What did Pelican do to my transfer?
A Monitoring Story

Jason Patton, CHTC

HTC 24

A day in the life of an OSPool user

- I submit a batch of jobs which use the OSDF (Pelican) plugin
  - transfer_input_files = osdf:///ospool/ap21/data/jcpatton/mydata.tar
- Most of my jobs run and finish successfully, yay!
- Some of my jobs go on hold with cryptic error messages, boo…

  -
    ```
     Error: Pelican Client Error: Attempt #3: from dtn-
    pas.cinc.nrp.internet2.edu:8443: failed connection setup: Get
    "https://dtn-pas.cinc.nrp.internet2.edu:8443(...Path...)": read tcp
    172.16.33.8:46386->163.253.29.17:8443...
    ```
  - Whatever, I release the held jobs and they complete, ok!
- My advisor finds out that I’m using data caching infrastructure
  - ”Why did you have to babysit (i.e. release) your jobs?”
  - “Who has access to our data now?”

Additional metadata added to Pelican plugin output

- Pelican client version
- Number of transfer attempts
- Per attempt:
  - Endpoint (i.e. cache)
  - Endpoint server version (i.e. XrootD version)
  - Time to first byte
  - Transfer duration
  - Transfer end time
  - Bytes transferred
  - Error message (if encountered)
- But where do these data go?

What did the file transfer plugin do to my transfer?
A Monitoring Story

Jason Patton, CHTC

HTC 24

CHTC HTCondor Software Suite PATh

The AP contains transfer plugin data

- Features added in 10.x and 23.x to improve visibility into job histories and file transfer histories
- Timeline
    - 10.3.0 - Per-execution *job epoch* history ClassAds added to the AP
    - 23.4.0 - Transfer plugin ClassAds added to job epoch history
    - 23.8.1 - HTCondor Python bindings support transfer plugin ads
    - 23.x (soon!) - condor_adstash supports transfer plugin ads
- Take home message: **File transfer plugins (e.g. Pelican) can emit arbitrary diagnostic data, this data is stored on the AP, and this data can be (remotely) queried.**

How to fetch transfer history ads

```
condor_history -epochs -type TRANSFER [options]
```

```
$ condor_history 785479.34 -epochs -type TRANSFER -limit 1 -long
ClusterID = 785479
EpochWriteDate = 1720189895
MachineAttrGLIDEIN_ResourceName0 = "UColorado_HEP"
MachineAttrGLIDEIN_Site0 = "Colorado"
MachineAttrName0 = "slot1_1@glidein_813411_196236950@lnxfarm205.colorado.edu"
NumShadowStarts = 2
OutputPluginResultList = { [ TransferUrl = "osdf:///ospool/ap40/data/redacted"; TransferType =
"upload"; DeveloperData = [ Attempts = 1; DataAge0 = 0.0; Endpoint0 = "ospool-
ap2140.chtc.wisc.edu:8443"; TransferTime0 = 17.333; ServerVersion0 = "XrootD/v5.6.9";
TimeToFirstByte0 = 0.2; TransferEndTime0 = 1720189895; TransferFileBytes0 = 181643160;
PelicanClientVersion = "7.9.2" ]; TransferEndTime = 1720189895; TransferSuccess = true;
TransferFileName = ”redacted"; TransferProtocol = "osdf"; TransferFileBytes = 181643160;
TransferStartTime = 1720189867; TransferTotalBytes = 181643160 ] }
ProcID = 34
```

How to fetch transfer history ads

Python binding: `Schedd.jobEpochHistory(**kwargs, ad_type="TRANSFER")`

- Returns an *iterator* that doesn’t contact the AP until consumed

```
>>> import htcondor
>>> schedd = htcondor.Schedd()
>>> hist_iter = schedd.jobEpochHistory(constraint="ClusterId == 785479 &&
ProcId == 34", projection=[], match=1, ad_type="TRANSFER")
>>> next(hist_iter)
[ ProcID = 34; ClusterID = 785479; EpochWriteDate = 1720189895; NumShadowStarts = 2; MachineAttrName0 =
"slot1_1@glidein_813411_196236950@lnxfarm205.colorado.edu"; OutputPluginResultList = { [ TransferUrl =
"osdf:///ospool/ap40/data/redacted"; TransferType = "upload"; DeveloperData = [ Attempts = 1; DataAge0 = 0.0; Endpoint0 =
"ospool-ap2140.chtc.wisc.edu:8443"; TransferTime0 = 1.733300000000000E+01; ServerVersion0 = "XrootD/v5.6.9"; TimeToFirstByte0 =
2.000000000000000E-01; TransferEndTime0 = 1720189895; TransferFileBytes0 = 181643160; PelicanClientVersion = "7.9.2" ];
TransferEndTime = 1720189895; TransferSuccess = true; TransferFileName = ”redacted"; TransferProtocol = "osdf";
TransferFileBytes = 181643160; TransferStartTime = 1720189867; TransferTotalBytes = 181643160 ] }; MachineAttrGLIDEIN_Site0 =
"Colorado"; MachineAttrGLIDEIN_ResourceName0 = "UColorado_HEP" ]
```

How to decipher a transfer history ad

```
        'ClusterID': 785419,
        'ProcID’: 1
        'NumShadowStarts': 1,
        'EpochWriteDate': 1720019431,
        'InputPluginResultList’: [{
                'TransferEndTime': 1720019430,
                'TransferFileBytes': 5875073024,
                'TransferFileName': ‘mydata.tar’,
                'TransferProtocol': 'osdf’,
                'TransferStartTime': 1720019430,
                'TransferSuccess': True,
                'TransferTotalBytes': 5875073024,
                'TransferType': 'download’,
                'TransferUrl': 'osdf:///ospool/ap40/data/jcpatton/mydata.tar’
                'DeveloperData’: {
                        'PelicanClientVersion': '7.8.1’,
                        'Attempts': 1,
                        'Endpoint0': 'sdsc-cache.nationalresearchplatform.org:8443’,
                        'ServerVersion0': 'XrootD/v5.6.9’,
                        'TimeToFirstByte0': 0.018,
                        'TransferEndTime0': 1720019430,
                        'TransferFileBytes0': 5875073024,
                        'TransferTime0': 9.8
                        },
                }]
```

**Job epoch identifying information
in top-level attributes**

**Common transfer plugin information
listed per plugin
in InputPluginResultList
or OutputPluginResultList**

**Plugin-specific information
and per-attempt information
in DeveloperData**

Admins: consider using condor_adstash

- condor_adstash is a tool that fetches certain (e.g. job) ClassAds, converts them to JSON, and optionally pushes them to a search engine database like Elasticsearch or OpenSearch
- **DeveloperData broken into JSON doc per transfer attempt**
- Coming soon (bug fixing!), but contact me for a container image that you can try today
- Consider the amount of metadata your users' transfers are creating when creating aliases and lifecycle policies!

Explore transfer data using Kibana

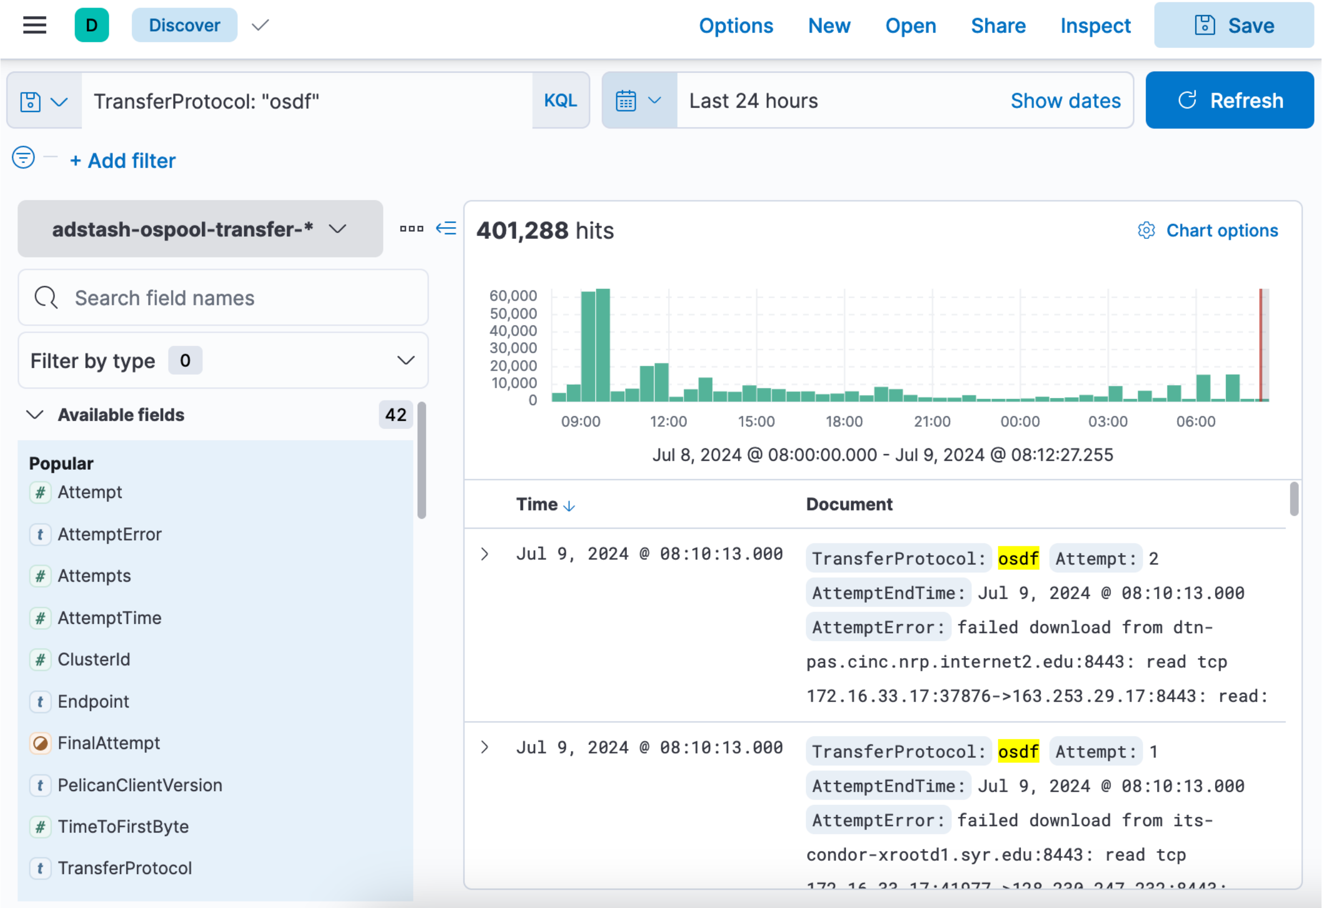

TransferProtocol: "osdf"

401,288 hits

Jul 8, 2024 @ 08:00:00.000 - Jul 9, 2024 @ 08:12:27.255

```
Jul 9, 2024 @ 08:10:13.000   TransferProtocol: osdf  Attempt: 2
                             AttemptEndTime: Jul 9, 2024 @ 08:10:13.000
                             AttemptError: failed download from dtn-
                             pas.cinc.nrp.internet2.edu:8443: read tcp
                             172.16.33.17:37876->163.253.29.17:8443: read:

Jul 9, 2024 @ 08:10:13.000   TransferProtocol: osdf  Attempt: 1
                             AttemptEndTime: Jul 9, 2024 @ 08:10:13.000
                             AttemptError: failed download from its-
                             condor-xrootd1.syr.edu:8443: read tcp
                             172.16.33.17:41077->128.230.247.232:8443:
```

CHTC HTCondor Software Suite PATh

Explore transfer data using Kibana

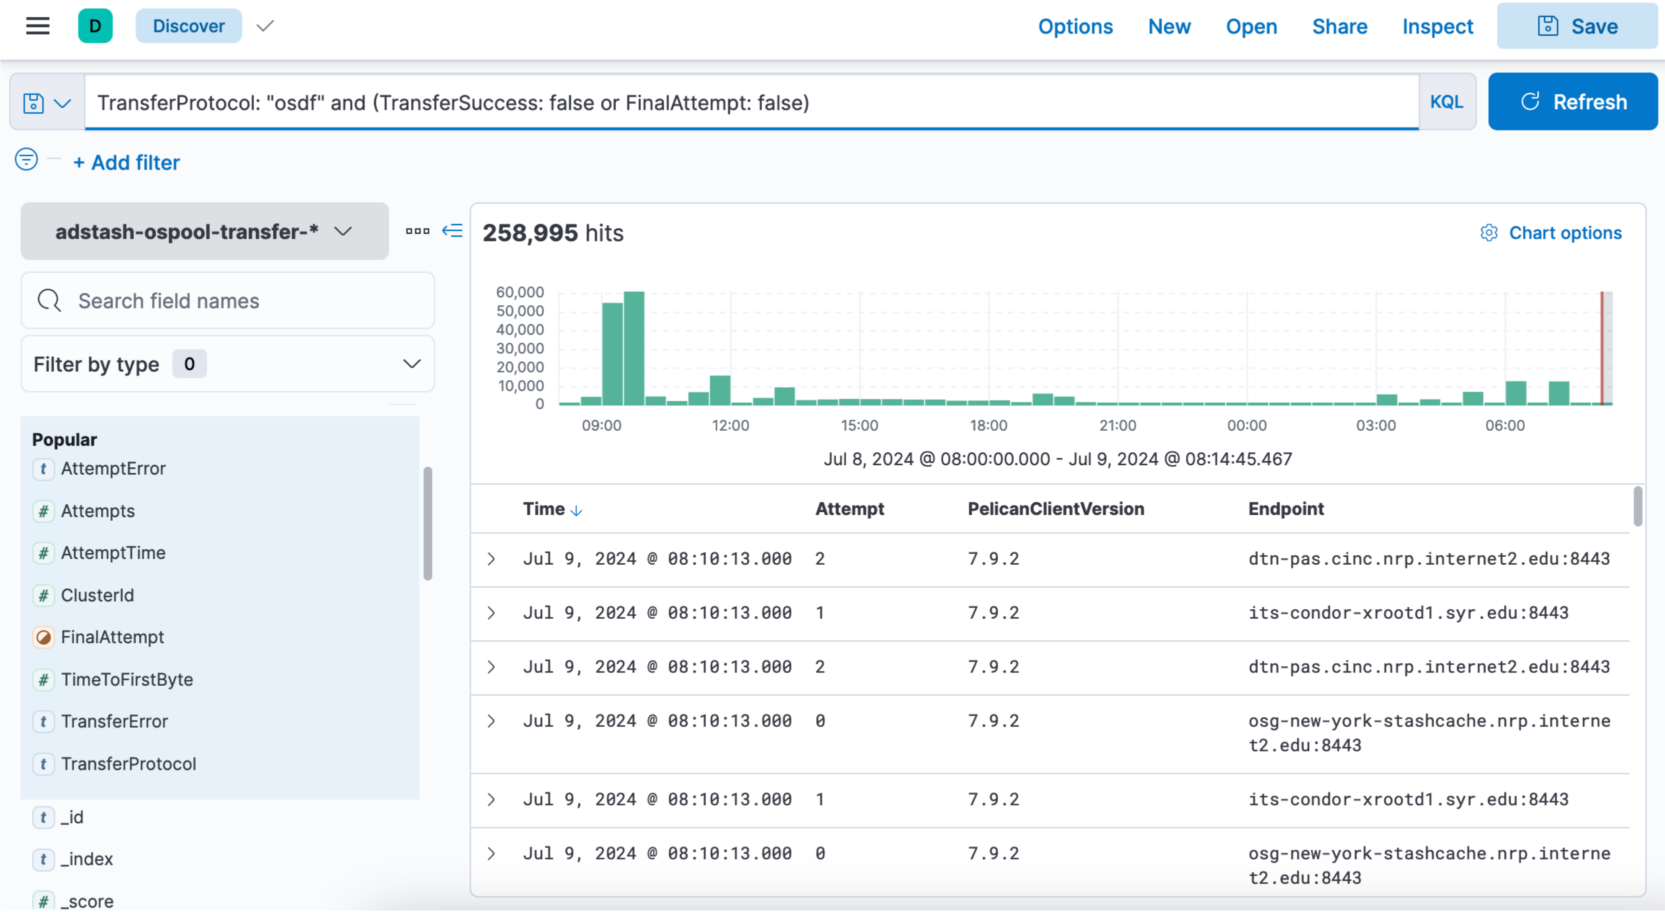

TransferProtocol: "osdf" and (TransferSuccess: false or FinalAttempt: false)

258,995 hits

Jul 8, 2024 @ 08:00:00.000 - Jul 9, 2024 @ 08:14:45.467

Popular
- AttemptError
- Attempts
- AttemptTime
- ClusterId
- FinalAttempt
- TimeToFirstByte
- TransferError
- TransferProtocol

- _id
- _index
- _score

Time ↓	Attempt	PelicanClientVersion	Endpoint
Jul 9, 2024 @ 08:10:13.000	2	7.9.2	dtn-pas.cinc.nrp.internet2.edu:8443
Jul 9, 2024 @ 08:10:13.000	1	7.9.2	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:10:13.000	2	7.9.2	dtn-pas.cinc.nrp.internet2.edu:8443
Jul 9, 2024 @ 08:10:13.000	0	7.9.2	osg-new-york-stashcache.nrp.internet2.edu:8443
Jul 9, 2024 @ 08:10:13.000	1	7.9.2	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:10:13.000	0	7.9.2	osg-new-york-stashcache.nrp.internet2.edu:8443

Explore transfer data using Kibana

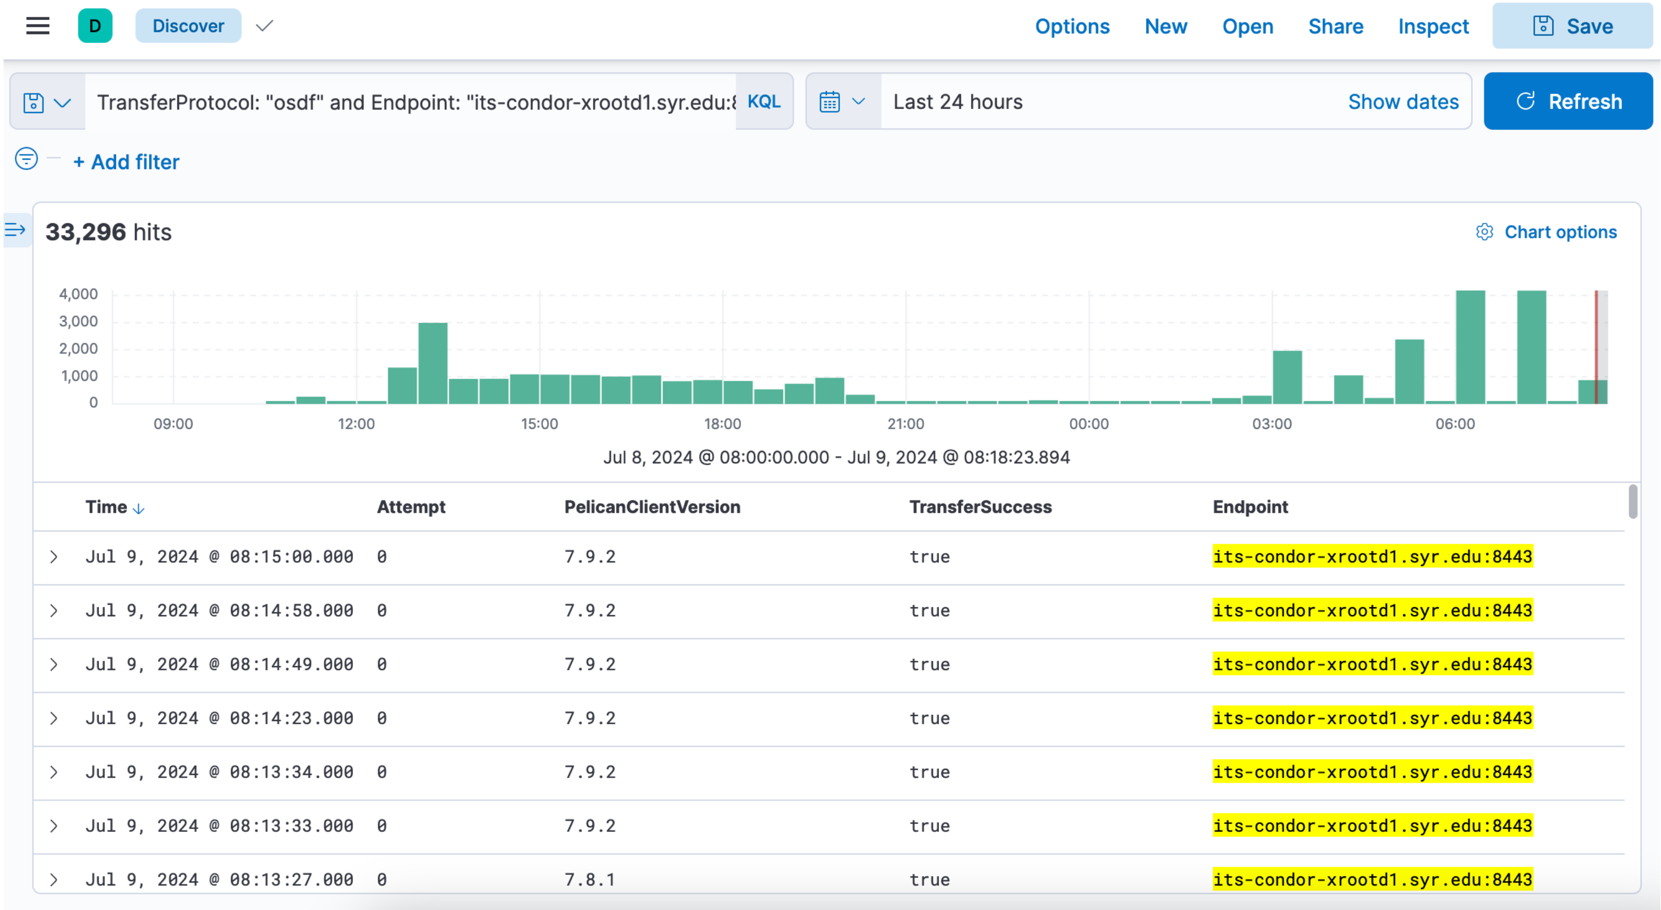

TransferProtocol: "osdf" and Endpoint: "its-condor-xrootd1.syr.edu:8

Last 24 hours

33,296 hits

Jul 8, 2024 @ 08:00:00.000 - Jul 9, 2024 @ 08:18:23.894

Time ↓	Attempt	PelicanClientVersion	TransferSuccess	Endpoint
Jul 9, 2024 @ 08:15:00.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:14:58.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:14:49.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:14:23.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:13:34.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:13:33.000	0	7.9.2	true	its-condor-xrootd1.syr.edu:8443
Jul 9, 2024 @ 08:13:27.000	0	7.8.1	true	its-condor-xrootd1.syr.edu:8443

Explore transfer data using Kibana

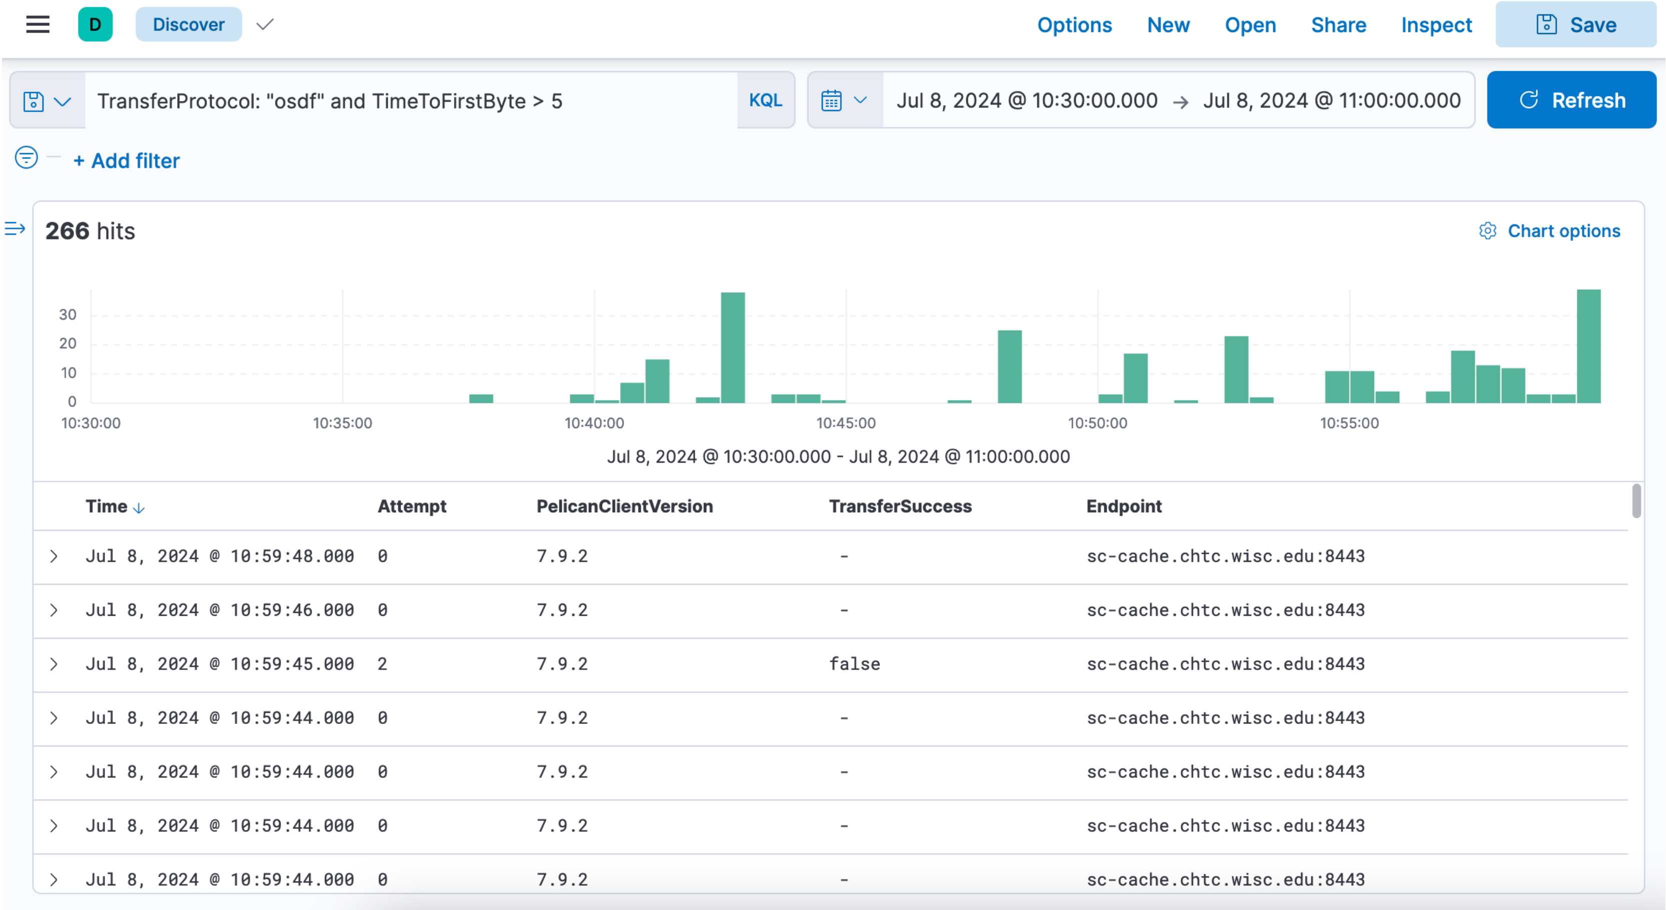

```
TransferProtocol: "osdf" and TimeToFirstByte > 5
```

Jul 8, 2024 @ 10:30:00.000 → Jul 8, 2024 @ 11:00:00.000

266 hits

Jul 8, 2024 @ 10:30:00.000 - Jul 8, 2024 @ 11:00:00.000

Time ↓	Attempt	PelicanClientVersion	TransferSuccess	Endpoint
Jul 8, 2024 @ 10:59:48.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:46.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:45.000	2	7.9.2	false	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:44.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:44.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:44.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443
Jul 8, 2024 @ 10:59:44.000	0	7.9.2	-	sc-cache.chtc.wisc.edu:8443

Explore transfer data using Kibana

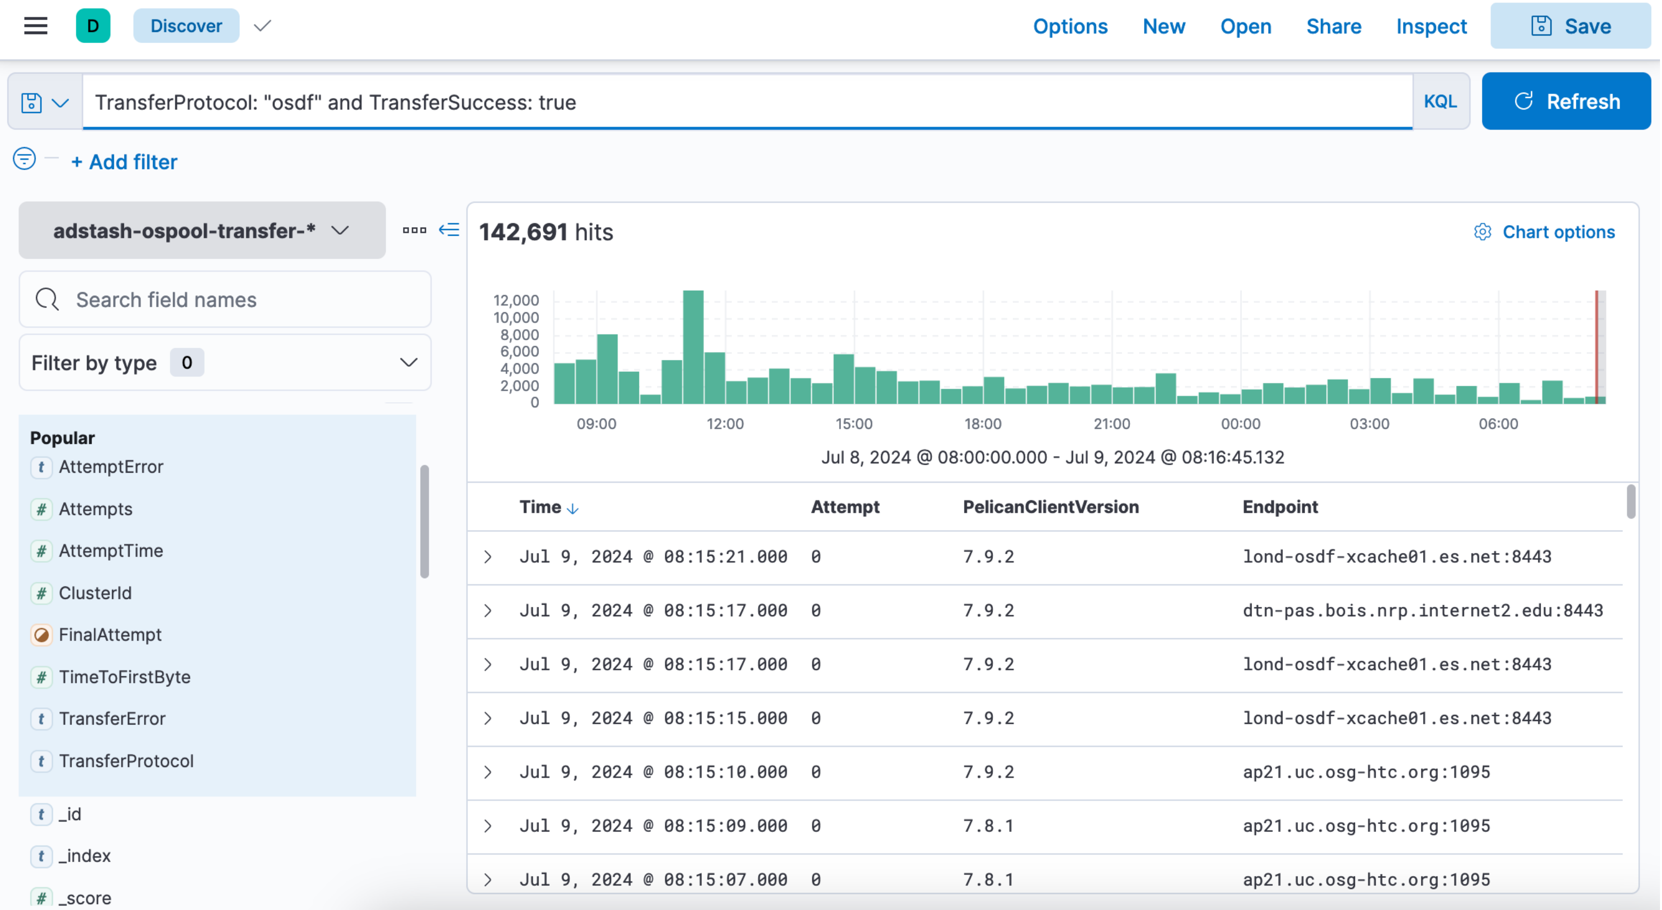

TransferProtocol: "osdf" and TransferSuccess: true

142,691 hits

Jul 8, 2024 @ 08:00:00.000 - Jul 9, 2024 @ 08:16:45.132

Time ↓	Attempt	PelicanClientVersion	Endpoint
Jul 9, 2024 @ 08:15:21.000	0	7.9.2	lond-osdf-xcache01.es.net:8443
Jul 9, 2024 @ 08:15:17.000	0	7.9.2	dtn-pas.bois.nrp.internet2.edu:8443
Jul 9, 2024 @ 08:15:17.000	0	7.9.2	lond-osdf-xcache01.es.net:8443
Jul 9, 2024 @ 08:15:15.000	0	7.9.2	lond-osdf-xcache01.es.net:8443
Jul 9, 2024 @ 08:15:10.000	0	7.9.2	ap21.uc.osg-htc.org:1095
Jul 9, 2024 @ 08:15:09.000	0	7.8.1	ap21.uc.osg-htc.org:1095
Jul 9, 2024 @ 08:15:07.000	0	7.8.1	ap21.uc.osg-htc.org:1095

Make a Grafana dashboard

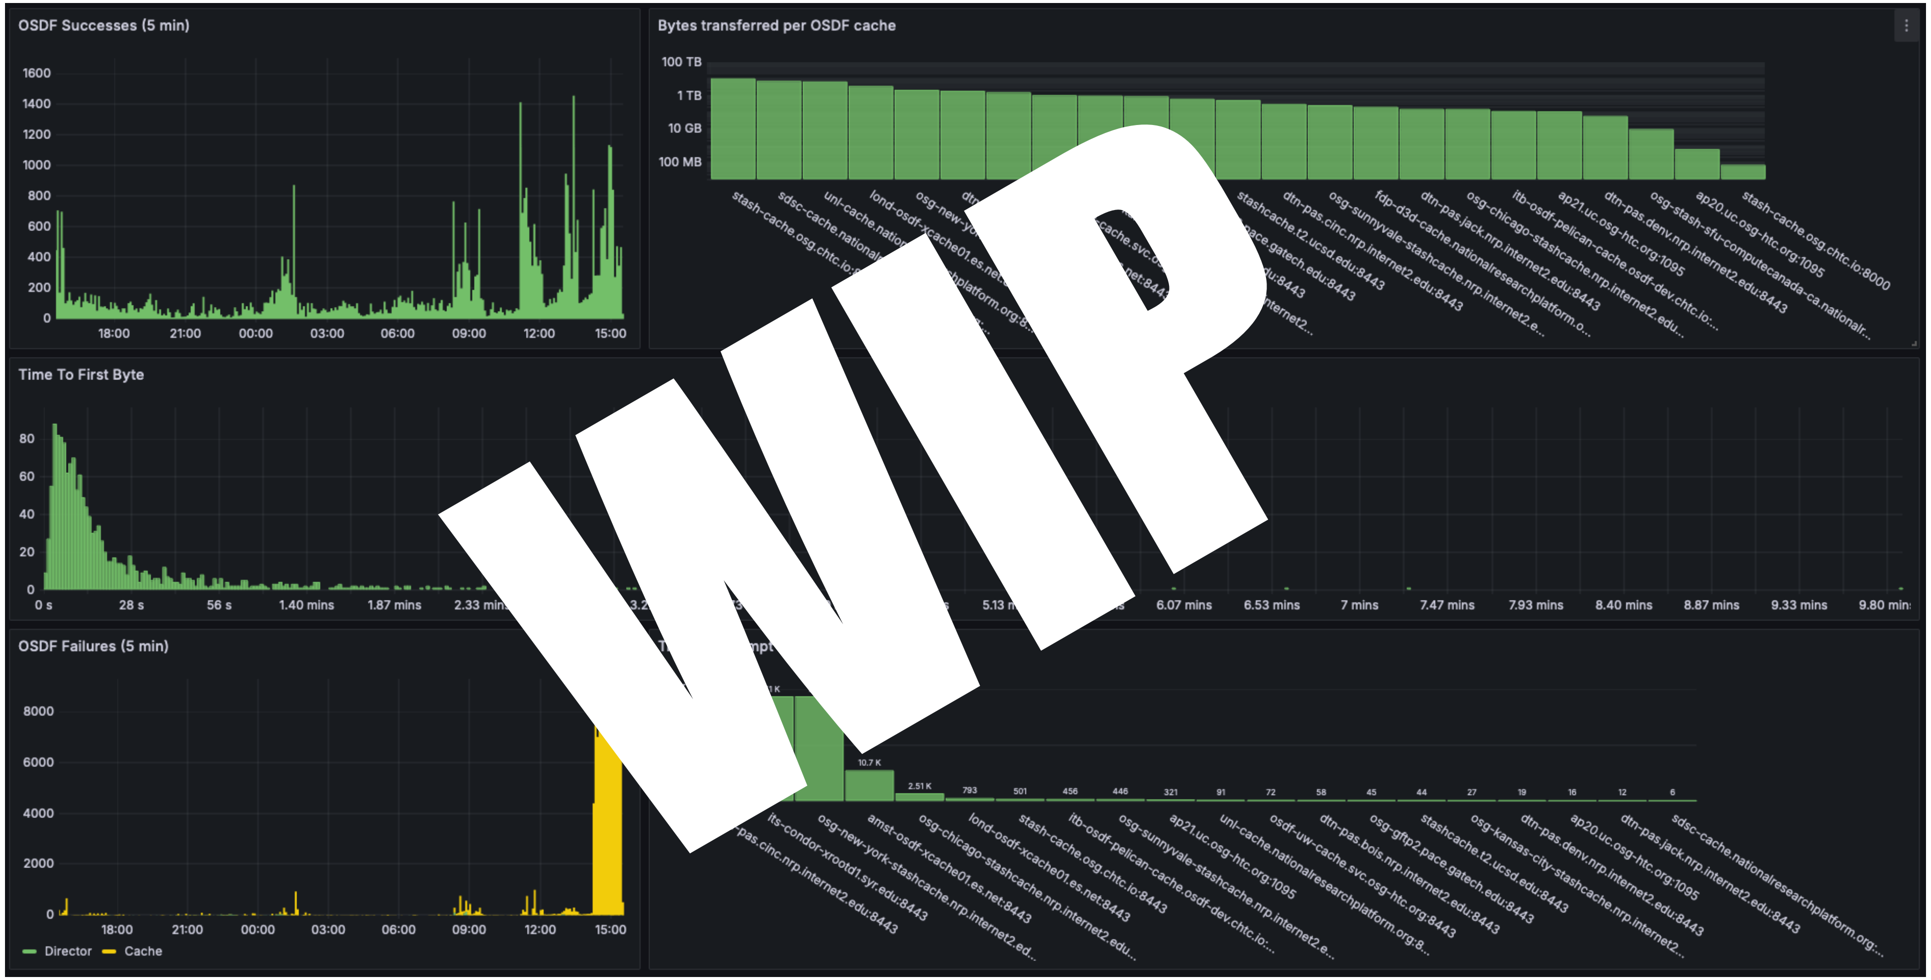
OSDF Successes (5 min)
1600
1400
1200
1000
800
600
400
200
0
18:00
21:00
00:00
03:00
06:00
09:00
12:00
15:00
Bytes transferred per OSDF cache
100 TB
1 TB
10 GB
100 MB
stash-cache.osg.chtc.io:8...
sdsc-cache.nationalr...
unl-cache.nation...
lond-osdf-xcache01.es.net...
osg-new-yo...
dtn...
stashcache.t2.ucsd.edu:8443
dtn-pas.cinc.nrp.internet2.edu:8443
osg-sunnyvale-stashcache.nrp.internet2.e...
fdp-d3d-cache.nationalresearchplatform.o...
dtn-pas.jack.nrp.internet2.edu:8443
osg-chicago-stashcache.nrp.internet2.edu...
itb-osdf-pelican-cache.osdf-dev.chtc.io:...
ap21.uc.osg-htc.org:1095
dtn-pas.denv.nrp.internet2.edu:8443
osg-stash-sfu-computecanada-ca.nationalr...
ap20.uc.osg-htc.org:1095
stash-cache.osg.chtc.io:8000
Time To First Byte
80
60
40
20
0
0 s
28 s
56 s
1.40 mins
1.87 mins
2.33 mins
5.13 mins
6.07 mins
6.53 mins
7 mins
7.47 mins
7.93 mins
8.40 mins
8.87 mins
9.33 mins
9.80 mins
OSDF Failures (5 min)
8000
6000
4000
2000
0
18:00
21:00
00:00
03:00
06:00
09:00
12:00
15:00
Director
Cache
10.7 K
2.51 K
793
501
456
446
321
91
72
58
45
44
27
19
16
12
6
...pas.cinc.nrp.internet2.edu:8443
its-condor-xrootd1.syr.edu:8443
osg-new-york-stashcache.nrp.internet2.ed...
amst-osdf-xcache01.es.net:8443
osg-chicago-stashcache.nrp.internet2.edu...
lond-osdf-xcache01.es.net:8443
stash-cache.osg.chtc.io:8443
itb-osdf-pelican-cache.osdf-dev.chtc.io:...
osg-sunnyvale-stashcache.nrp.internet2.e...
ap21.uc.osg-htc.org:1095
unl-cache.nationalresearchplatform.org:8...
osdf-uw-cache.svc.osg-htc.org:8443
dtn-pas.bois.nrp.internet2.edu:8443
osg-gftp2.pace.gatech.edu:8443
stashcache.t2.ucsd.edu:8443
osg-kansas-city-stashcache.nrp.internet2...
dtn-pas.denv.nrp.internet2.edu:8443
ap20.uc.osg-htc.org:1095
dtn-pas.jack.nrp.internet2.edu:8443
sdsc-cache.nationalresearchplatform.org:...
WIP

Interesting findings so far

Director outages

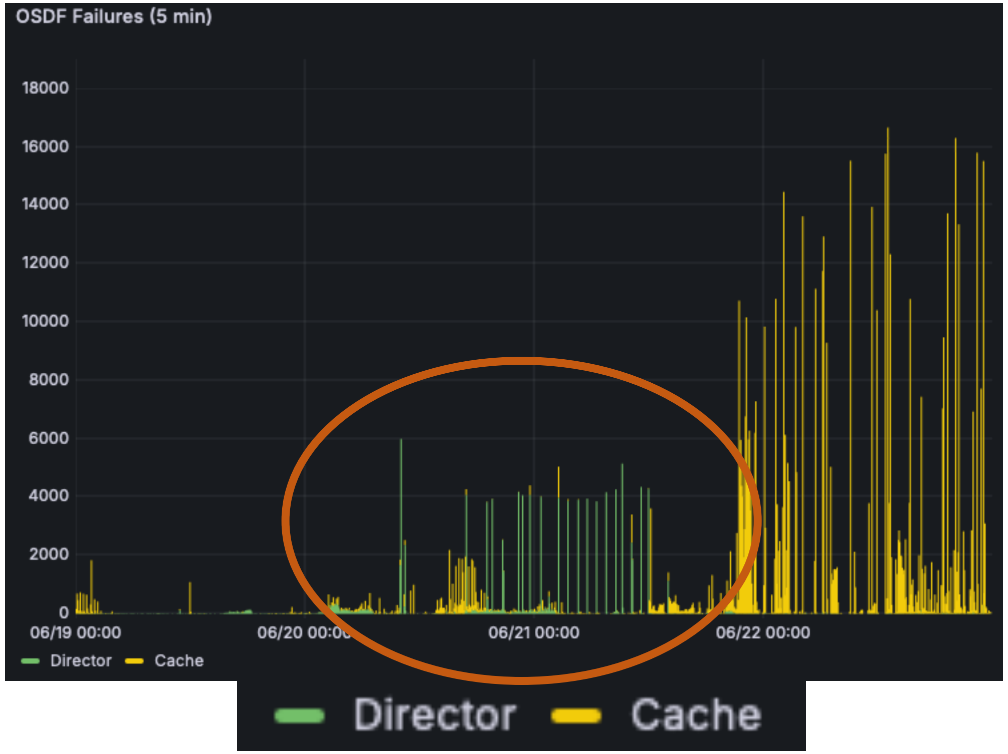

Grafana histograms seem a little off

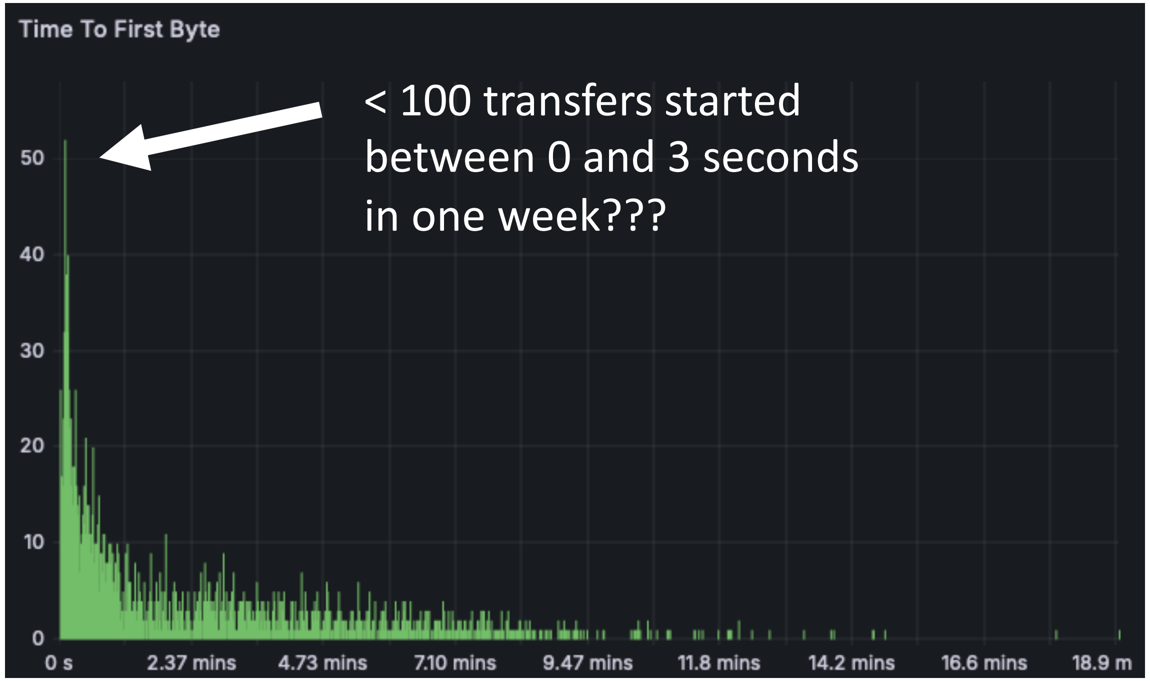

< 100 transfers started
between 0 and 3 seconds
in one week???

800,000 transfers in < 2 seconds

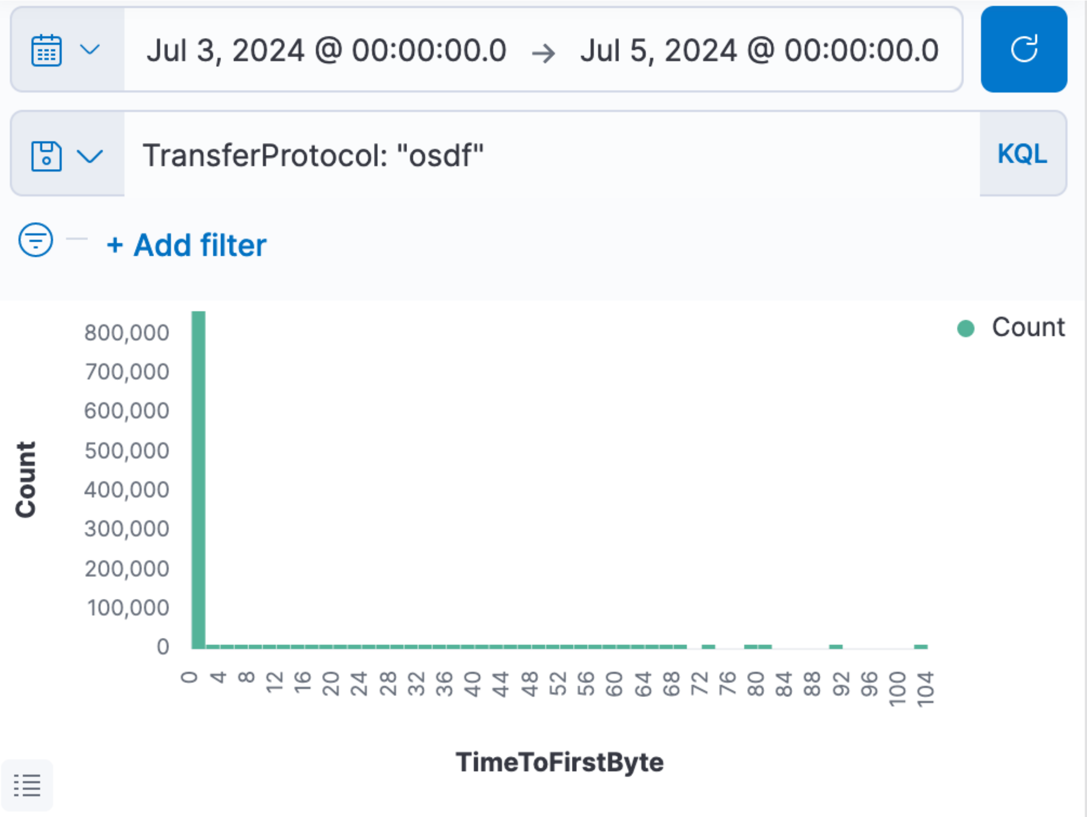
Jul 3, 2024 @ 00:00:00.0 → Jul 5, 2024 @ 00:00:00.0

TransferProtocol: "osdf"

~30 transfers in < 2 seconds

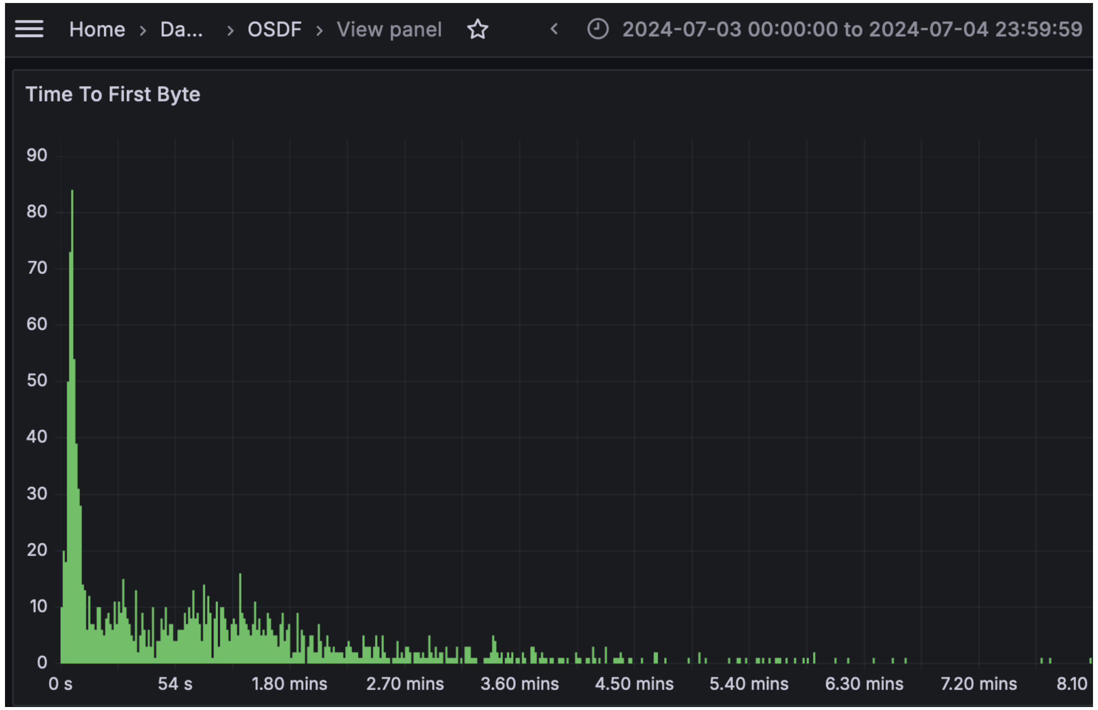
Home › Da... › OSDF › View panel

2024-07-03 00:00:00 to 2024-07-04 23:59:59

Next steps

- Fixing adstash bugs!
- Tweaking and validating Grafana dashboard metrics
- Detect trends and alert OSDF operators?
- Provide easier history access (htcondor CLI tool?) to answer researcher questions about errors and where their files are?
- Let us know what else you would like to see!

What if I administer an origin?

- Metrics pushed to Prometheus, which is also compatible with Grafana
- See Patrick’s lightning talk later today!

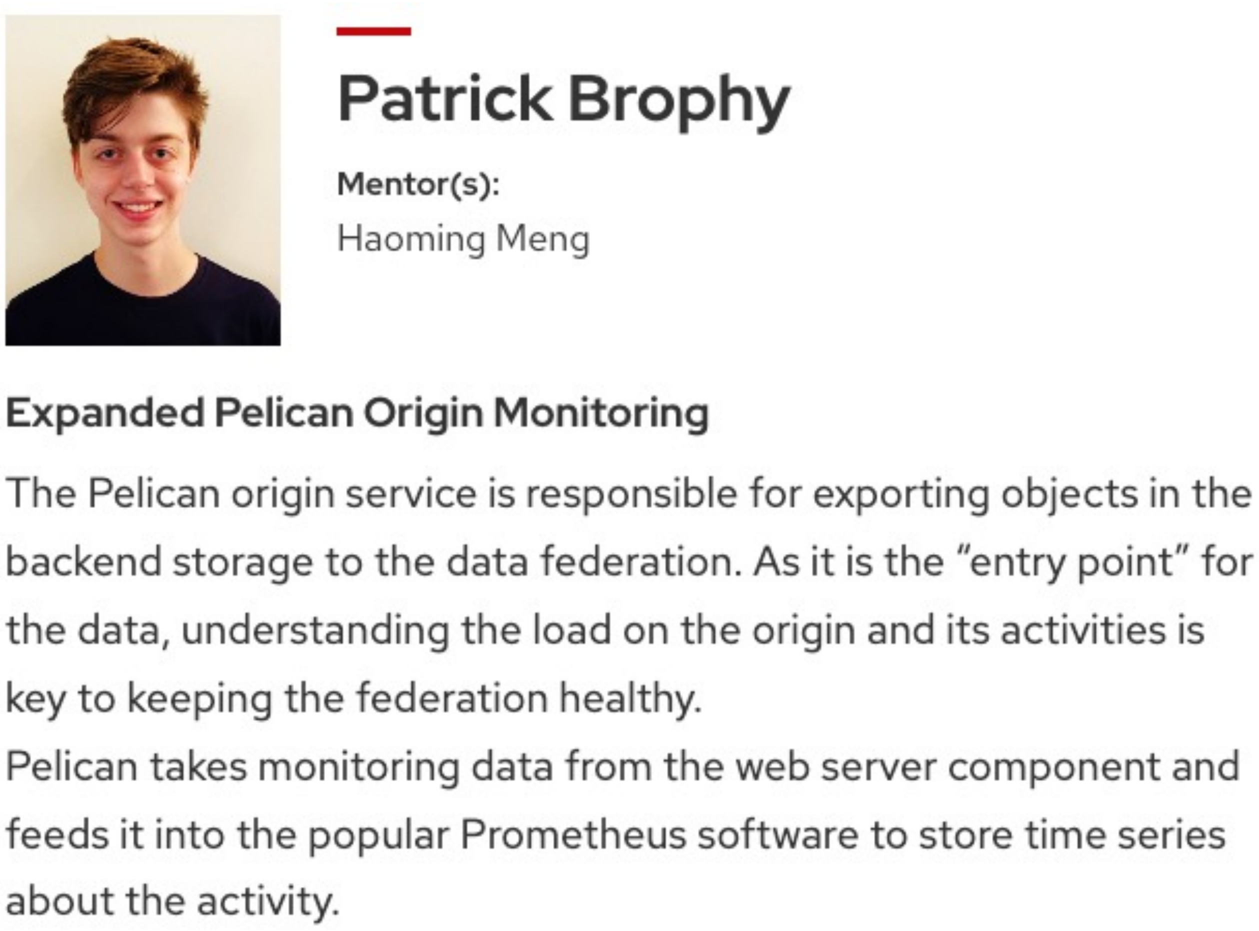

Patrick Brophy

Mentor(s):
Haoming Meng

Expanded Pelican Origin Monitoring

The Pelican origin service is responsible for exporting objects in the backend storage to the data federation. As it is the “entry point” for the data, understanding the load on the origin and its activities is key to keeping the federation healthy.

Pelican takes monitoring data from the web server component and feeds it into the popular Prometheus software to store time series about the activity.

"Data in Flight - Delivering Data with Pelican"

Wednesday, "Complex Workflows Track", 1:30 PM

Interested in contributing to the OSDF? This tutorial will guide you through how to use Pelican to connect your data to an OSDF-like system.

To participate in the hands-on portion of the tutorial, you must register at **`go.wisc.edu/cfs143`** ***before*** **end-of-day Tuesday**

Registration link and more information is available in the "session details" page in the schedule.

Questions?

jpatton@cs.wisc.edu

This project is supported by the National Science Foundation under Cooperative Agreements OAC-2030508 and OAC-2331480. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

CHTC HTCondor Software Suite PATh

Bonus slide!

- Turn on job epoch history

```
JOB_EPOCH_HISTORY = $(SPOOL)/epoch_history
```

- If you are administering a glidein pool, consider additional attributes in your job and transfer plugin ClassAds:

```
SYSTEM_JOB_MACHINE_ATTRS = $(SYSTEM_JOB_MACHINE_ATTRS), Name,
GLIDEIN_ResourceName, GLIDEIN_Site
```

```
TRANSFER_JOB_ATTRS = $(TRANSFER_JOB_ATTRS), MachineAttrName0,
MachineAttrGLIDEIN_ResourceName0, MachineAttrGLIDEIN_Site0
```