

CMS Diskless T2s

Daniele Bauer, **Chris Brew**, Dave
Colling, Simon Fayer



- Diskless Tier 2s are an important idea for simplifying smaller sites
- But are they fully useful or would we need to restrict the workloads to low bandwidth ones
- CMS jobs at Oxford reading from were an ideal place to test this.

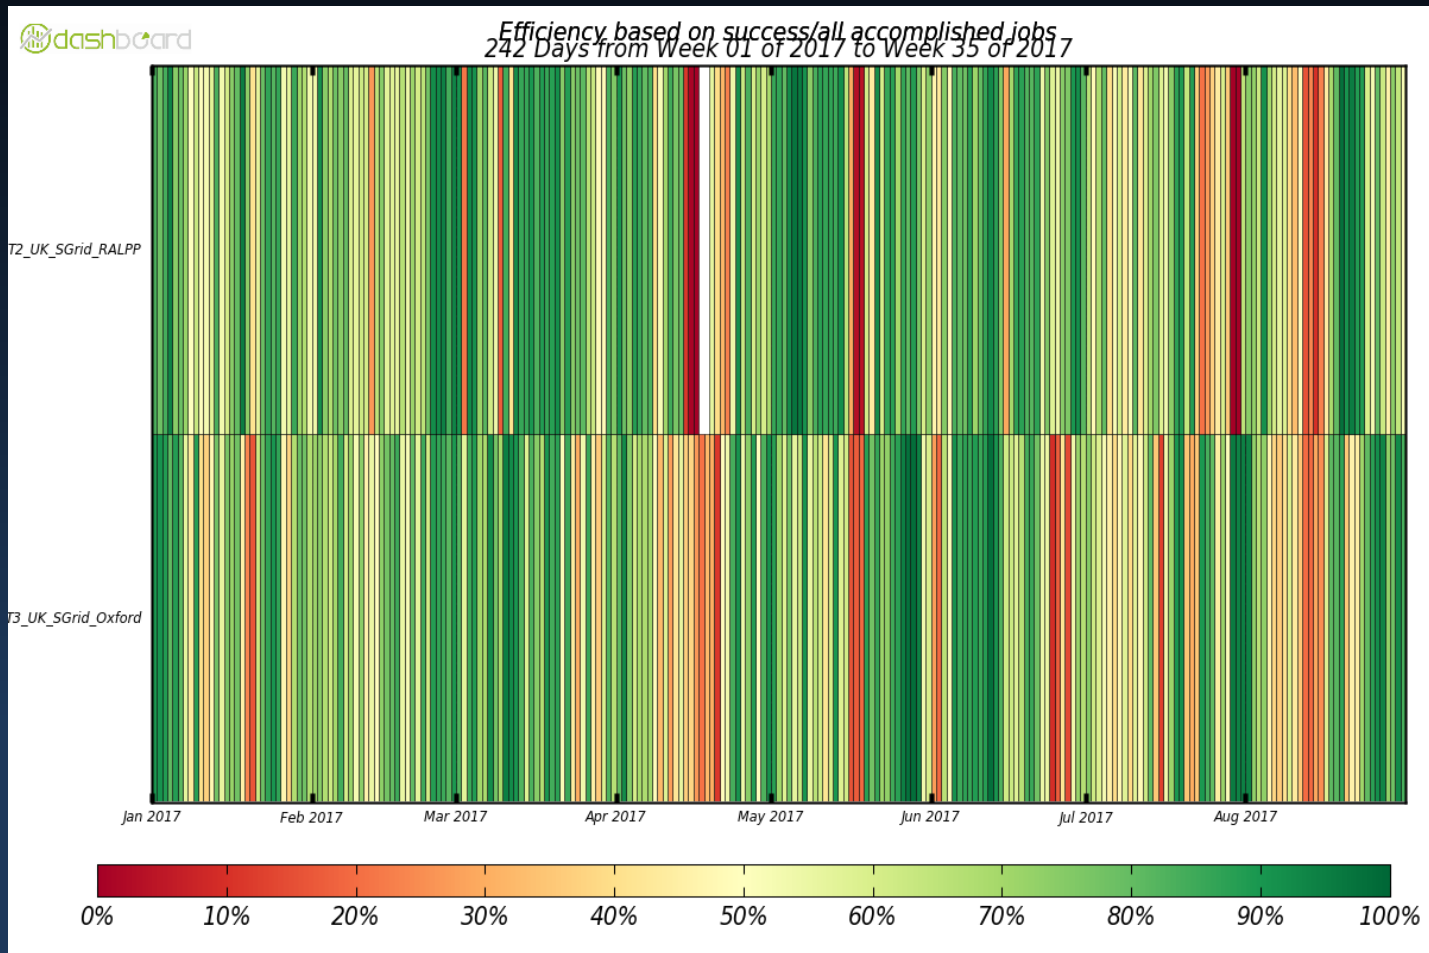


Why CMS jobs at Oxford?

- CMS central config already supports the idea for diskless sites for special opportunistic sites
- CMS also already had the AAA infrastructure for WAN reads of data from within the code
- Oxford was already registered and set up as a CMS T3 site but was under utilised
- December 2016, we stopped CMS work at Oxford, decommissioned the storage and configured CMS jobs at the site to read from and write to Oxford



Can we run jobs?




Can we run them efficiently?


(or as efficiently as CMS ever does)



Number and Efficiency of Running Jobs at Oxford




Efficiency against Number Running for Jobs at Oxford






Transfer Throughput

2017-01-01 00:00 to 2017-06-30 00:00 UTC



Destinations

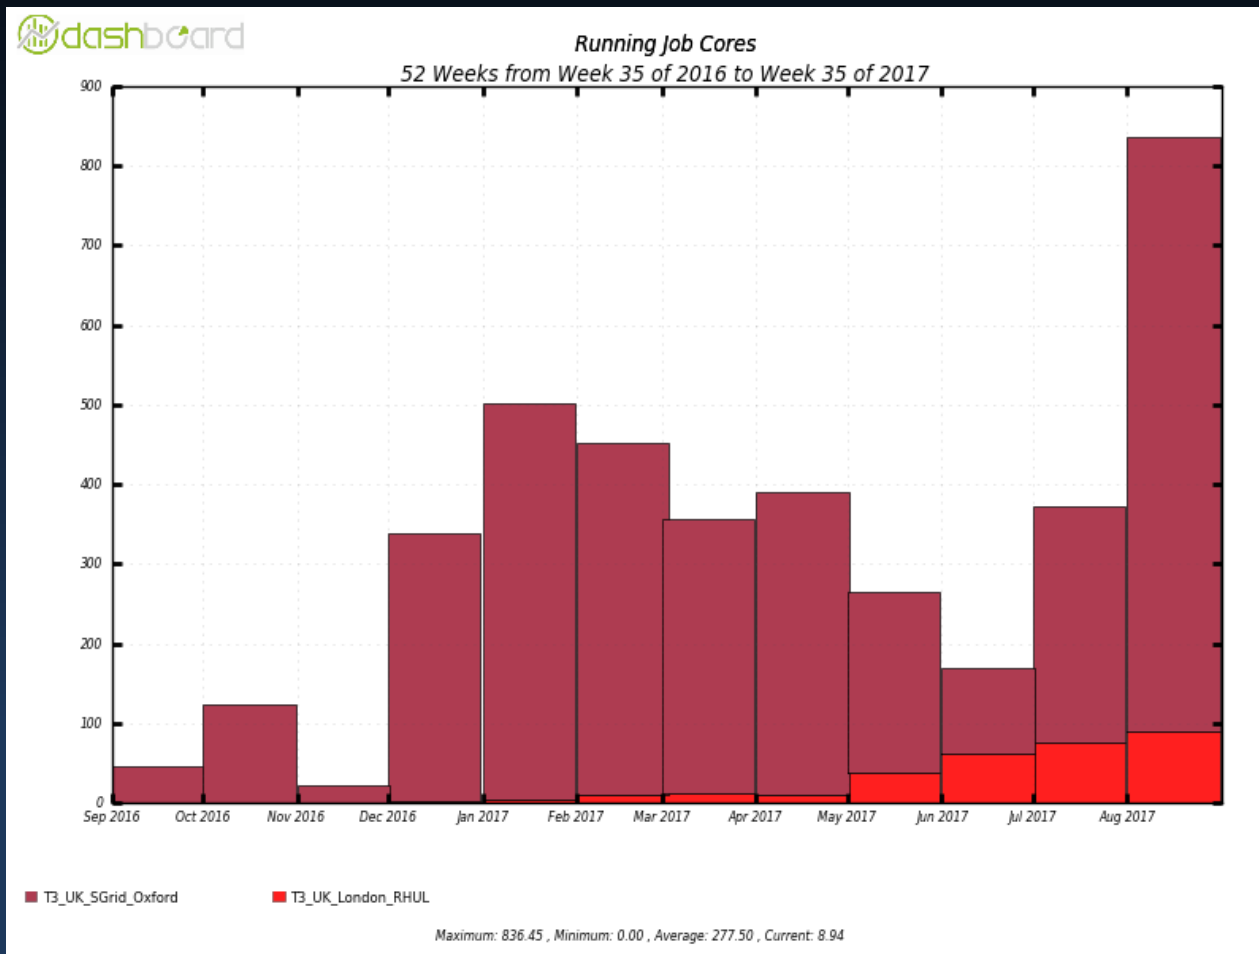


Next Steps

- CMS jobs with remote IO work well, we continued with the setup after the tests
- We then set up RHUL as a satellite of IC



Oxford and RHUL Usage

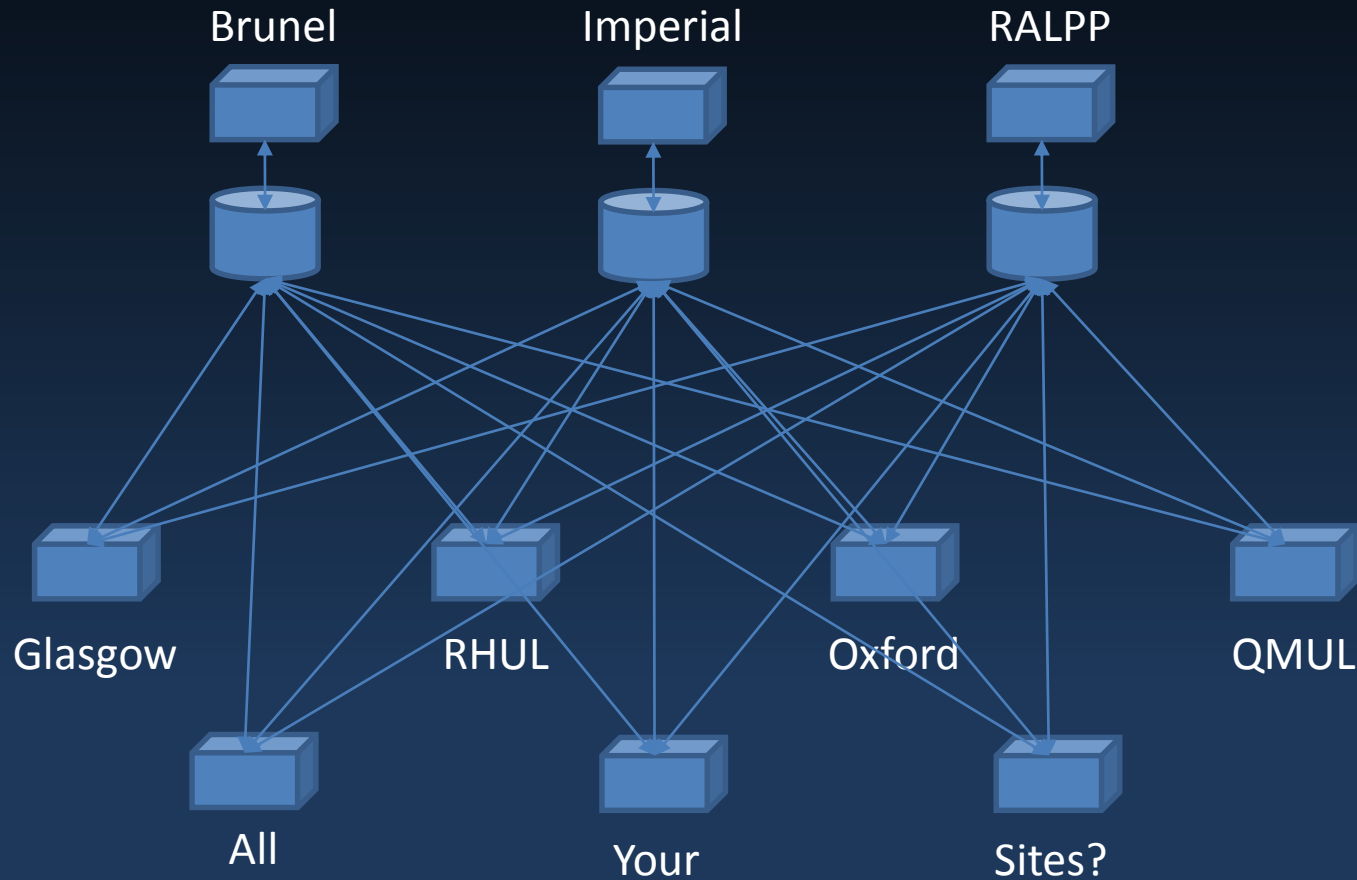


But...

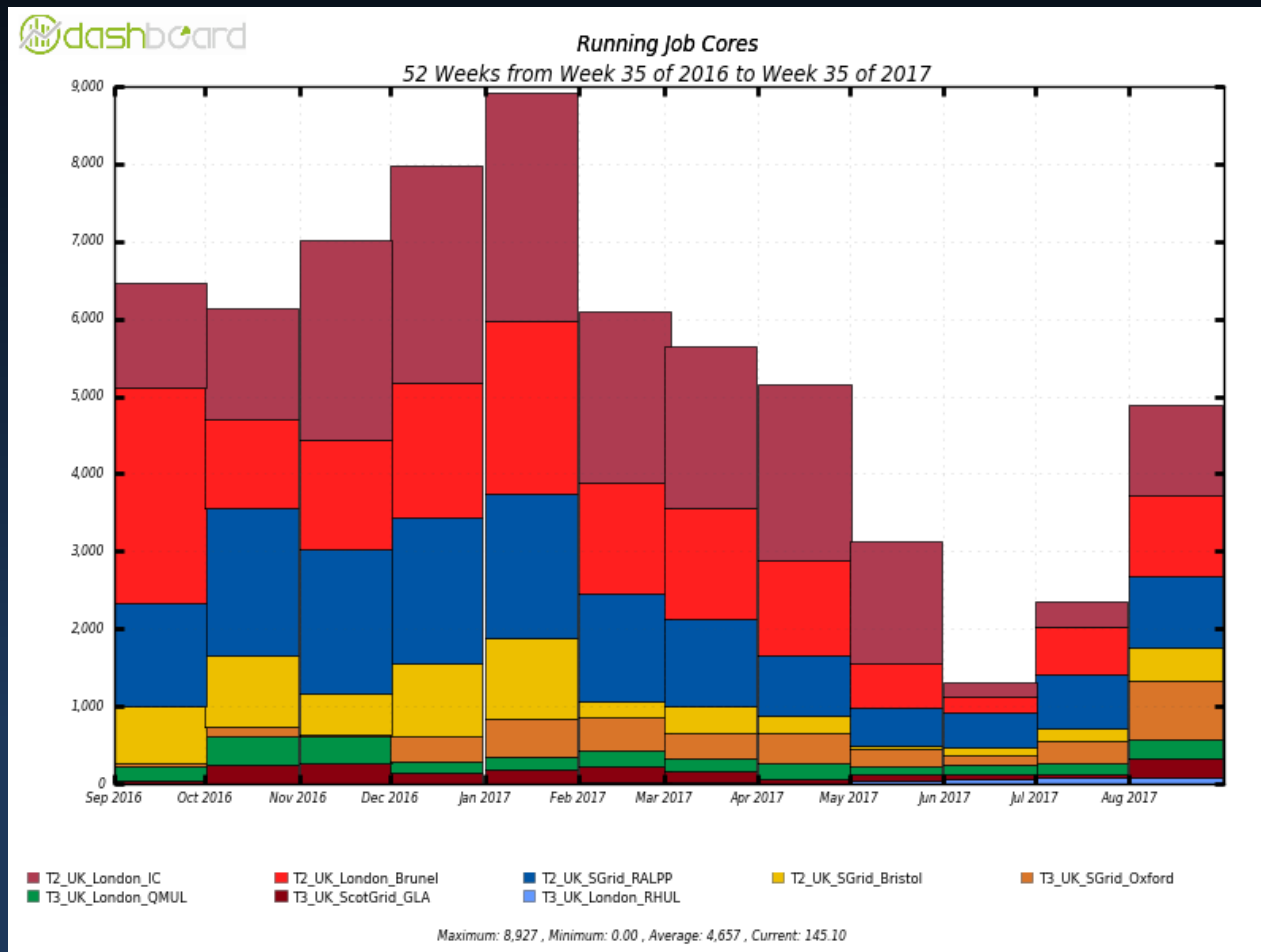
- CMS site config supports a many-to-many mapping between storage and compute as long as the software can resolve the data location
- So we set up a UK xrootd redirector and registered the RALPP, IC and Brunel AAA endpoints with it




CMS UK T2/T3 Mesh



CMS UK T2 and T3 Usage



CMS Testing Plan for Echo



- | | | | | |
|----------------------|-------------------|---------------------|---------------------|-------------------|
| T2_AT_Vienna | T2_ES_IFCA | T2_BE_UCL | T1_US_FNAL_Disk | T1_RU_JINR_Disk |
| T2_CN_Beijing | T2_FI_HIP | T2_DE_DESY | T2_DE_RWTH | T1_ES_PIC_Buffer |
| T1_FR_CCIN2P3_Disk | T2_ES_CIEMAT | T2_BE_IHHE | T2_CH_CSCS | T2_EE_Estonia |
| T0_CH_CERN_Disk | T1_DE_KIT_Disk | T2_UK_London_Brunel | T2_CH_CERN | T1_RU_JINR_Buffer |
| T1_FR_CCIN2P3_Buffer | T2_FR_CCIN2P3 | T1_IT_CNAF_Buffer | T1_US_FNAL_Buffer | T2_BR_SPRACE |
| T1_IT_CNAF_Disk | T2_UK_SGrid_RALPP | T1_ES_PIC_Disk | T2_HU_Budapest | T2_IT_Pisa |
| T1_DE_KIT_Buffer | T2_IN_TIFR | T2_UA_KIPT | T2_US_Purdue | T2_IT_Legnaro |
| T2_US_Wisconsin | T2_US_Nebraska | T2_FR_GRIF_LLRL | T2_US_Caltech | T2_US_Florida |
| T2_US_UCSD | T2_RU_IHEP | T2_US_Vanderbilt | T2_US_Ioannina | T2_PL_Swierk |
| T2_IT_Bari | T2_KR_KNU | T2_TR_METU | T2_UK_SGrid_Bristol | ... plus 11 more |

Maximum: 948.22 MB/s, Minimum: 0.16 MB/s, Average: 245.95 MB/s, Current: 1.42 MB/s

