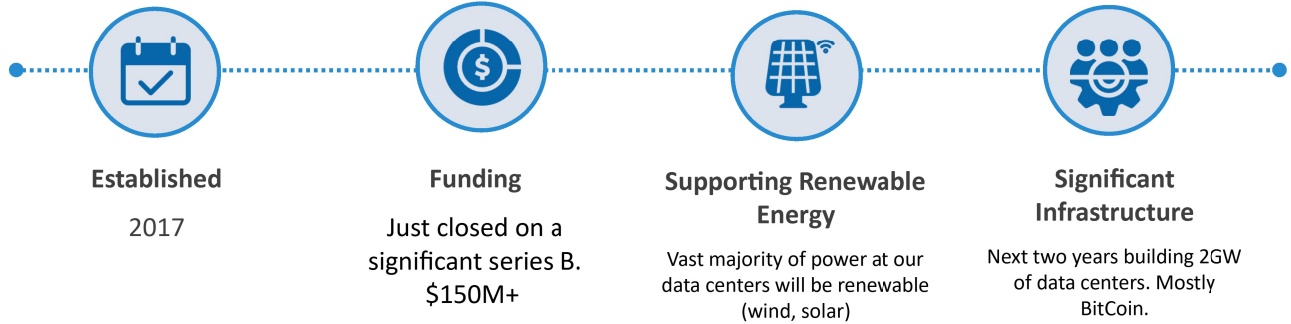


Lancium


Software and technical solution that enable the faster decarbonization of the electrical grid



LANCIUM COMPUTE CAMPUSES

Our Data Centers

- Located in high renewable density zones where price of energy is often negative.
- Rampable .. We can very rapidly drop and add load so as to help stabilize the grid. Revenue from providing these “ancillary services” can be a significant fraction of the energy cost.
- All-in electrical rate 4-6 cents per KWh depending on QoS.



West Texas Congestion



2013

Competitive Renewable Energy Zone reactive compensation plan was designed to accommodate **14 GW** of West Texas renewable generation



2021

>25 GW of renewable generation is expected to be connected in West Texas

2012 8,220 MW

2014 8,634 MW

2016 9,842 MW

2018 11,606 MW

2020 13,335 MW

2021 18,687 MW

West Wind

Y/Y
+40%

2012 10 MW

2014 32 MW

2016 296 MW

2018 1,325 MW


2020 4,246 MW

2021 7,010 MW

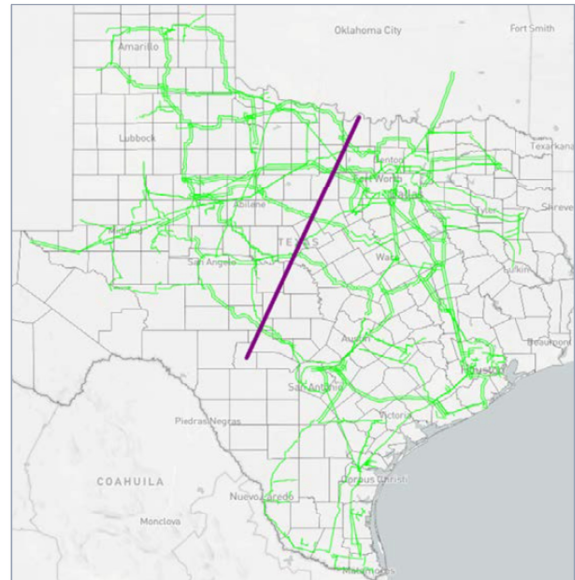
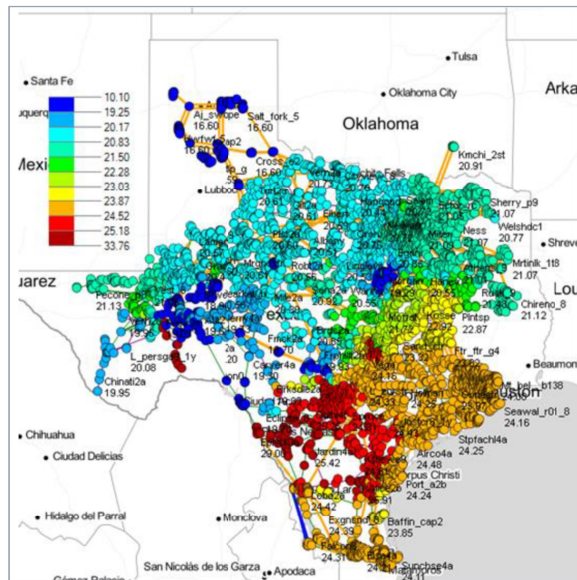
West Solar

Y/Y
+65%

~20 GW of renewable energy



Annual Average Nodal Heat Map - 2019

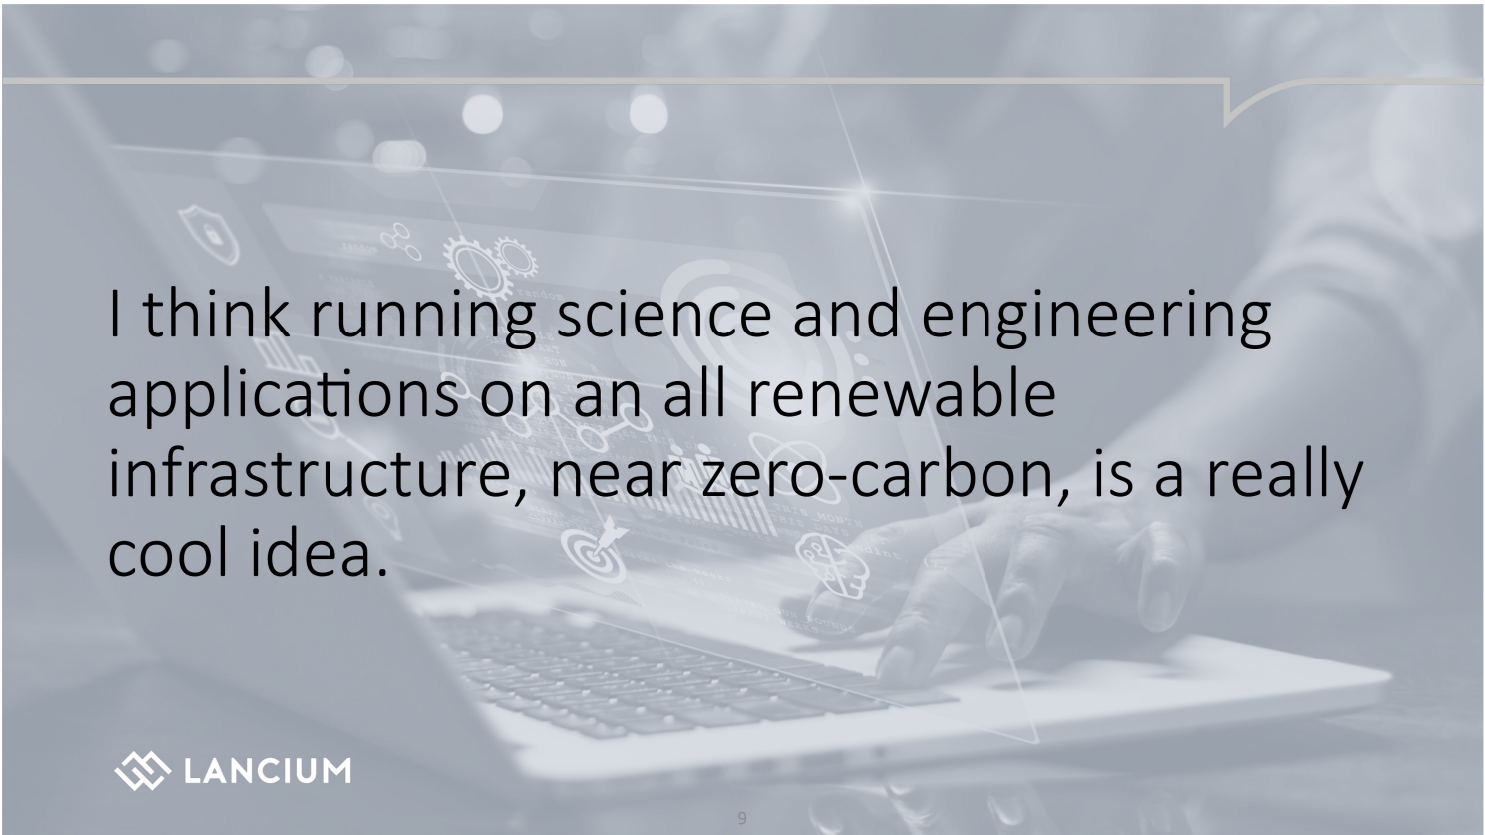


Interested in how computation can accelerate the renewable transformation: google uva cs grimshaw renewable

Lancium Renewable Sites Start Coming Online 2022-2023

- Ft Stockton, Abilene, Childress
- TTU/Lubbock (Research facility)
- The new sites (except TTU) will have 100 Gbs networks.
- Ft. Stockton will have 73 MW in late 2022. Abilene will have 48 MW early 2023. Increasing in 48 MW chunks thereafter.





I think running science and engineering applications on an all renewable infrastructure, near zero-carbon, is a really cool idea.



9



**LANCIUM COMPUTE
SERVICES**

10



Lancium Compute Overview



Containers - RCLS - **Lancium Remote Command Line Services** provide a Linux command line service using Singularity containers that allows users to run thousands and thousands of computations remotely on LCI resources. The command line executes in the context of a Singularity container of the users choice.



DS - Lancium **Data Services** provide customers with a global-scale distributed file system in which data stored at Lancium sites, customer sites, and customer-partner sites can be accessed securely from anywhere with a network connection.

11



Lancium Compute – Compute Services

- Designed for high throughput computing and low-degree-parallel applications.
- Singularity container-based computation model
 - Users register their own container images or use pre-configured images for popular applications such as Amber, Tensorflow, etc.
- There is both a CLI and a REST API.
- Users specify the job resource requirements, e.g., 8 GB of memory, 8 vCPUs, and the command line to run, files to stage in and out, e.g.,

NEWS FLASH
Lancium now
supports MPI
applications with up
to 90 nodes.

```
lcli job run --name "BLAST Tutorial job" --command "blastp -  
query \  
cow_db_small/cow.1000.protein.faa -db  
human_db/human.1.protein.faa -evalue 1e-5 \  
-max_target_seqs 1 -num_threads 4" --image lancium/blast --cores  
4 --mem 8 \  
--input-file cow_db_small.tar.gz --input-file human_db.tar.gz  
- returns (among other things) a job number  
lcli job show <job number> -- gives an update on the job  
lcli job terminate <job number> -- terminates the job
```



Lancium Compute – Data Services

- Staging – data is staged from/to customer sites to run jobs
- Lancium storage – each user has a directory structure, akin to home, where they can put/get data, and that can be a source/target for job data.
 - Campus file system data resources can be securely linked (subject to access control) into user and organizational directories.
- Read Only Cache – users can specify a sub directory of their storage area that can be replicated mounted as a read only file system for their running jobs.
- SubCluster and node storage.
 - Job input data is staged into a job working directory on a subCluster head node and mounted in /jobDir in the running Singularity container. /jobDir is the current working directory when the job is started.
 - Job files/directories to be staged out must be in /jobDir when the job terminates.
- Each node has ~860GB of local scratch storage mounted into /tmp in the running Singularity container.
- /cvmfs mounted in each Singularity container
- Squid/Frontier services also available.

13



Lancium Compute – Business Aspects

- No ingress/egress fees
- Firm contracts. Lancium takes the over-use risk.
- Rack Rate: \$0.01/core hour, \$0.10/K80 hour, \$10 TB/month for persistent data.
- Quality of Service model
 - Current
 - On-demand | Best Effort
 - Real Soon now
 - On-demand (aka 100%) | 50% | 25% | Best Effort
- Significant QoS, volume and prepay discounts.
 - Up to 40% off.

14



OSG AND LANCIUM


OSG

The OSG job mix is a “poster child” for the type of load we want to run. High throughput, interruptible.

We run squid servers. All nodes mount /cvmfs.

“Production” version OSG Docker backfill containers run in 48 vCPU Singularity containers with just shy of 96GB memory. /cvmfs mounted.

We donate ~1M vCPU/hrs a week



CAMPUS COMPUTING AND THE COMPUTING CONTINUUM

17

LANCIUM



CC* Part 4, “Campus Computing and the Computing Continuum”

- Working off last years RFP
- \$400,000 over two years
- “Proposals are required to commit a minimum of 20% shared time on the cluster and describe their approach to making the cluster available as a shared resource external to the campus. One possible approach is ... the Open Science Grid.”
- “Campus Computing solutions that integrate cloud services are particularly encouraged.”

18

LANCIUM

Open Science Grid

- We have a good working relationship with the OSG and can handle providing OSG with the compute resources to satisfy your commitment to contribute to a shared resource.
 - Eliminates your need to set up and manage OSG software on your equipment.

19

LANCIUM

Lancium as a Cloud Provider

RFP explicitly allows the use of cloud resources.

believe we are the lowest cost, and most environmentally friendly, cloud provider.

on-demand rack rate, \$400,000 is 40M core hours, 80M vCPU hours.

Far more at lower qualities of service.

Even more with our "OSG special pricing".

and advantages:

You can run far more jobs requiring more resources at any given instant than you could get with \$400,000 worth of hardware.

When you are not using resources, you are not paying for resources.

You can add GPUs when needed.

No system administrator time required, use your precious human resources to work with end users, not keeping the HW/SW stack operational.

20

LANCIUM

Lancium "OSG Special"

	On-demand Rack Rate (cents/hr)	On-demand OSG Special	Best-effort OSG-Special
Core hour, 2-vCPU/hr	1	0.8	0.5
K80	10	8	4
A6000	100	80	40

These terms are good both for proposals to NSF you might make for CC*, or in general for the next six months if you ask us for the "OSG-Special."

CONCLUSION

Lancium & OSG



There is a massive renewable energy transformation underway.



Computation is one of the industries that can most easily exploit the massive amounts of stranded energy in the Great Plains.



Lancium is building renewably powered data centers in the Great Plains to leverage the on-going transformation.



Lancium has demonstrated the ability to run high-throughput OSG jobs.



You won't find a better price for HTC cycles than Lancium.



Thank You

Andrew.grimshaw@lancium.com

