



ADVANCED DATA ANALYSIS

Exploring job and transfer data in Elasticsearch via Kibana, directly, and through the Lens API

July 24, 2024

ELASTICSEARCH CONCEPTS

- Currently the primary store in Landscape for non-metric data – logs, events, job details, etc.
- Data model and interchange format is JSON documents composed of multiple key:value fields.
- Data storage engine is Apache Lucene.
- Data is stored in indices, distributed between one or more shards - each shard is a Lucene database.
- Commonly data is stored in time-based index patterns for partitioning and scalability (e.g. fifebatch-logs-YYYY.MM.DD). Clients usually support wildcards (e.g. fifebatch-logs-*), some “smarter” than others.

ELASTICSEARCH CONCEPTS

- An index has defined field mappings that are set when data is added, based on a template.
 - The mapping defines the type of field (string, float, etc) and how it is analyzed.
 - Fields can be given specific static types, or can be dynamically assigned based on the first value seen.
 - Changing a field mapping requires re-indexing the data!
- By default string fields are analyzed by lowercasing and splitting on word boundaries (space, dash, etc). Great for ordinary text search, not great for identifying strings (e.g. host names).
- Most strings we work with are “keyword” type that is not analyzed.

KIBANA

- Kibana is the native interface to Elasticsearch
 - explore raw data (“discover”)
 - Create ad-hoc visualizations (“visualize”)
 - Make dashboards, but that’s admins/by request only currently. You can usually make the same dashboard in Grafana.

Discover

Visualize

Dashboard

The screenshot shows the Kibana interface with the following components highlighted:

- Discover:** The top navigation menu and the main search area. The search bar contains the query "Search... (e.g. status:200 AND extension:PHP)". The left sidebar shows the "Discover" tab selected.
- Visualize:** The second navigation menu item.
- Dashboard:** The third navigation menu item.
- Search/filter:** A callout box pointing to the search bar and the "Add a filter" button.
- Toolbar:** A callout box pointing to the top right navigation buttons: "New", "Save", "Open", "Share", and "Last 24 hours".
- Timepicker:** A callout box pointing to the time range selector, which is currently set to "April 18th 2018, 11:04:24.140 - April 19th 2018, 11:04:24.140" with a "Time" dropdown menu.

The main content area displays a bar chart with a y-axis labeled "Selected Fields" and a time series. Below the chart, a list of fields is visible:

- sourcepath
- tags
- transfer_rate
- transfer_time
- user
- @timestamp

The bottom right of the screenshot shows a log entry for "April 19th 2018, 11:04:21.000" with the following details:

```
bytes_transferred: 0 process_id: 2789 ifdh_version: v2_1_0 n
ode_name: gm2samgpvm01.fnal.gov ifdh_use: interactive message
: 2018-04-19T16:04:21Z gm2samgpvm01.fnal.gov gm2pro/gm2/v2_1_0/g
m2[2789]: ifdh: ifdh cp: transferred: 0 bytes in 13.1062 seconds
type: ifdh tags: transferred, interactive @timestamp: April 1
```

KIBANA TIPS

- Everything is interactive! Click on a legend value or in a table to filter.
- Filters can be “pinned” to carry between views
- Kibana supports powerful full-text search queries using Elasticsearch [query string query](#), which are in turn based on Lucene expressions.
 - Specify field to search against as “field:term”
 - Combine terms with boolean operations (OR, AND, NOT - caps are important!)
 - Group terms with parens “name:(foo bar)” (OR is implicit)
 - Match wildcards “name:foo*” and regex “name:/foo.*”
 - Numeric comparisons “Memory_mb:>2000”



LET'S LOOK AROUND

<https://landscape.fnal.gov/kibana>

ELASTICSEARCH DATA

Key indices and fields

fifebatch-history: historical job data

fifebatch-events: Job state changes

hepcloud-classads-slots: current glidein data

fife-dh: IFDH transfer events

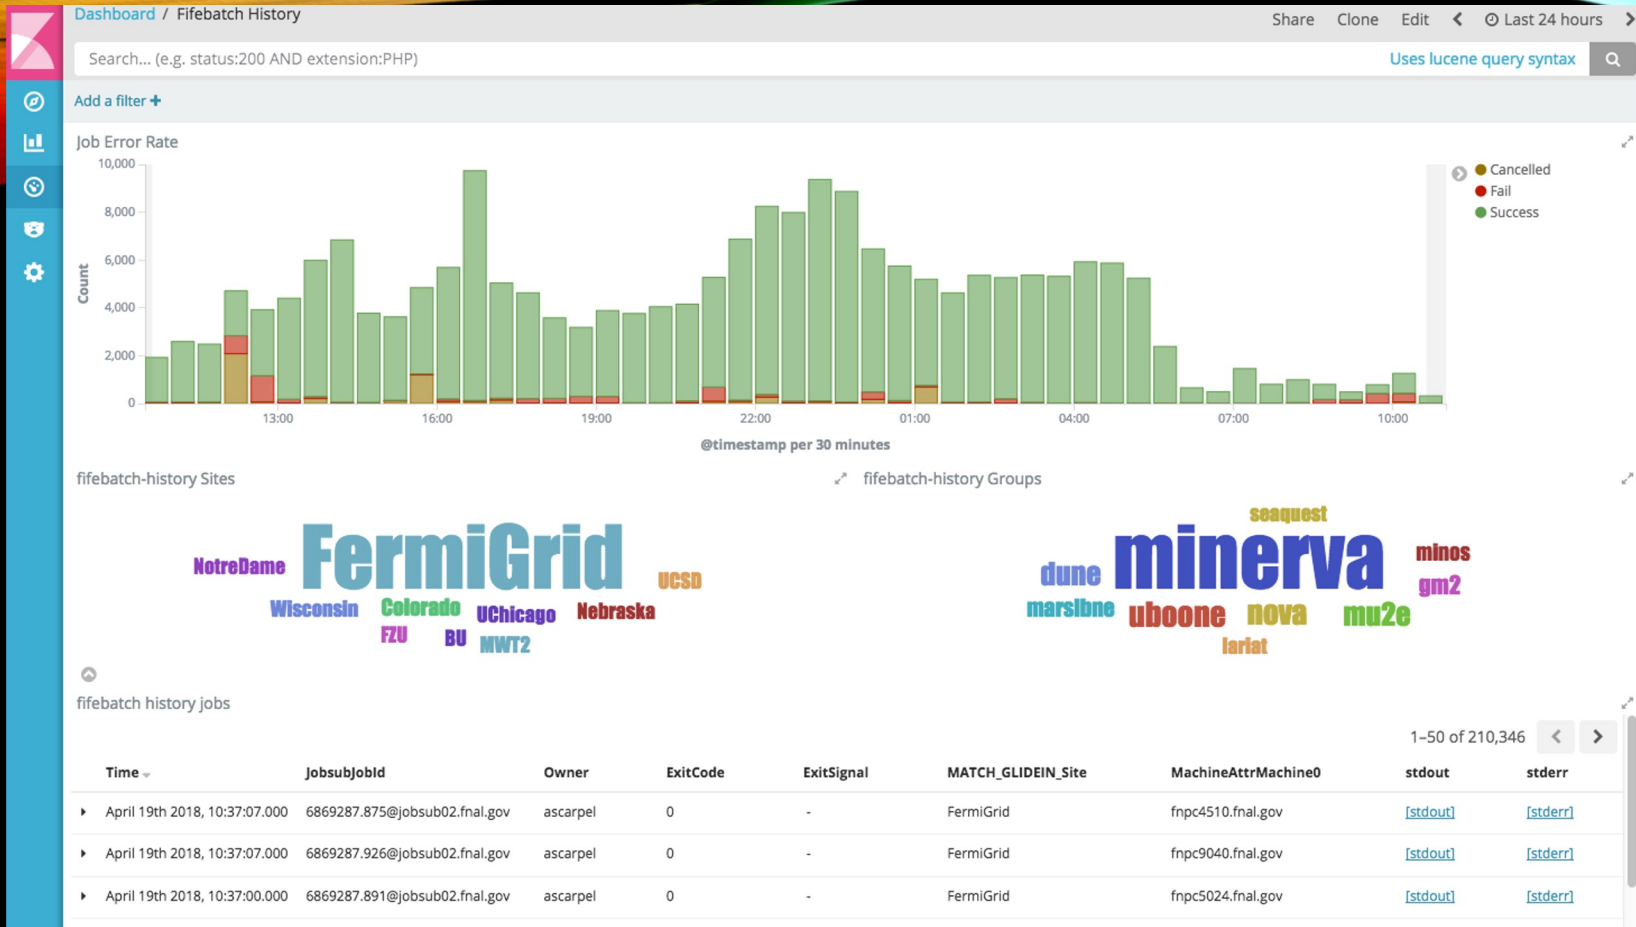
dcache-billing-events: dCache transfers

(Naming things is hard, “fife” doesn’t mean “fife only” and probably has CMS data in it and more.)

FIFEBATCH-HISTORY

Complete job attributes (ClassAds) representing final state of jobs.

- Daily indices (e.g. fifebatch-history-2018.04.19)
- Keep two years
- Complete “raw” HTCCondor ClassAds



<https://landscape.fnal.gov/kibana/app/kibana#/dashboard/ba047b90-b8ca-11e7-989a-91951b87e80a>

Similar: <https://landscape.fnal.gov/monitor/d/000000167/user-batch-history?orgId=1&refresh=5m>

HISTORY FIELD REFERENCE

Key Attributes

JobsubJobId "123456789.0@fifebatch2.fnal.gov"

ClusterId "123456789"

Procid "0"

host "jobsub01.fnal.gov"

JobStatus

1: idle

2: running

3: cancelled

4: completed

5: ~~held~~

Owner "novapro"

JobsubClientKerberosPrincipal

"amoren@FNAL.GOV"

Jobsub_Group "nova"

Final Attributes

ExitCode/ExitSignal

Success: "ExitCode:0 AND JobStatus:4"

Fail: "NOT ExitCode:0 AND JobStatus:4"

MachineAttrMachine0 "fnpc9060.fnal.gov"

MachineAttrGLIDEIN_Site0 "FermiGrid"

CommittedTime "3065"

RemoteWallClockTime is similar, but includes evictions

RemoteUserCpu "2396" (s)

RemoteSysCpu "12" (s)

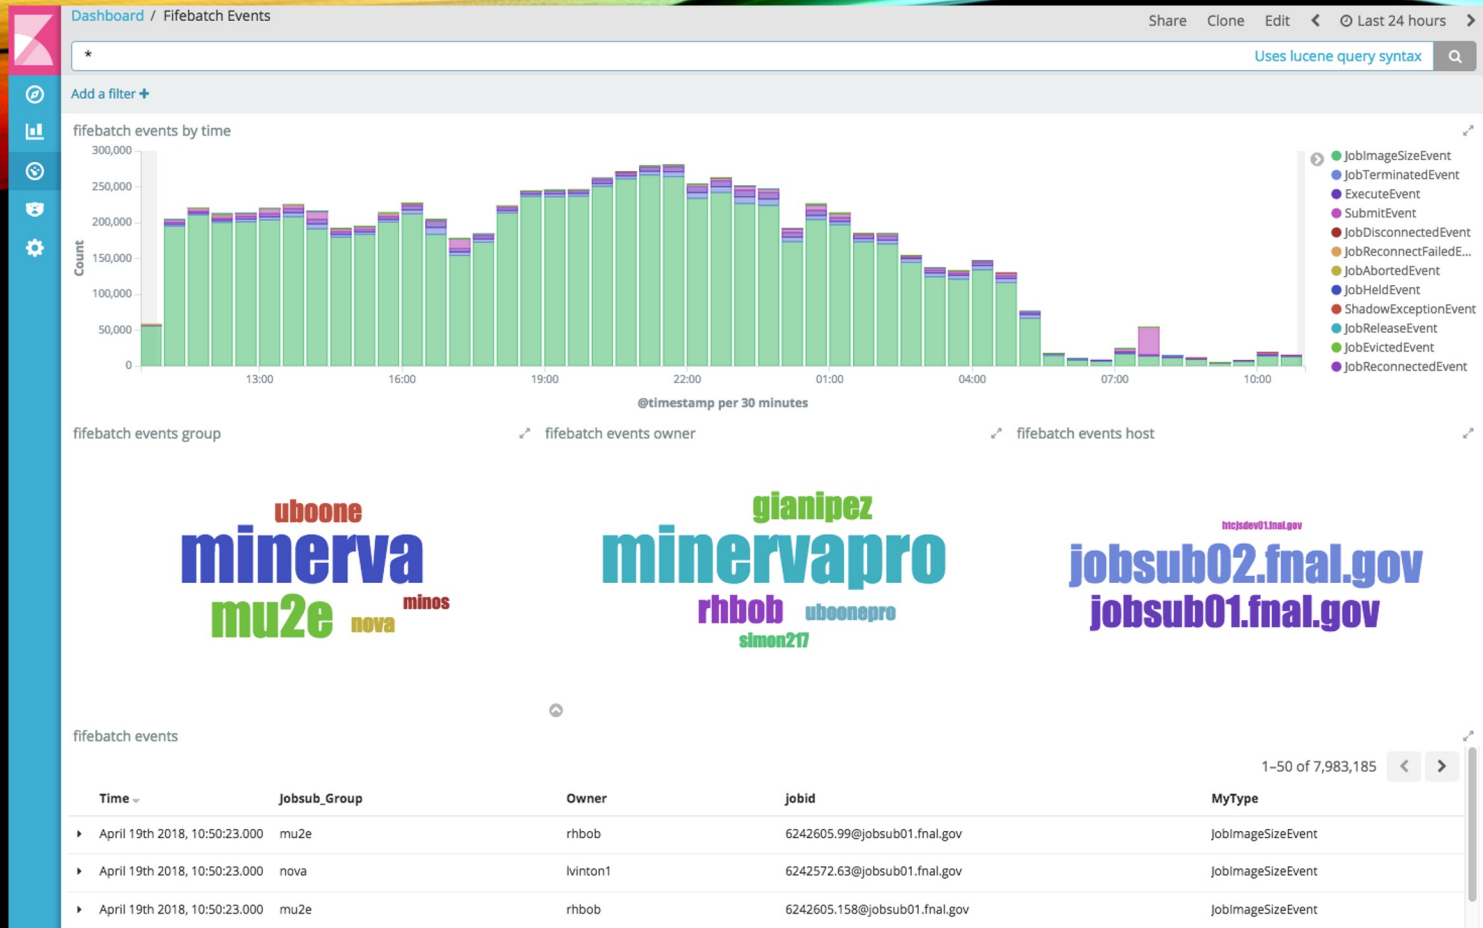
ResidentSetSize_RAW "1,343,928" (KB)

DiskUsage_RAW "110" (KB)

FIFEBATCH-EVENTS

HTCondor event logs (job submitted, job started, etc) collected in “real time”

- Daily indices (e.g. fifebatch-events-2018.04.19)
- Keep 180 days



<https://landscape.fnal.gov/kibana/app/kibana#/dashboard/fifebatch-events>

Similar: <https://landscape.fnal.gov/monitor/d/000000167/user-batch-history?orgId=1&refresh=5m> (again)

EVENTS FIELD REFERENCE

Key Attributes

jobid "123456789.0@fifebatch2.fnal.gov"

cluster "123456789"

process "0"

host "fifebatch2.fnal.gov"

Owner "novapro"

JobSub_Group "nova"

MachineAttrGLIDEIN_Site0 "FNAL"

MachineAttrGLIDEIN_ResourceName0
"GPGGrid"

MachineAttrMachine0 "fnpc9060.fnal.gov"

Event Attributes

TriggerEventTypeName "5"

TriggerEventName

"ULOG_JOB_TERMINATED"

MyType "JobTerminatedEvent"

Key Events

1	SubmitEvent
2	ExecuteEvent
4	JobEvictedEvent
5	JobTerminatedEvent
6	JobImageSizeEvent
9	JobAbortedEvent
12	JobHeldEvent
13	JobReleasedEvent
22	JobDisconnectedEvent
23	JobReconnectedEvent

HEPCLOUD-CLASSADS-SLOTS

Select slot/machine attributes (ClassAds) updated every five minutes from fifebatch HTCondor collectors.

SLOTS FIELD REFERENCE

Key Attributes

Name

"slot1_3@glidein_1045375_1086175008@fnpc4201.fnal.gov"

GLIDEIN_Site "FNAL"

GLIDEIN_ResourceName "GPGrid"

SlotType

"Partitionable": pilot

"Dynamic": claimed

State "Claimed"

Resources

For "Partitionable" slot this is what's remaining unclaimed, for "Dynamic" this is what's claimed by the job.

Cpus "35GB"

Disk "35GB"

Memory "3GB" (raw in bytes)

Memory_mb "3072"

Memory_gb "3"

FIFE-DH

IFDH event logs (start transfer, end transfer, POMS data) collected in “real time”

- Daily indices (e.g. fife-dh-2018.04.19)
- Keep 180 days

Search... (e.g. status:200 AND extension:PHP)

Uses lucene query syntax

Add a filter +


fifebatch-history Groups




history owner tag cloud




Job Error Rate




IFDH Transfer Events



history walltime percentiles



ifdh transfer time percentiles



failed jobs per site

MachineAttrGLIDEIN_Site0: Descending	Count
FermiGrid	37,866
Colorado	149
UChicago	50
Wisconsin	27

failed jobs per host

MachineAttrMachine0: Descending	Count
fnpc9044.fnal.gov	224
fnpc9057.fnal.gov	211
fnpc9058.fnal.gov	204
fnpc17135.fnal.gov	198

failed jobs per owner

Owner: Descending	Count
rhbob	31,528
gianipez	2,000
cbsykes	1,797
minervapro	723

FIFE-DH FIELD REFERENCE

Key Attributes

jobid/JobsubJobId

"123456789.0@fifebatch2.fnal.gov"

cluster "123456789"

process "0"

user/Owner "novapro"

experiment/Jobsub_Group "nova"

ifdh_event_type "starting_transfer"

Transfer Attributes

node_name "fnpc6013.fnal.gov"

sourcepath

"/pnfs/GM2/scratch/.../gm2offline_reco_15395.00068.root "

destpath

"/pnfs/GM2/scratch/.../gm2offline_reco_15395.00068.root "

bytes_transferred "2,016,223"

transfer_time "4.78027" (s)

Event Types

starting_transfer	"Ifdh cp" start
transferred	"Ifdh cp" finished
starting_file_transfer	Start of single file transfer
finished_file_transfer	End of single file transfer
begin_execution	Job started
finished	Job finished
failed_transfer	Failed transfer
poms_data	POMS metadata for job
ifdh_log_msg	Catch-all

DCACHE-BILLING-EVENTS

dCache event logs (“door” and “pool” events) collected in “real time”

- Daily indices (e.g. dcache-billing-events-v1-2024.07.22)
- Keep 90 days

> Search... (e.g. status:200 AND extension:PHP)

Options Update

Add a filter +


dcache_metric_transfers_write

dcache_metric_transfers_read

5.569TB
Transfers (Write)

23.755TB
Transfers (Read)

dcache_datehistogram_read



dcache_popular_pools_read

dcache_popular_area_read

v-stkendca1908-3 v-stkendca2209-7
v-stkendca2022-2
v-stkendca2207-7
 w-cmsstor203-disk-disk4

nova.scratch@enstore GM2.scratch@enstore
cms.cms11@enstore
 dune.scratch@enstore
 minerva.mc_reconstructed@enstore

DCACHE FIELD REFERENCE

Key Attributes

cellType “door” or “pool”

UserNamePrincipal “novapro”

(see also: subject, UidPrincipal, GidPrincipal)

storage_group “nova”

billingPath “/pnfs/...”

transferPath “/pnfs/...” (usually the same)

isWrite “read” or “write”

status.code “0” for success

status.message (for failures)

transferSize “1234” (bytes)

transferTime “1234” (seconds)



ACCESSING DATA

Kibana

Grafana

cURL

Client libraries

Lens API

GRAFANA

- Elasticsearch index patterns get added as data sources in Grafana, usually with a similar name.
 - Notable exception: “fifebatch-events” -> “fifebatch logs” (sorry)
 - “fife events” is outages
 - Demonstrate “Explore”
- Often easiest to start with a similar dashboard or graph
 - Demonstrate how to copy dashboards to scratch area
 - Demonstrate how to copy panels to new dashboard
- Demonstrate different panels
 - Time series and aggregations
 - Table

CURL

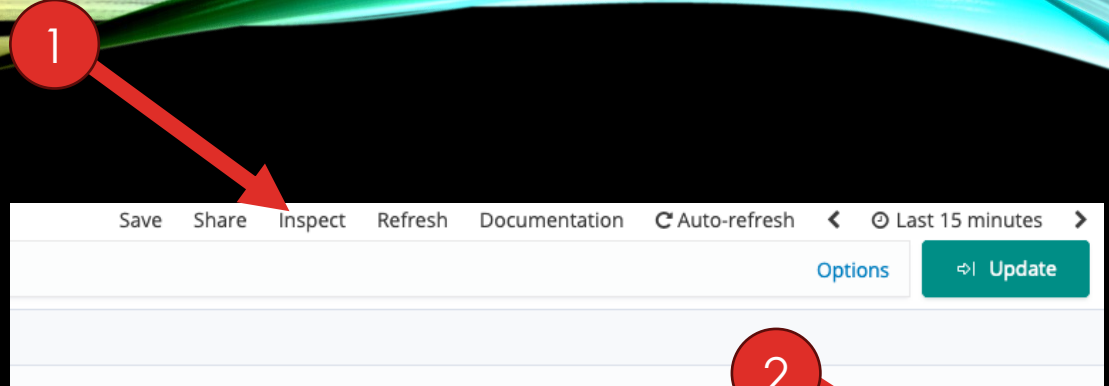
Elasticsearch uses a RESTful HTTP interface, we expose that read-only at <https://fifemon-es.fnal.gov>

The Elasticsearch [query DSL](#) is quite complex JSON-based language, but it's very powerful.

In addition to searching for documents, it can do aggregations on numeric fields, group results into nested buckets, and more.

PRO TIP:

Click “Inspect” at the top of a Kibana visualization (1) then select “View: Requests” (2) then “Requests” (3) to see the raw query that Kibana is making to Elasticsearch.



GET A SINGLE DOCUMENT BY ID

```
$ curl 'https://fifemon-es.fnal.gov/fifebatch-jobs/job/6185988.0@jobsub01.fnal.gov?pretty'  
{  
  "_index" : "fifebatch-jobs.01",  
  "_type" : "job",  
  "_id" : "6185988.0@jobsub01.fnal.gov",  
  "_version" : 576,  
  "found" : true,  
  "_source" : {  
    "cluster" : 6185988,  
    "Owner" : "adi",  
    "RequestDisk" : 35000000,  
    "NumJobStarts" : 1,  
    "HoldReasonCode" : 26,  
    "RemoteUserCpu" : 4982.0,  
    "time_ratio" : 0.0,  
    "type" : "job",  
    ...  
  }  
}
```

SEARCH FOR DOCUMENTS

```
$ curl 'https://fifemon-es.fnal.gov/fifebatch-history-*/_search?pretty' -d '{
  {
    "query": {
      "query_string": {
        "query": "Owner:novapro AND NOT ExitCode:0 AND EnteredCurrentStatus:[now-12h TO now]"
      }
    },
    "size": 1,
    "_source": [
      "JobsubJobId"
    ]
  }
}'

{
  "took" : 96,
  "timed_out" : false,
  "hits" : {
    "total" : 15,
    "hits" : [
      {
        "_index" : "fifebatch-history-2018.04.19",
        "_type" : "condor_history",
        "_id" : "AWLc4W_HXgSn6-mb-4kT",
        "_score" : 6.5376205,
        "_source" : {
          "JobsubJobId" : "6239238.0@jobsub01.fnal.gov"
        }
      }
    ]
  }
}
```

AGGREGATE

```
$ curl 'https://fifemon-es.fnal.gov/fifebatch-history-*/_search?pretty' -d '{
  "query": {
    "query_string": {
      "query": "Owner:novapro AND NOT ExitCode:0 AND EnteredCurrentStatus:[now-12h TO now]"
    }
  },
  "size": 0,
  "aggs": {
    "walltime": {
      "sum": {
        "field": "CommittedTime"
      }
    }
  }
},
{
  "took" : 138,
  "timed_out" : false,,
  "hits" : {
    "total" : 15,
    "max_score" : 0.0,
    "hits" : [ ]
  },
  "aggregations" : {
    "walltime" : {
      "value" : 353307.0
    }
  }
}
```

CLIENT LIBRARIES

There are many client libraries to help communicate with and query Elasticsearch.

For example, [elasticsearch](#) and [elasticsearch-dsl](#) for Python.

Example: <https://landscape.fnal.gov/docs/using/#elasticsearch>

LANDSCAPE API

- Landscape has a GraphQL API for current and historical job data.
 - Documentation at <https://landscape.fnal.gov/docs/using/#graphql>
 - Explore data and schema at <https://landscape.fnal.gov/lens>