

xAOD Analytics

Ilija Vukotic

September 2015

Acknowledgements

Results here reflect contributions from a number of people and groups:

ASG: Attila Krasznahorkay

Rucio: Mario Lassnig, Thomas Beermann

Analysis Support: Doug Benjamin

Infrastructure: CERN IT & Lincoln Bryant (MWT2 Chicago)

Questions and Goals

Which data formats are used for analysis?

Which branches are used?

Difference between on-grid/off-grid?

Which versions of ROOT are used?

Measure IO (read) performance

Small recap on data sources ...

1. Panda Job Archive.



1. xAOD monitoring information.



We collecting usage data from both grid and off-grid jobs.

Only user jobs were monitored - still investigating why no production jobs.

Small recap ... continued

Starting 08/07/2015 we collect from 100% jobs. Still there are jobs not using the instrumented code version so absolute numbers are still not there.

We collect ~ 1GB/day:

accessedFiles - full path for all the files in the job

accessedBranches - for each branch that was accessed a number of accesses/events

accessedContainers - for each container that was accessed a number of accesses/events

ReadCalls

ReadSize

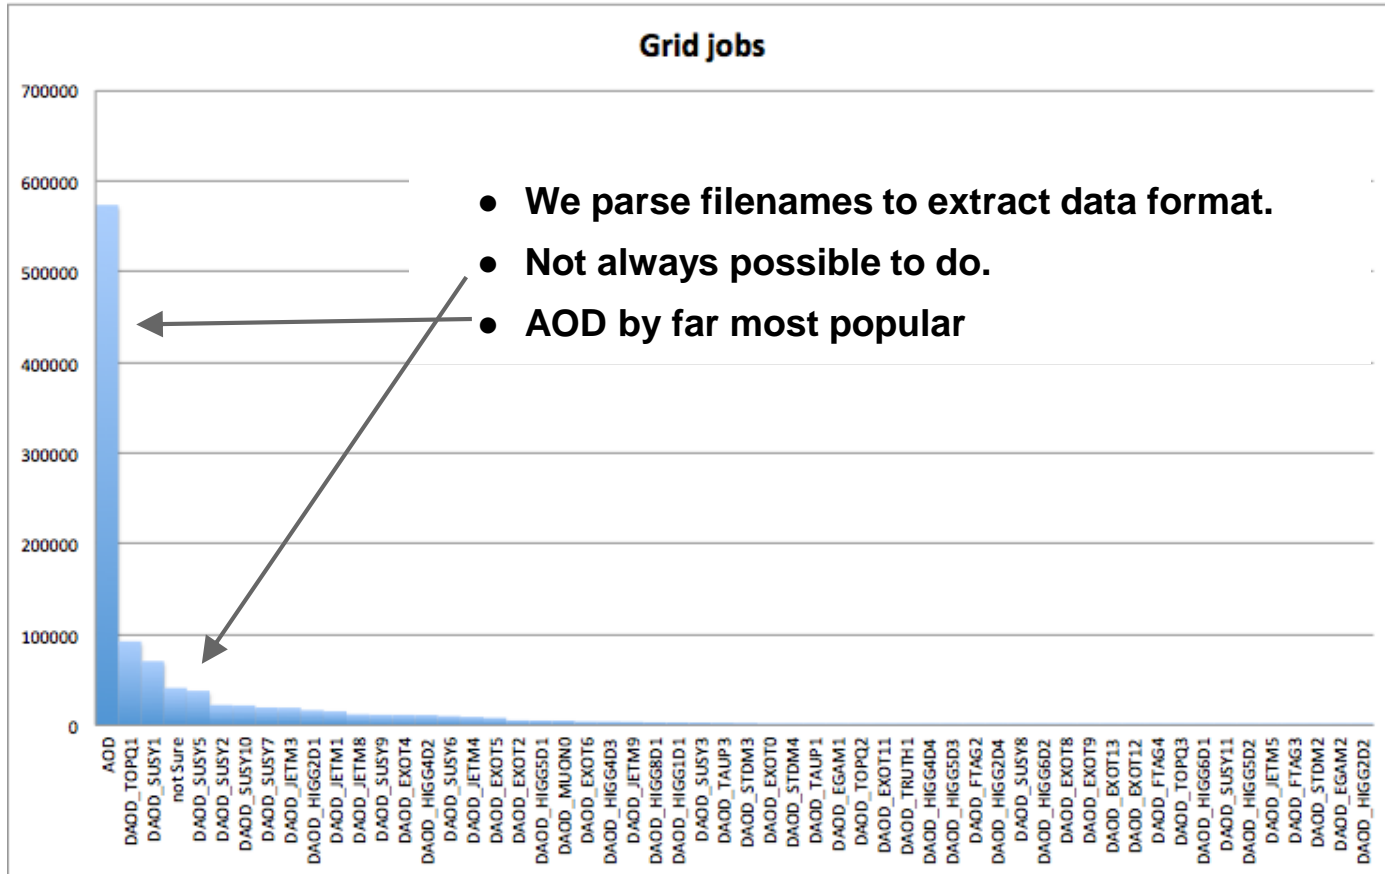
CacheSize

ROOT_RELEASE

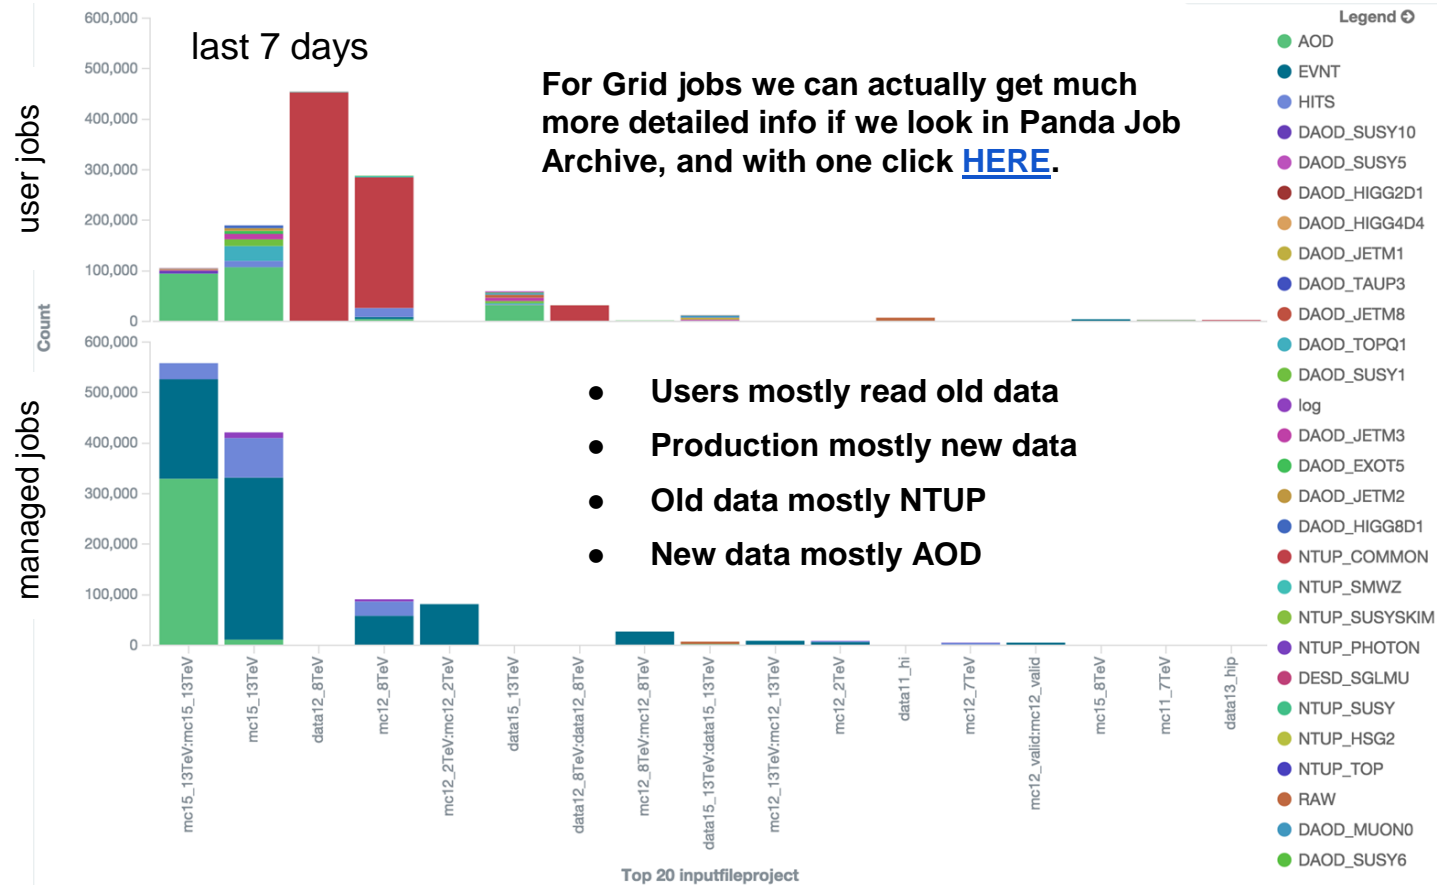
pandaID, taskID - if grid job

Here we use all the data collected in August and first week of September.


Q1. Which data formats are used



Q1. Which data formats are used



Q1. Which data formats are used



Q2. What is the read performance

To first order depends on the average read size.

Very important for local, essential for remote reading.

Q2. What is the read performance

Grid


access type	jobs	read calls	read size [kB]	cache size	files	accessed branches	accessed containers
		(averaged over all jobs)					
<i>xrootd</i>	493244	11678	9407	2.24E+07	4.05	186.66	26.53
<i>dcap</i>	167265	28910	13	2.15E+07	4.86	189.93	29.66
<i>local</i>	409150	37716	102	2.09E+07	3.38	186.24	27.55

Off-Grid


access type	jobs	read calls	read size [kB]	cache size	files	accessed branches	accessed containers
		(averaged over all jobs)					
<i>xrootd</i>	78788	8924	27481	30394336	2.66	182.52	25.20
<i>dcap</i>	39098	70487	8	10283204	7.12	224.07	35.04
<i>local</i>	657262	111118	20	8545653	2.63	210.28	27.30

Q2. What is the read performance

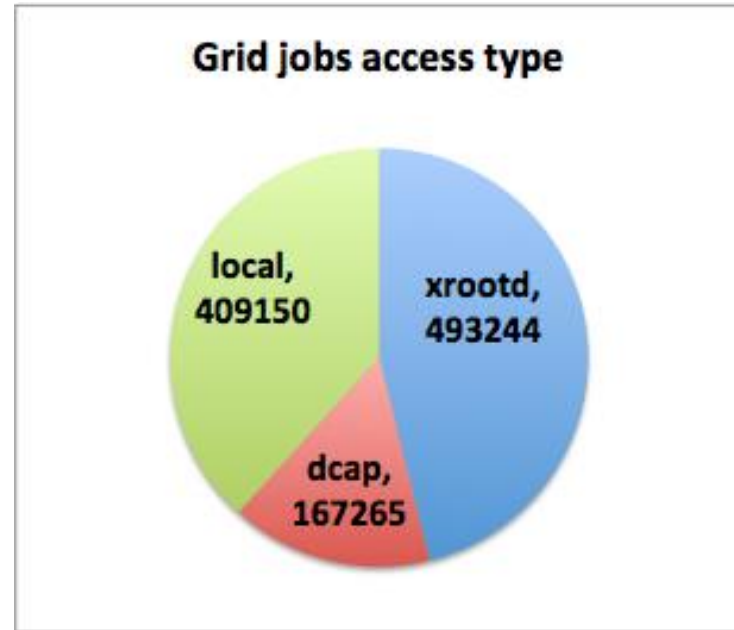
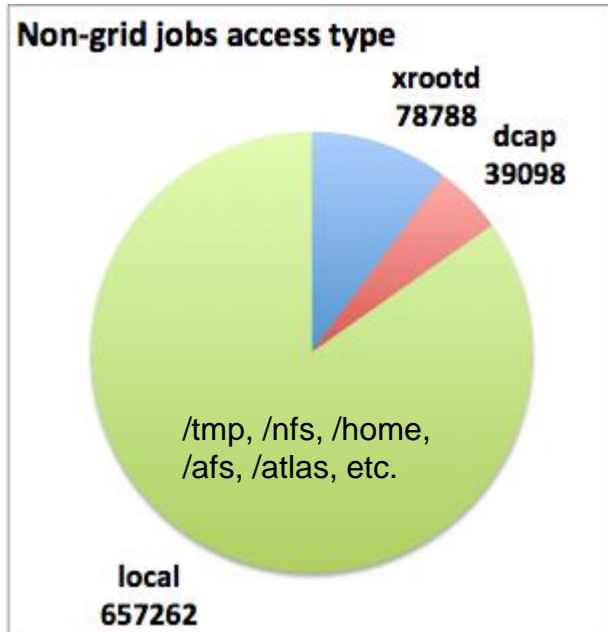
Non-grid reading



Grid jobs reading



Q2. What is the read performance

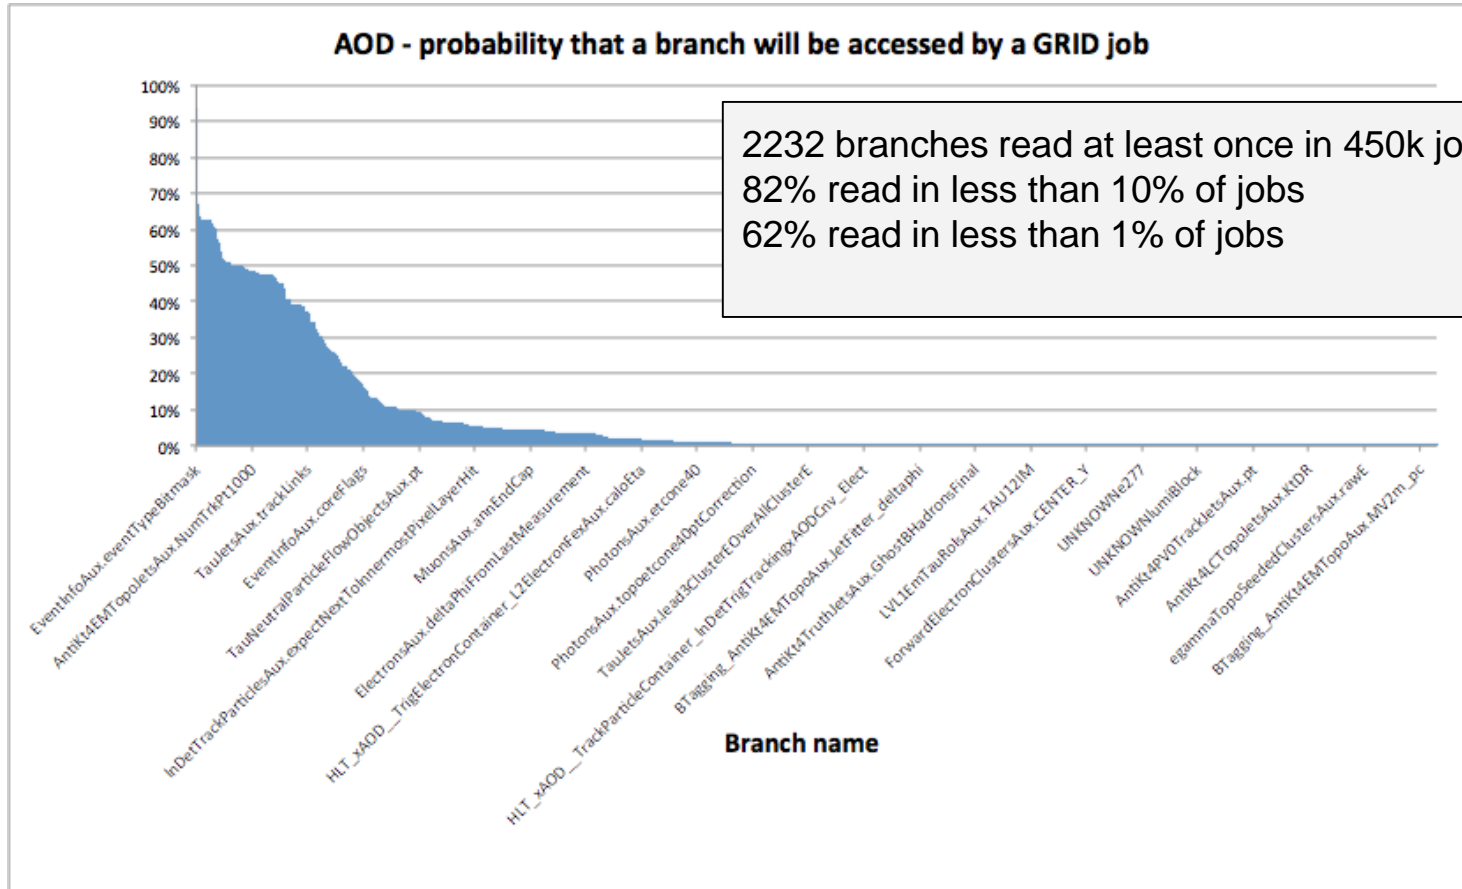


in conclusion: most of the difference comes from different mix of protocols used.

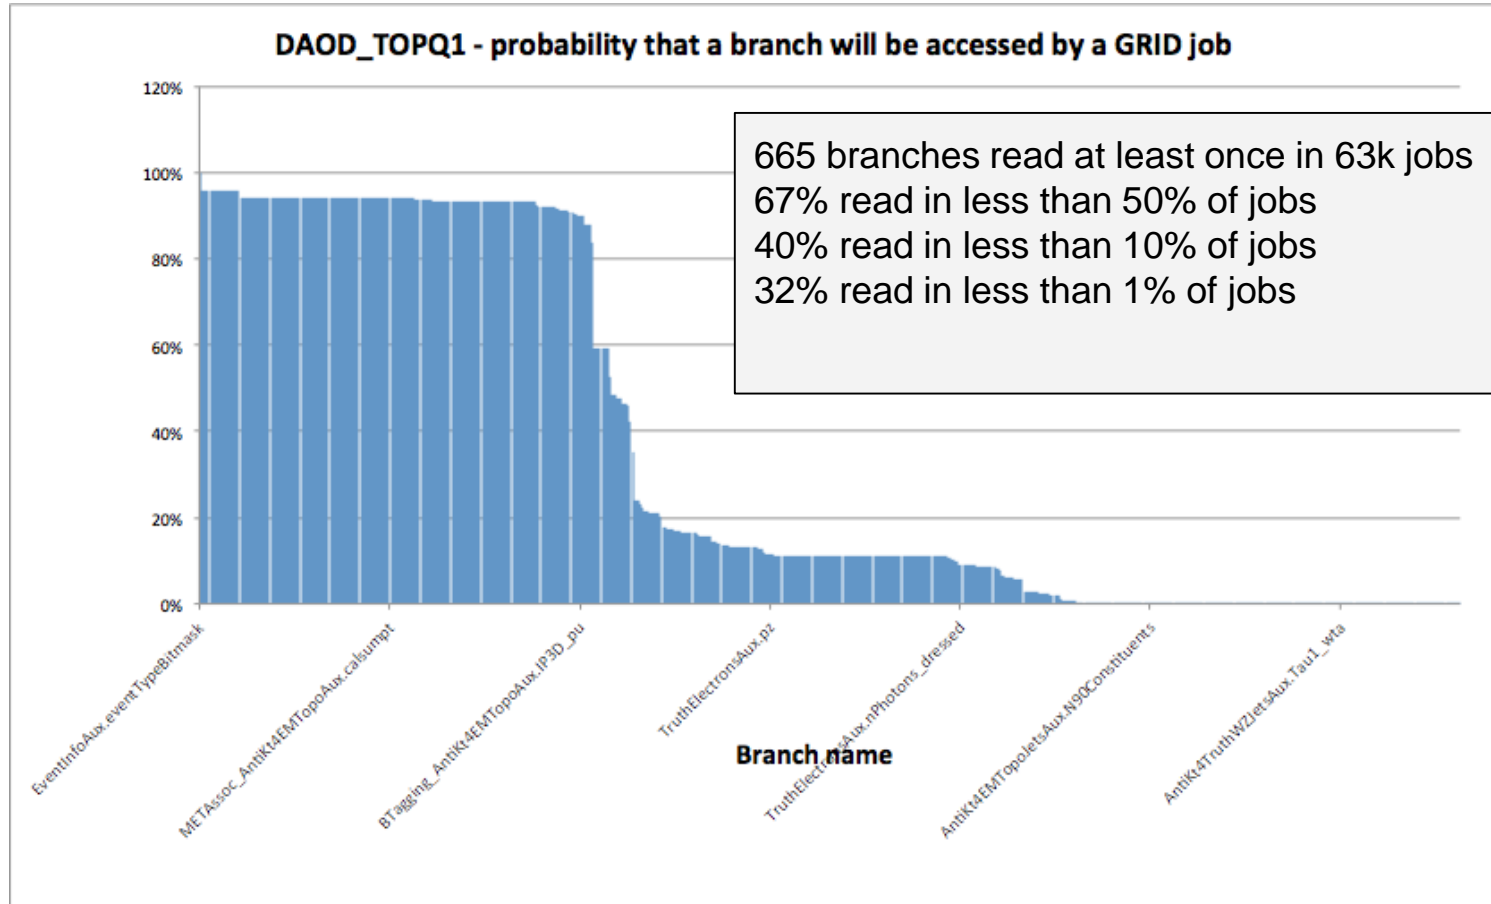
Q3. Which branches are read

It would be great if we could give a feedback to people doing derivations on what branches are really read.

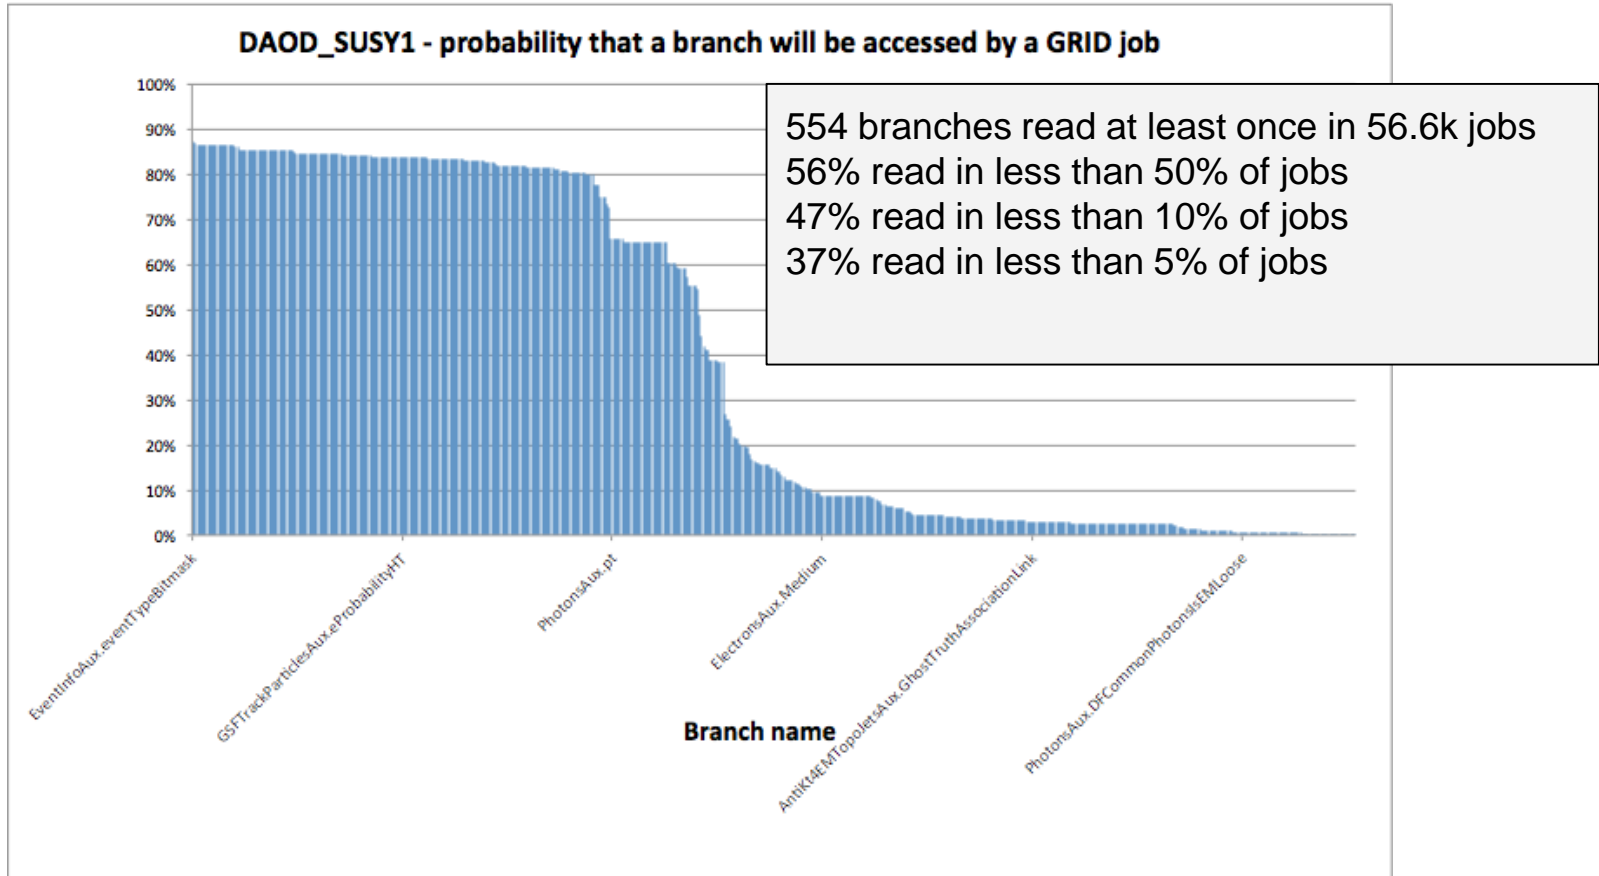
Q3. Which branches are read



Q3. What branches are read

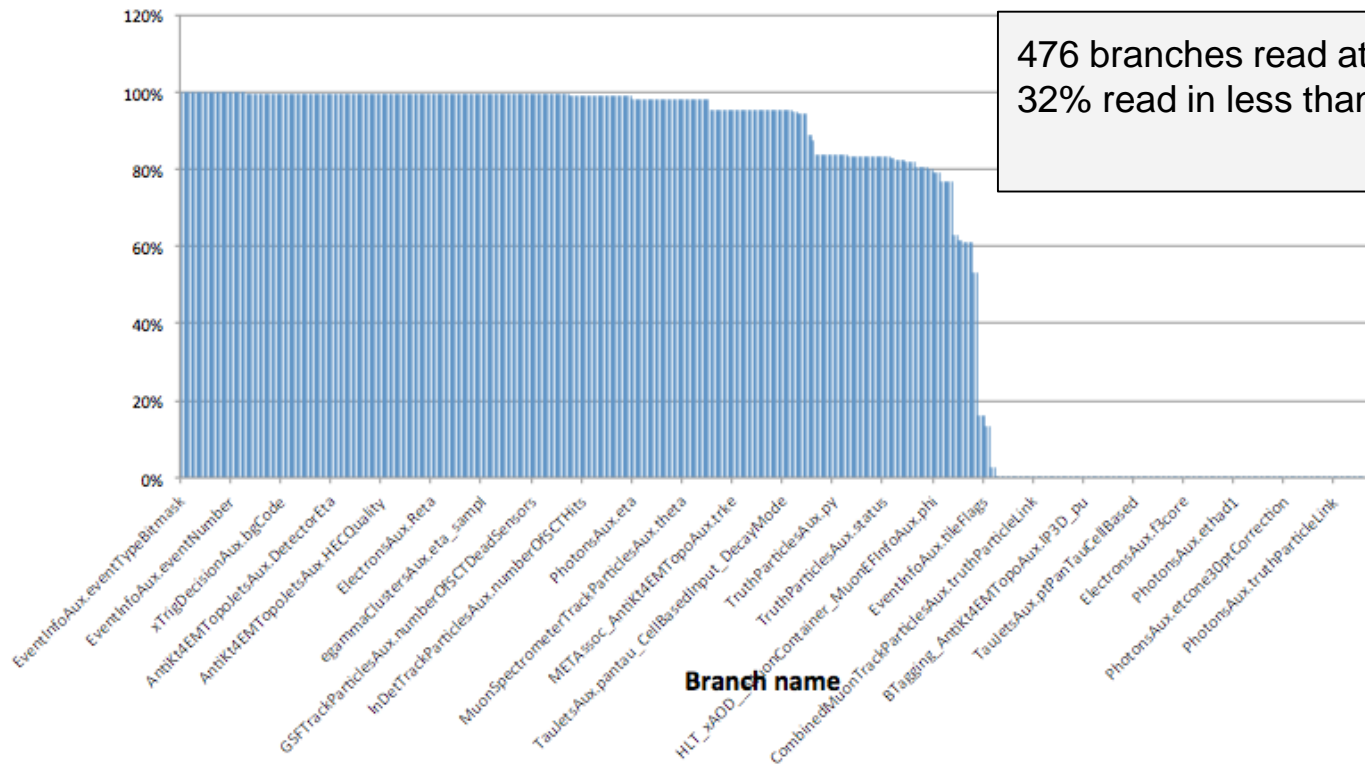


Q3. What branches are read



Q3. What branches are read OFF GRID


DAOD_HIGG4D2 - probability that a branch will be accessed by a OFF - GRID job



476 branches read at least once in 96k jobs
32% read in less than % of jobs

Q3. What branches are read ON GRID

DAOD_HIGG4D2 - probability that a branch will be accessed by a GRID job



Q4. What ROOT versions are used

	JOBS	
ROOT version	Grid	Off-Grid
6.02/10	15378	16860
6.02/12	1054281	757480
6.02/05		77
6.04/00		17
6.04/02		693
6.02/05		77

Only September data

To Do

Ilija - Make cron jobs to index summed up info once a day into ElasticSearch. Create two dashboards. Fix cache calculations.

Attila - change code so it does not re-learn for each file (medium priority)

Doug - feedback analysis results to ASG, derivations people

Attila - change code so it uses knowledge on branches used by previously finished jobs (low priority)

Attila - add sending of the TaskID equivalent for off-grid jobs running in Condor, PBS, Torque.

Attila - calculate and send a branch fingerprint.