



Contribution ID: 1

Type: **Presentation**

Decarbonizing Scientific Computing

Scientific computing uses a tremendous amount of energy, and given the location of most HPC centers, results in a similarly large amount of CO₂ emissions. In the US, for example, in 2019, every MWh of generated power on average led to 0.7 metric tons of CO₂ emissions. To address the huge carbon footprint of computing Lancium Compute is building, low carbon, renewable-energy-driven data centers in the Great Plains of the United States.

Because the wind does not always blow, and the sun does not always shine, our data centers must be able to rapidly ramp our computing and electrical load up and down in order to balance the electrical grid. Not all applications are suitable for this sort of load management. However, many batch based scientific jobs, e.g., High Throughput Computing jobs, are ideal.

In 2021 we began to work with the Open Science Grid to support their HTC load. As part of that integration effort we added support for CVMFS for all containerized jobs. Later work with the US CMS and US ATLAS team led us to further deploy a hierarchical Squid architecture to support Frontier.

In this talk I will briefly present electrical grid basics, and explain how the characteristics of renewables make them difficult to integrate into the grid. I follow with a discussion on how controllable loads can solve these problems, and how computing can be an excellent controllable load. I will then describe our quality of service model and multi-site system architecture with CVMFS to support high throughput jobs single node jobs and low-degree parallel jobs at our clean compute campuses in Texas.

Primary author: GRIMSHAW, Andrew (Lancium Compute)

Presenter: GRIMSHAW, Andrew (Lancium Compute)