



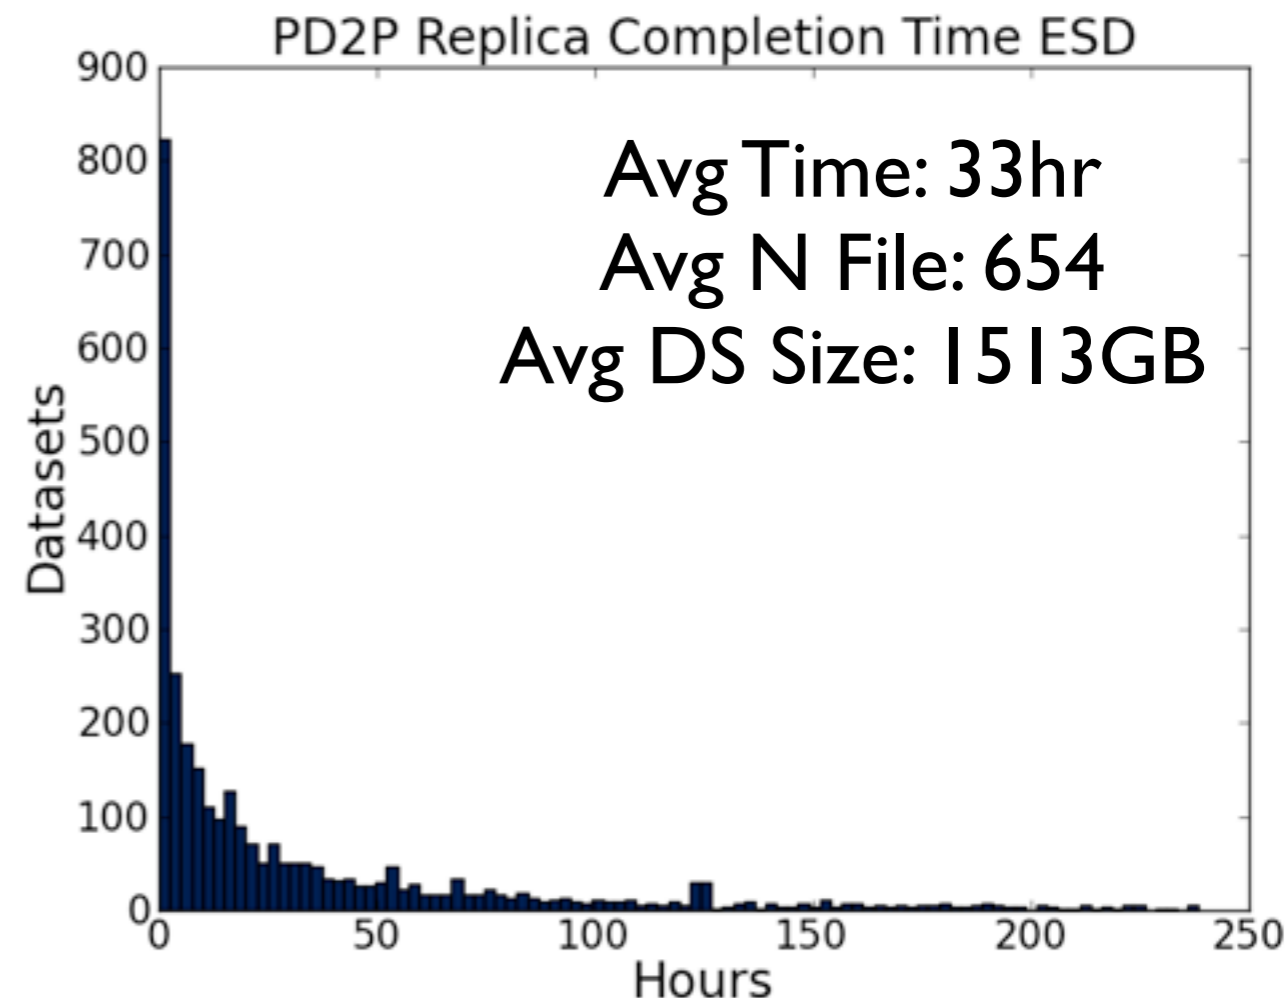
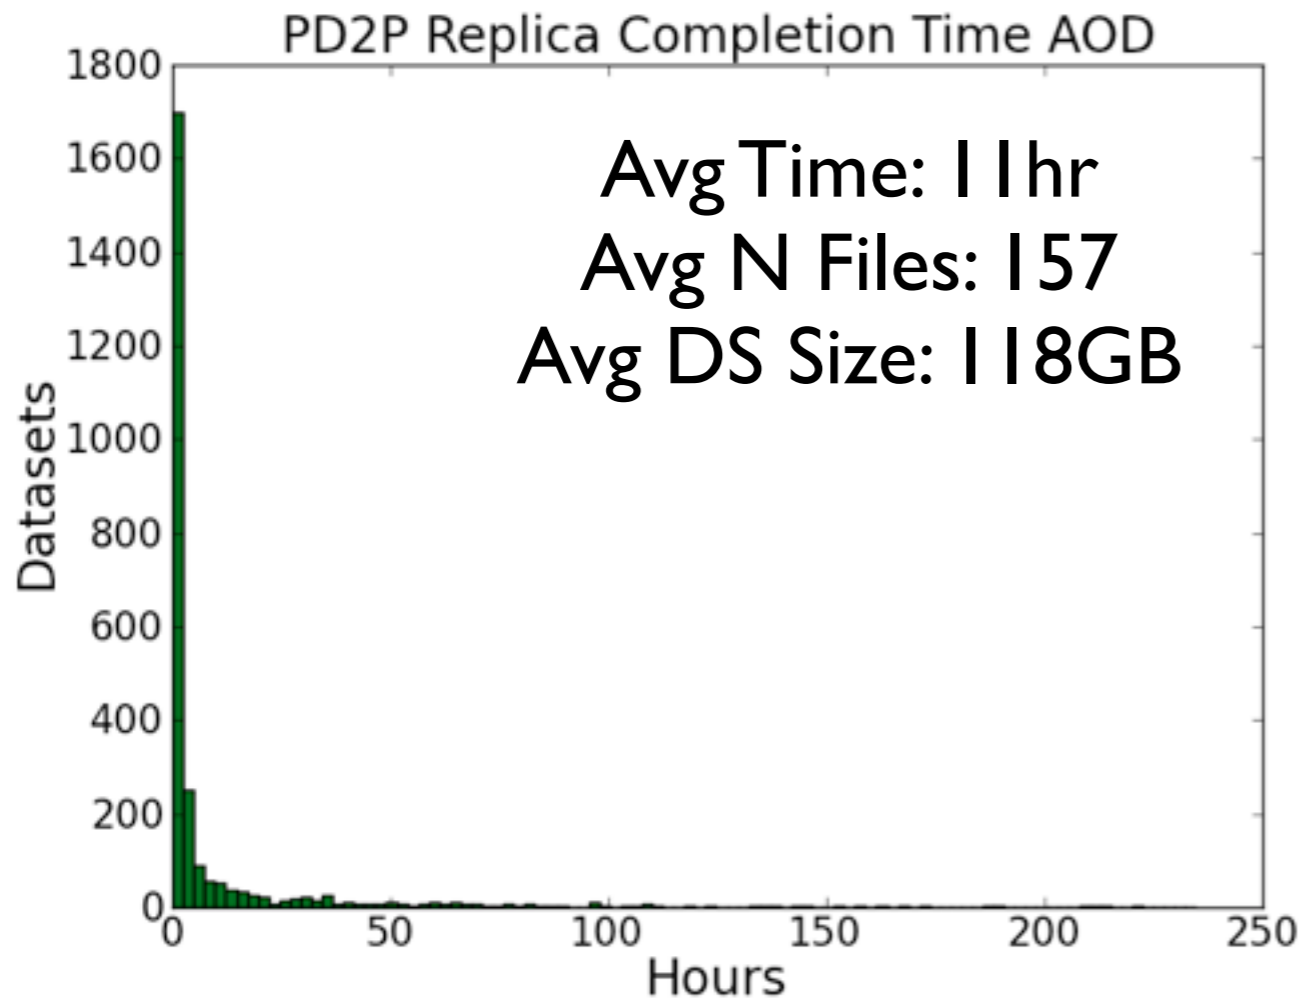
PD2P: The Executive Summary

Graeme Stewart

(Selected highlights from the ADC Dev meeting
26/1/2011)


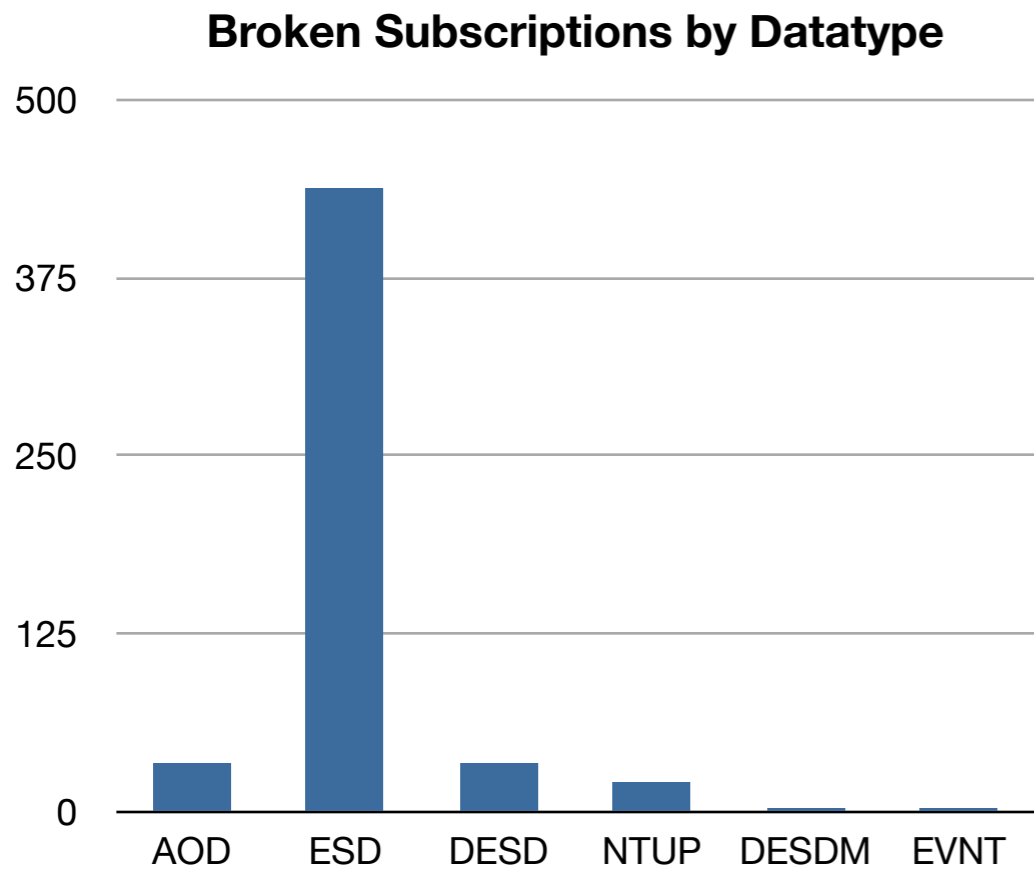


ESD vs. AOD



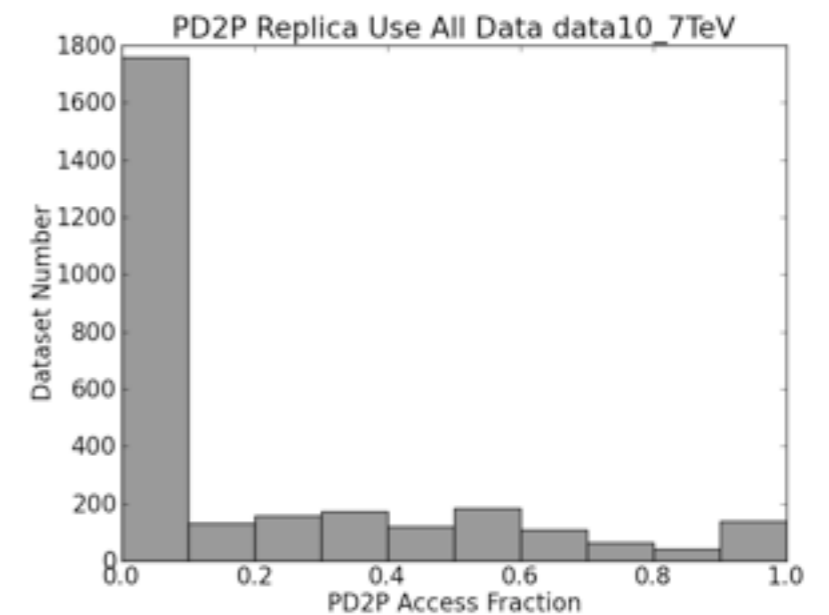

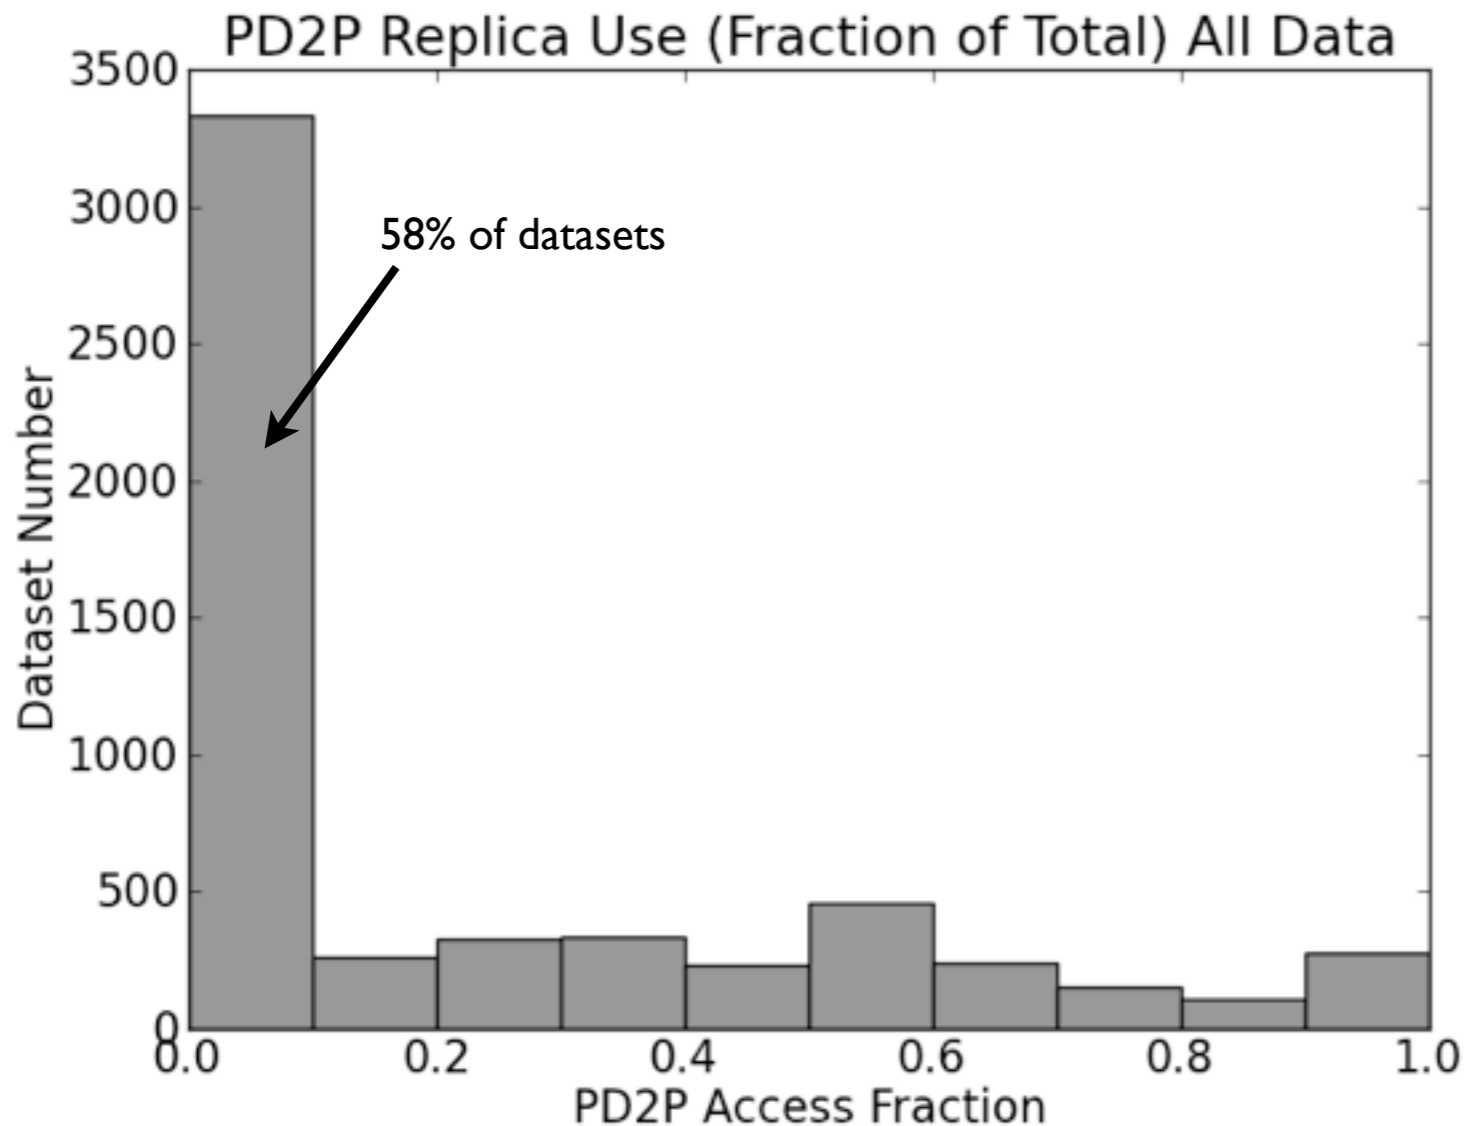
- Small datasets complete better and faster!

What Breaks





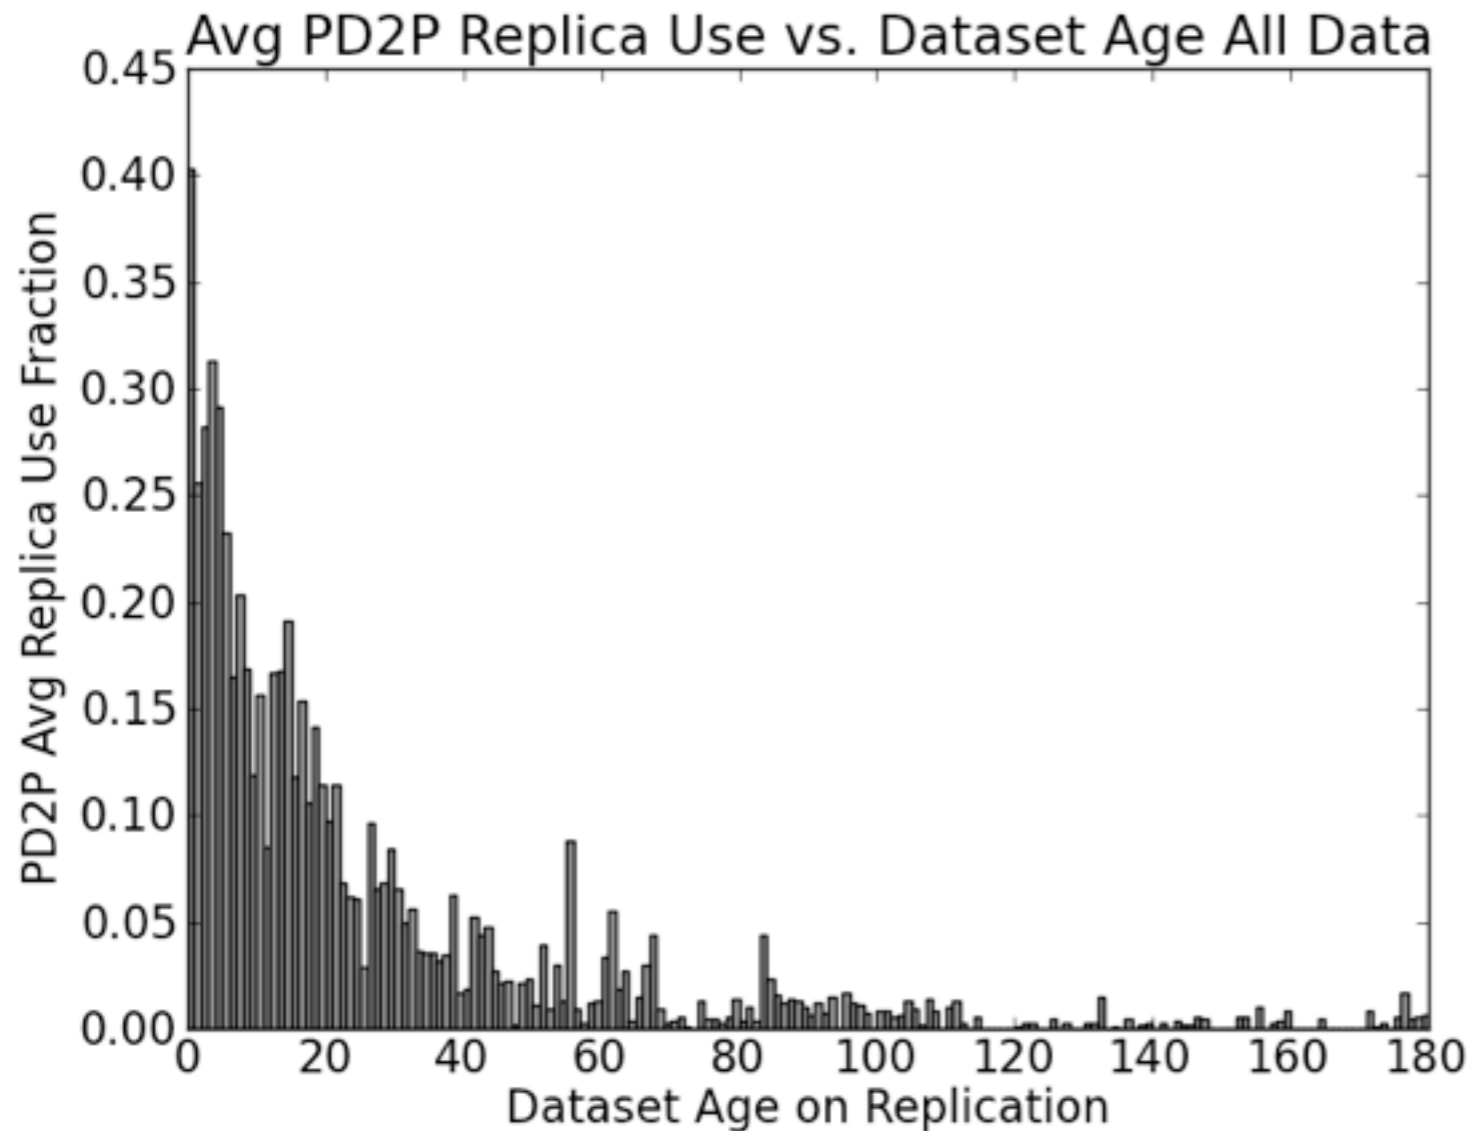
- Mainly ESD, biased towards data10_7TeV
 - Broken ESD subscriptions are for larger datasets: Avg N Files, 1057; Avg DS Size 2624GB
 - This is ~twice the ESD average

PD2P Replica Use



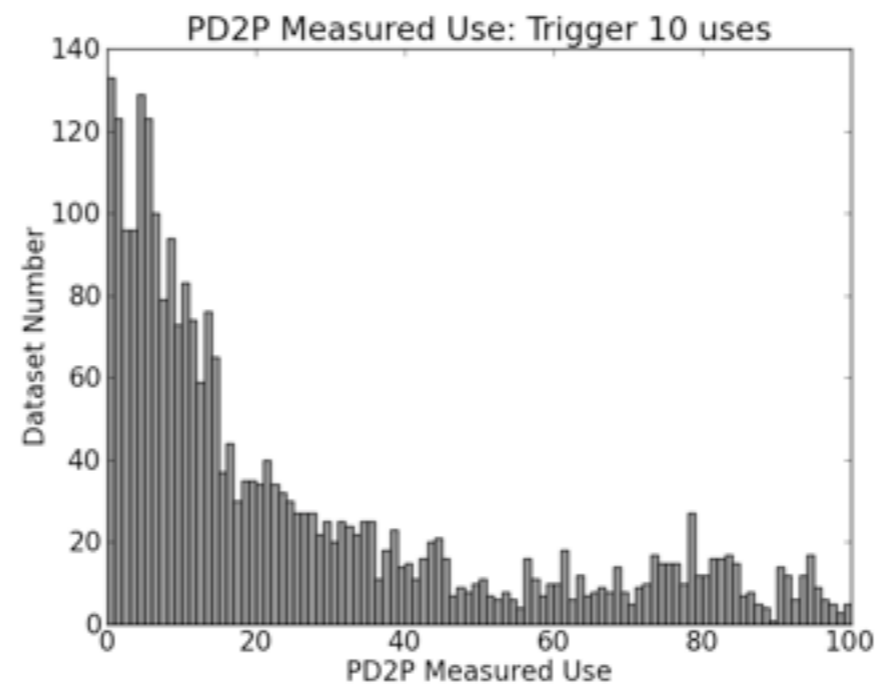


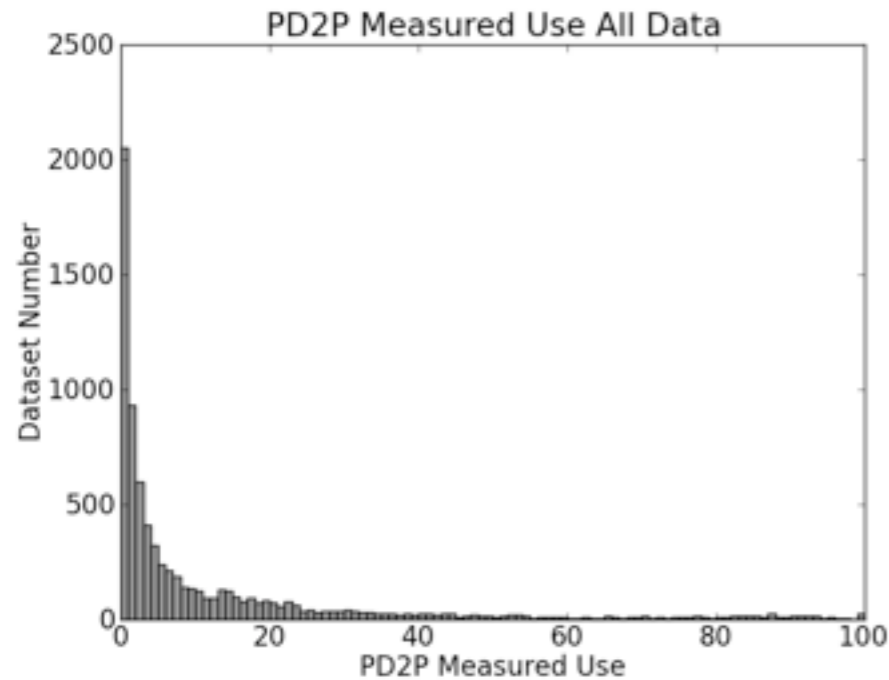
- Lots of data unused
- No strong relationship with data type or project

Reuse vs. Age II



- Strong correlation between PD2P replica use and dataset age

Apply Different Replication Triggers



$\Sigma \text{ datasetNumber} = \text{replicationLoad}$
 $\Sigma \text{ datasetNumber} \times \text{Accesses} = \text{usefulness}$

- Seems that a higher trigger does help PD2P 'efficiency'
- But would it use T2s and network capacity most efficiently?
- Some level of speculative replication is a good thing

Discussion

- Larger datasets break more often
 - Subscribing dataset shards (`_dis`) for large datasets might make sense for faster, more reliable completion
 - Some technical problems to consider
- Filtering based on data type or project would seem to be of little value
- Filtering based on dataset age looks like a very promising optimisation
- Raising the replication trigger should make PD2P more 'efficient' (`datasetUse/replicationLoad` goes up)
 - But need to optimise for best use of network and Tier-2 resources
- Other 'Open Questions':
 - Feedback between DDM and PanDA for site selection: to be discussed
 - More aggressive rebrokering would help PD2P reuse
 - Dan assures us this is fine for outputs
 - PD2P@TIs - wait for Kaushik?

Proposal

Dataset Age	Threshold
< 1 week	1
1 week - 1 month	2
1 month - 3 months	3
3 months - 9 months	5
> 9 months	10

- Modify PD2P based on dataset age
- Reduce rebrokering time to 12 hours
- Use length of subscription queue as negative weight for T2s