

---

# Benchmarking discussion

Fengping Hu, Lincoln Bryant, Ilija Vukotic  
Enrico Fermi Institute  
University of Chicago



---

IRIS-HEP AGC workshop 2023  
05.05.2023





# Target

---

- Process 200TB in 20 minutes
  - Assume read 10% →  $20\text{TB}/20*60\text{s} = 133\text{Gbps}$
  - Assume each event is 2KB and 10% is streamed →  $133*10^6/8*0.2 = 83.5\text{Mevents/s}$
  - Assume 25 kevents/core/second => needs 3340 cores
  - $133\text{Gbps}/35 \text{ nodes} \rightarrow 3.8\text{Gbps per node}$
  - Each core  $133 \text{ Gpbs}/3340 = 40\text{Mbps}$

# Sanity Check



	Target	UChicago AF
<b>Disk</b>	200 TB	>1 PB
<b>Network</b>	133 Gbps	200 Gbps WAN >200 Gbps LAN
<b>CPU</b>	3340 CPU Cores	4560 CPU Cores (K8S)

- The target is well-scoped for the UChicago AF



# Throughput test setup

---

- Coffea-casa Dask workers read data from ~800 files across 15 fast compute nodes with 25Gbps links
- Each Dask worker requests a few selected columns via:  






```
uproot.open(filename) ["Events"].arrays(arrays_to_read)
```
- Compared three file access methods:
  - Locally, via CephFS shared filesystem:
    - 19 hyperconverged nodes with spinning disks and 10Gbps links
  - Remotely, via XRootD at UNL:
    - <https://xrootd-local.unl.edu:1094//store/user/AGC/nanoAOD/>
  - Remotely, via XRootD with a local XCache at UChicago





# Local filesystem throughput scaling with number of workers

- From left to right the number of workers:  
50 (~1GB/s), 100 (~2GB/s),  
200 (~3GB/s), 400 (5GB/s),  
800 (5GB/s)
- Scales almost linearly in the beginning
- 10–20MBps per worker







# Remote files throughput scaling with number of workers

- Scales linearly from 50, 100 and 200 workers, bottlenecked somewhere from 200 workers to 400 workers
- Total is  $40\text{MBps/node} * 15\text{ nodes} = 600\text{MBps}$
- Per workers is  $600\text{MBps} / 200\text{ workers} = 3\text{MBps/worker}$


c022.af.uchicago.edu




c023.af.uchicago.edu



c026.af.uchicago.edu



c027.af.uchicago.edu





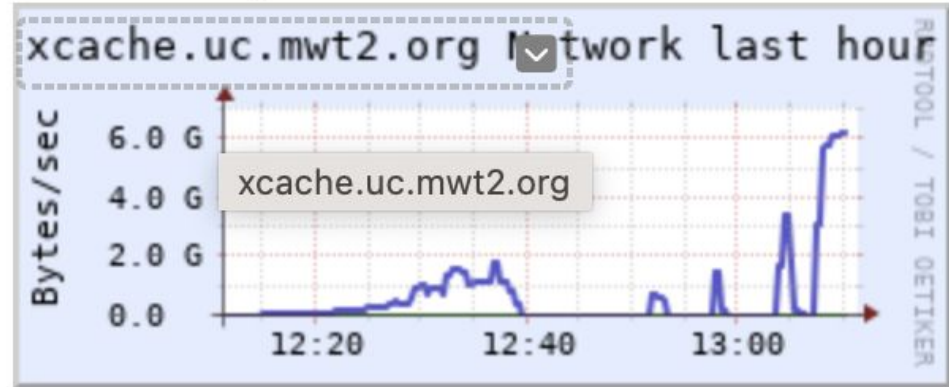




# Remote files throughput with xcache

- xcache.af.uchicago.edu:1094
- Xcache at UChicago AF: 2\*25Gbps links

xcache.uc.mwt2.org





# Misc data points

---

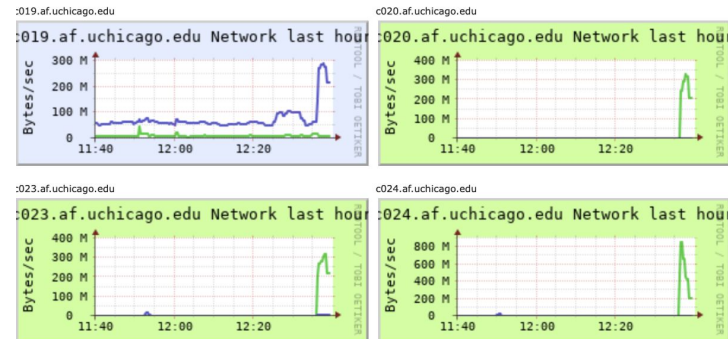

- Cat file to /dev/null – 113MBps
- Cat file to /dev/null after os cache – 3300MBps
- `uproot.iterate(f"{fname}:Events", expressions = arrays_to_read)`
  - Read a few columns – 50MBps(filesize/time)
  - Read a few columns with os cache – 823MBps
  - Read all the columns – 35MBps
  - Read all the columns with os cache – 113MBps





# Running the benchmark as htcondor job directly vs in dask

- Dask worker are also Docker universe HTCondor jobs, we can control the number of workers via the scaling parameters
- With direct HTCondor jobs, we run one task as one job, so there will be 787 jobs running concurrently, this will be similar to the 800 Dask worker case
- The data throughput graph shows the similar results indeed



# Meeting the Challenge

---



- What needs to be improved to hit our target of 200TB in 20min?
  - Identify and resolve bottlenecks in the Dask pipeline



# ServiceX data Access

---

Tested current production version (1.1.4).

Two instances deployed on UC AF (xAOD and Uproot).

Allocated 1k cores for the tests.

Each dataset run concurrently (using ServicexDataBinder).

Datasets:

- Uproot – 6 datasets – 3TB in 21k files
- xAOD – 9 datasets – 136TB in 117k files

# ServiceX – Uproot



Dataset	Files	Size[GB]
single_top_tW_nominal	50	8
single_top_s_nominal	114	10
t_chan_nominal	2506	365
ttbar_scaleup	917	178
ttbar_PS_var	443	93
ttbar_nominal	7066	1355
wjets__nominal	10199	1048
<b>Sum:</b>	<b>21295</b>	<b>3057</b>

Reading a single variable.

When reading remotely, it takes hours.

All requests stay at 10 transformers – default minimal scale.

Reading fully cached data: 16 minutes.

CPU utilization never goes above 15%, so horizontal autoscalers never trigger (default is 30%).

Had to manually lower HOA to 10%, then it scales up but not very fast.

Title	Submitted by	Start time	Finish time	Status	Files completed
ttbar_nominal	Ilija Vukotic	2023-05-01 22:05:54	2023-05-01 22:22:30	Complete	7065 of 7065
single_top_tW_nominal	Ilija Vukotic	2023-05-01 22:05:53	2023-05-01 22:06:23	Complete	50 of 50
ttbar_PS_var	Ilija Vukotic	2023-05-01 22:05:54	2023-05-01 22:14:30	Complete	443 of 443
ttbar_scaleup	Ilija Vukotic	2023-05-01 22:05:53	2023-05-01 22:15:43	Complete	917 of 917
single_top_t_chan_nominal	Ilija Vukotic	2023-05-01 22:05:53	2023-05-01 22:13:33	Complete	2506 of 2506
single_top_s_chan_nominal	Ilija Vukotic	2023-05-01 22:05:53	2023-05-01 22:06:43	Complete	114 of 114
wjets_nominal	Ilija Vukotic	2023-05-01 22:05:53	2023-05-01 22:21:25	Complete	10199 of 10199





# ServiceX Uproot

Never scales up to use all cores.

XCache and S3 endpoints basically idle.

Not all requests start right away – some more optimization possible in file path finding.

Factor 4–5 speedup probably possible.



# ServiceX xAOD – single dataset



Reading a single dataset (30885 files ~30TB ).

Single Jet collection, simple cut selection.

Selection fully cached in xcache.


Transformer CPU utilization at 70–80%.

Transformers scaled up to the limit of 750 cores.

We can do much better with startup – write a custom HPA.

Once that's done, the next limiting factor will be number of CPUs available (our xCache can support up to ~3500 transformers like this).

Plato at 750 transformers reached.



# ServiceX xAOD



Period	Files	Size[GB]
K	9312	11704
M	18747	24008
F	5416	5800
I	801	1000
C	9567	12544
L	30885	31006
D	16966	23311
Q	11164	12153
O	14079	17228
<b>Sum:</b>	<b>116937</b>	<b>138755</b>

Processing fully cached data: 1.5 hours.

Took ~40 min to get all the datasets looked up and started.

Scaled up to ~900 cores.

Sometimes transformers would suddenly scale down. It appears not all the input data was in the XCache.

Reading a single Jet collection, simple cut selection.

When reading remotely, it takes whole morning, several retries.

When cached CPU utilization at 70–80%.

At this scale path lookups take considerable time (some of the lookups expired from the cache – TTL: 1 day).

Title	Submitted by	Start time	Finish time	Status	Files completed	Workers
period_M	Ilija Vukotic	2023-05-02 16:53:51	-	Submitted NaN%	0 of 0	-
period_Q	Ilija Vukotic	2023-05-02 16:53:51	-	Running 1%	2872 of 22328	74
period_D	Ilija Vukotic	2023-05-02 16:53:51	-	Submitted NaN%	0 of 0	-
period_F	Ilija Vukotic	2023-05-02 16:53:51	-	Running 69%	3762 of 5416	79
period_C	Ilija Vukotic	2023-05-02 16:53:51	-	Submitted NaN%	0 of 0	-
period_K	Ilija Vukotic	2023-05-02 16:53:51	-	Running 82%	7649 of 9312	156
period_O	Ilija Vukotic	2023-05-02 16:53:51	-	Running 34%	4850 of 14079	107
period_L	Ilija Vukotic	2023-05-02 16:53:51	-	Running 94%	29242 of 30885	443
period_I	Ilija Vukotic	2023-05-02 16:53:51	-	Submitted NaN%	0 of 0	-

# ServiceX xAOD

Only a little bit of space to optimize transformers.  
Requests take time to ramp up.  
x 3 improvement possible.

