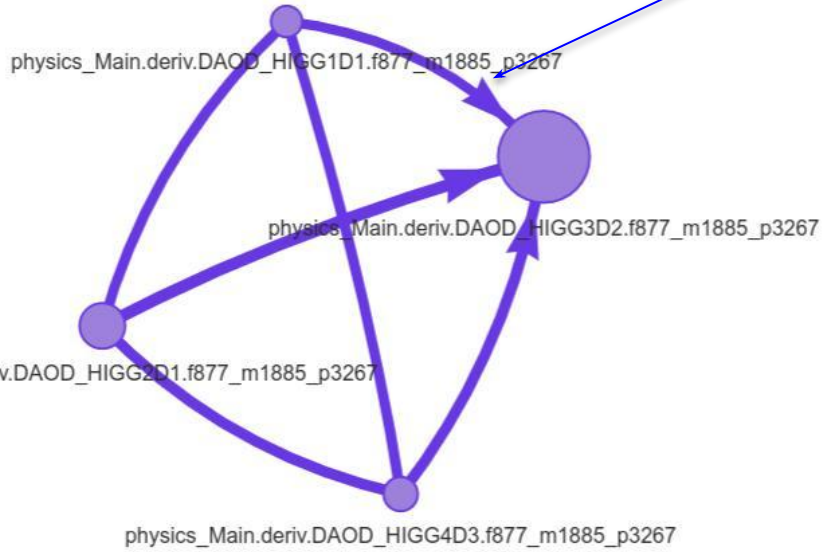





Subsets



subset




*identity
(mutual subset)*





Unique Events Details



[click here](#)

03_m1912_p3336

13_m1912_p3402

physics_MinBias.merge.AOD.f903_m1912

physics_Main.merge.AOD.f903_m1912

physics_Background.merge.AOD.f903_m1912

physics_CosmicCalo.merge.AOD.f903_m1912

physics_ZeroBias.merge.AOD.f903_m1912

data17_13TeV.00341419.physics_Background.merge.AOD.f903_m1912:
unique = 262834
non-unique = 20118
uniqueness = 92.89%

none



All Overlaps of AODs

show all overlaps within a tag

show all overlaps between tags

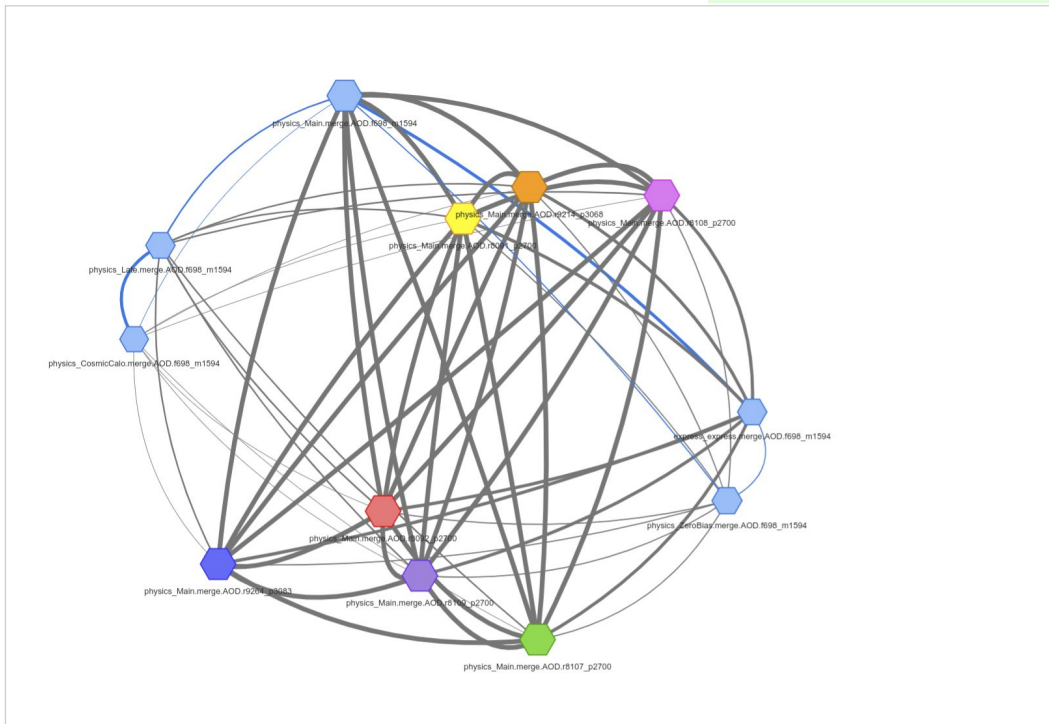

E116.1/00299184: live Help

Cluster by AMI Tag Cluster by group size Expand all clusters

overlap thresholds:


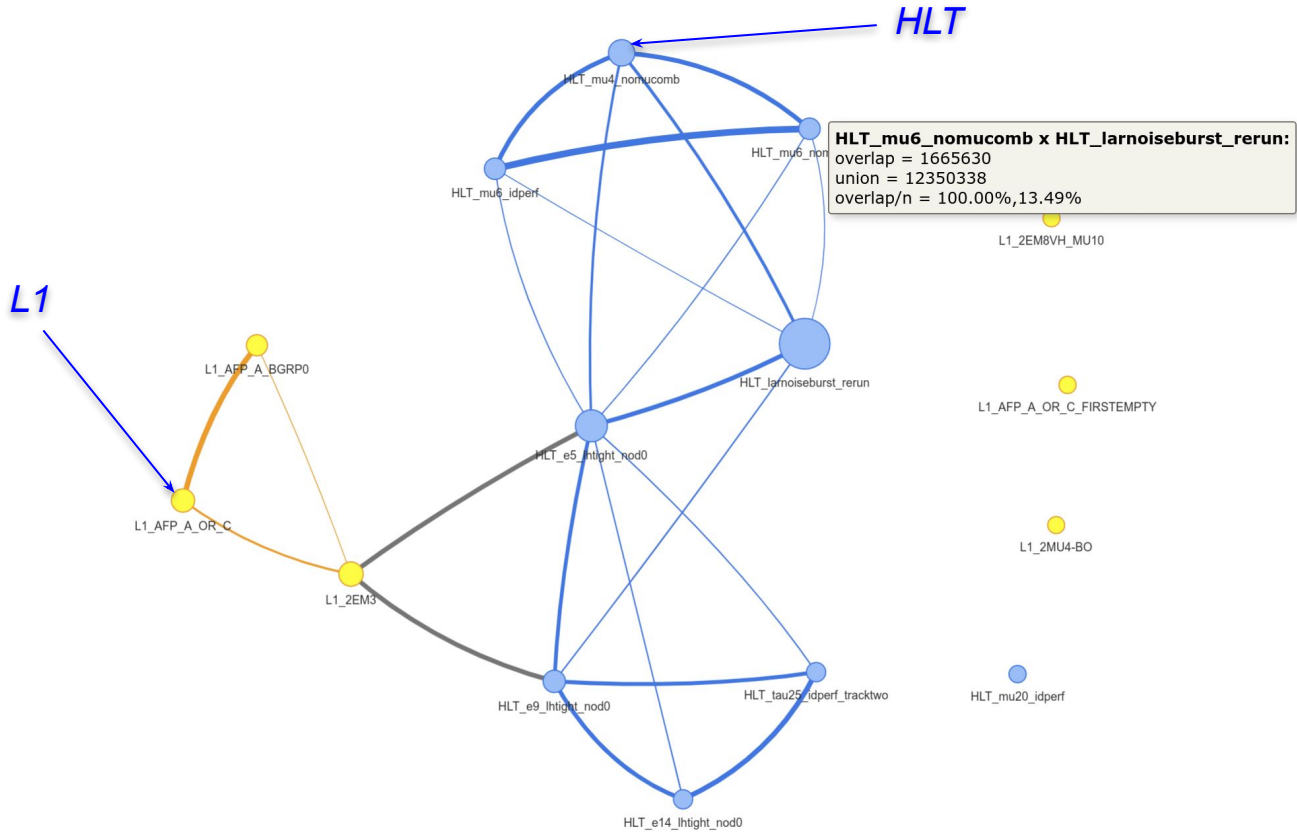
tag level: 1 target: filter: .AOD. Recreate

Context-sensitive menu will be here.





Trigger Overlaps





Clusters

E116.1/00298690:


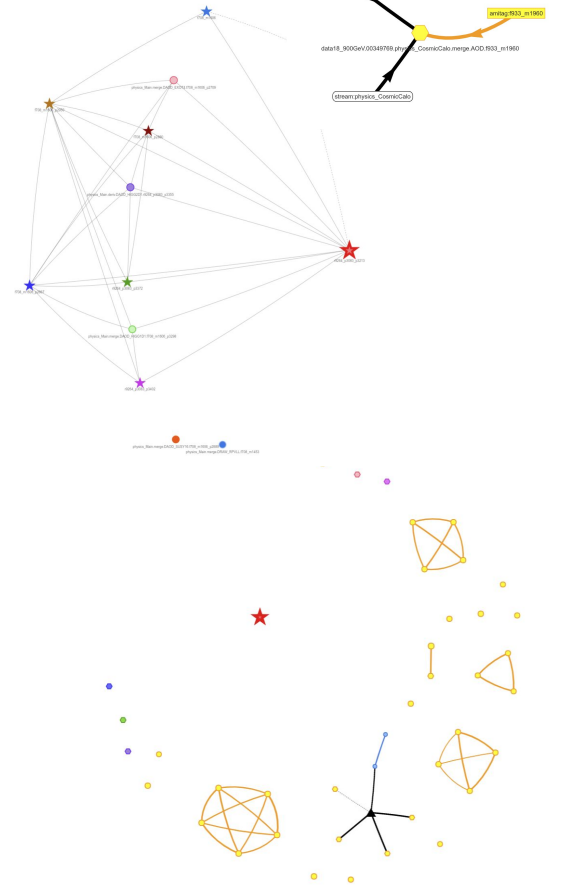


live [Help](#)

Cluster by AMI Tag Cluster by group size Expand all clusters

overlap thresholds:

tag level: target: filter: [Recreate](#)


Context-sensi





Web Client - Correlogram

Various types of data views



Atlascope 02.00.00+ [02/Oct/2021 at 23:13:57 CEST by atlevind for CERN] Reset

GERN-Proxy Add canonical:DAOD_SUSY12

Search Show - Table

Execute -

Customize the interactions with the graph.

Cluster by group type
 Cluster by group size
 Expand all clusters
 Show all edges
 hierarchical (Lup/tr |size/hierarchy) live

clusterize zoom cluster stabilize get children get parents remove old

filter: Apply select: limit(100)

Results Table Image Plot SkyView

A = data15_13TeV_28082 physics_Main merge DAOD_EXOT12_r7562_p2521_p2950_10332160
B = data15_13TeV_28082 physics_Main merge DAOD_EXOT3_r7562_p2521_p2950_10552452

| | |
|--------------|---------------|
| A | 128979 (94%) |
| B | 106537 (78%) |
| A^B | 98308 (72%) |
| A-A^B | 30671 (22%) |
| B-A^B | 30671 (6%) |
| A^vB | 137208 (100%) |

Sending Gremlin request to //atlas-event-index.cern.ch/Atlascope/Proxy.jsp?server=http://atlas-event-index.cern.ch:8183: g.E["3yrd-36w-6u@-3a0"].id().next().toString().replaceFirst(".*", "Edge.jsp?id=")

Sending Gremlin request to //atlas-event-index.cern.ch/Atlascope/Proxy.jsp?server=http://atlas-event-index.cern.ch:8183: g.E["3yrd-36w-6u@-3a0"].id().next().toString().replaceFirst(".*", "Edges.jsp?id=")

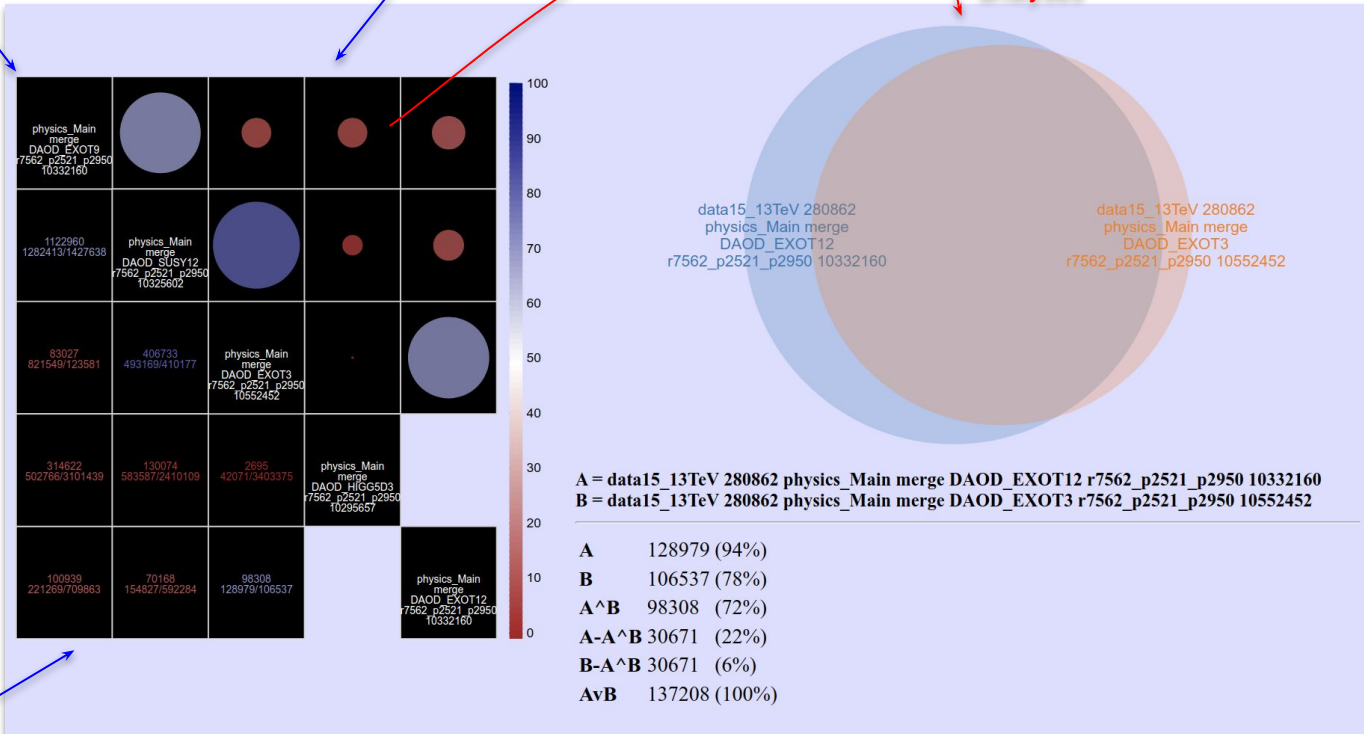
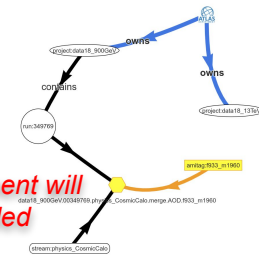


Web Client - Overlaps between Datasets

Dataset names

Overlaps as circles

Hovering over an overlap element will give a Venn diagram with detailed analyses

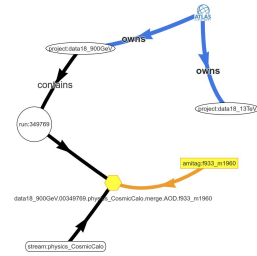


Overlaps as numbers



Web Client - Plots

Table columns can be presented in various plot types



ATLAS AtlasScope 02.00.00 - [03/Oct/2021 at 10:49:34 CEST by atevind for CERN] [Reset]

CERN-Proxy [Add] Search ATLAS dataset: DAOD_EXOT8 [Search] [Show - Table -]

Execute `g V().has('tbl', 'dataset')`

Customize the interactions with the graph. [Cluster by group type] [Cluster by group size] [Expand all clusters] [Show all edges] [Hierarchical (L up/R C size/hierarchy)] [Live]

filter: [Apply] [select] [limit(10)]

| version | runno | project | streamname | prodstep | datatype | dsid | dstypeid |
|------------------|--------|--------------|--------------|----------|-------------|----------|----------|
| r264_p3083_p4077 | 300800 | data16_13TeV | physics_Main | deriv | DAOD_BPHY21 | 11950848 | 10387 |

Evolution Plot Scatter Plot SkyView

| version | runno | project | streamname | prodstep | datatype | dsid | dstypeid | is_deleted | status | has_raw | has_trigger | prev_seen | tid | lbl | phoenix | fullfil | importDate |
|-------------------|---------|------------------|--------------|-----------|-------------|-----------|-----------|------------|-----------|---------|-------------|-----------|----------|---------|---------|---------|---------------|
| undefined | 2048 | r264_p3083_p4077 | undefined | undefined | undefined | undefined | undefined | undefined | AVAILABLE | true | false | 8162 | 20406081 | dataset | true | true | Mon Sep 27 18 |
| r264_p3083_p4077 | 302393 | data16_13TeV | physics_Main | deriv | DAOD_BPHY21 | 44456704 | 10387 | 0 | 2996970 | W | 11 | | | | | | |
| r264_p3083_p4077 | 302393 | data16_13TeV | physics_Main | deriv | DAOD_BPHY21 | 24009472 | 10387 | 0 | 4412611 | TL | 22 | | | | | | |
| r7562_p2521_p2950 | 280862 | data16_13TeV | physics_Main | merge | DAOD_EXOT12 | 42328576 | 10274 | 0 | 304311 | SL | 22 | | | | | | |
| r264_p3083_p4077 | 300655 | data16_13TeV | physics_Main | deriv | DAOD_BPHY21 | 940800 | 10387 | 0 | 1469531 | TL | 21 | | | | | | |
| r264_p3083_p4077 | 3022919 | data16_13TeV | physics_Main | deriv | DAOD_BPHY21 | 26108624 | 10387 | 0 | 425607 | TL | 22 | | | | | | |

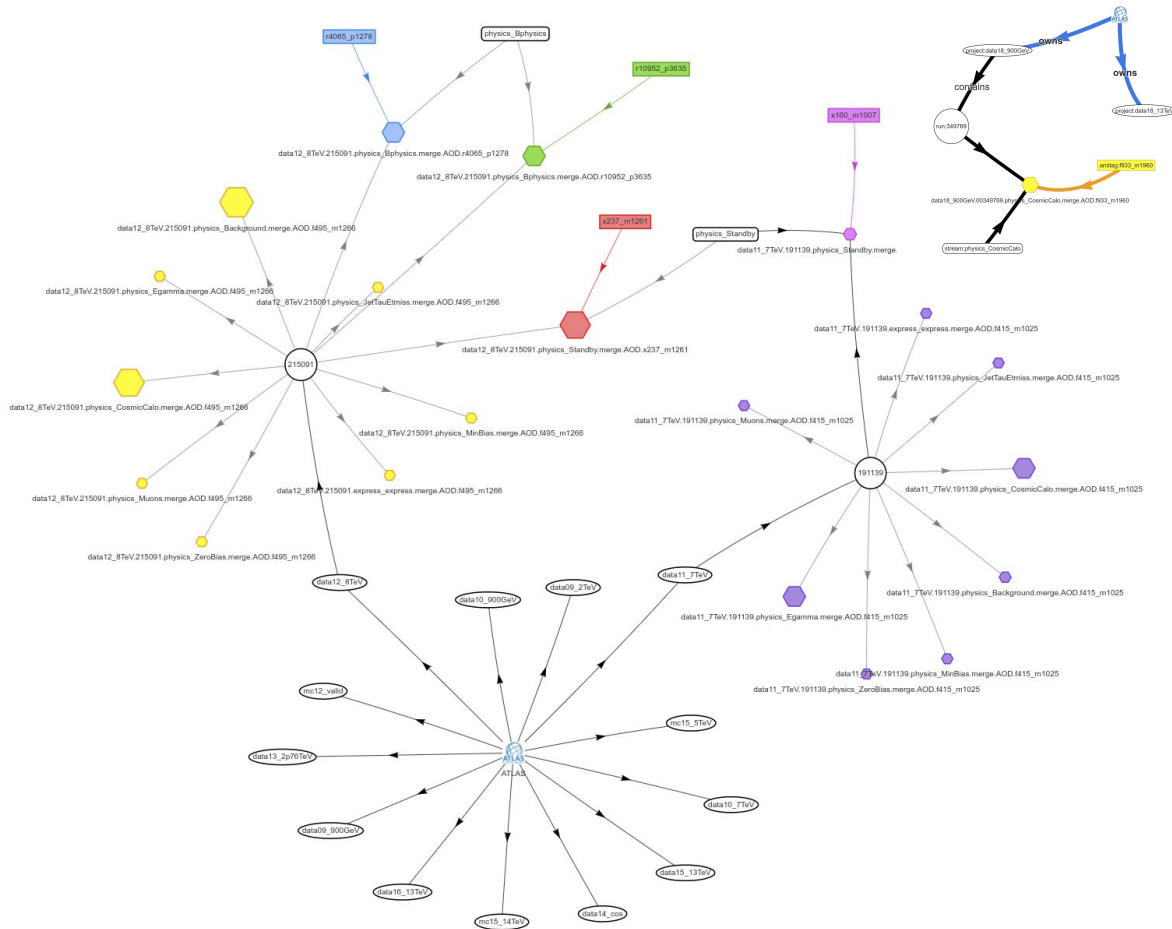

Select graph server and initial graph, then select an element to see possible actions.

Sending Gremlin request to //atlas-event-index.com.ch/AtlasScopeProxy.jsp?server=http://atlas-event-index.com.ch:8163: g V().has('tbl', 'dataset').limit(10)

Showing 10 new elements

Results Table Image Plot SkyView

Events dataset



Event Index Home:

<https://atlas-event-index.cern.ch/>



Presentation to CHEP'2019:

<https://docs.google.com/presentation/d/12UnR3iDWmYYKUQuZ5Hu0RdKbZAXC5uc3-LFKvgBbehS/edit?usp=sharing>