
Impacts of Cloud, HPC, and Opportunistic on Facilities

Gathering questions & areas to investigate to manage “on-prem” vs cloud resource provisioning

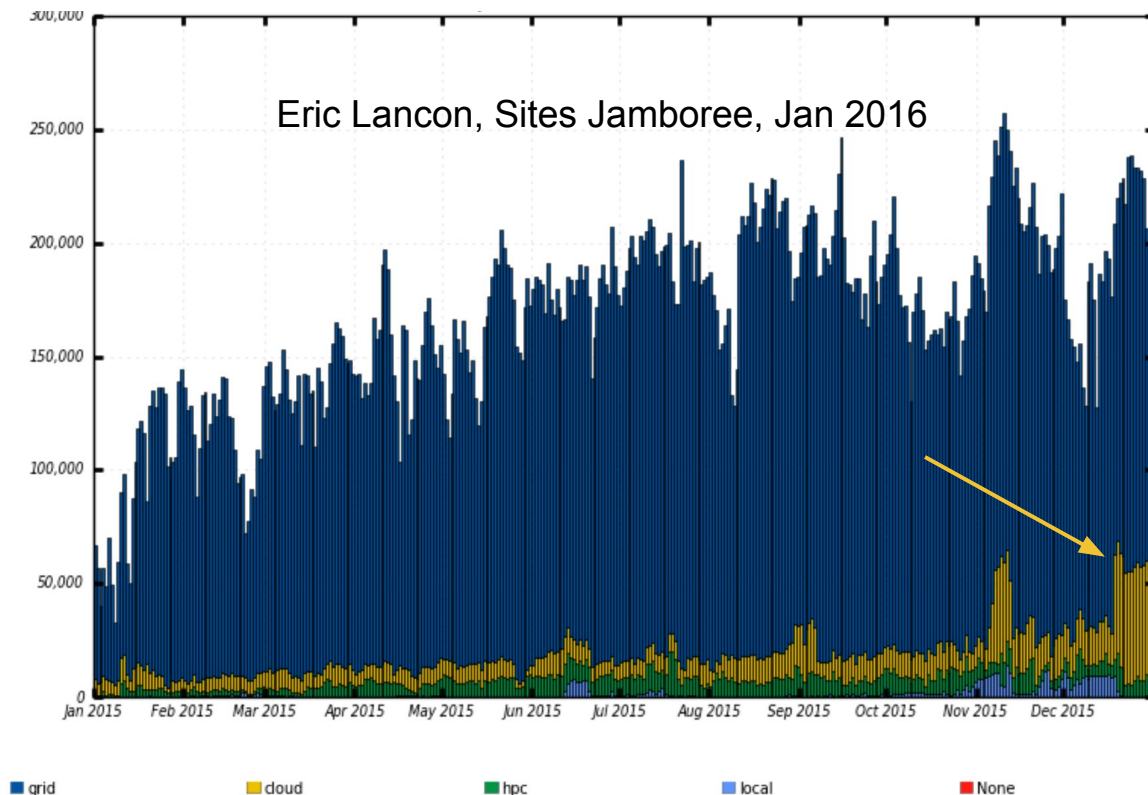
Rob Gardner • University of Chicago

US ATLAS Distributed Facilities Workshop, Clemson, March 14, 2016



Background: growth of non-grid

- Many years of cloud R&D in ATLAS
- In the past year have seen more efforts to access resources beyond the dedicated facilities with much success
- HPC - allocated - additional resource opportunity
- Economies of scale making cloud more competitive for some workloads



Questions for Tier1, Tier2 Facilities

- When do cloud economies of scale render traditional deployments obsolete?
- How to make accurate comparison of total costs given the many (local) factors that vary site to site?
- How to include estimates for workload types?
- Figuring in bandwidth and storage costs.

Questions, II

- Time dependencies?
 - Procurement decisions are made once or twice a year
 - Equipment of varying vintage and useful life
 - Dropping instance costs
 - Spot versus on-demand
- Vendor dependence of cost model?
- Storage capacity versus compute as a service
- Specifying what we **must** own versus lease