Lancium leverages four mechanisms and a computational philosophy to simultaneously reduce costs and dramatically reduce green-house gas emissions as compared to traditional data centers.

Balance the grid. Lancium acts as a Controllable Load Resource (CLR) by being able to rapidly adjust the load on the grid. This enables the grid to take more renewable energy, with dynamic loads that can take more load during periods of over generation resulting in less curtailment and shedding load quickly during periods of reduced renewable power generation. This makes it possible use more renewable energy (less CO2) and balance the grid without the need for dirty, carbon-based "peaker" power plants. As a bonus we are able to reduce the total cost of our power!

Locate near wind farms. Lancium is locating many of our data centers near wind farms where electricity is both cheap and emission free. This allows us to buy electricity, often cited as 40% of the cost of operating a data center, at a significantly lower price. As a bonus, for every MW of wind energy used we reduce CO2 emissions by 0.7 metric tons.

Use older gear. Lancium uses older generation, more cost-effective nodes and processors. Per core performance improvements have slowed to 2.5-3.5% a year. This means that a 3-5 year old processor is perhaps 85%-90% as fast as the latest and greatest processor. Yet it costs less than half of a current generation machine. Similarly a 5-8 year old processor is 50%-80% as fast as the latest and greatest, yet costs 15%-25% as much. The downside is the older machines are not as dense or power-efficient. That is not a problem for Lancium as our electricity is very cheap and green, and our data centers are very inexpensive. As a bonus, we give new life to old gear and reduce the impact on the environment caused by disposal of old equipment and manufacturing new gear.

Inexpensive data centers. Lancium data centers are far less expensive than traditional data centers. Lancium data centers eliminate traditional HVAC, using forced air to cool equipment. Our data centers are more akin to warehouses, containers, or pole buildings than traditional data centers. This is possible because modern processors can operate at significantly higher temperatures than older processors. We have found via experimentation that machines can operate without significantly higher failure rates at input air temperatures up to 115F. Ambient temperatures rarely get over 100F at our data center locations. If it gets so warm that we need to reduce operations at one site, we simply migrate jobs elsewhere.

Going wide. Our computational philosophy is to deal with slower machines and intermittent power availability by "going wide". Going-wide recognizes that in capacity-oriented, high throughput computing, what the customer cares about is the rate at which we complete jobs, not the time for any particular job to complete. For example, if we are without power at any given site 10 percent of the time we can achieve the same throughput by increasing the resource count by 10/9=1.11, 11%. In a similar manner, we deal with slower CPUs and GPUs by adding more equipment. We can add more and more equipment because the older gear is far cheaper than new gear, because our data centers are inexpensive per square foot, and our electricity costs are so low.