

Enabling Microsoft OneDrive Integration with HTCondor

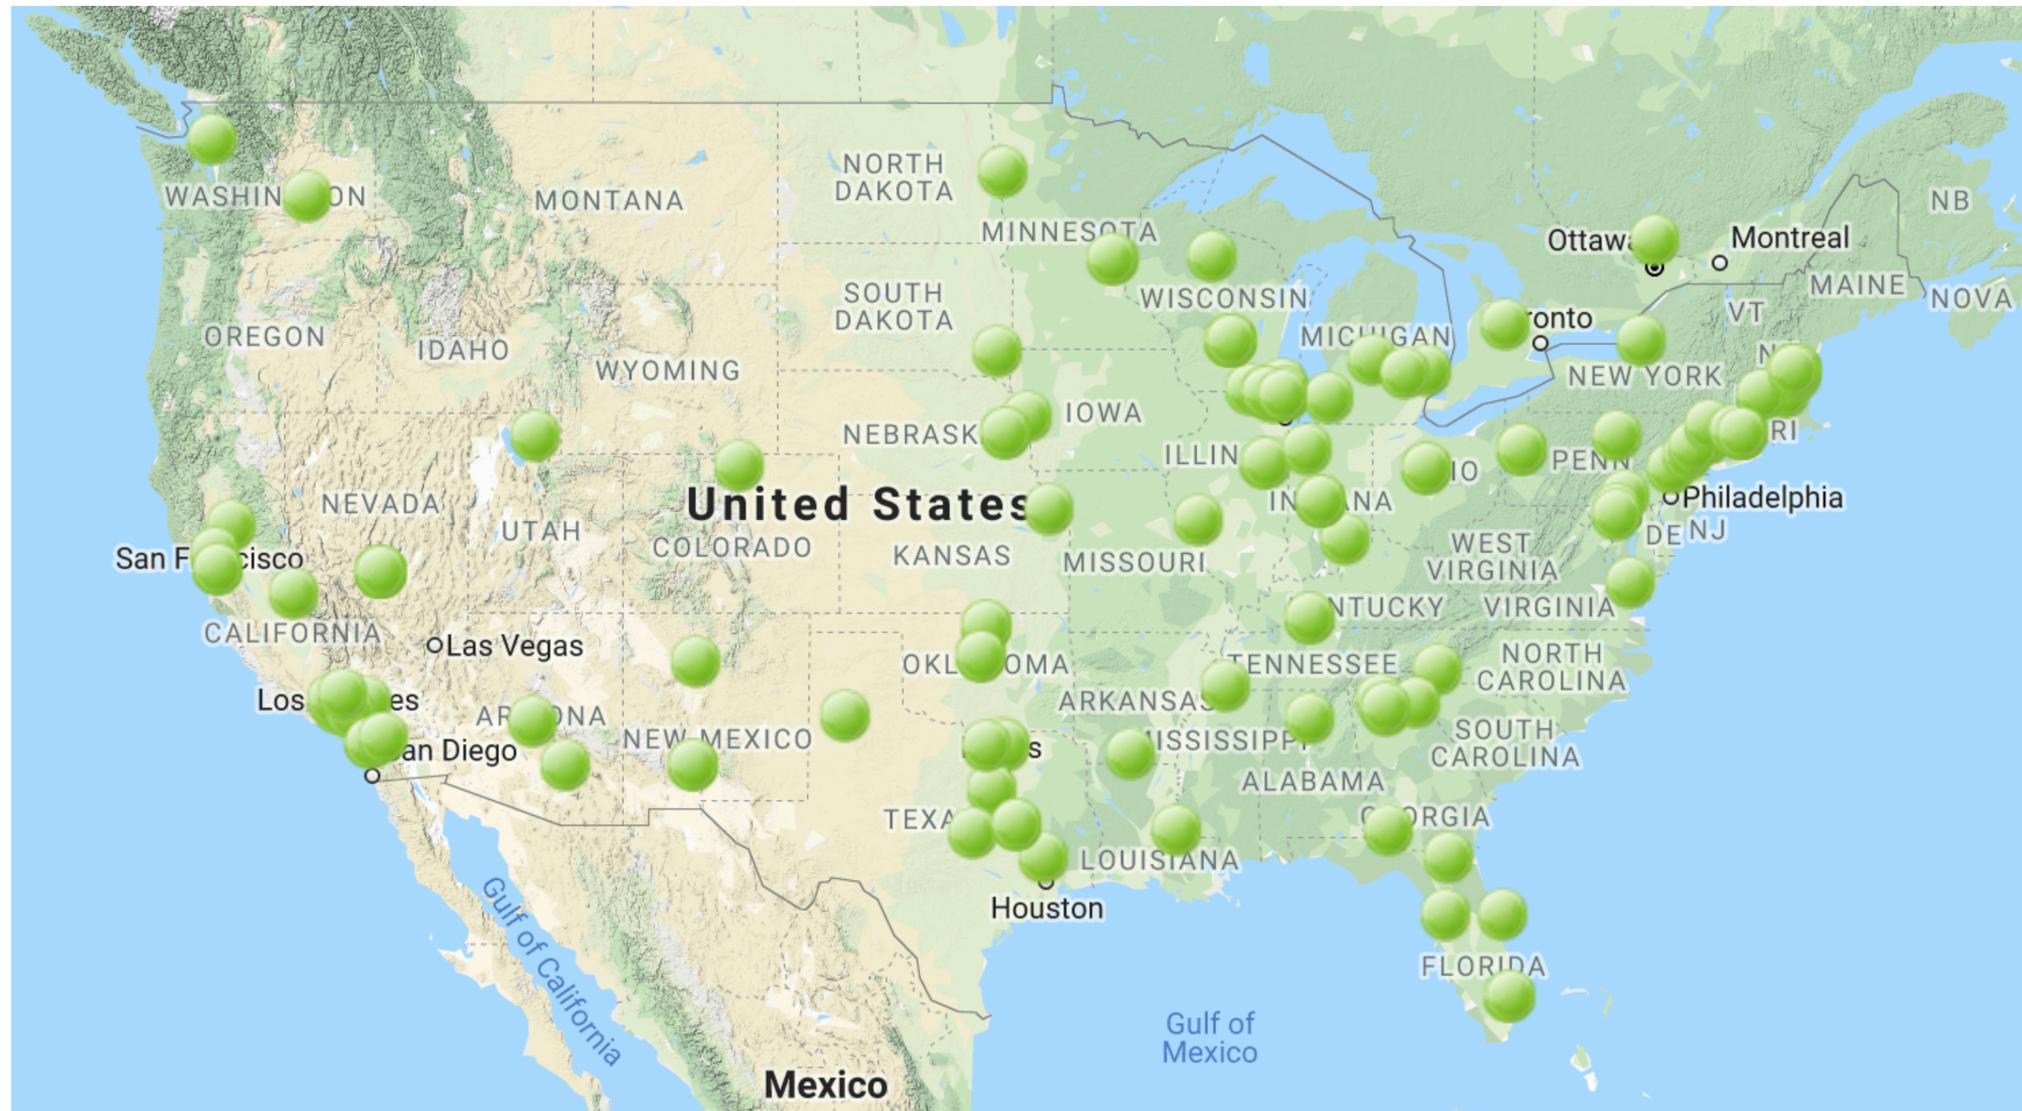
Derek Weitzel
University of Nebraska—Lincoln

Imagine a Scenario

- Jane the biologist wants to run BLAST
- She can run it on your laptop, but it'll take weeks to complete the queries.
- She chooses to run the jobs on cyberinfrastructure like the Open Science Grid

DHTC on the OSG

- The Open Science Grid provides DHTC services to hundreds of researchers across 100+ sites



Data...

- Jane's data is on her laptop
- She could use several different tools to transfer the data to the submit node
- But every time she updates her queries or database, she will have to re-transfer

Data Transfers to the Worker Node

- Transfer from the submit host to the WN is simple, but is bottlenecked
- Many services provide data transfer capabilities
 - Globus provides data server to data server transfers
 - Aspera, also data server to data server
- Neither provide data distribution to the execution hosts
- Need a CDN like service to deliver data to WN

Comparison with HPC

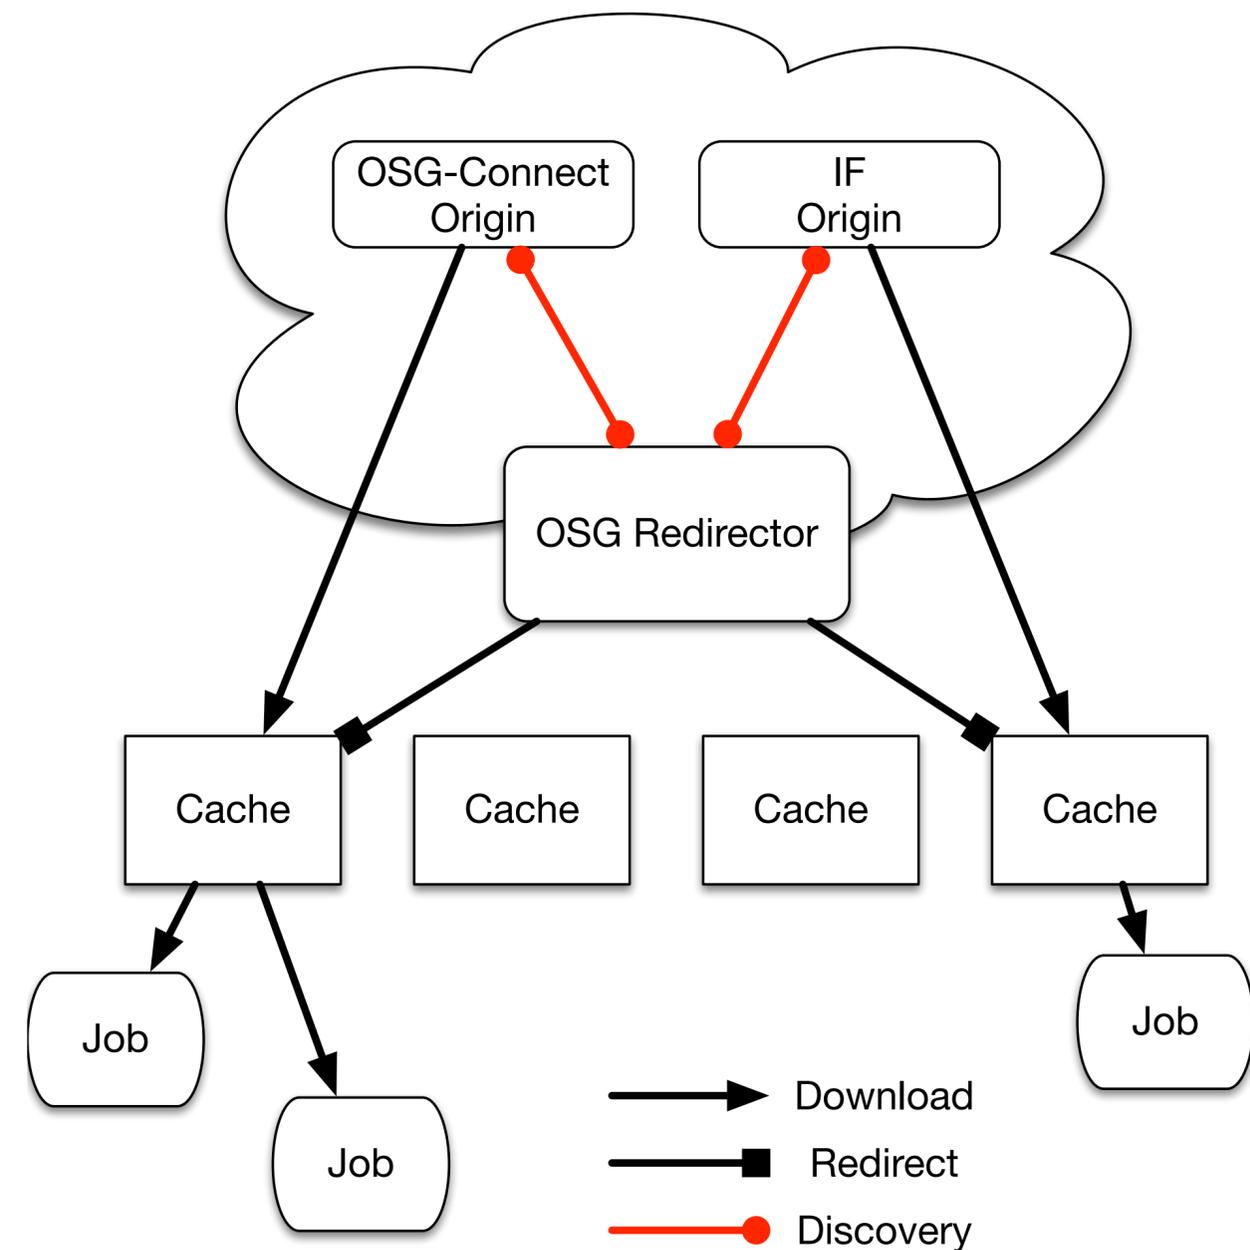
- With HPC, there is (usually) a shared filesystem to distribute data
- Transfer data from your laptop to the cluster, then use the filesystem to distribute the data.
- With distributed high throughput computing (DHTC), no shared filesystems are available.

OSG Methods for Data Distribution

- StashCache
 - Regionally distributed Caches spread throughput the OSG
- Proxied HTTP
 - Each site has a HTTP cache for requests

StashCache

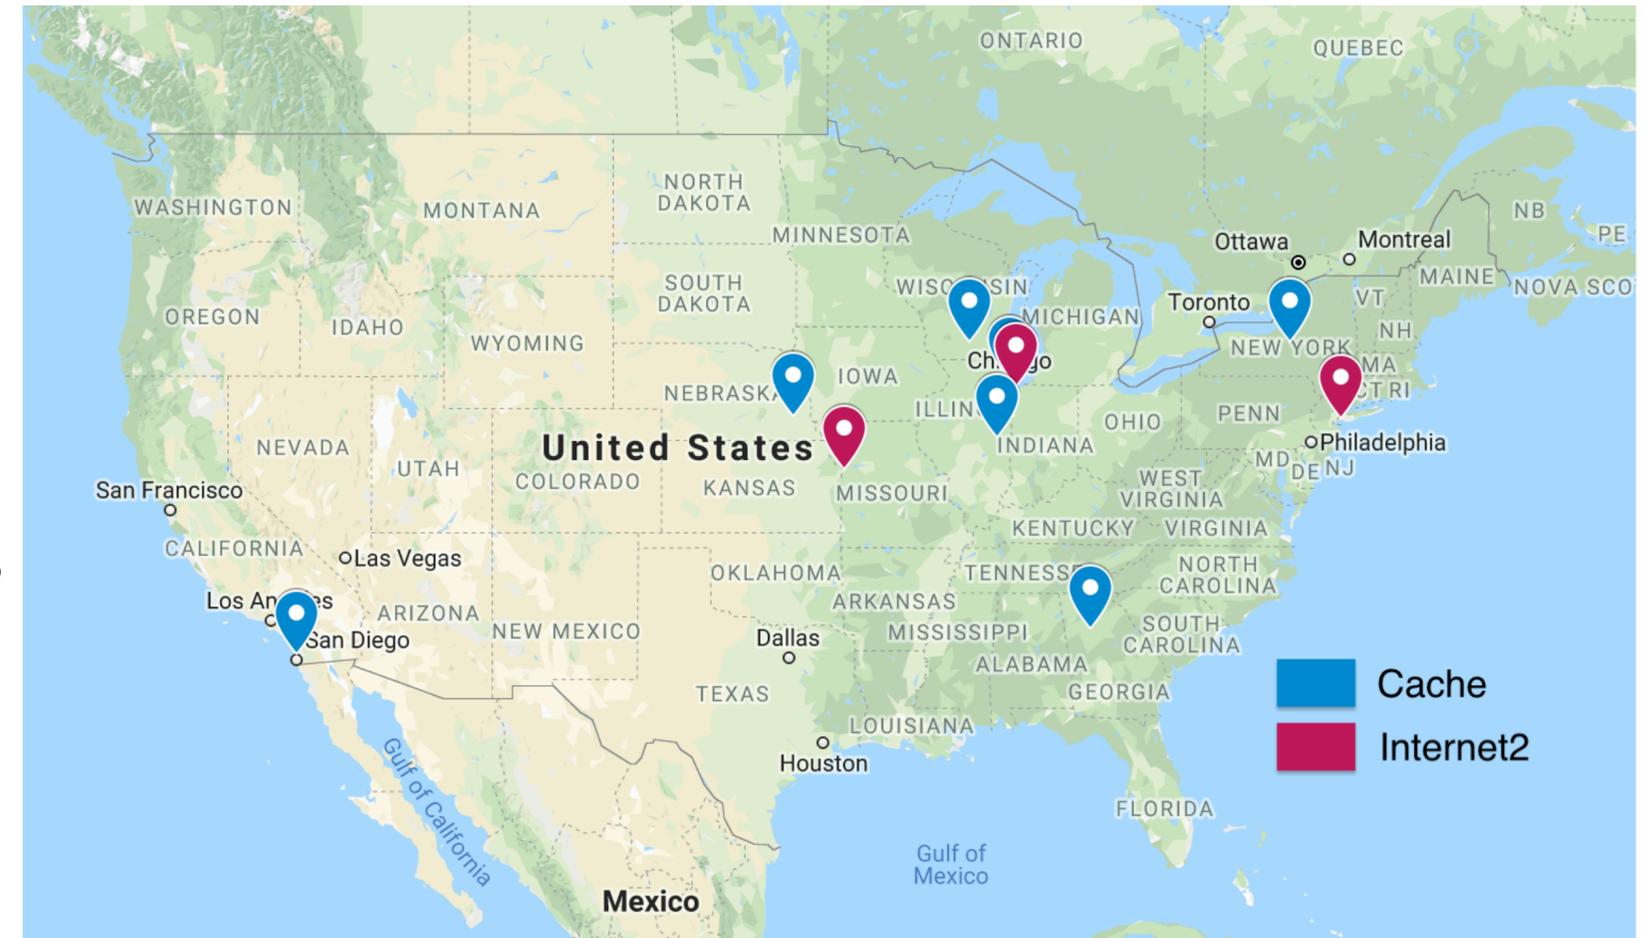
- Copy data to an Origin
- Everything is public, no credentials
- See StashCache paper here at PEARC19



Weitzel, D., Zvada, M., Vukotic, I., Gardner, R., Bockelman, B., Rynge, M., Hernandez, E.F., Lin, B. and Selmecci, M., 2019. StashCache: A Distributed Caching Federation for the Open Science Grid. In *Practice and Experience in Advanced Research Computing (PEARC 19)* (accepted) doi:< 10.1145/3332186.3332212

StashCache

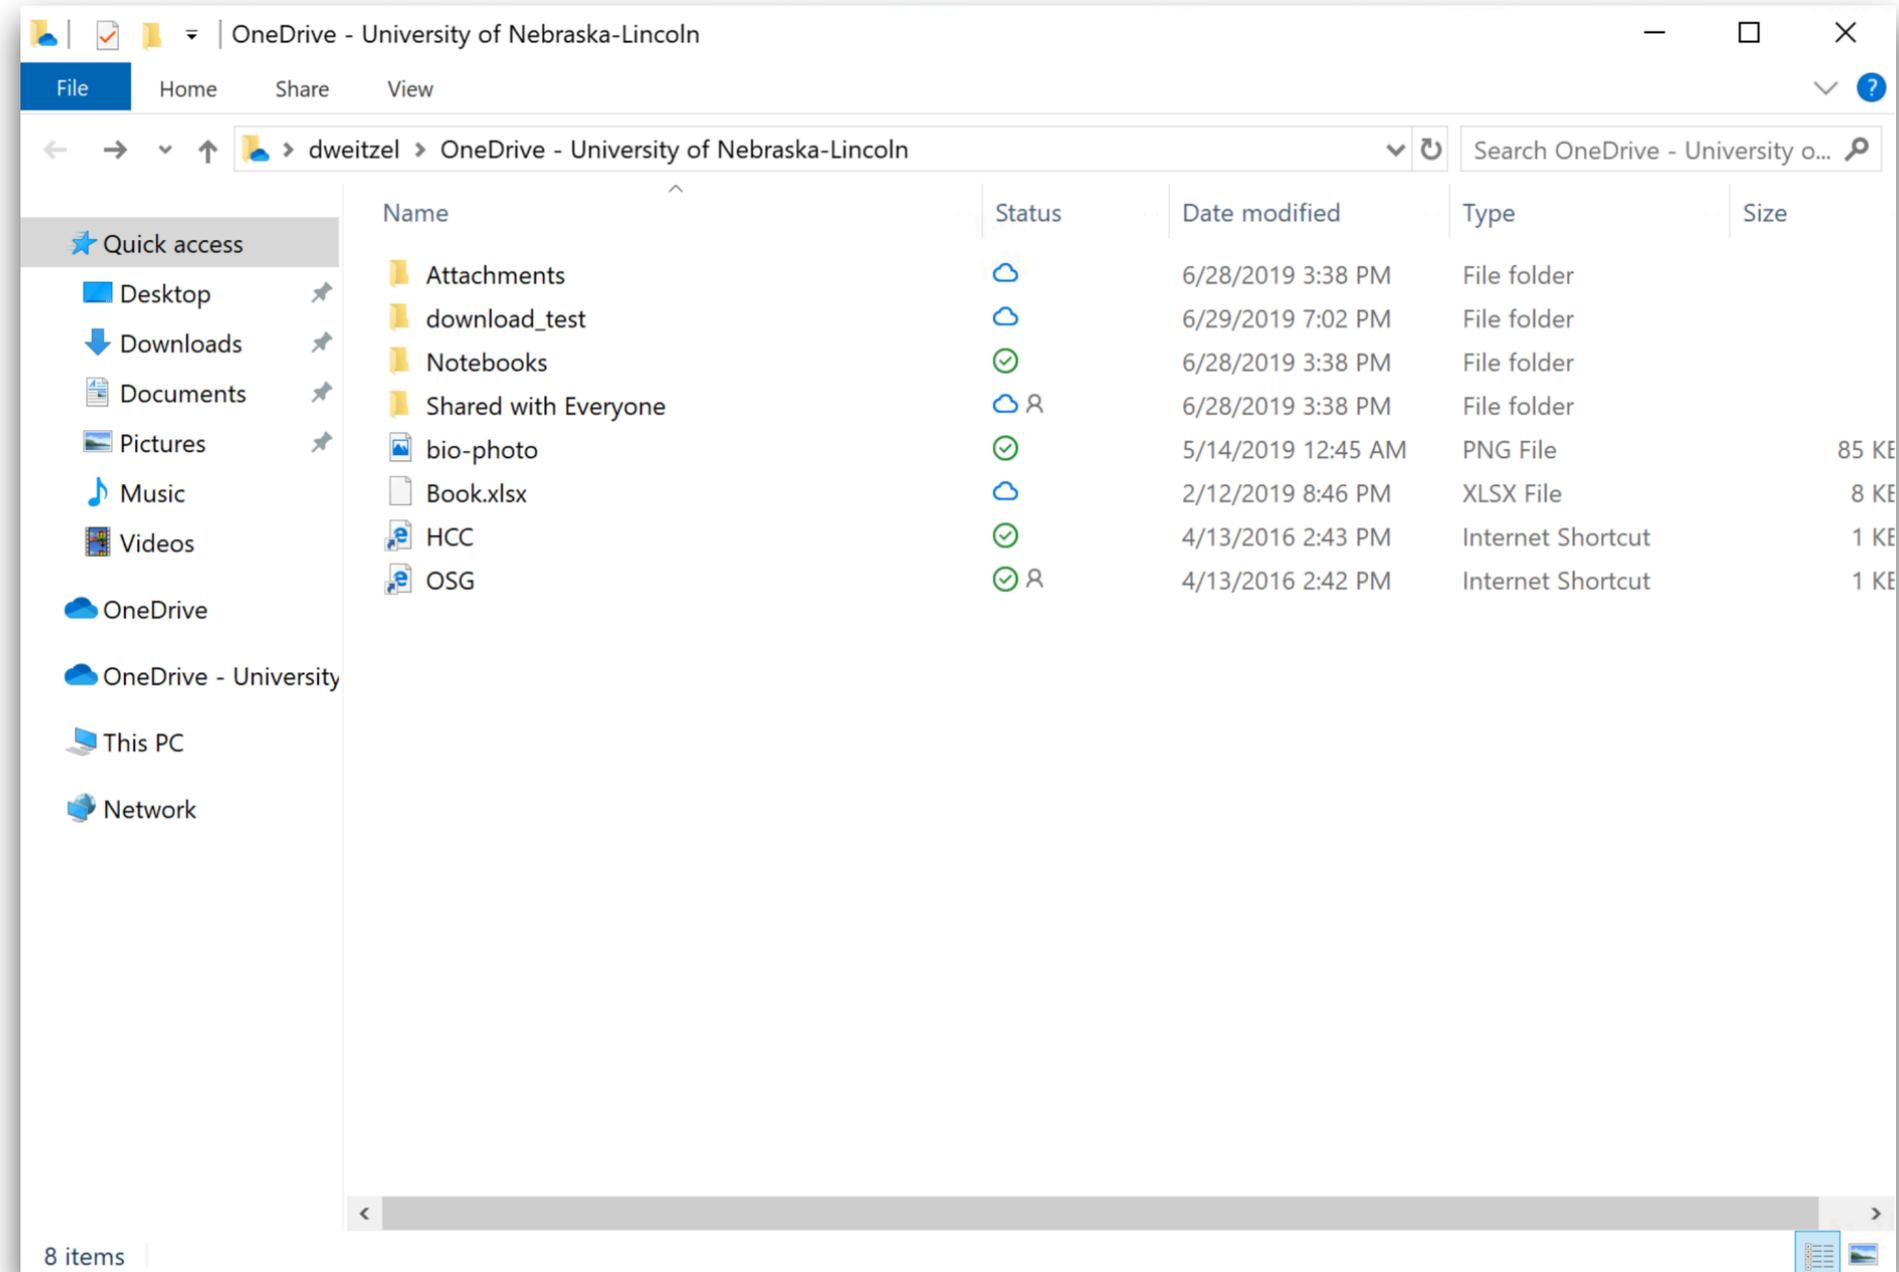
- Copy data to an Origin
- Everything is public, no credentials
- See StashCache paper here at PEARC19



Weitzel, D., Zvada, M., Vukotic, I., Gardner, R., Bockelman, B., Rynge, M., Hernandez, E.F., Lin, B. and Selmecci, M., 2019. StashCache: A Distributed Caching Federation for the Open Science Grid. In *Practice and Experience in Advanced Research Computing (PEARC 19)* (accepted) doi:< 10.1145/3332186.3332212

OneDrive

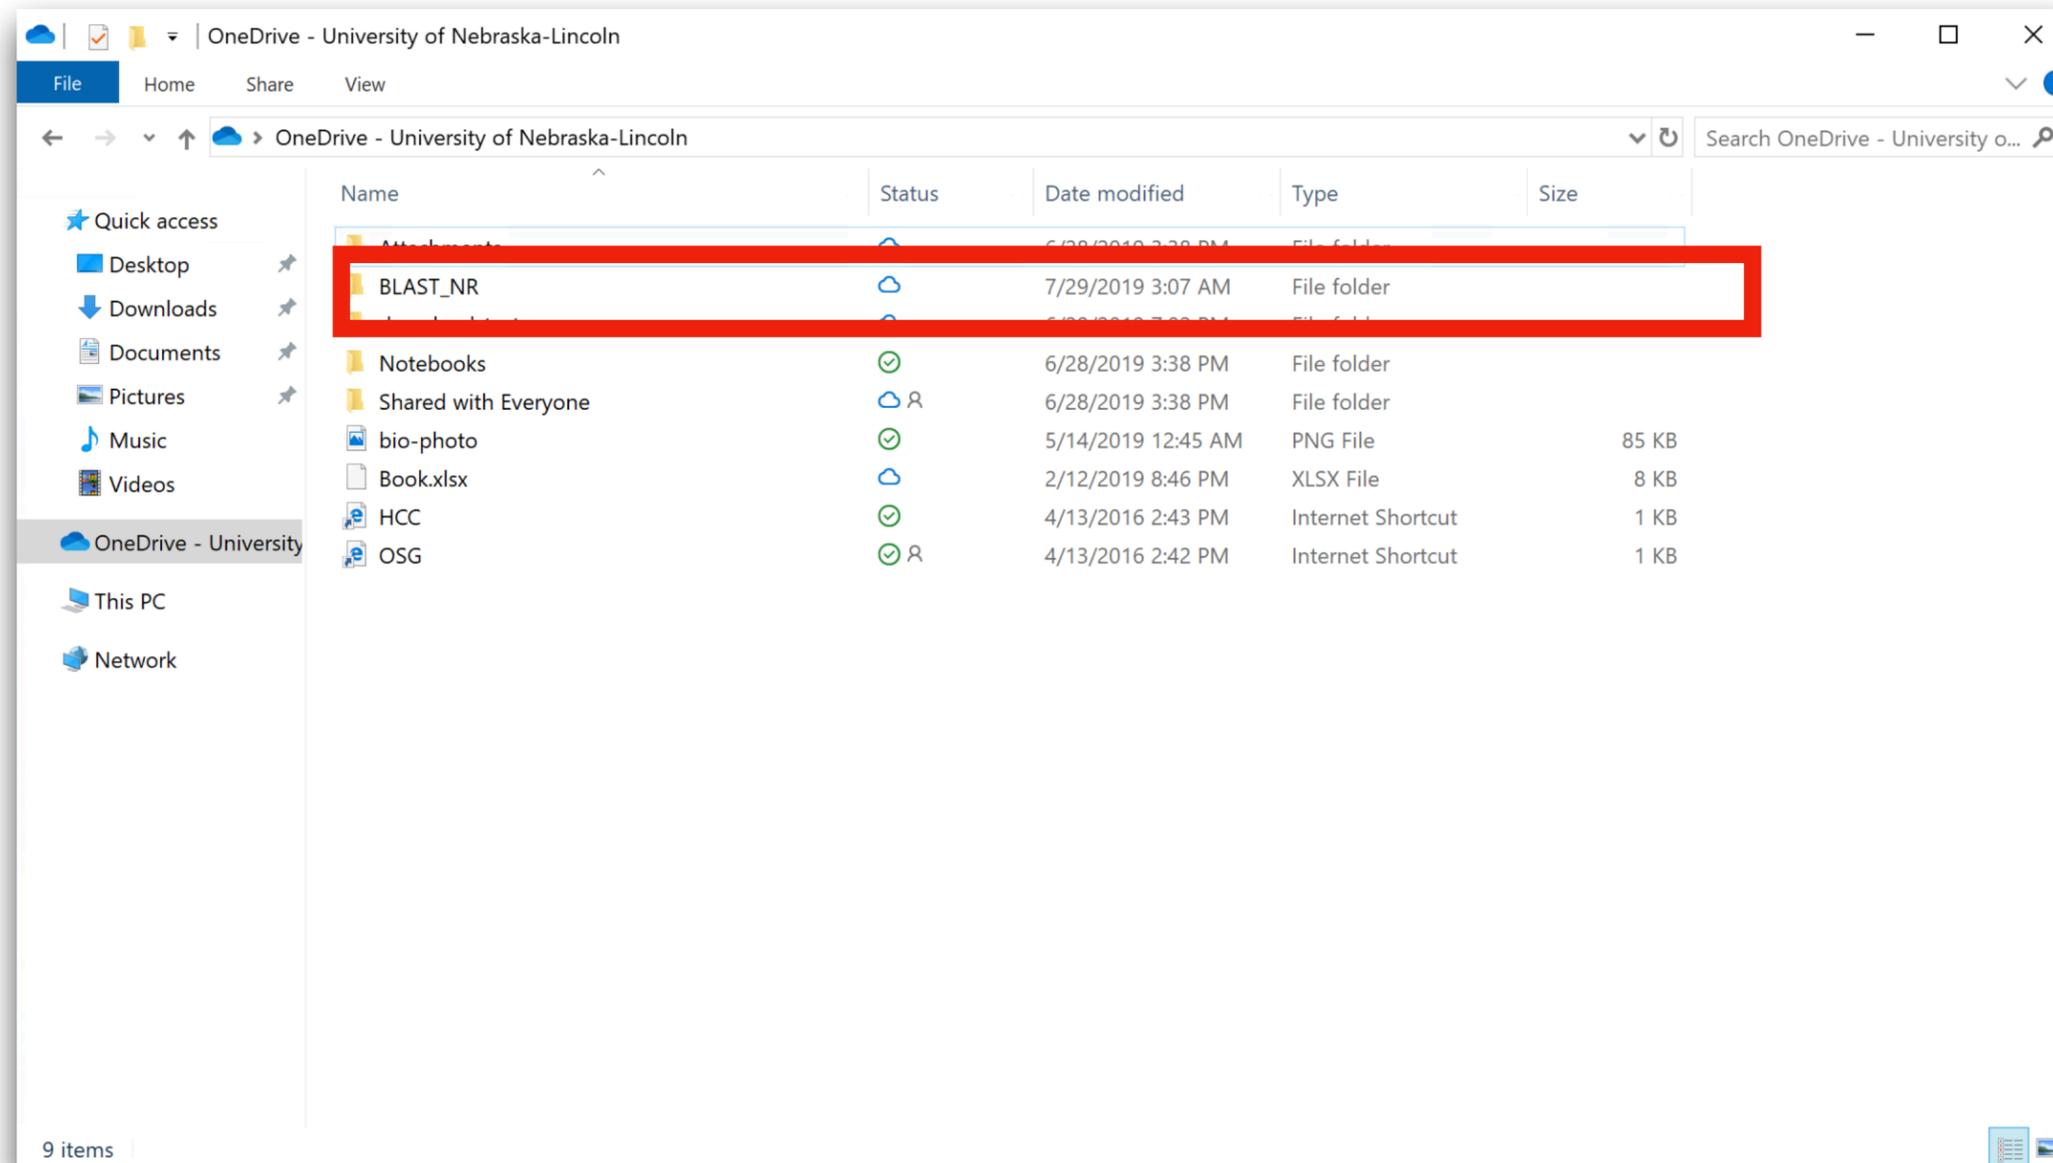
- Cloud Storage folder similar to DropBox, Box...
- Presented as a directory on your laptop (here's mine!)



OneDrive

OneDrive

- Jane can upload BLAST DB to OneDrive



HTCondor Integration

- HTCondor handles OneDrive credentials
- Transparently transfers data from OneDrive for the job
- Submit file has special syntax to specify OneDrive requirement

```
executable = run_blast.sh
```

```
output = out
```

```
error = err
```

```
log = log
```

```
should_transfer_files = YES
```

```
when_to_transfer_output = ON_EXIT
```

```
transfer_input_files = onedrive:///blast_nr.tar.gz
```

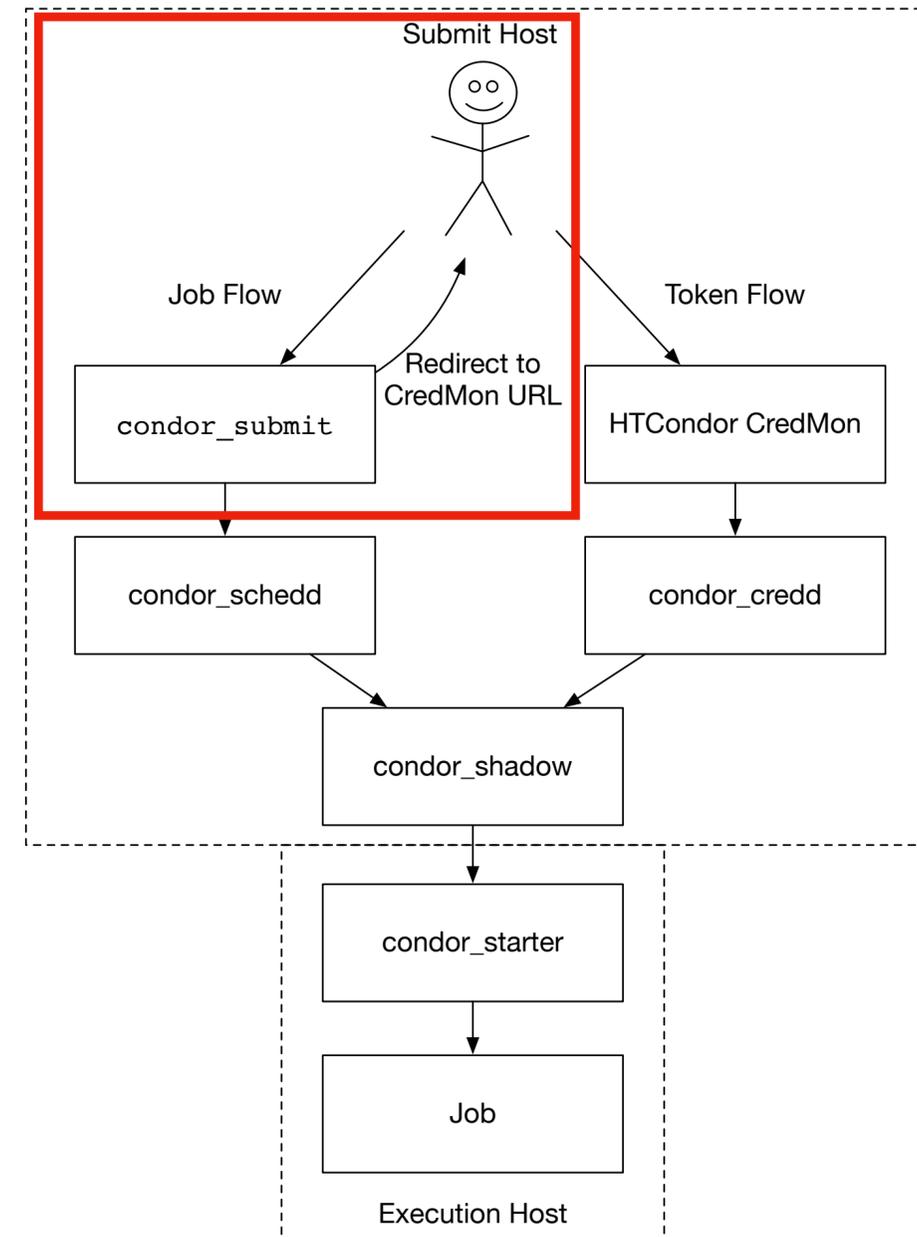
```
use_oauth_services = onedrive
```

```
onedrive_oauth_permissions = User.Read
```

```
queue
```

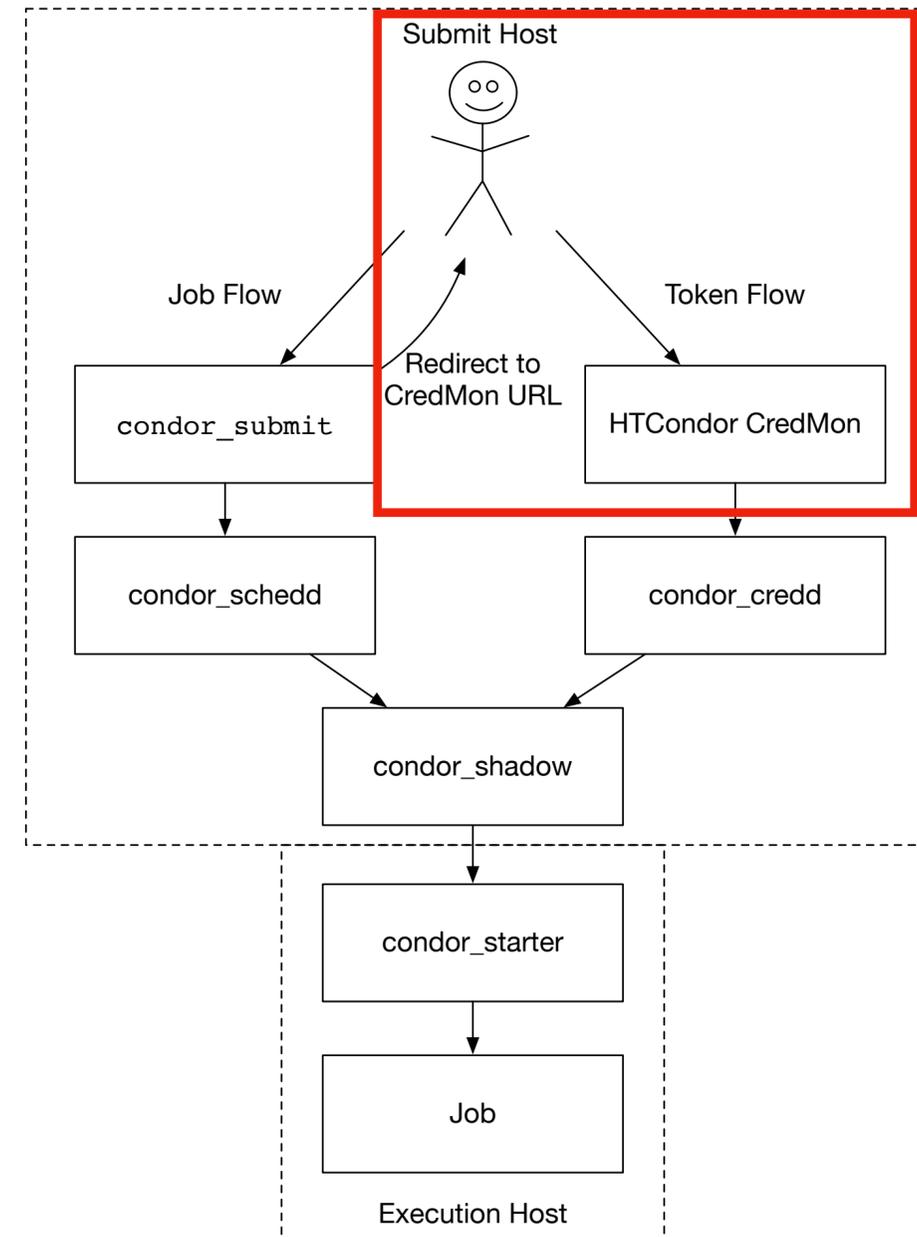
Credential Handling

1. User Submits a job
2. First Time:
 1. It returns a special URL to connect to the HTCCondor CredMon
 2. User goes to the HTCCondor CredMon



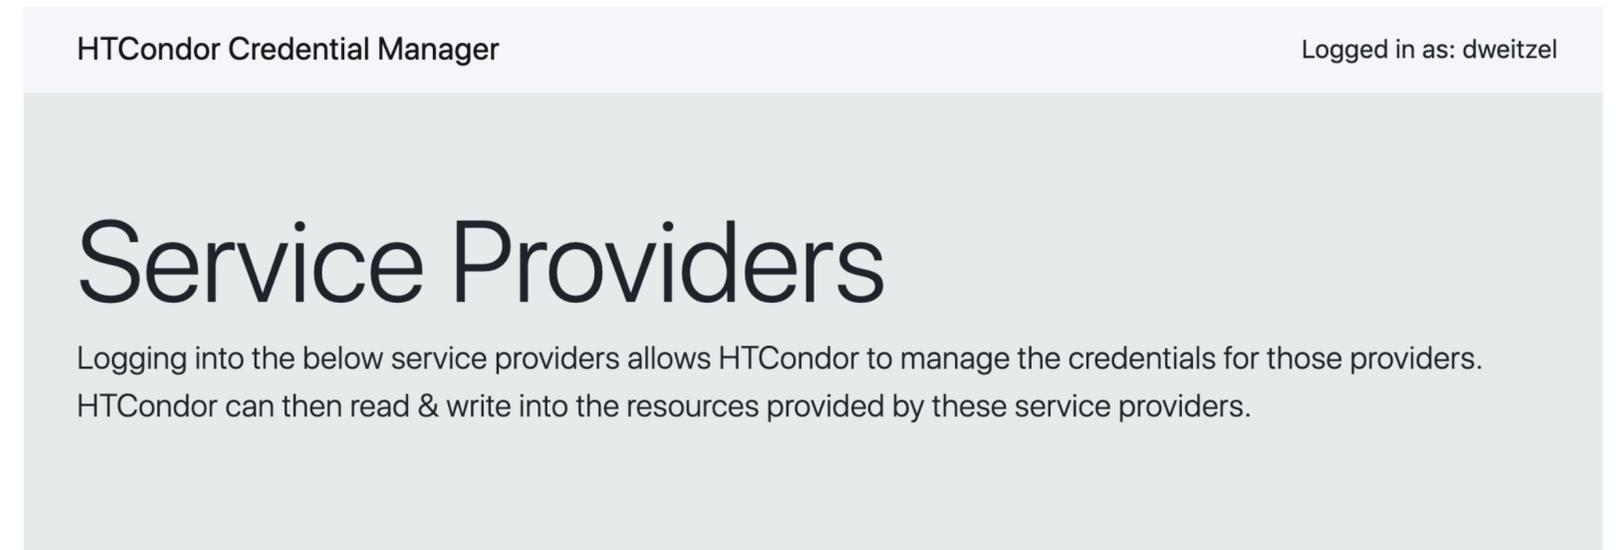
Credential Handling

1. User Submits a job
2. First Time:
 1. It returns a special URL to connect to the HTCCondor CredMon
 2. User goes to the HTCCondor CredMon



HTCondor Credential Manager

- Login with your OneDrive account
- Goes through the OAuth flow to get a Token
- The Token is short lived (20 min)
- The Credential Manager refreshes the token



HTCondor Credential Manager Logged in as: dweitzel

Service Providers

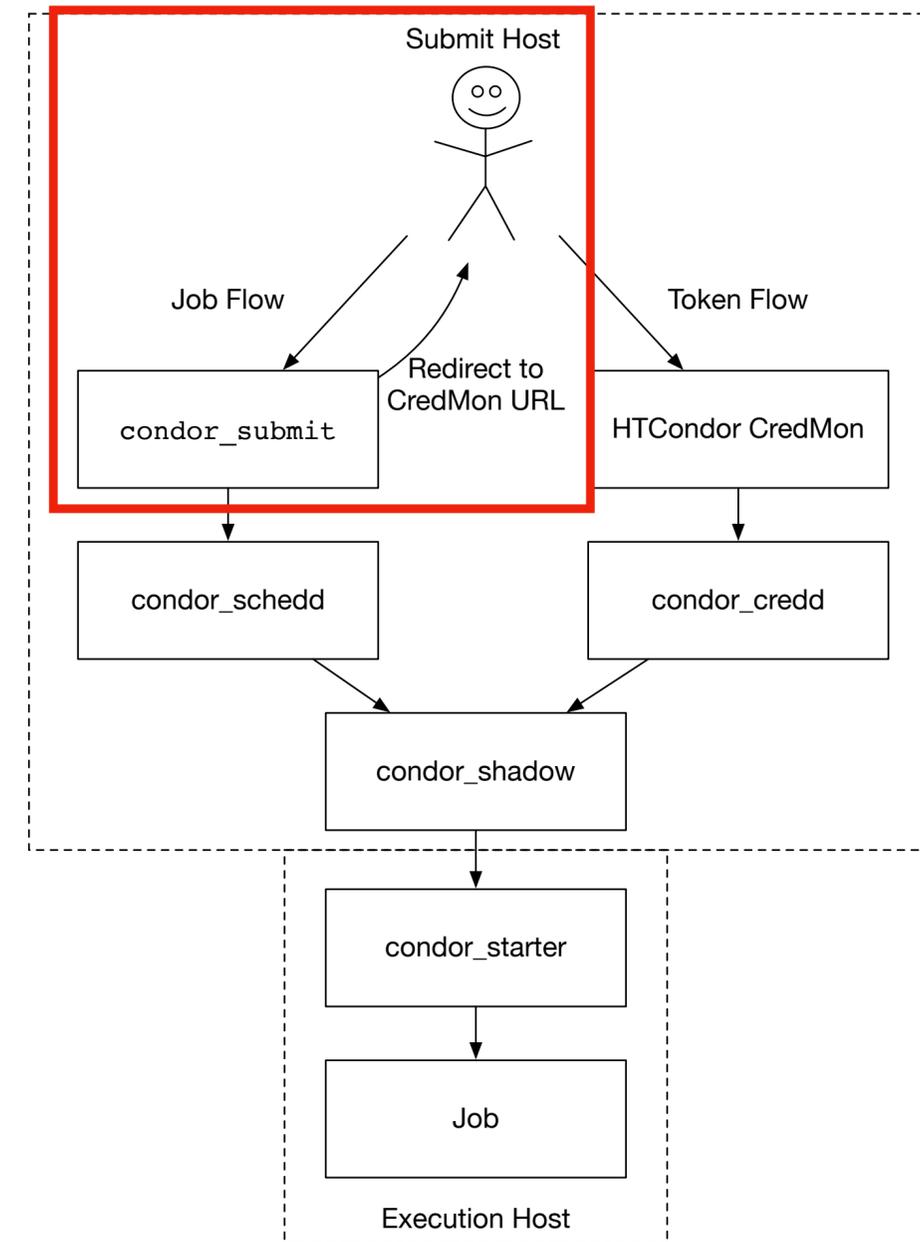
Logging into the below service providers allows HTCondor to manage the credentials for those providers. HTCondor can then read & write into the resources provided by these service providers.

onedrive Login

Login

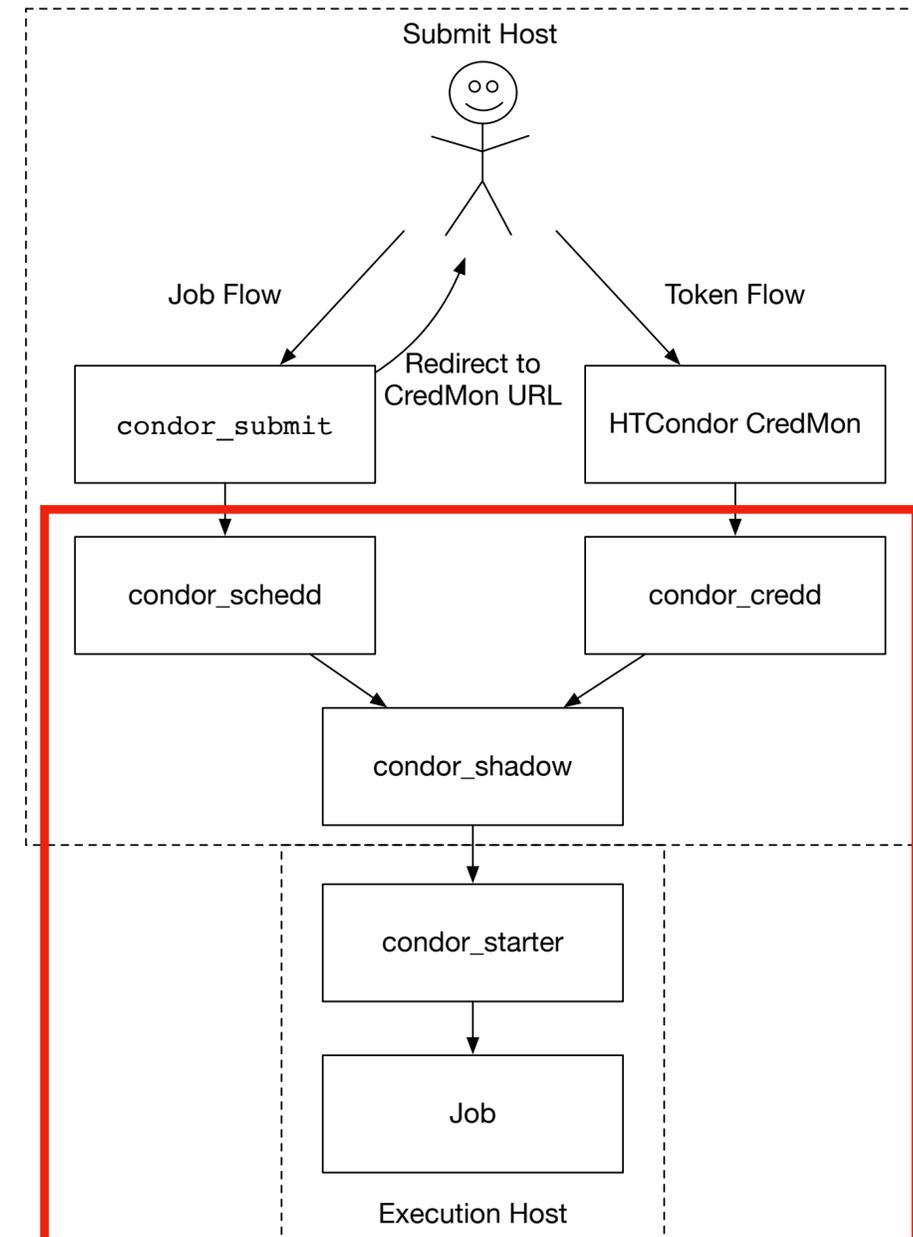
Credential Handling

1. User Submits a job
2. First Time:
 1. It returns a special URL to connect to the HTCCondor CredMon
 2. User goes to the HTCCondor CredMon
3. User submits the job



Credential Handling

1. User Submits a job
2. First Time:
 1. It returns a special URL to connect to the HTCCondor CredMon
 2. User goes to the HTCCondor CredMon
3. User submits the job
4. The credential is sent with the job to the worker node



Experiments

- Compare OneDrive to existing transfer mechanisms to worker nodes.
- Run on multiple real sites on the OSG, not a laboratory setting
- Use the top 4 sites of opportunistic usage, Colorado, Bellarmine, Syracuse, Chicago

Experiments

- Create test framework to download multiple sized files
- Test file sizes are from the percentiles of file sizes used in StashCache by users.

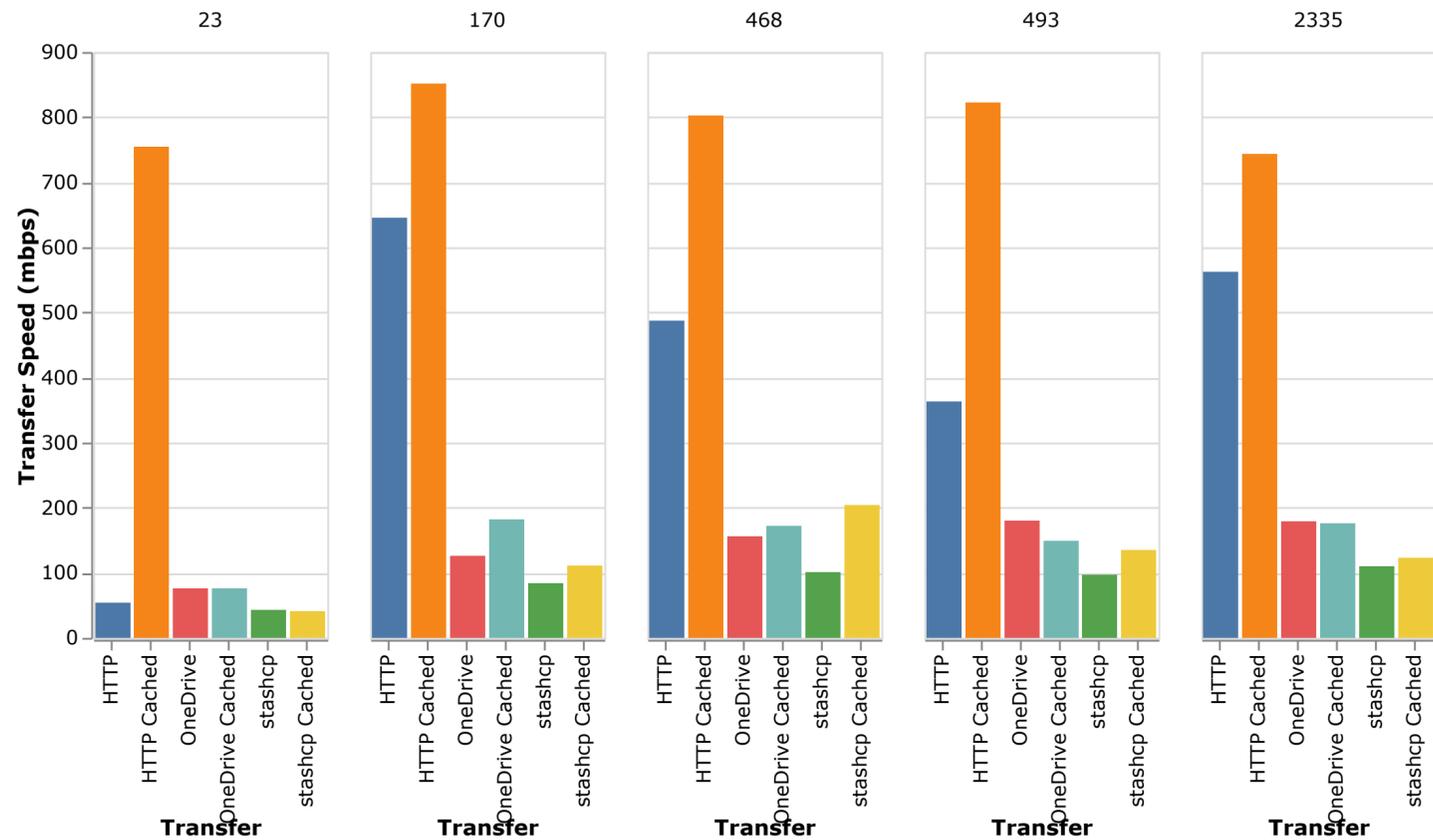
Percentile	Filesize
1	5.797 KB
5	22.801 MB
25	170.131 MB
50	467.852 MB
75	493.337 MB
95	2.335 GB
99	2.335 GB

Results

- Different sites perform significantly different. As an example, here is Colorado and Bellarmine

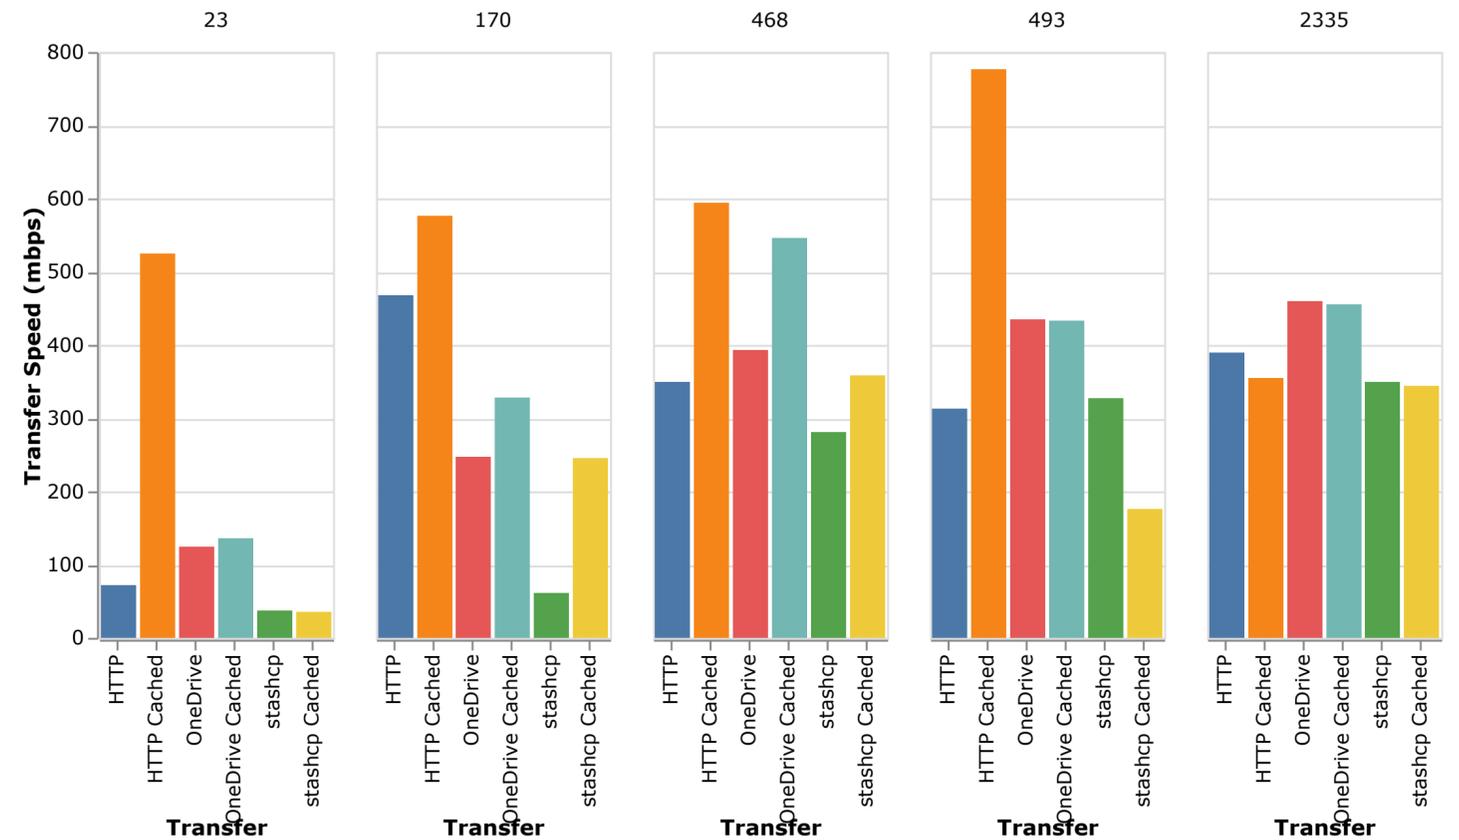
Colorado

File Sizes (MB)



Bellarmino

File Sizes (MB)



Results

- Measure performance. Positive percentage means OneDrive performed faster transfer than the test.

Site	HTTP 22.8MB	HTTP 2.3 GB	StashCache 22.8MB	StashCache 2.3GB
Syracuse	-80.00%	-77.80%	7.30%	-77.60%
Colorado	-89.90%	-75.90%	89.40%	46.10%
Bellarmine	-76.40%	29.70%	248.10%	33.80%
Chicago	-85.50%	4.10%	-41.60%	36.00%
Average	-82.90%	-29.90%	75.80%	9.50%

Results

- In General, OneDrive is slower than HTTP

Site	HTTP 22.8MB	HTTP 2.3 GB	StashCache 22.8MB	StashCache 2.3GB
Syracuse	-80.00%	-77.80%	7.30%	-77.60%
Colorado	-89.90%	-75.90%	89.40%	46.10%
Bellarmine	-76.40%	29.70%	248.10%	33.80%
Chicago	-85.50%	4.10%	-41.60%	36.00%
Average	-82.90%	-29.90%	75.80%	9.50%

Results

- But, OneDrive is faster than StashCache in most cases

Site	HTTP 22.8MB	HTTP 2.3 GB	StashCache 22.8MB	StashCache 2.3GB
Syracuse	-80.00%	-77.80%	7.30%	-77.60%
Colorado	-89.90%	-75.90%	89.40%	46.10%
Bellarmine	-76.40%	29.70%	248.10%	33.80%
Chicago	-85.50%	4.10%	-41.60%	36.00%
Average	-82.90%	-29.90%	75.80%	9.50%

Experience Working with OneDrive

- OneDrive interface can be complicated with Personal Accounts, Institutional Accounts, ...
- OneDrive is included in *most* Office365 agreements with Microsoft (Does your university have one?)
- Although I did not hit any download limits for OneDrive during the test, full production may hit limits
- OneDrive has file size limits of 15GB



OneDrive

Conclusions for Jane

- She can transfer the data to her OneDrive folder
- Data is instantly available for her workflows on the OSG
- Transferring files is easy with HTCondor managing the OneDrive credentials and transfers

Dissemination

- The OneDrive download capabilities are under review for inclusion in HTCondor
- Once included, the capability will be available on every HTCondor installation
- Will be included along with a Box and Google Drive plugins to download files.



Conclusions

- OneDrive provides an easy-to-use interface for researchers
- OneDrive's performance for data distribution is comparable to current data distribution methods