

CMS Thoughts on Squid Caches

- ▶ CMS has no problem with using a common set of squid caches between ATLAS and CMS
 - ▶ Mostly an issue at Tier-1s (fewer Tier-2s are shared)
 - ▶ Currently CMS sees about 10GB used at the launch pad
 - ▶ Current recommendation for the deployed squid cache is 100Gb, so there is enough space
 - ▶ Important that information intended to be served from frontier is reasonably designed not to flush the whole cache
 - ▶ There isn't a concept of quota, ATLAS and CMS both have reasonable queries
- ▶ Monitoring of squids has up to now generally been done by the CMS team
 - ▶ Should think about how this evolves with a more general service
- ▶ There has been some discussions about using squids to support opportunistic VOs (at least on the OSG side)
 - ▶ This is interesting technology, but we should keep the critical information like conditions from opportunistic access
 - ▶ If someone wants to pursue squids for community activities, we should deploy 1 for critical activities and one for community