




Status of E1 Core & U1




- Status
- MC Dataset Overlaps
- New Aux Accessors
- Graph DB
 - Database
 - Web Service
 - Migration plans




Status - Import


More than 200 000 000 000 events available !



Events (200819704088 @ 2018-09-30)




Datasets (163033 @ 2018-09-30)






Status - Access



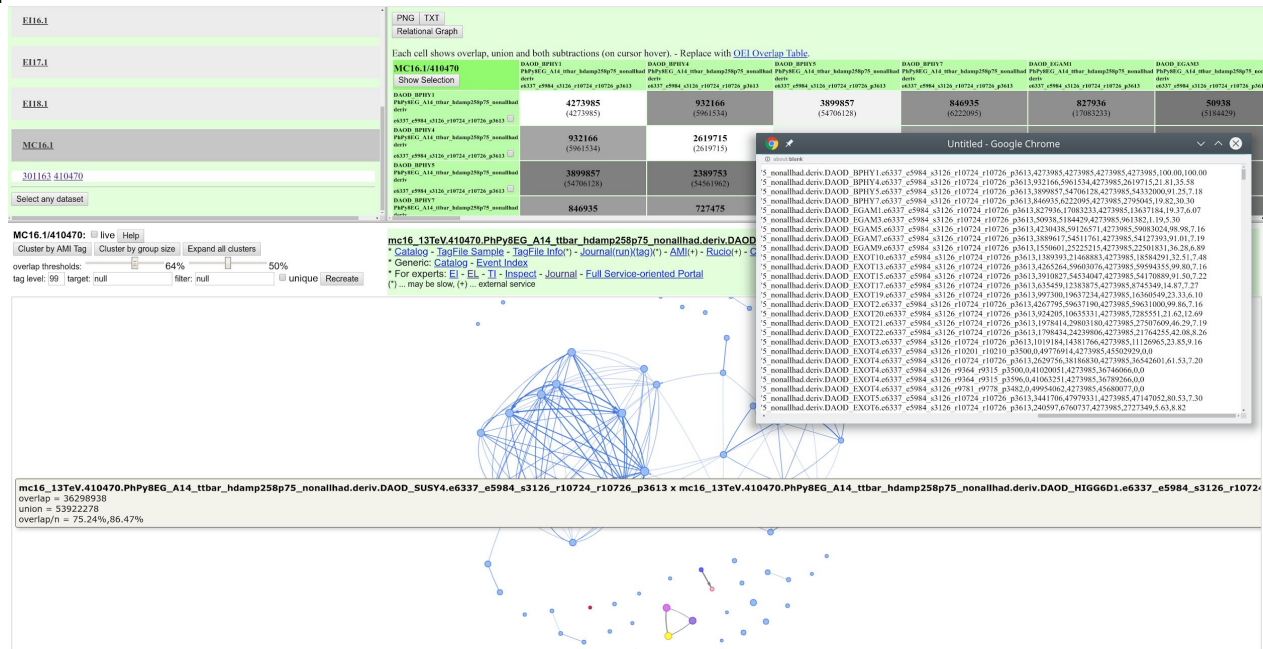
Daily Access statistics (3232909 @ 2018-09-30)






MC Dataset Overlaps

- Implemented on request for run MC16.1/410470
 - Different handling of RunNumbers
- Triggered some improvements to general interface
- New possibility: TXT export from WS





New Aux Accessors




- **net.hep.atlas.Database.EIHadoop.Accessor.Aux.Filler**
 - Fill TagFile with missing columns from another TagFile
 - For example: fill trigger info from AOD to DAOD
- **net.hep.atlas.Database.EIHadoop.Accessor.Aux.Unique**
 - Extract only unique events from a set of TagFiles

\$ ei -query ... -mr ... -aux ...




Graph Databases



- Traditional data structures in HEP:
 - tuples (tables)
 - trees
 - nested tuples (trees of tuples)
 - relational (SQL-like)
- Schemafull or schemaless
- **But many of our data are graph-like & schemaless**
 - **Entities with relations**
 - $G = (V, E)$ # *graph* = (vectors, edges)
- Not well covered by relational (SQL) databases
 - We don't need only a possibility to add new data with pre-defined relations
 - We need to add new relations
- Graph databases exist since a long time
 - Matured only recently thanks to Big Data & AI
 - Very good implementations & (de-facto) standards available



Graph Database for E1




- **Large parts of E1 data are graph-like**
 - Catalog:
 - Various relations between dataset
 - Dataset overlaps
 - Events:
 - Versions of the same event in different places



Graph Databases

Standards & Choices




- De-facto standard language/api: **Gremlin**
 - Gremlin is a functional, data-flow language to traverse a property graph. Every Gremlin traversal is composed of a sequence of (potentially nested) steps. A step performs an atomic operation on the data stream. Every step is either a *map*-step (transforming the objects in the stream), a *filter*-step (removing objects from the stream), or a *sideEffect*-step (computing statistics about the stream).
 - Gremlin supports transactional & non-transactional processing in declarative or imperative manner.
 - Gremlin can be expressed in all languages supporting function composition & nesting.
 - Supported languages: Java, Groovy, Scala, Python, Ruby, Go, ...
- Commonly used framework: **TinkerPop**
- Leading implementation: **JanusGraph**
 - Supported storage backends: Cassandra, HBase, Google Cloud, Oracle BerkeleyDB
 - Supported graph data analytics: Spark, Giraph, Hadoop
 - Supported searches: Elastic Search, Solr, Lucene
- Chosen visualisation: **visj.org**



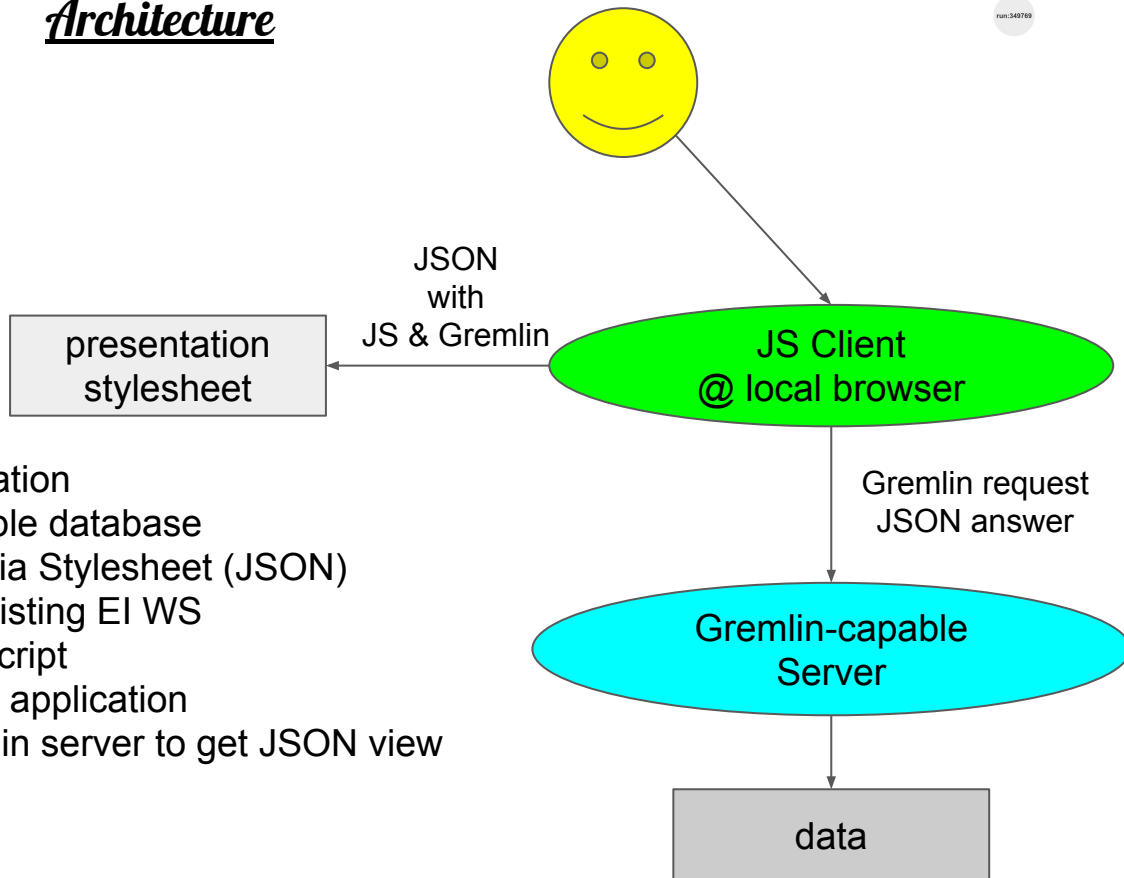


Graph Database for E1

Architecture



- Using standard Graph DB
 - Importing data from E1
- Generic Web Service graphical visualisation
 - Can display any Gremlin-compatible database
 - Visualisation can be customised via Stylesheet (JSON)
 - To give the same L&F as existing E1 WS
 - Implemented completely in JavaScript
 - So doesn't need server-side application
 - Connects to standard Gremlin server to get JSON view of data

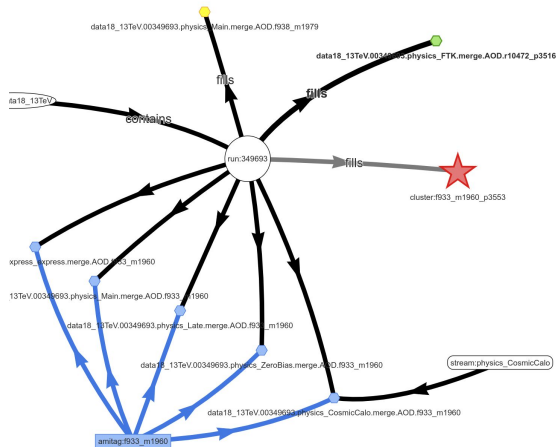
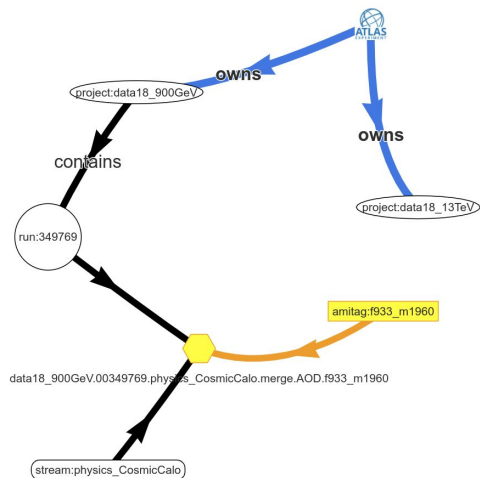





Graph Database for E1

Status

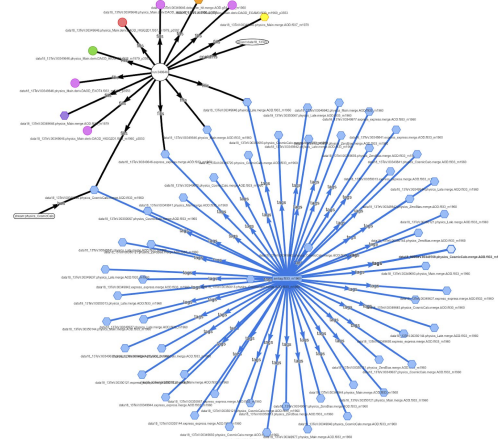
- Currently using standard installation (with Oracle BerkeleyDB backend)
 - Will migrate to HBase backend
- Snapshot of E18 Catalog with a subset of overlap tables imported
 - All Catalog & all overlap tables can be easily imported
- Most of the graphical part implemented
 - By standalone JS implementation



Customize the interactions with the graph.
Cluster by group type ... Expand all clusters
Catalog - Sample - Info +
live expand children expand parents remove old

data18_900GeV/00349769.physics_CosmicCalo.merge.AOD.933_m1960
Catalog - Sample - Info +

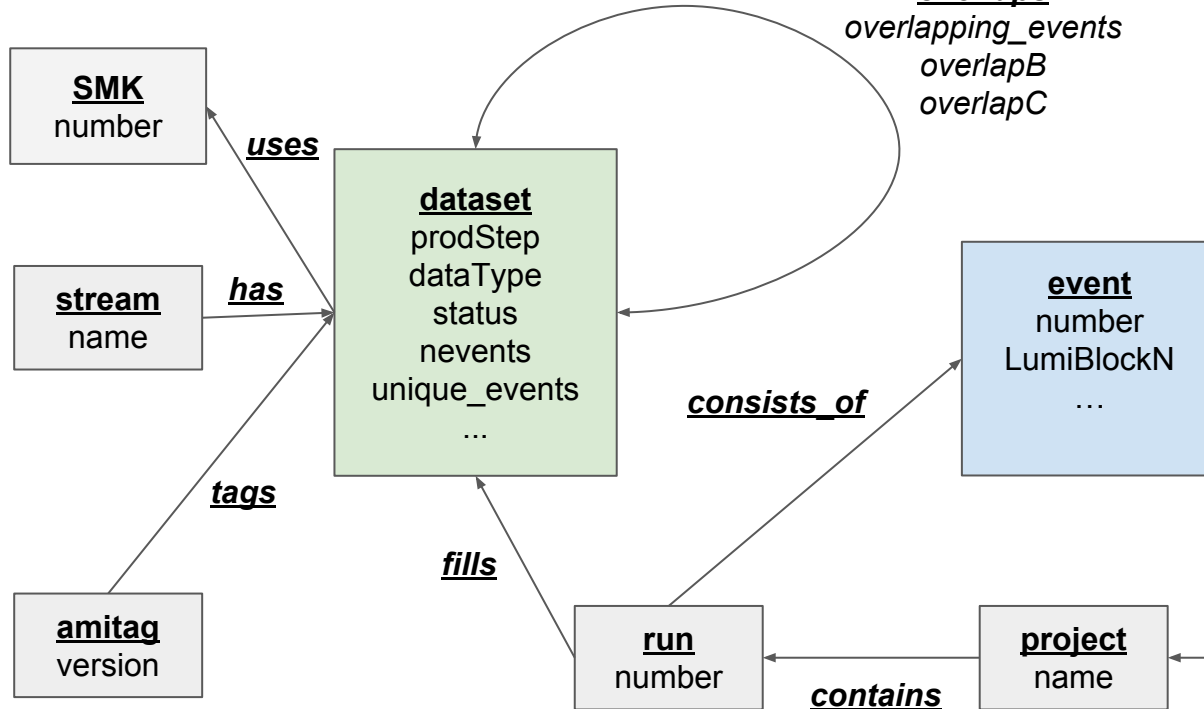

933_m1960 # Showing 0 new elements # Showing 0 new elements # Showing 55 new elements # Showing 55 new elements #



--- commands output ---



GraphDB Schema




- Should decide what is a feature and what is a relation
- Arrows have only logical meaning, they can be navigated equally from both sides
- Relations and features have defined multiplicities (not shown here) and types (int, string, date,...)
- Defined entities (and combinations of them) can be indexed for fast search
- Other features and relations can be freely added to any entity (vertex or edge)



Graph Database for E1

Plan - Datasets




- Install JanusGraph on Hadoop server (aiatlas016,54)
 - With data stored in HBase
 - [J.G. is already installed and functional on aiatlas016, using BerkeleyDB]
- Replicate all Catalog & Overlap data
 - Scripts exist, just to run them
- Migrate Catalog WS, CLI, API (the same API, new implementation)
 - Catalog & SimpleCatalog classes



Graph Database for E1

Plan - Events

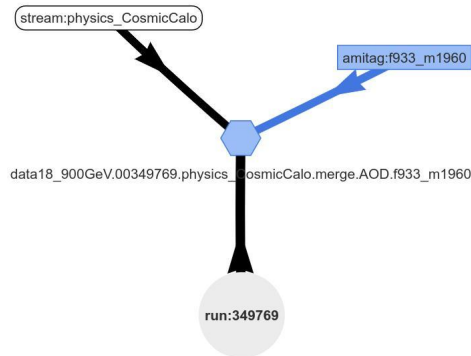


- **Structure of HBase backend of JanusGraph is very similar to our EventLookup HBase table (designed by Rainer)**
 - **So we can expect similar performance (i.e excellent performance)**
 - **Much simpler API: standard Gremlin instead of Rainer' script & quite complex access code**
- **Design & Create JanusGraph HBase for events**
 - Including data, which are currently only in HDFS files
 - Also trigger
 - While GraphDB doesn't require schema, we should define a reasonable structure
 - And formal schema/constraints (for some entities) can speed up searching
- **Replicate all data**
 - Connected to Catalog (already in JanusGraph)
- **Migrate WS, CLI, API (the same API, new implementation)**
- **Setup new workflow**
- **Switch to new implementation**
 - With existing CLI, WS



Graph Database for E1

Gremlin examples (Groovy style)



```
# add a vertex 'experiment' with the name 'ATLAS'
g.addV('experiment').property('name', 'ATLAS')
# add edges 'owns' from all vertices 'project' to vertex 'experiment' 'ATLAS'
g.V().hasLabel('project').addE('owns').from(g.V().hasLabel('experiment').has('name', 'ATLAS'))
# a function deriving a dataset name (which is not stored as such)
# from existing dataset relations by traversing the graph
def datasetName(d) {
    return d.sideEffect(values("prodStep").store("4"))
        .sideEffect(values("dataType").store("5")).in()
        .sideEffect(hasLabel("run").values("name").store("2"))
        .sideEffect(hasLabel("amitag").values("version").store("6"))
        .sideEffect(hasLabel("stream").values("name").store("3"))
        .sideEffect(hasLabel("run").in().hasLabel("project").values("name").store("1"))
        .cap("1", "2", "3", "4", "5", "6").next().values().join().toString()
}
```


There is also SQL API to Gremlin

in a similar way, one can navigate from an event to its dataset of certain dataType and amitag



Graph Database for E1


Summary



- Existing implementation provides
 - Excellent performance
 - Flexible architecture - new requirements are implemented within hours
- Most data (Catalog + EL) are stored in HBase tables
 - Using de-facto graph structure (entities with relations)
- Will migrate storage to GraphDB on HBase
 - i.e. keep storage & migrate structure into standard implementation
- Will keep (and partially re-implement) API and UI (CLI+WS)
- Adiabatic change
 - All clients keep working with the same interface
 - Minimal perturbation
- Will lose possibility to (directly) use Hadoop M/R
 - Those tasks should be re-developped
- Generic browser for all GraphDB data
 - Atlas has many domains with graph-like data
 - Graph databases are often used in AI (machine learning)



Info



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

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

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

GIT: <https://gitlab.cern.ch/atlas-event-index/Atlas-Event-Index-Core.git>

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

Questions & Comments: <mailto:Julius.Hrivnac>