Sustainability, Lancium, and HEP

Dr. Andrew Grimshaw
Lancium Compute
Emeritus Professor of Computer Science, University of Virginia

September 5, 2022

LANCIUM
BALANCING ENERGY

I am not a HEP person ….

- I’ve been working with computational scientists for three decades.
- 2013, my sons and nephews in the LHC
- Only time my youngest ever said something related to my job was “cool”.

LANCIUM
BALANCING ENERGY

Agenda

- The message
- The renewable transformation of the Texas electricity markets
- The challenges of renewable power
- Addressing the challenges
- Lancium Clean Campuses – 2 GW of renewable power
- Conclusion

LANCIUM
BALANCING ENERGY

The Message

- Renewables are the least expensive way to produce power in the world today.
- The pace of renewable deployment is staggering.
  - And gated only by the ability of electrical grids to safely absorb the power
- Controllable loads, particularly loads that can be rapidly reduced, are critical to maintaining grid stability in the presence of non-constant solar and wind power.
- Computation can be used as the controllable load.
- Lancium is building GW scale clean campuses that provide very inexpensive, mostly carbon-free, electricity for computation.

1: “Renewables 2020: Analysis and forecast to 2025”, International Energy Agency
2: Lazard, Levelized Cost of Energy and Levelized Cost of Storage – 2020

LANCIUM
BALANCING ENERGY

The Renewable Transformation - Texas

LANCIUM
BALANCING ENERGY

Texas - ERCOT

- Renewables are exploding in Texas
- Partly because of Texas's wind and solar assets: in other words, the wind blows a lot in the plains, and the sun shines most days.
- Partly because of their market structure that incentivizes producers to produce; in Texas you can charge what you want, and the market clears to the lowest bidders.
    - Renewables are cheaper than the fuel price of natural gas, oil and coal.
    - Subject to the constraint that there must be a portion of generation that is "synchronous".
- Let's look at Texas/ERCOT in a bit more detail

LANCIUM
BALANCING ENERGY

ERCOT WIND & SOLAR – ANNUAL ADDITIONS

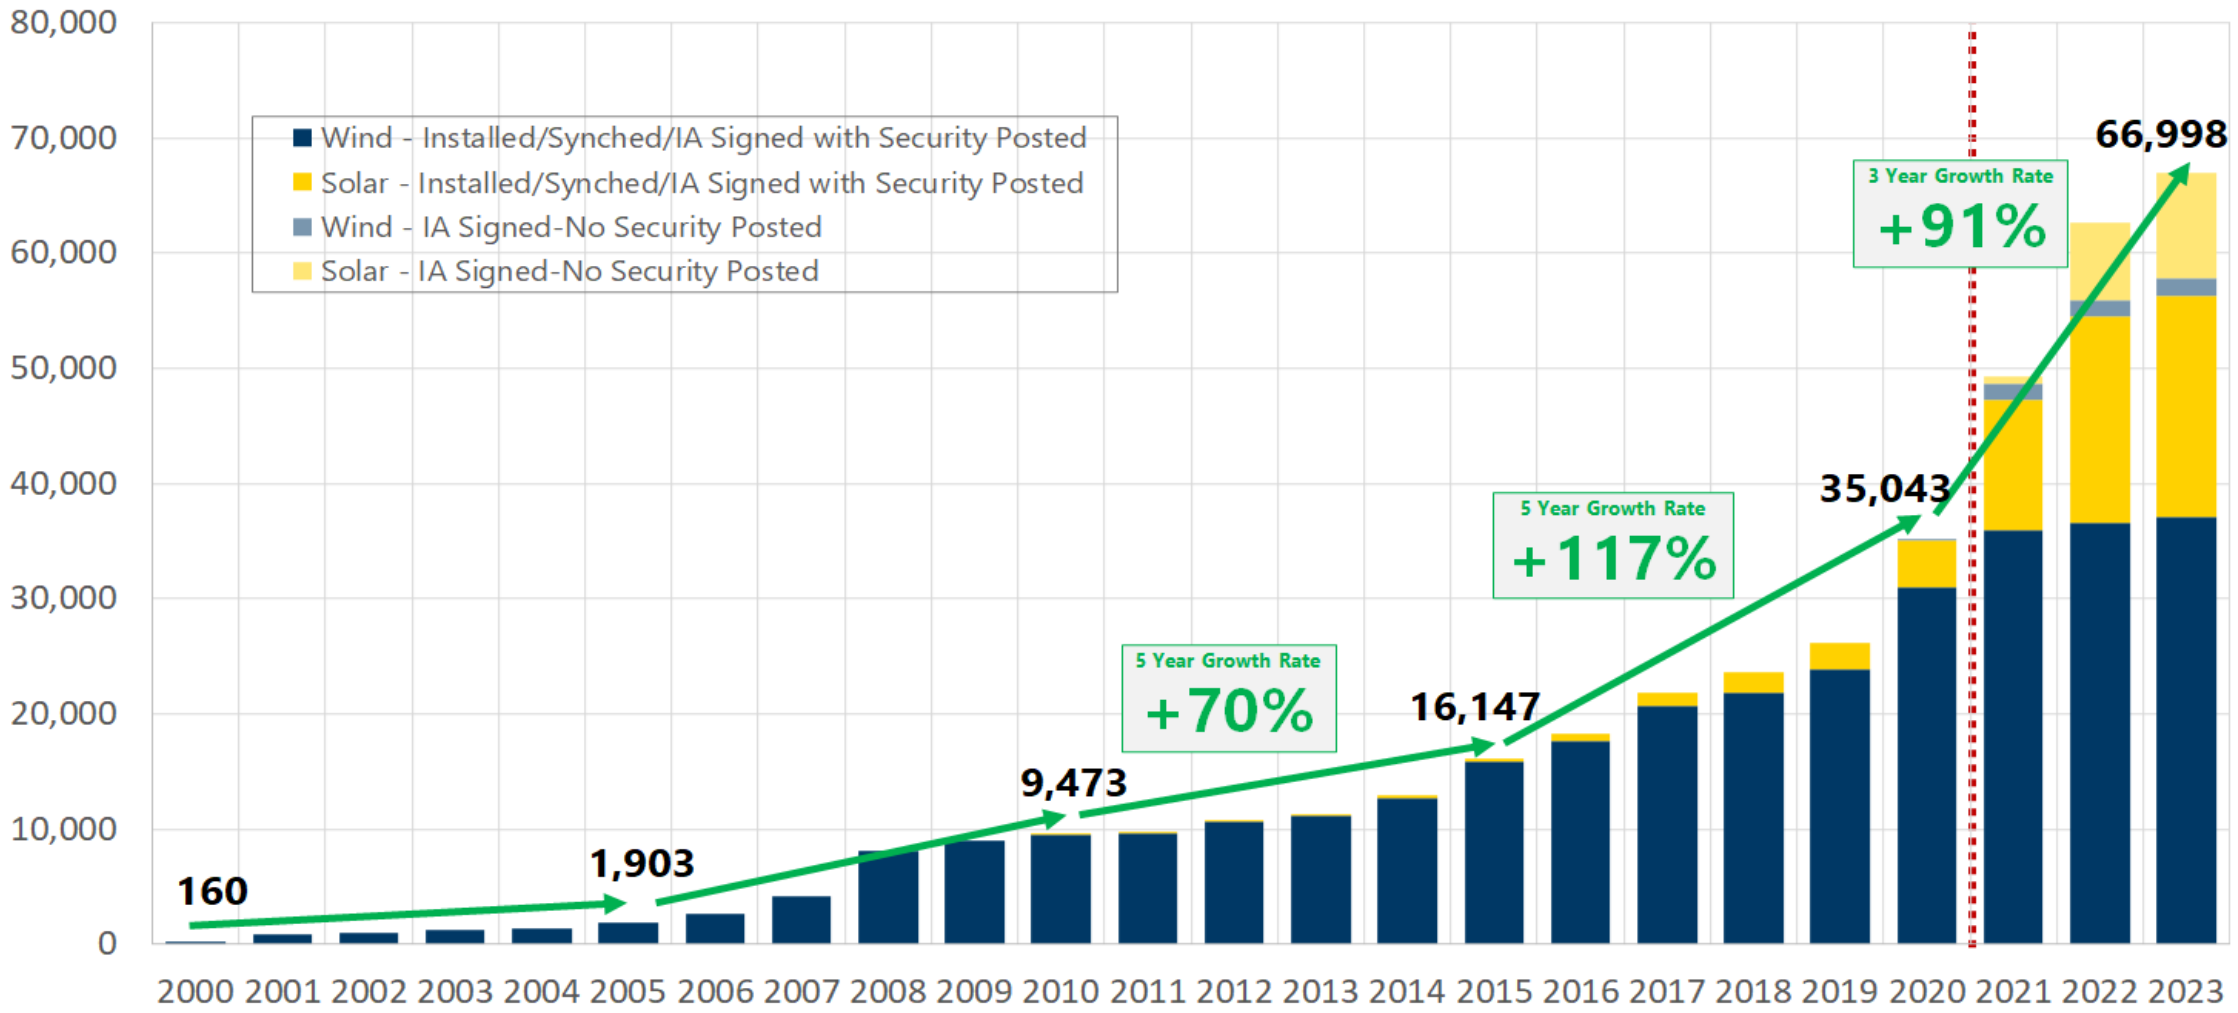

ERCOT's Rapid Renewable Generation Growth...

Percent Energy (MWh) Generated in ERCOT by Fuel Type

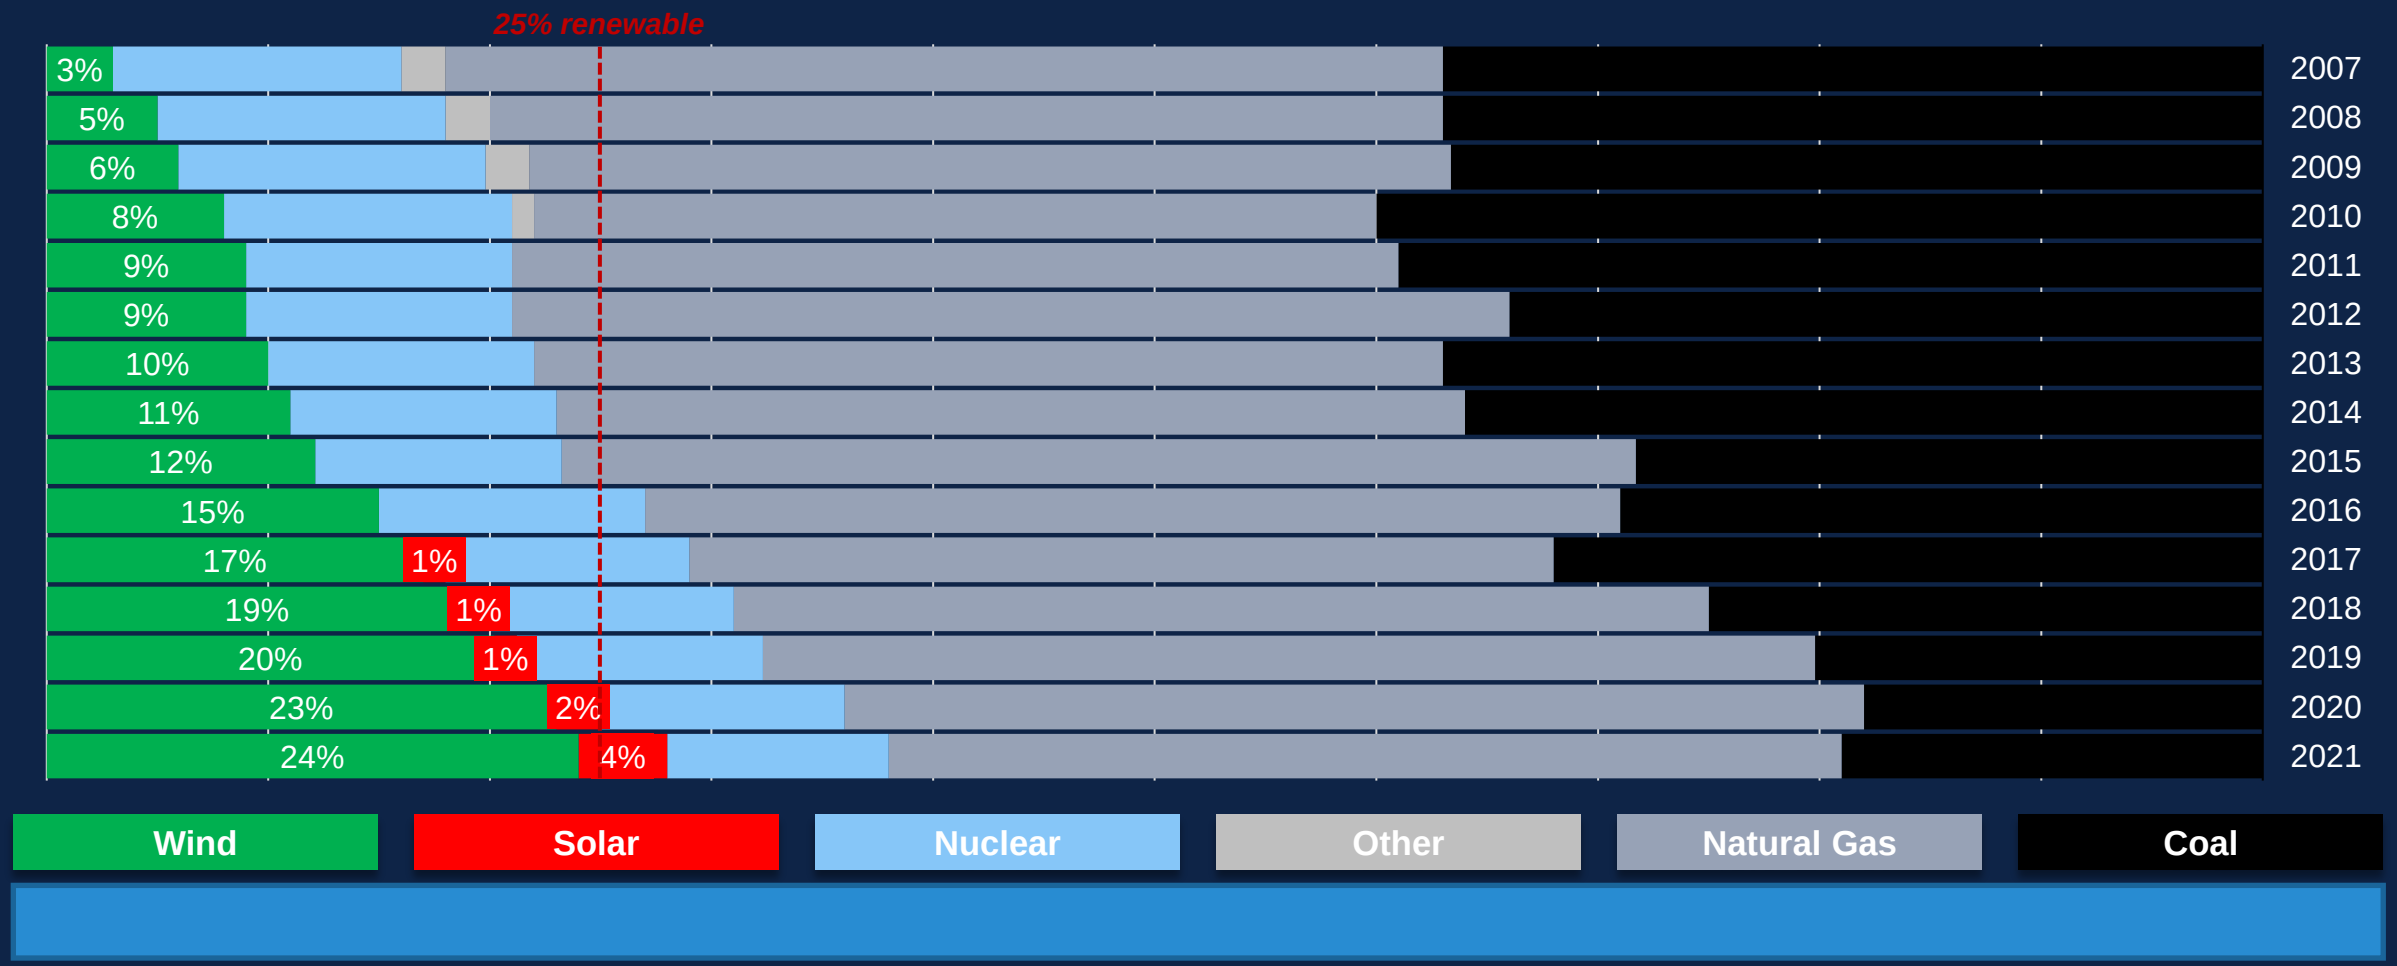

LANCIUM
BALANCING ENERGY

Source: ERCOT - Market Data

West Texas Load vs Installed Capacity by Technology

- West Texas renewables growth is outstripping load growth

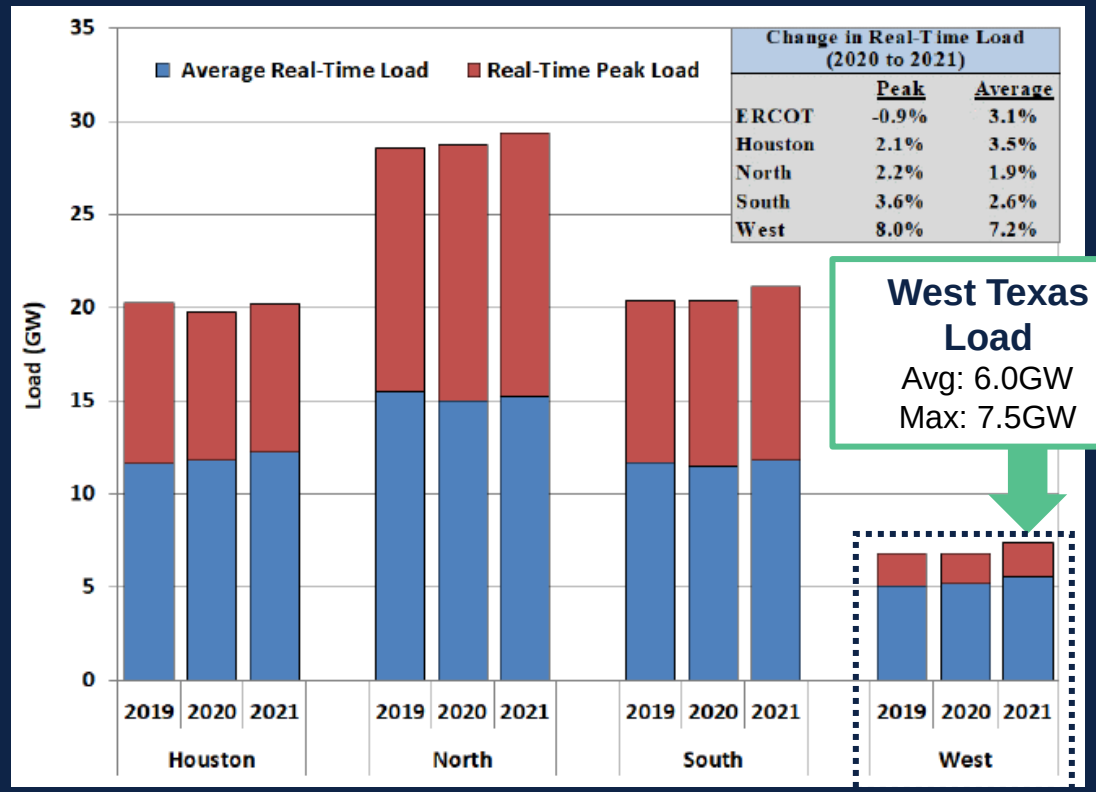

Change in Real-Time Load (2020 to 2021)	Peak	Average
ERCOT	-0.9%	3.1%
Houston	2.1%	3.5%
North	2.2%	1.9%
South	3.6%	2.6%
West	8.0%	7.2%

West Texas Load
Avg: 6.0GW
Max: 7.5GW

West Texas Supply
Wind: 24GWs
Solar: 7GWs
NatGas: 4GWs

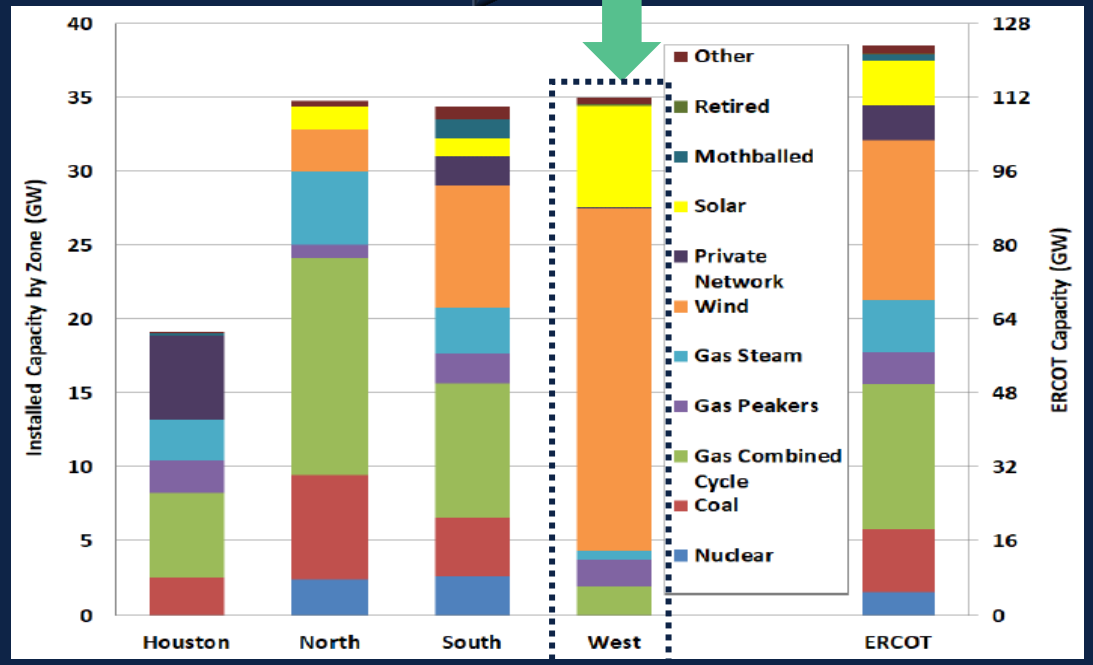

In 2021, West Texas load averaged 6GWs with installed capacity at 35GWs (24GWs Wind, 7GWs Solar, 4GWs NatGas)

LANCIUM
BALANCING ENERGY

Texas West Load Zone

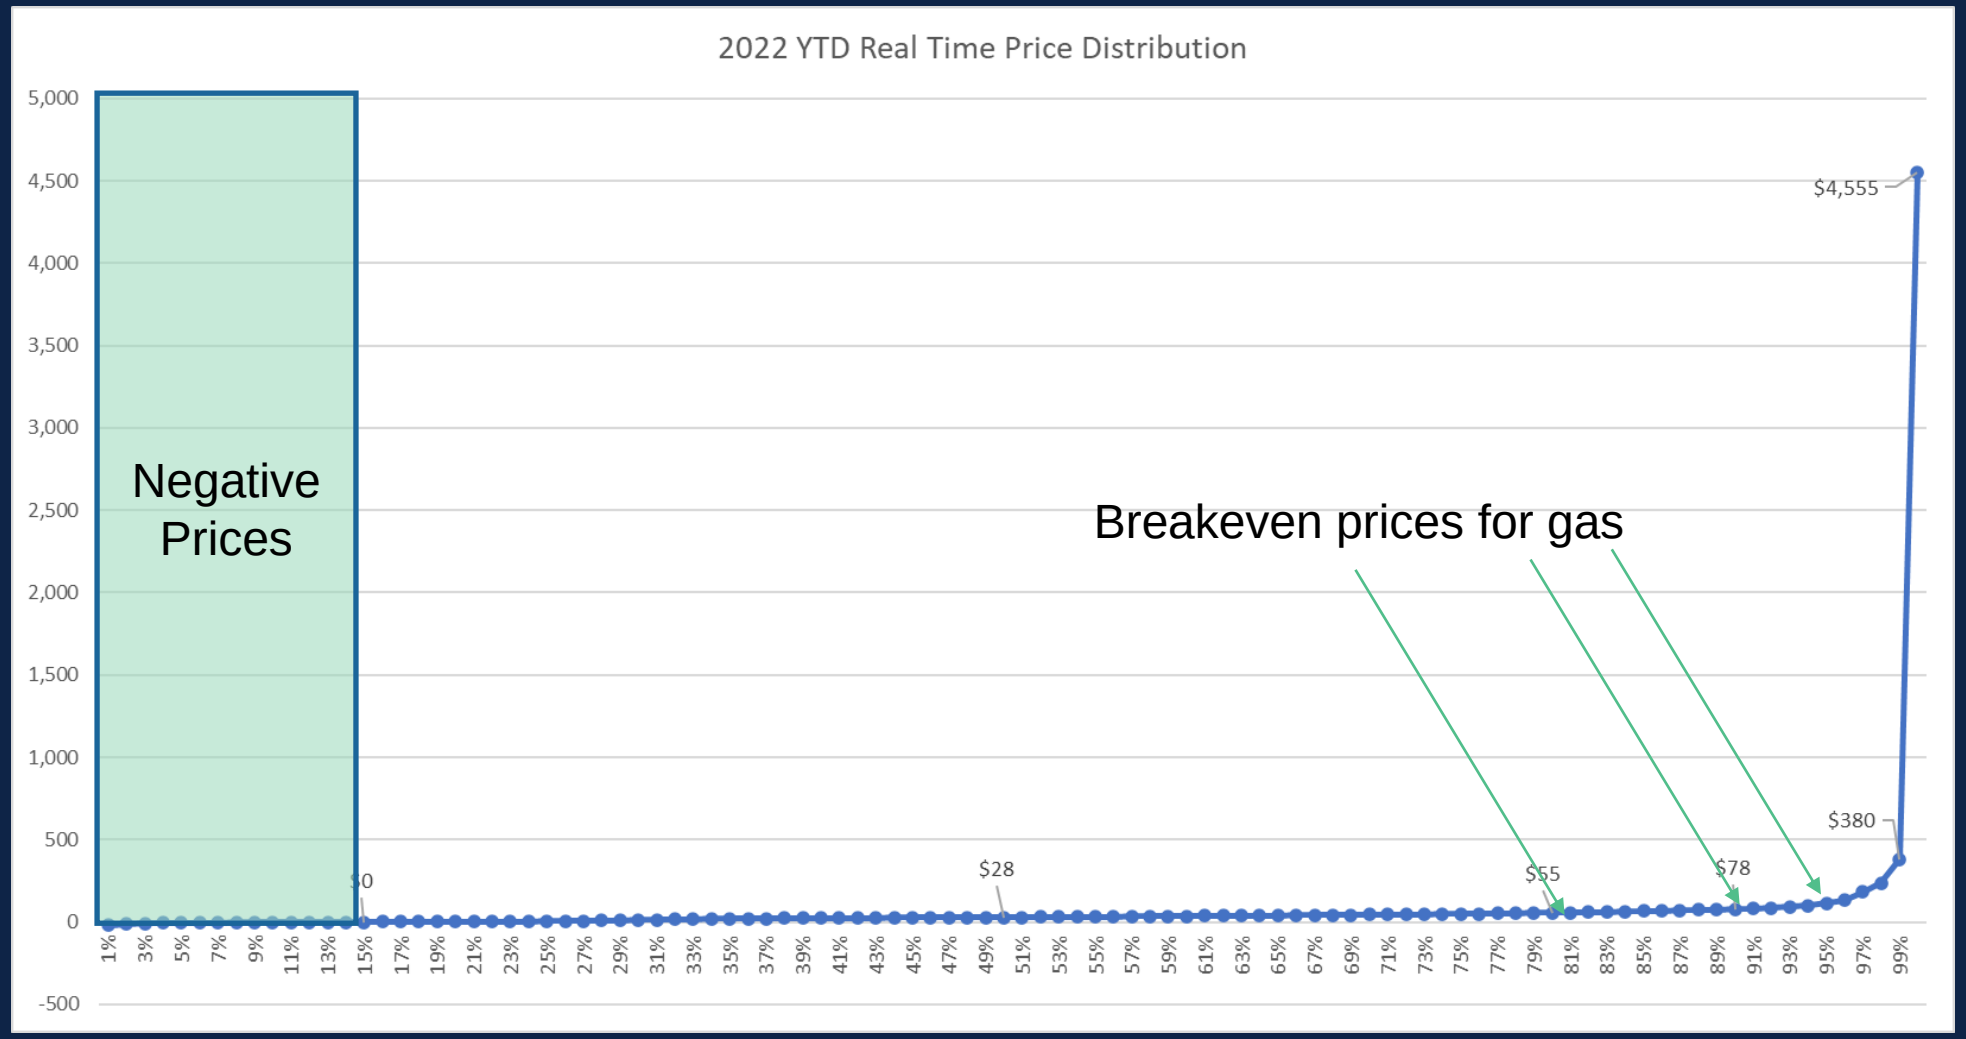

LANCIUM
BALANCING ENERGY

Challenges with Wind/Solar

LANCIUM
BALANCING ENERGY

Challenges

- Congestion – the load (people, industry) is far from the generation.
- They are variable and intermittent.
  - Solar is (obviously) diurnal, varies with the seasons, and can change rapidly due to clouds.
  - Wind varies constantly, though has trends:
    - Stronger in winter than summer. Stronger at night.
- As load approaches generation, frequency starts to drift (called an excursion). If load not reduced, or generation increased, generators can be destroyed. The grid operator will not let this happen. They are taken offline to prevent damage. This can lead to blackouts to reduce load.
  - Because of this grid operators (ERCOT) keep a significant amount of reserve capacity that can come on-line very quickly (~15s), quickly (~15 minutes), slowly (~3 hrs).
  - They pay for this .. It is called an “ancillary” service.
  - “Capacity” can be increases of generation OR decrease of load.
- In the absence of massive (and expensive) energy storage systems, or the ability to rapidly shed load, or increase production, the amount of renewable power the grid can absorb is limited without risking instability and blackouts.

Congestion

LANCIUM
BALANCING ENERGY

ANNUAL AVERAGE NODAL LMP HEAT MAP - 2019

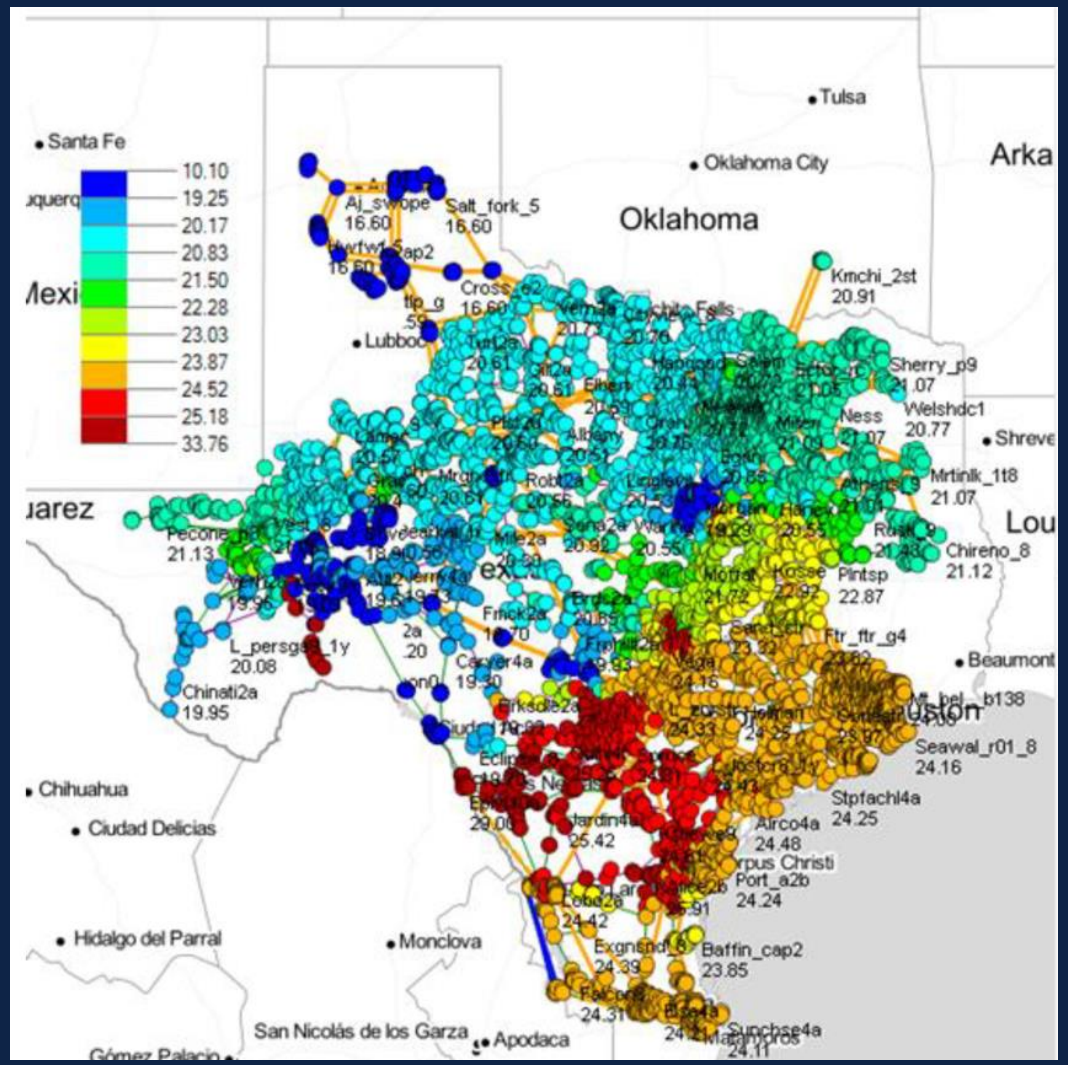

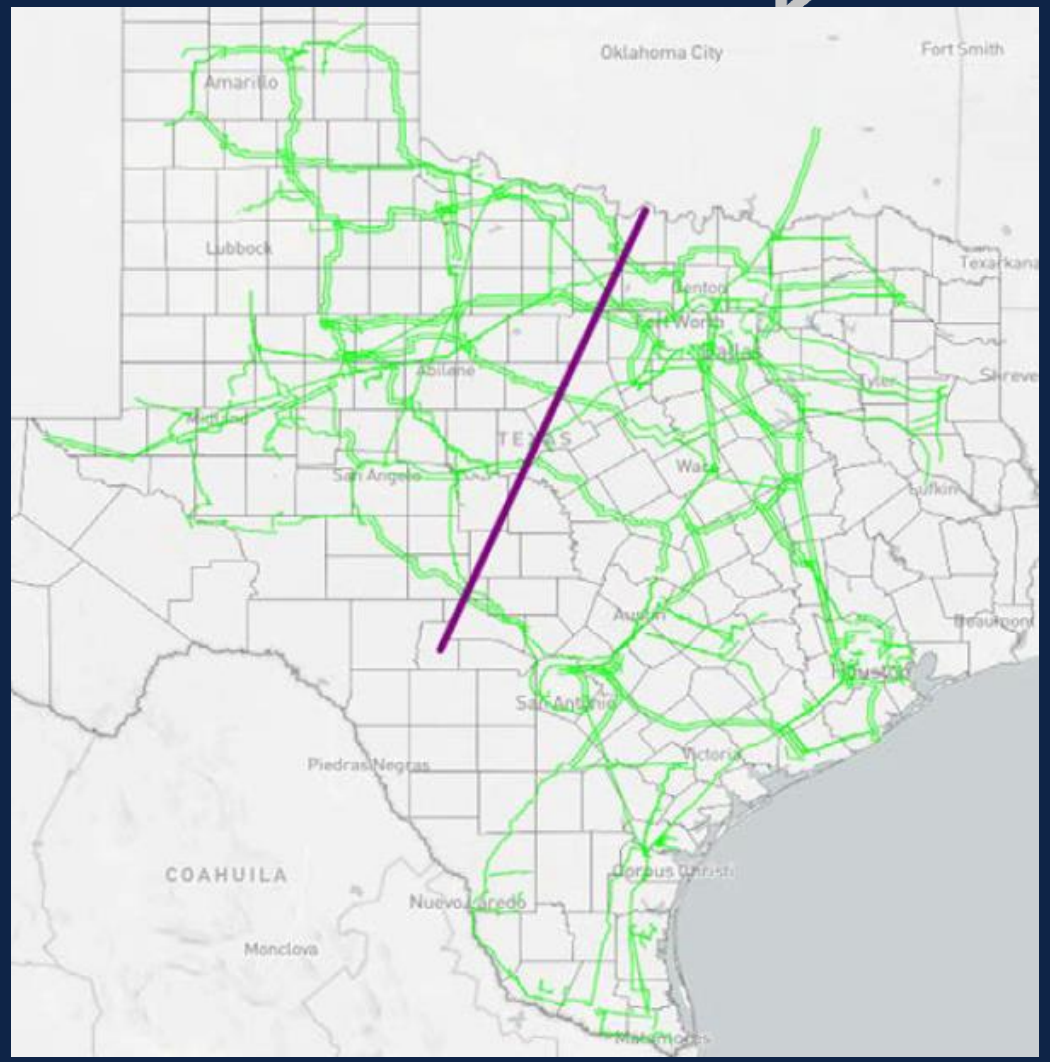

LANCIUM
BALANCING ENERGY

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

West Wind	
2012	8,220 MW
2014	8,634 MW
2016	9,842 MW
2018	11,606 MW
2020	13,335 MW
2021	18,687 MW

Y/Y +40%

West Solar	
2012	10 MW
2014	32 MW
2016	296 MW
2018	1,325 MW
2020	4,246 MW
2021	7,010 MW

Y/Y +65%

BALANCING ENERGY

Solar 7GWs | Wind 19GWs | NatGas 4GWs

West Supply: 30GWs

5GW

West Load 5GWs

Export Limit 15 GWs

West Texas Export Constraints
16 Tie Lines – Max 12-15 GWs

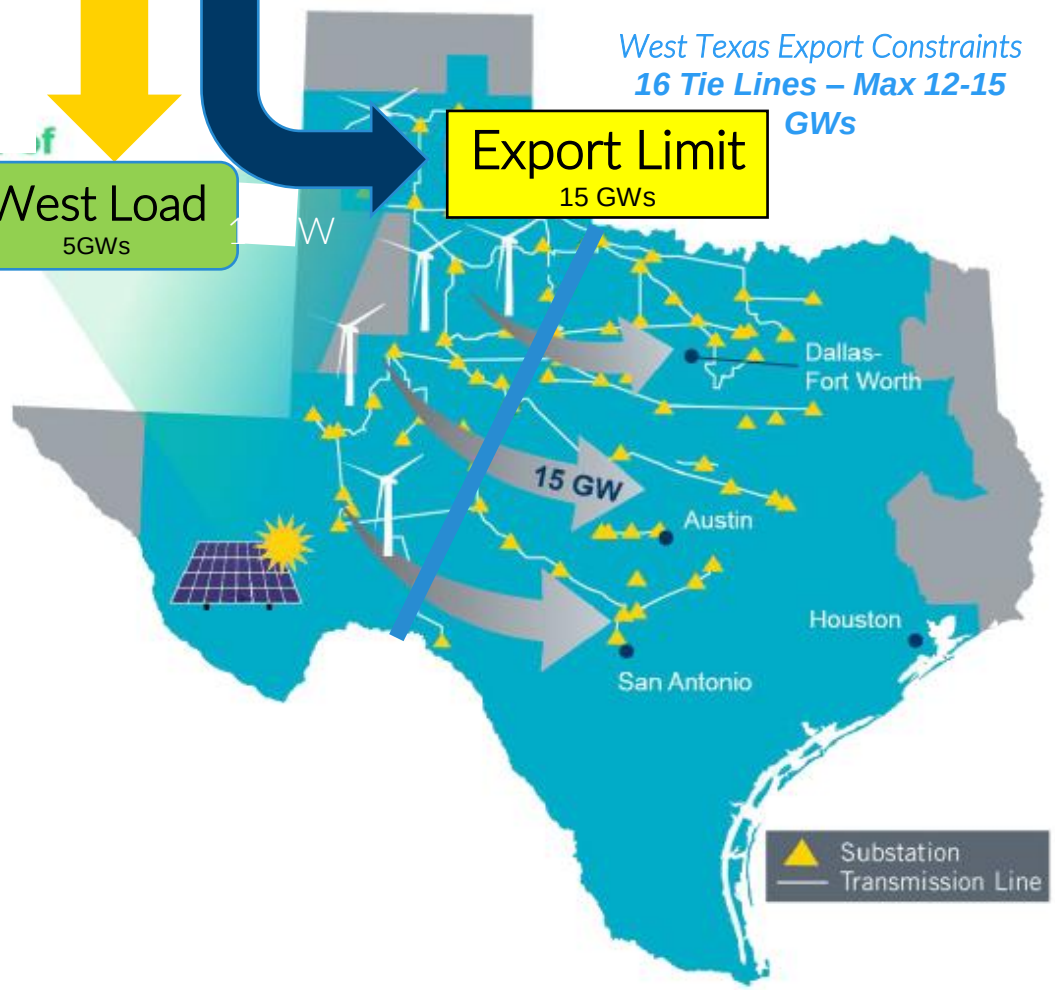

Congestion leads to curtailment

- When a generator is told not to produce energy it is being curtailed.
- This happens all the time
  - That is why you often see a wind farm with only one or two turbines turning, the rest have been curtailed.
- Generators do not get paid for curtailed power.
  - Makes renewables less profitable … meaning less will be deployed.
  - It is also a waste of carbon-free power!
- Surely there is something we can do with that curtailed power? ☺

LANCIUM
BALANCING ENERGY

Intermittent Output

Sometimes the wind does not blow, and the sun does not shine so much.

LANCIUM
BALANCING ENERGY

Wind power production functions

- On the right is data from a small, 150MW, facility in the panhandle near Lubbock.
- One rarely gets the full capacity out of a wind system.
- Below that is a graph that shows what % of time you get a certain amount of power.
- Wind power can vary very rapidly.

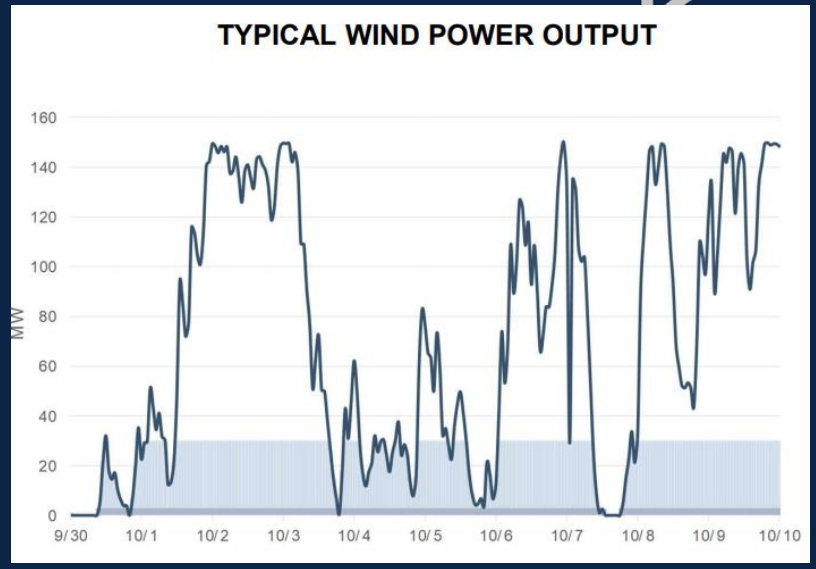

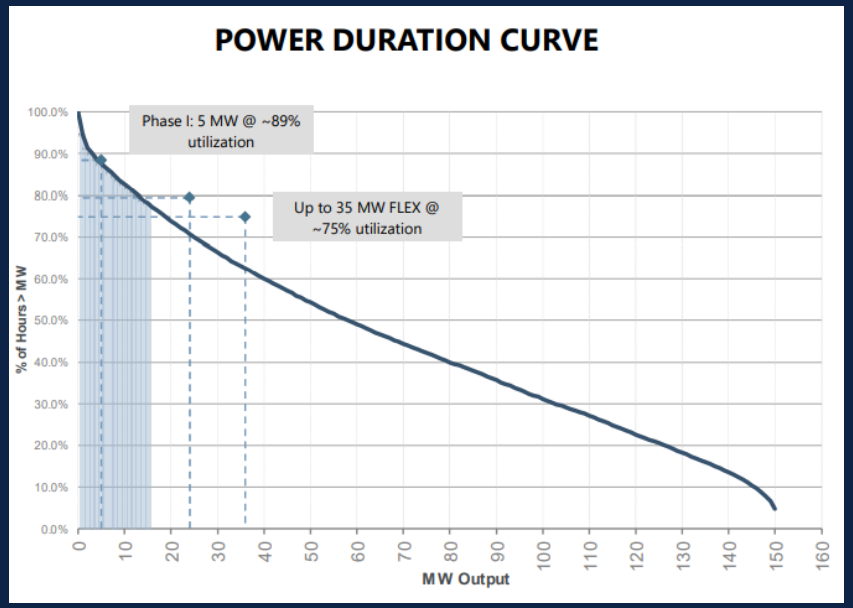

LANCIUM
BALANCING ENERGY

Wind and Solar are complimentary

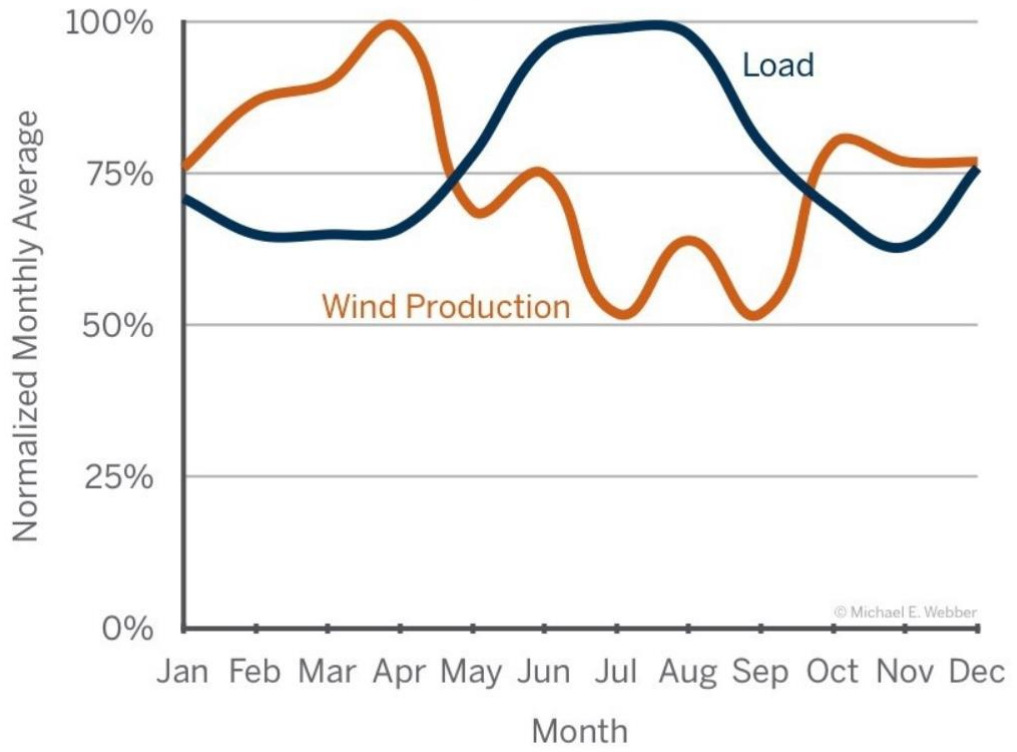

Figure 119: Continental wind production is lowest during the summer months, when demand for electricity is at its greatest.

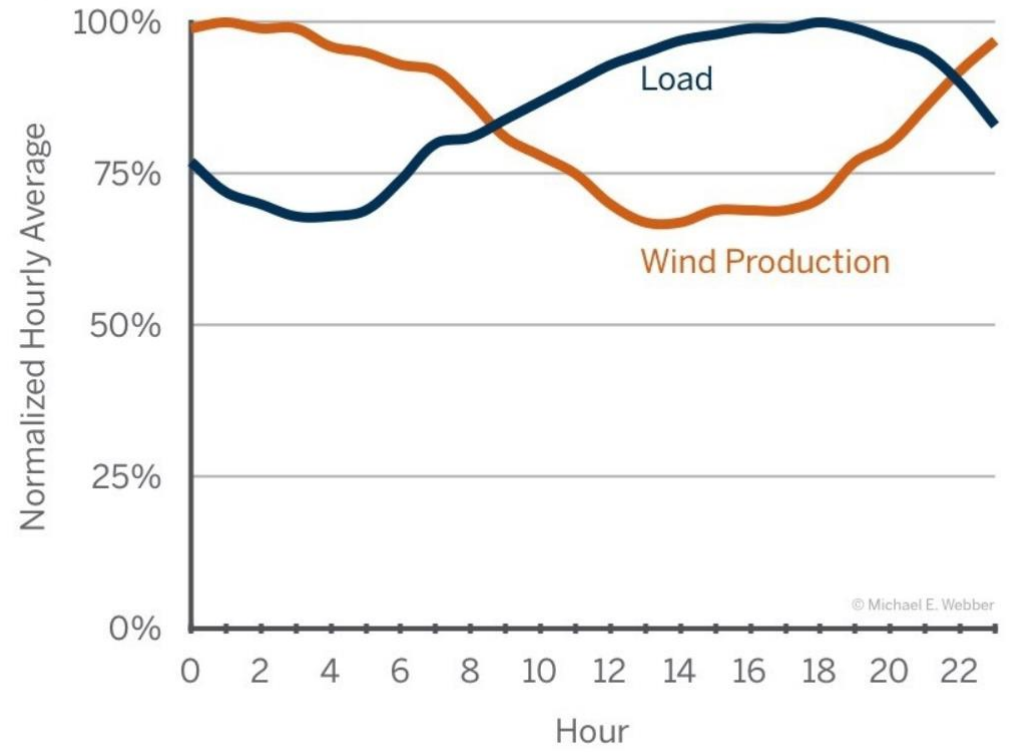

Figure 120: Continental wind is also out of alignment with demand during the day, with production being lowest during the times with the most demand.

LANCIUM
BALANCING ENERGY

Hourly balancing challenge for ERCOT – March 18, 2022 – March 20, 2022

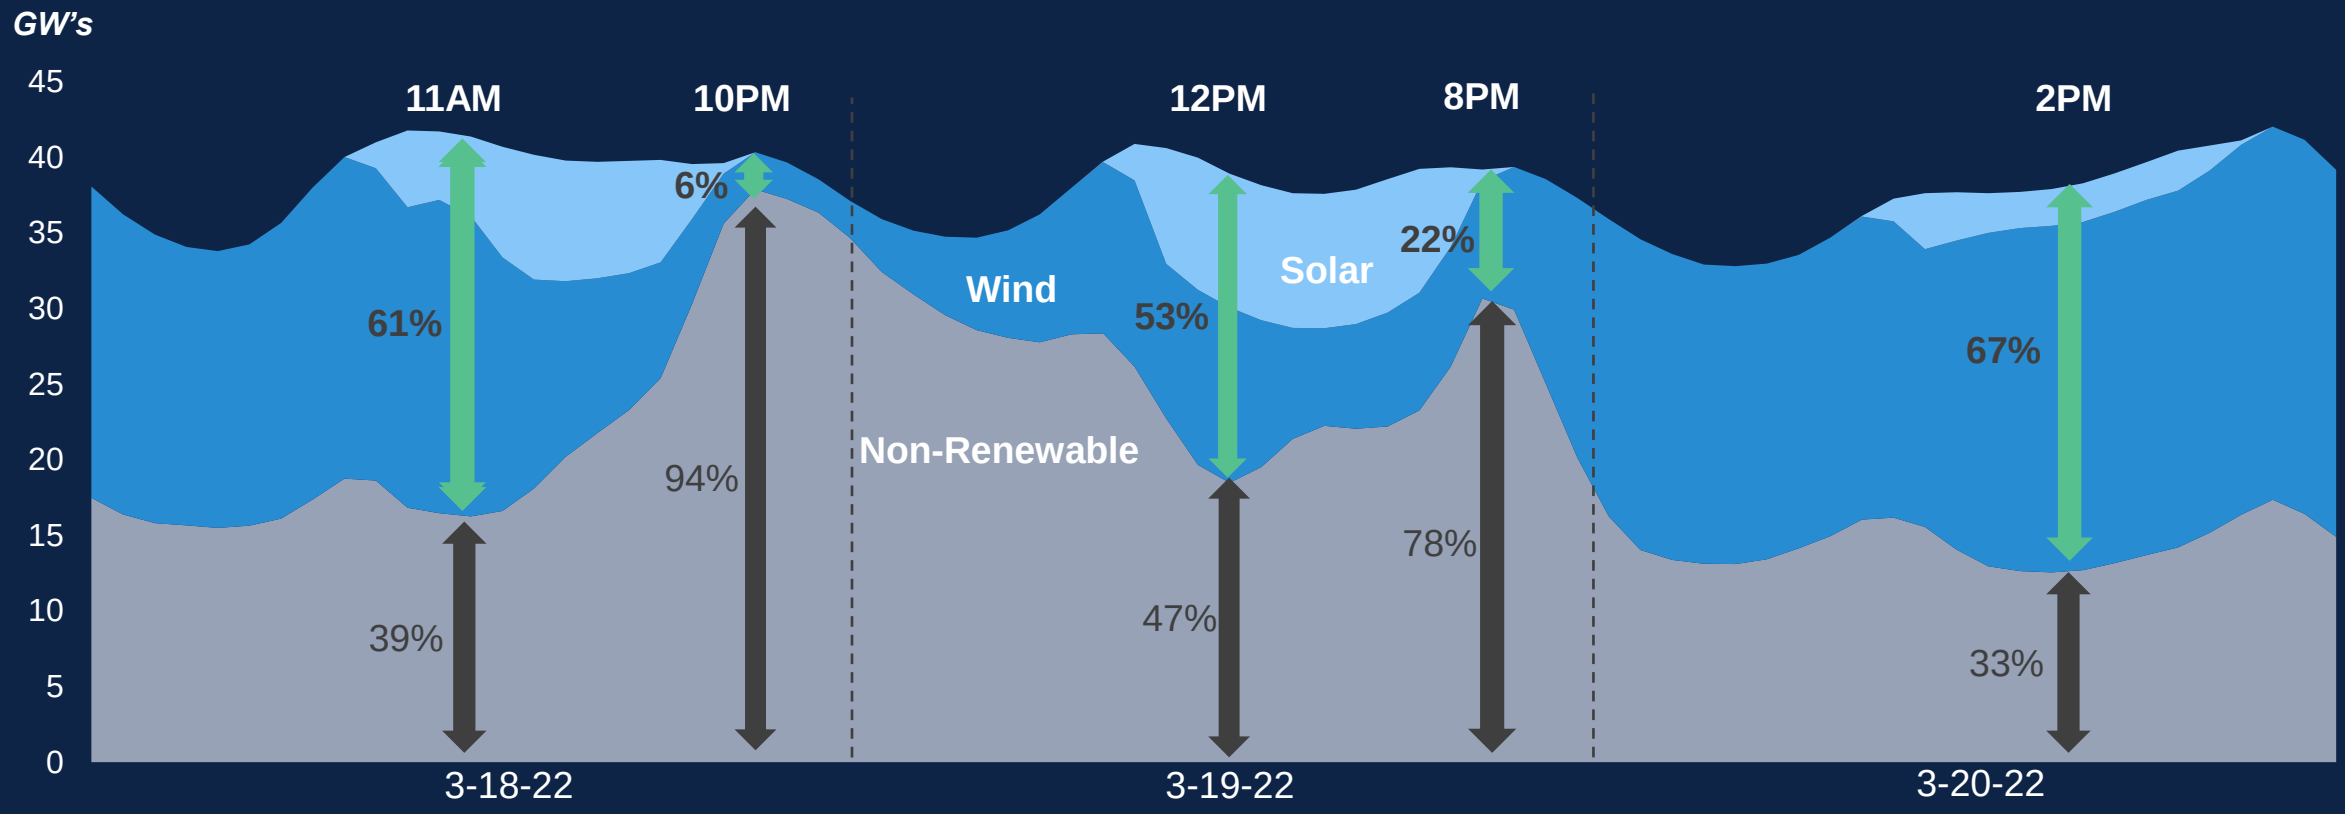

Batteries and load flexibility can provide short-term balancing throughout the day

LANCIUM
BALANCING ENERGY

Sources and Notes: ERCOT

Addressing the Challenges

LANCIUM
BALANCING ENERGY

To address congestion and variability we need:

- To move load close to production.
  - There are many industries that consume huge amounts of power: aluminum, steel (arc furnaces), heavy manufacturing, computing (think about training large models).
  - But most require significant infrastructure and logistics to operate. Computation does not.
  - We're going to focus on computation in close proximity to the power generators.
- A controllable load. We need to be able to ramp the controllable load down when the renewable resources are not producing enough – or are forecast to be unable to produce.

LANCIUM
BALANCING ENERGY

Controllable load

- A controlable load is a load that can ramp up and down on command.
- We can use a large (multi GW) compute load to soak up excess power when renewable energy is plentiful and cheap, and reduce load when renewables are scarce and electricity is more expensive.
- By using the excess energy at a **nonzero** price renewable generation becomes more profitable, leading to more renewable generation.
  - Thus, when the region is at, say 10% of generation capacity, it is 10% of a larger base, meaning less brown energy is needed.
- It will also lead to less carbon-based generation!

One possible use:
World's best glider ride!

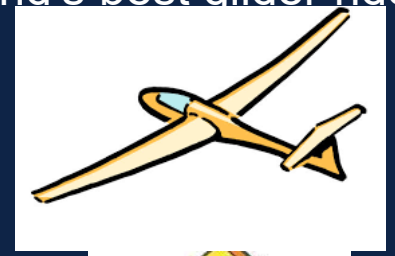

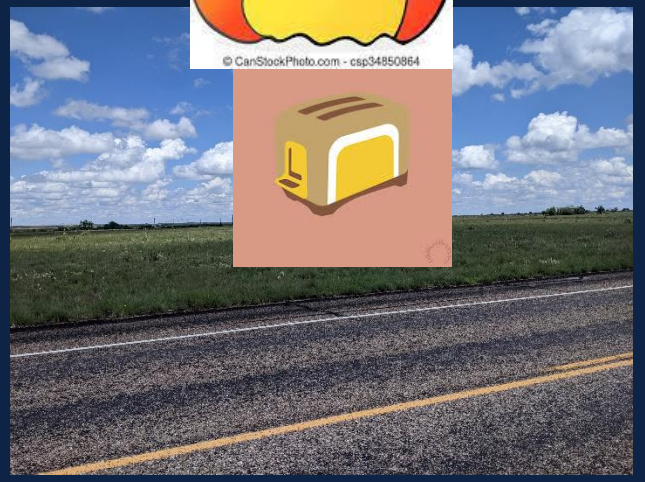

LANCIUM
BALANCING ENERGY

Lanncium

Mission: To decarbonize the grid.

LANCIUM
BALANCING ENERGY

Who we are

Lancium is a leading energy technology and infrastructure company ***dedicated to accelerating the energy transition.***

What we do

We create technologies and build infrastructure that enable more clean energy production while balancing and stabilizing the power grid.

LANCIUM

- Founded in 2017, HQ : The Woodlands, TX
- Two-large scale Clean Campuses (500 MW renewable energy) under construction in Ft. Stockton & Abilene, TX
- First company to commercialize a Controllable Load Resource (CLR) in ERCOT
- Industry-leading Lancium Smart Response™ technology protected by a large patent estate
- Lancium Compute – 100% renewable energy powered HPC Platform-as-a-Service (PaaS) and Custom HPC Cloud

Properties of end market applications

Common traits make certain applications ideal users of Lancium's technology

1 Large Consumers of Power	Energy intensive applications in large addressable markets, where power costs are a large component of the consumers total cost
2 Enabling Reduced Emissions Footprint	End markets currently utilizing significant amounts of carbon-based energy with room to reduce emissions footprint
3 Individually Controllable	Applications where consumption can be turned up or down individually, or in aggregate
4 Consistently Use Power	Steady power draw means it can respond to grid signals in real-time
5 Interruptible	Applications can be turned down without losing significant value from work-in-process

Lancium Footprint: 5,000 acres with ~7.5 GW Potential Capacity

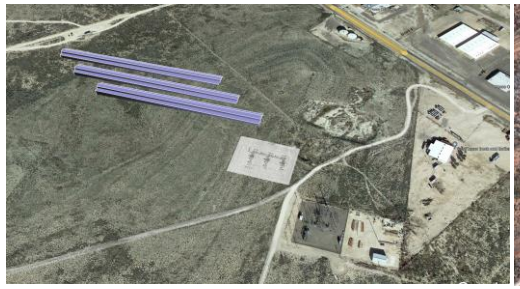


Name	1 Fort Stockton	2 Abilene	3 Childress	4 – 8 Various
Description / Status	Lancium Clean Campus 1 Under Construction 111 Acres	Lancium Clean Campus 2 Under Construction 875 Acres	Lancium Clean Campus 3 In Planning 1,400 Acres	Lancium Clean Campus 4-8 In Planning 1,861+ Acres
Target Date	Q2-2022 (Phase 1) Q1-2023 (Phase 2)	2Q-2023 (Phase 1) Q2-2024 (Phase 2)	Q1-2025	2024+
Installed Capacity	25 MW (Phase 1) 300 MW (Phase 2)	200 MW (Phase 1) +1,000 MW (Phase 2)	1,000 MW	5,000+ MW
Key Stats	Utility's and Lancium construction started	Utility's and Lancium construction started	Interconnection approval expected 6/1/22	Under Lancium control
Interconnection Status	TNMP 25MW Load Request Contracted & Paid TNMP 300MW Load Request Contracted & Paid	AEP 200MW Load Request Contracted & Paid AEP 1,000MW Approved	AEP 1,000MW Load Request & Under Study	NA

Lancium's three primary end markets

Large energy consumers whose work can be interrupted on demand

Bitcoin Mining

Overview

- A "miner" is a shoebox sized computer that hosts an Application Specific Integrated Circuit (ASIC)
- Bitcoin miners compete to solve complex math problems
- The first computer to solve the problem is rewarded
- Bitcoin miners use significant amounts of power to conduct their work

High-Performance Computing

Overview

- HPC jobs can include:
  - *Drug discovery*
  - *Scene rendering for entertainment*
  - *Aircraft / vehicle design*
  - *Machine learning*
  - *High energy particle physics*
- **Computing tasks can be moved to another site or restarted later based on grid conditions without affecting the results**

Hydrogen Production

Overview

- Hydrogen is a clean, versatile fuel
  - *Transportation*
  - *Industry*
  - *Power generation*
- Near-zero emissions when produced with renewables and when consumed in any application
- Texas is well positioned to become the central hydrogen hub in Texas

Lancium Compute

Unhook Computation from fossil fuels by dancing with the grid.

LANCIUM
BALANCING ENERGY

Lancium Compute Offerings

HPC PaaS

1. **HPC PaaS Offering**: Pre-configured HPC Clusters to focus your efforts on HTC and loosely-coupled parallel jobs
2. **Just Bring your Workloads and Data:** We make it easy to launch your workloads and retrieve the results for analysis

Custom HPC PaaS

1. **Customize HPC PaaS**: Modify the HW infrastructure of our HPC PaaS to suit your workload needs (and leverage our SW Stack)
2. **Just Bring your Workloads and Data:** Leverage our state-of-the-art HPC Software Stack to obtain the PaaS benefit on custom HW Config

Custom HPC System (BYODC)

1. **Customize HPC:** A data center to suit the specific requirements of your HPC environment and business model (HW, SW, cooling, space design etc..)
2. **Carbon-Free Environment**: Our clean campuses are **powered 100% by renewable energy** to host and maintain your HPC Infrastructure

200 MW Renewable Power at Lancium's Clean Campus in Abilene, TX in May 2023

LANCIUM
BALANCING ENERGY

How do we do it?

LANCIUM
BALANCING ENERGY

Computation as a flexible load

- Energy costs are very low when renewables are providing power.
- Three regimes of interest
  - Plenty of power … power almost free … go go go
  - On the edge of load approaching max generation ..
    - Need to rapidly ramp power consumption up and down for primary frequency response, and ramp power consumption downwards in response to price signals.
  - Power shortage … characterized by high power prices … reduce power consumption to minimum levels so that population centers don't have blackouts and we don't have to spin up brown-energy Peaker plants.
- Note that the ability to *accurately* predict the power regime for a site in the future can be quite useful. ☺

LANCIUM
BALANCING ENERGY

To drop load

- We migrate jobs in space and time to manage load.
- Migration “in time” means to pause the job (eliminating its power consumption) and restart it at a later time (when power is more plentiful.)
- Migration “in space” means to move the job from one site to another.
- To quickly (single digit seconds) drop load in half, “freeze” computations.
- To drop load further:
    1. Persist jobs to disk.
    2. As jobs are persisted, shutdown their nodes.
    3. If more power savings required, “unplug” nodes.

LANCIUM
BALANCING ENERGY

Lancium Compute offers three services

- 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.

- *VMS* - Lancium Virtual Machine Services provides customers the ability to define, instantiate, and manage Linux KVM virtual machines on LCI resources.

- *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.

LANCIUM
BALANCING ENERGY

Remote Command Line Services

- Containerization is a trend sweeping computing.
  - Containerization cleanly addresses the software configuration problem.
  - Containers “jail” the application, preventing it from “getting out” and attacking the host operating system or other container instances.
- We use Singularity with a tool to automatically generate a Singularity image from a Docker image.
- Users can upload their own images or use provided images.
  - Users specify the job resource requirements, e.g., 8GB of memory, 8 vCPUs, and the command line to run, files to stage in and out, e.g.,

```
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
```

LANCIUM
BALANCING ENERGY

Lancium Compute Today

- Providing ~ 1M vCPU hrs/week for Open Science Grid .. 42M+ vCPU hours so far.
- Have run CERN/CMS jobs for US CMS.
- Running low-degree parallel MPI CFD jobs for commercial customer– less than 100 nodes.
- Rendering for Grid Markets.
- Financial & business analytics – essentially ML/AI.

LANCIUM
BALANCING ENERGY

Vis a vis HEP and situation in Europe

- In the short term (less than one year) we could take up to ~2MW of compute load; ~100,000 cores.
  - A member of the US CMS team has told us though that we would need to place significant extra storage at our site and run a software stack to manage it. We are willing to do this at cost.
  - This is clearly just a drop in the bucket compared to the challenge. But lots of small measures add up. (1 MW == .8760 GWh over a year.)
- In the longer term we can take as much computing as desired either as PaaS or via some negotiated mechanism, BYODC, whatever.
- Physics experiments certainly welcome!
- We could even swing 1.21 GW for a time portal a la “Back to the Future.” ☺

LANCIUM
BALANCING ENERGY

Conclusion

What I hope you learned.

LANCIUM
BALANCING ENERGY

The Message

- Renewables are the least expensive way to produce power in the world today.
- The pace of renewable deployment is staggering
  - And gated only by the ability of electrical grids to safely absorb the power.
- Controllable loads, particularly loads that can be rapidly reduced, are critical to maintaining grid stability in the presence of non-constant solar and wind power.
- Lancium is building GW scale data centers that provide very inexpensive, mostly carbon-free, electricity for computation. Moving computation from brown energy regions to Lancium data centers reduces overall CO2 due to computational science.
- This in turn can increase the development of new renewable resources, leading to more available renewable power even on cloudy, low-wind days.

LANCIUM
BALANCING ENERGY

Our vision is to build geographically dispersed data centers.

- In the “wind belt”.
- When the wind does not blow in Texas, we’ll migrate north to the Dakota’s, or somewhere else where the wind is blowing, the sun is shining, and the electricity is cheap and carbon-neutral.
- Further, with decent forecasting we can begin migrations before the power levels shift, particularly with jobs that are being scheduled for the first time.
- Every MW of computation we shift from brown energy regions into the wind belt reduces overall CO2 by up to 5K mT/yr.

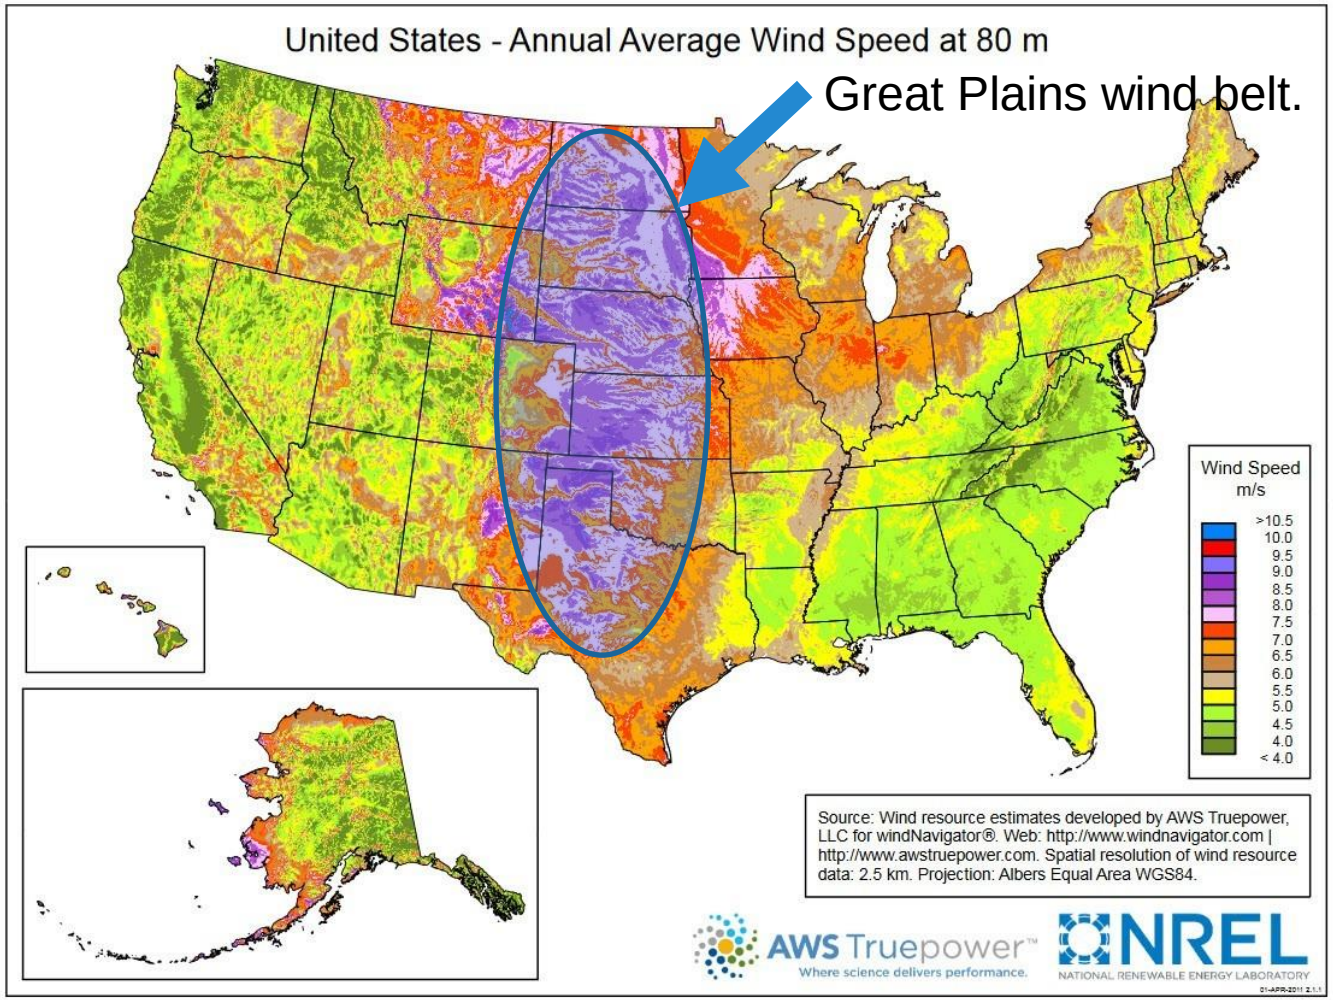

LANCIUM
BALANCING ENERGY

I believe we are at the beginning of a new industrial revolution; just as oil and gas replaced coal, water and animal power as the power sources of industry, wind and solar will gradually yet surely replace oil and gas as the primary power source.

And just as in the first industrial revolution, it is the efficiency of the new technology that is the main driver in the transition. While it is true that it is in our collective best interest to turn away from carbon-emitting energies, it is the fact that electricity is simply cheaper when generated from renewables that is securing a carbon-free, or carbon-reduced, future.

Lancium is at the center of this transformation.

END

LANCIUM
BALANCING ENERGY

Frequency Stability

LANCIUM
BALANCING ENERGY

Frequency & Inertia

- As demand approaches supply
  - Frequency starts to drop
  - As frequency starts to drop current starts to rise on spinning generators.
  - As current starts to rise, so too does temperature and R
  - If it goes to far, generators “trip-off”.
  - This causes supply to drop – potentially causing a chain reaction.
  - To prevent collapse, load is shed by turning off load zones, i.e. blackouts.
- Immediate response to stabilize is “Primary Frequency Response”, must be done very rapidly. Typically within 15 seconds.
- Longer term response, e.g., within 15 minutes is called “Load Response” or LR. This is addressed by adding more generation or reducing load. Increases and decreases are controlled by the grid operator using computer generated “set points”.
- If you can vary your load rapidly, and in response to grid operator signals or price signals, you can be paid for it in the “ancillary markets”.
  - This can significantly reduce the price you pay for electricity.

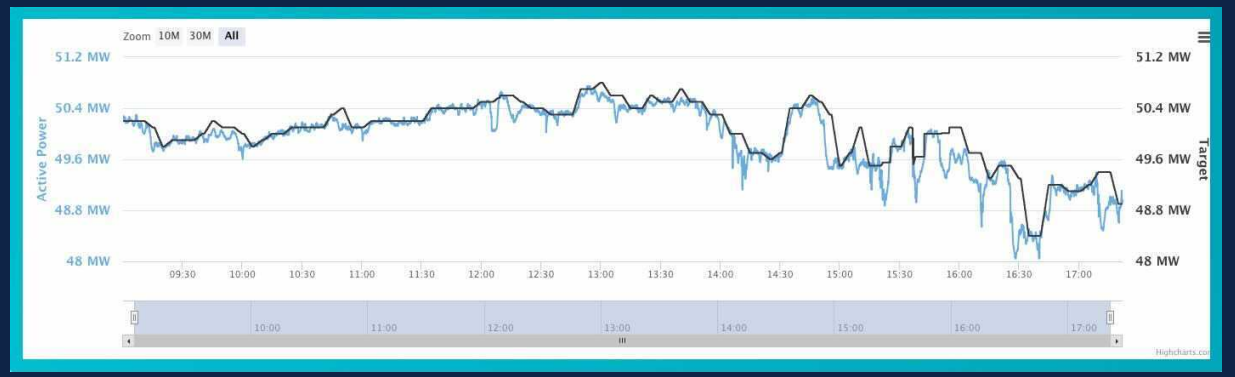

LANCIUM
BALANCING ENERGY

….Causing Increased Prevalence of Negative Power Prices…

- Depressed power prices can inhibit the buildout of new renewable generation

Frequency of negative real-time pricing

2013

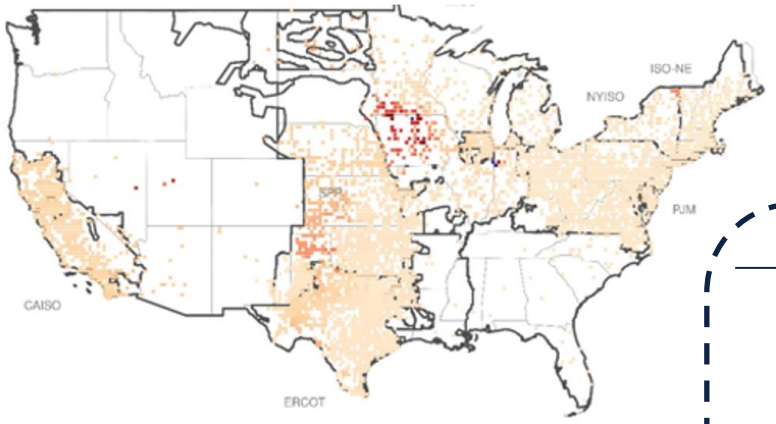

Lancium founded

2017

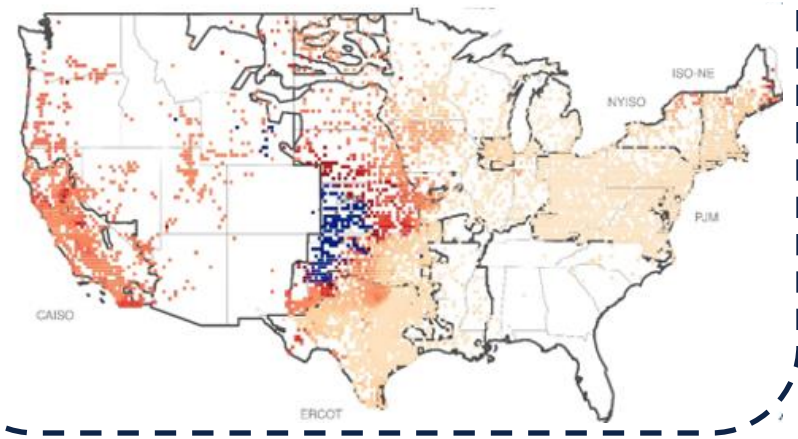

2021

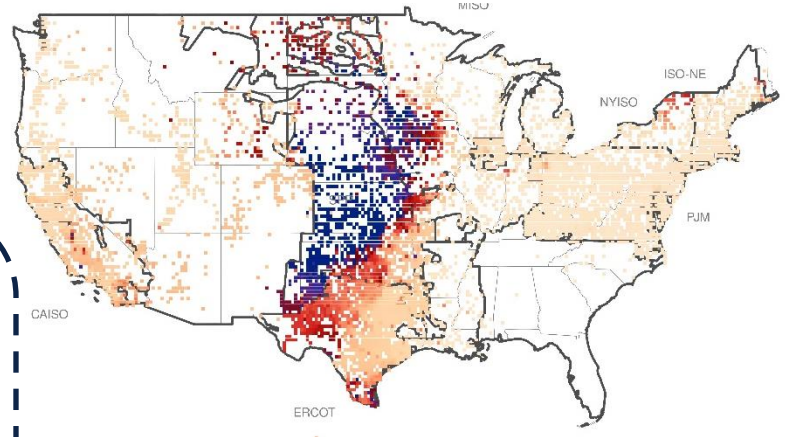

0% **10%** **20%**

Source: Lawrence Berkeley National Laboratory – Energy Markets & Policy – Frequency of negative locational marginal prices at nodes in the seven organized wholesale markets in the United States

Controllable Loads Are the Best Application of Grid Flexibility

- Technological options and solutions for flexibility

		Flexibility			Economy	Quantity	Emissions
Resource designation	**Resource type**	**Start-up time**	**Run-time**	**Rampability**	**Start-up cost (USD/MWh)**	**Minimum load** (% of nominal power)	**Emissions intensity**
Baseload fossil generation	Generation	1–10 h	Continuous *(Subject to frequent outages)*	✗	55 - > 100	25–50%	High
Natural Gas Peaker	Generation	5 min	~ Days	✓	< 1 *(excludes degradation)*	20% per unit	High
Batteries	Storage	Seconds	< 4 h	✓	0 *(excludes degradation)*	0%	Varies
Load - Traditional demand response	Load	Minutes	Continuous	~	N/A	5–40%	N/A
LANCIUM	**Load**	**Seconds**	**Continuous**	✓	**0**	**0%**	**Low-Negative**

Source: International Renewable Energy Agency

2021 YTD REAL-TIME PRICE DISTRIBUTION

SmartResponse Value Proposition - Economic Dispatch in West Texas (LZ_West – 01/01 to 06/28)

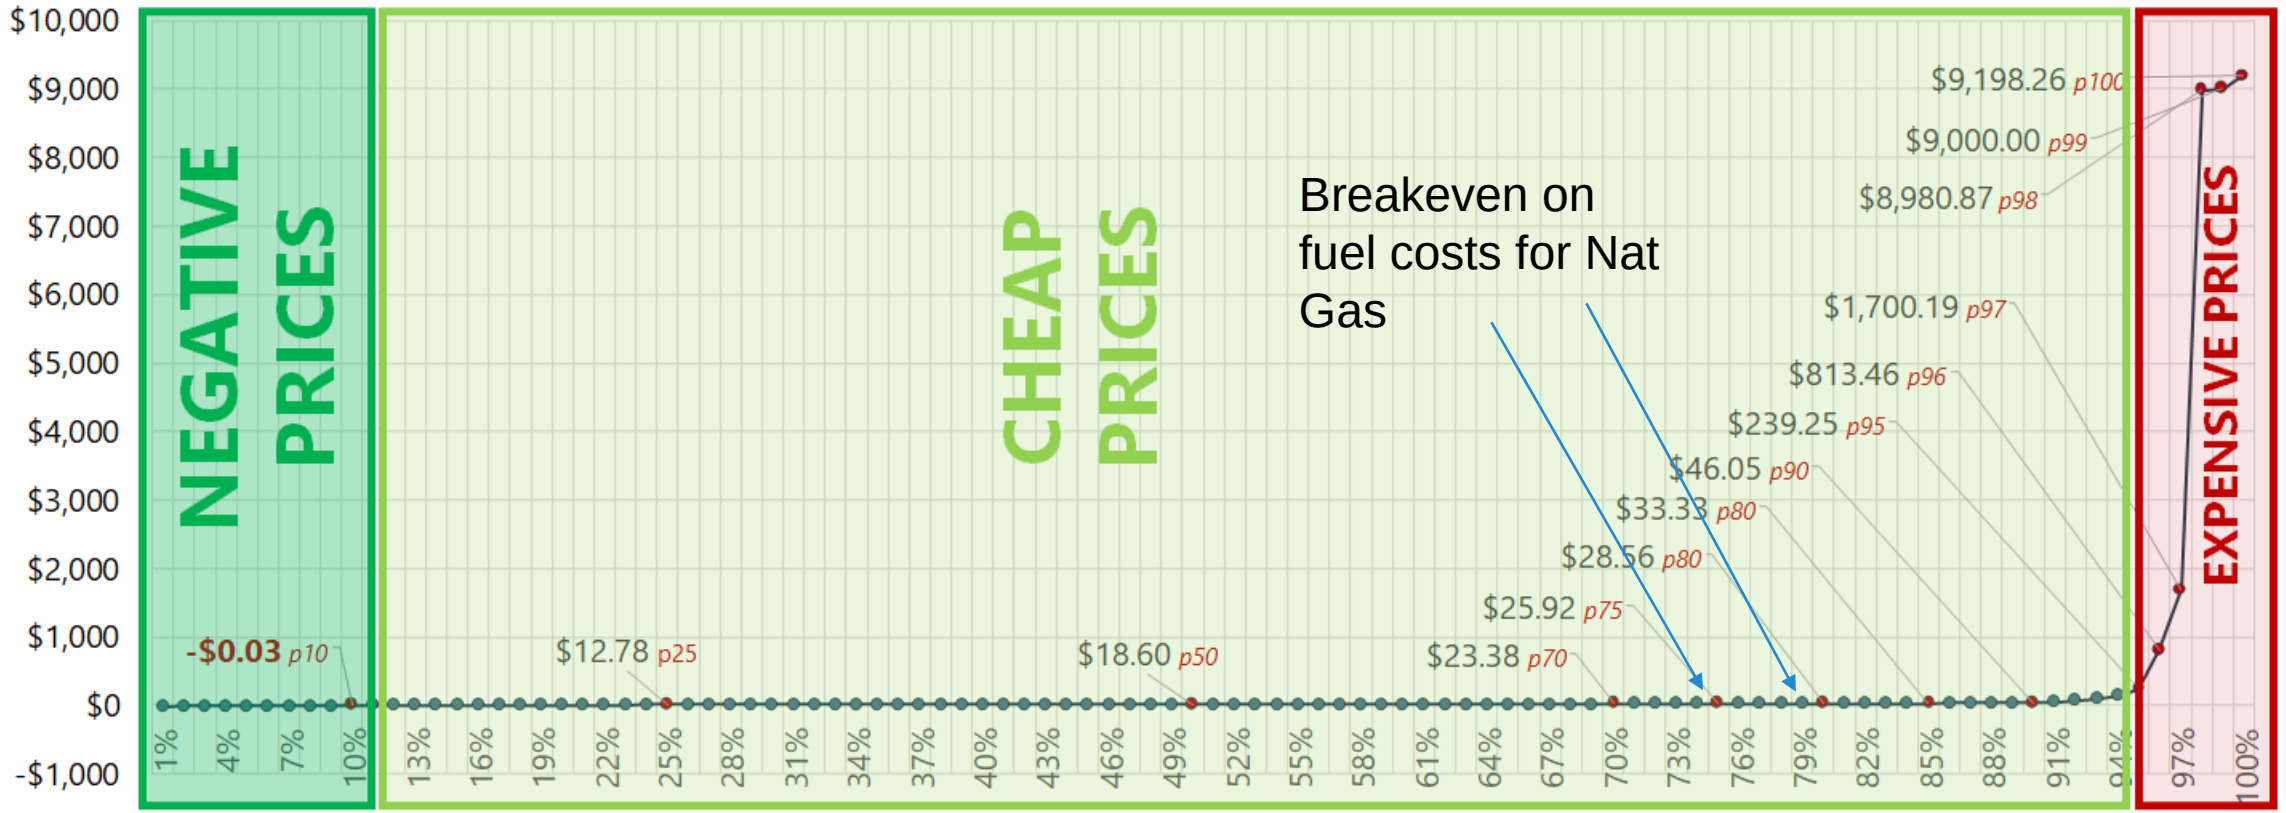

YTD - 10% of all hours settled negative

Lithium is not the answer

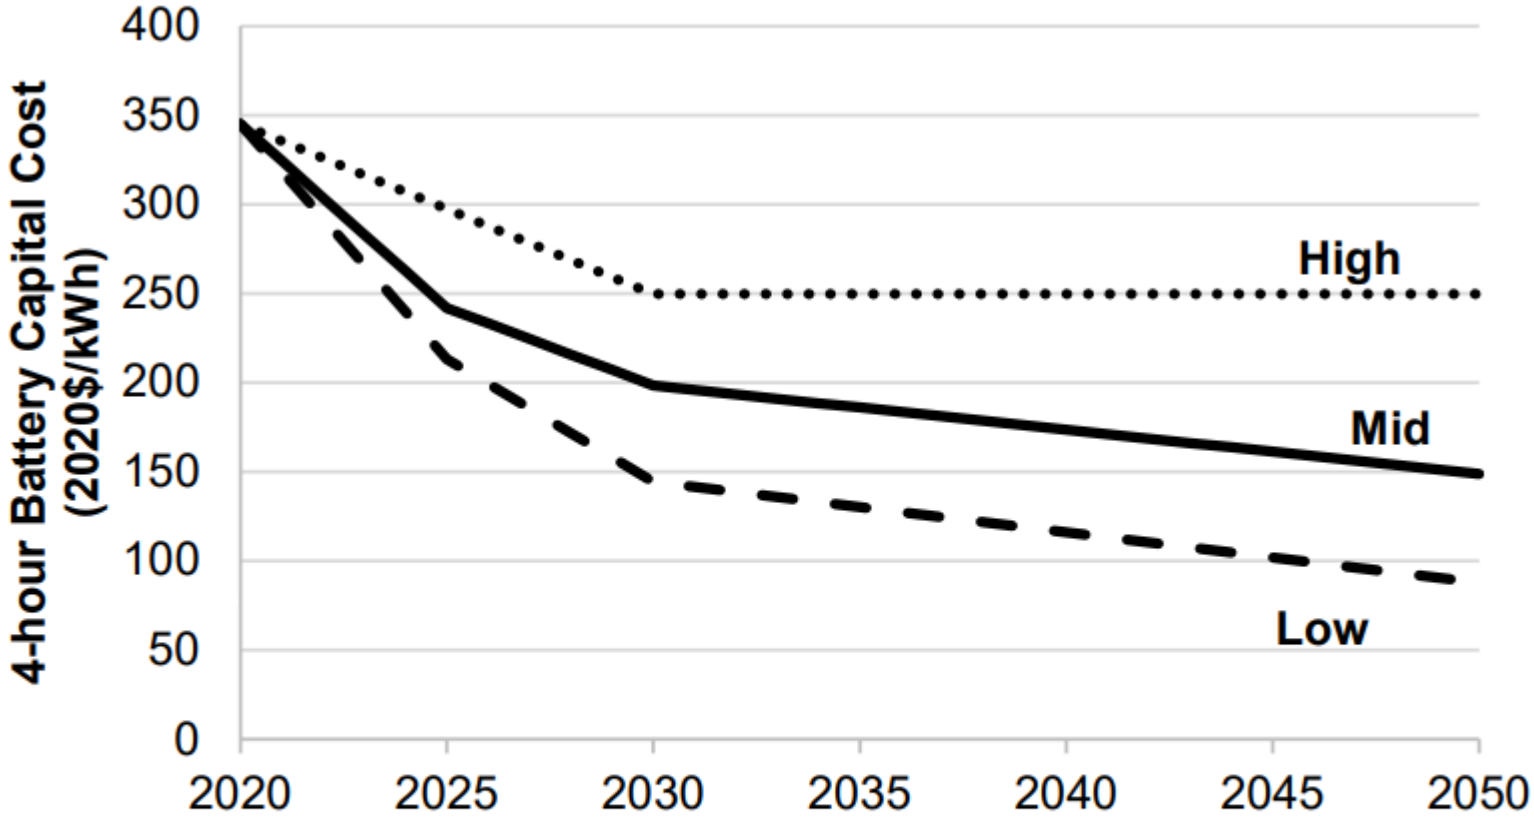

Figure ES-2. Battery cost projections for 4-hour lithium ion systems.

Cost Projections for Utility-Scale Battery Storage: 2021 Update Wesley Cole, A. Will Frazier, and Chad Augustine, National Renewable Energy Laboratory, NREL/TP-6A20-79236 June 2021. For non Li, see https://ambri.com/technology/