


WN Xrootd Configurations

- Effects of various WN proxy-cache configurations tested (v4, Xcache-v5, Mem-proxy-v5);
 - Using speed transfer tests of ~ 1GB files, and from ATLAS job results data
- Between v4 and Xcache, job success rates are ~ similar.
 - For the Memcache; small increase in direct-io failures; larger effect for staged jobs.
- (Especially) for the 2019 tranche, ML and Multicore Personalities give different responses to the speed tests for Xcache)


| cache | Count | direct-io | Success[%] | |
|-------|--------|-----------|------------|-----------|
| | staged | | staged | direct-io |
| 4 | 138583 | 184026 | 95.3 | 90.0 |
| Mem | 6093 | 6367 | 89.1 | 88.0 |
| X | 7112 | 8930 | 96.8 | 89.6 |

Wn-2019-dell

Personality = ml

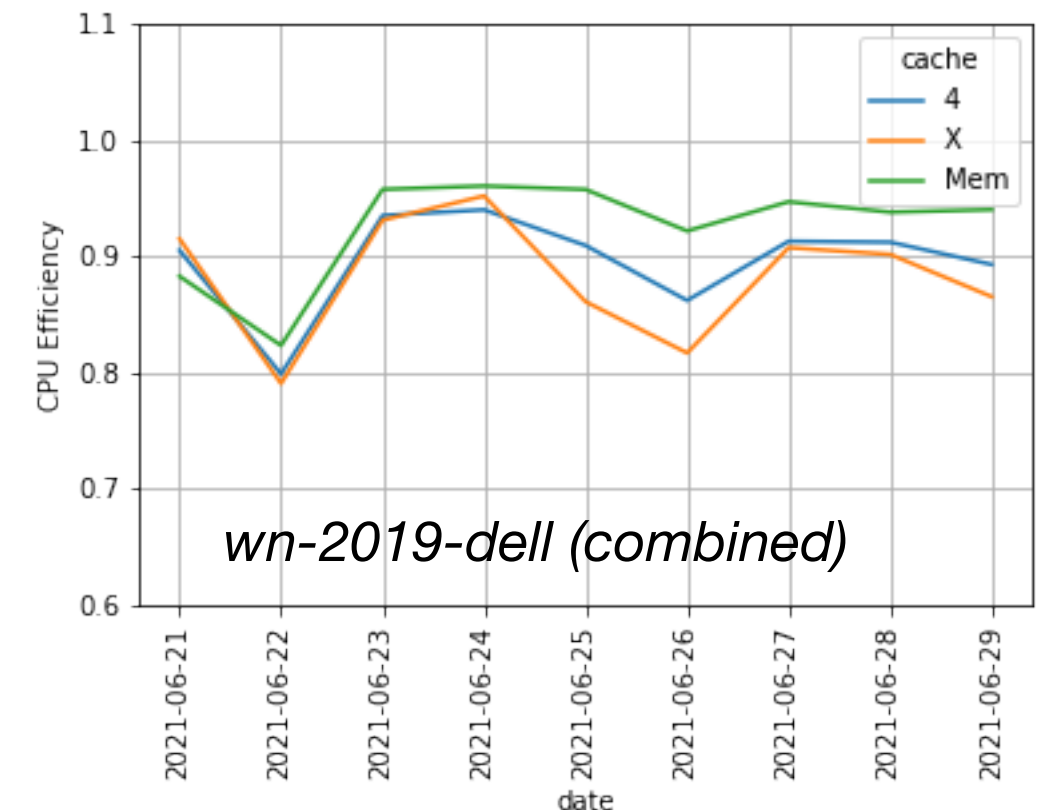


Personality = multicore



cache
 4
 X
 Mem

- For the 2019-tranche, while it appears that Memcache improves overall efficiency, much depends on streaming mode and personality/load.



- Memcache speeds appear more stable across different generations / personalities / loads
- Xcache performance more 'extreme' in (un)favourable conditions
-

Reduced Slots and ATLAS HS06

- Some discrepancy over HS06 reported by ATLAS and RAL/Vande
 - ATLAS sees currently ~ 121k HS06 (cf. 173k pledge)
 - Vande reports running at 102% of fairshare.
- ATLAS needs a site-averaged Corepower (but unlikely to be so wrong)
- Number of running cores ~ agree, so issue is elsewhere;
- Once we agree, can do something about the low running ?
- 2014,15 and 16 tranches recently updated to run at half the number of existing slots;
 - 2014 already was ~ best eff. Tranche (probably unneeded)
 - No clear improvement observed, but external factors (job mix, other VO's might have larger effect)

