



[00:00:00] Andrew Johnson: Hi, everyone. In this episode, I sit down with Chris Silvestre, an analyst on our U.S. equity team, and hear about some of his learnings from a recent research trip to northern Virginia, the data center capital of the world. We cover electricity demand in the U.S., how the critical themes of electrification and AI are influencing that demand, and the significant role that data centers play. It's packed with insights into the challenges and constraints in meeting these growing power needs and how businesses, regulators, and investors are responding. I hope you enjoy it as much as I did.

[00:00:37] Disclaimer: This podcast is for informational purposes only. Information relating to investment approaches or individual investments should not be construed as advice or endorsement. Any views expressed in this podcast are based upon the information available at the time and are subject to change.

[00:00:53] Andrew Johnson: Hey, Chris! Welcome back to the podcast.

[00:00:55] Chris Silvestre: Andrew, thanks so much for having me back.

[00:00:57] Andrew Johnson: Before we dive in, Chris, I think some background may be helpful for listeners today. This episode was born out of an internal memo that you sent around following a research trip you did a couple of months ago where you visited several data centers in Northern Virginia. This is one of the many examples of you and the rest of the research team getting out there, turning over stones, and really searching for businesses that may fit into the portfolio.

As is often the case, you come away with many more learnings that can help piece together the mosaic of risks and opportunities that exist within the investment universe. In this case, you shared learnings that touched on a few key big themes these days, one of them being AI and, more specifically, the power or the electricity that is needed to keep the AI models moving.

So with that, let's jump right into it. I'm thinking the best place to start is to provide some context on just electricity demand more generally speaking. What does that look like over the last decade or so? What are the key factors that have influence over that demand?

[00:02:02] Chris Silvestre: Great way to kick it off.

So over the past decade—just to be clear, we're talking about the U.S. specifically here; the rest of the world is a little bit different—but in the U.S., electricity demand has been pretty well flat for about a decade, and there are a couple of reasons for that, but basically think about electricity demand as having three key end-use sectors roughly split evenly.



So industrial, residential, and commercial. On the industrial side, industrial production has been pretty well flat in the U.S., so that mostly explains that. Then on the commercial and residential side, electricity use has also been pretty flat. Although we've had more structures, it's been offset by energy efficiency. So you put that all together, and pretty well, electricity use has been flat for a decade in the U.S.

[00:02:46] Andrew Johnson: Demand has been relatively flat for some time. What are the expectations for demand going forward? What would be the key drivers behind those expectations?

[00:02:55] Chris Silvestre: The source we typically look at is the IEA (International Energy Agency). They put out projections annually for what electricity consumption and production projections are going to be over the next couple of decades. The latest release they put out was in March 2023, and that release showed that they expected electricity demand to increase by about 25% in the U.S. through 2050. That's about a percent a year or less than a percent a year, and that's after stagnating for the better part of a decade.

The drivers of that are multifold. So, we talked about the end-use sectors, residential, and commercial. There you're getting a lot more electrification. You're getting things like natural gas; there's a push to shift that over to electricity. Think about natural gas burning stoves or natural gas for heating purposes. And the reason for that is to meet ESG goals. It's more environmentally friendly to heat a home with electricity than it is with natural gas. Same thing for cooking purposes. So that's one kind of driver of demand. And you're actually getting certain states introducing bans on new natural gas connections and forcing new homes to do electricity only.

So that's on the structure side. Then, of course, electric vehicles. Everyone knows about that. That's also expected to be a big driver of demand going forward. And that will, of course, take away from petroleum and add incremental demand for electricity. Then the last driver, which has really only transpired in the past year and a bit. And this is the most significant or projected to be the most significant is data centers. The driver of that is artificial intelligence. It's just far more energy-intensive than your typical data center and that's why this is a relatively new development.

Not that data centers are a new development. Of course, they've been around for decades, but it's really the AI-specific use case and demand that's projected to add a lot of incremental demand that wasn't forecast just 18 months ago.

[00:04:56] Andrew Johnson: I gather data centers are at the heart of AI and these large language models and, presumably, without enough power, we can't run these things. What role do data centers play in the broader electricity demand landscape?

[00:05:10] Chris Silvestre: We'll start with some context. Roughly, an AI-related data center could consume about 10 times the power of a typical data center. If you look at the power rating of an AI chip, NVIDIA's latest data chip is about 700 watts. A typical power capacity or power rating of a CPU chip pre-AI workload would be about 70. So that's a 10x increase in the power consumption from this new generation of GPUs that are for AI-specific or high-performance-computing specific workloads. That's really the key driver of the incremental demand. For the same footprint and the same data center, you need 10 times more power for these new workloads than you do for the old-style data center.



I have some specific examples. In the U.S. Equity strategy, we own Southern Company, and we own AEP (American Electric Power Company). I always find it helps to provide specific examples to really drive it home. Southern Company is a vertically integrated utility, so they manage the generation of the power, the transmission to the load center, and then the actual distribution of the customer. They create something called an Integrated Resource Plan every once in a while. They file that resource plan with the regulator, essentially requesting budget or permission to invest in generation, transmission, and distribution to meet the needs.

In their previous integrated resource plan, they projected something like 400 megawatts of incremental demand from data centers to 2030. In their most recent integrated resource plan, they projected 6600 megawatts of incremental demand. So that's a 17x increase in projected electricity demand just from data centers through 2030. So that's really what's driving the vast majority of the incremental, unexpected demand projections from AI data centers.

[00:07:09] Andrew Johnson: We've had this increasing electrification of everything in our lives, which has been a theme for some time. And now this relatively new technology in AI, these large language models that appear to have even greater need for power. What specific challenges does AI pose to the current infrastructure for electricity supplies that you were just alluding to?

[00:07:31] Chris Silvestre: We can start by going to the beginning of the genesis of data centers and, essentially, data center hubs. So in the early 2000s, they really started as buildings where carriers exchange information, literally your cell phone carrier or your broadband carrier. They would come together, and they have to exchange information. And then that naturally created a hub when the hyperscalers and the cloud companies wanted to connect to the Internet and wanted to provide their customers access to the Internet. So you have essentially a clustering effect.

It's an organic process. If you think about it like how a nation develops, settlers land in one city, and then the population starts growing there, and then maybe they branch out into another city, and that becomes the second biggest city and population center. That's kind of how the data center market works because it makes sense to be close to other data centers and other, essentially, providers of information or information storage centers so that you can talk to each other and minimize the costs related to transferring data and all be in the same place.

So that's how the data center market really developed in that organic process. And that led to a clustering of data centers. Interestingly, Virginia was the first data center market in the world, is still the largest data center market in the world, and is projected to be the largest for decades to come because of that process. The reason that happens is because it's very expensive to move data around. You actually have to get it from place to place.

If you can imagine having one data storage center in Virginia, let's say the first, it doesn't really make sense to build something in the middle of the U.S. You have to lay a fiber connection to actually transfer the data quickly, and then it's cost prohibitive. It just naturally follows a very organic process.

That leads me to the second part of your question, which is the challenges. The challenge is that you need power in certain spots and the right amount of power in those spots in order to actually support this process. What that results in is basically a clustering of power demand. So that's causing some challenges. It's not that the U.S. doesn't have power or the ability to generate power. It just may not have it in the right spots at the right time to



support this new incremental demand from data centers.

There are some questions about whether the traditional clustering paradigm is still applicable with AI and, without getting too technical, there are two key processes in high-performance computing. One is training models where you have a whole bunch of data, and you actually have to train a model. The second is making predictions on those models. There's a thought that maybe those can be divorced and the training can happen in more remote locations, and the inference making predictions on those models can happen in the traditional data center hubs. That's still to be determined, but the point is, most of the demand is still going to be in the traditional data center hubs, and that's where the power is actually needed. That is creating some challenges at the moment.

[00:10:37] Andrew Johnson: As you were explaining that and giving that background, I was just thinking how easy it is for us to take for granted the movement of data when we're just using our phones or our computers on a daily basis. But there's this entire infrastructure behind it, and with that comes constraints. And that's what you were just segueing into there. So what are the main constraints in meeting the current power needs of data centers?

[00:10:59] Chris Silvestre: One is land, the data centers themselves. I think what's important to understand is that data centers, the power capacity is pretty well immutable. Once it's built, you can't really change it, and that's because customers are really buying power capacity. They're not really buying floor space. If you look at all the pricing and how this product is quoted, it's by power capacity.

The reason is because customers can put whatever they want in their cage. The way a data center works is they build kind of a shell of a building and inside of that building, they build a bunch of different cages. I'm talking about a third-party data center where you could have multiple customers in the same data center, not a purpose built for one customer. The data center is responsible for providing everything to the cage, and the customer provides whatever's in the cage. What's interesting is a lot of the third-party data center builders and providers don't actually know what the customer has in the cage. That can be proprietary information, and it's not really their responsibility.

So the point of all of that is the data centers are really selling power capacity. In order to change that, it's very difficult. Because if you think about power coming in from a transmission source, maybe you have a substation that you have to build on the data center site, and then you have to run cabling to the actual building, and then you have to run cabling from the building to every single unit in that building. Once you've done that, it's pretty costly to rip that all out and redo all of that. A lot of the time, it's actually not possible. You actually don't have the space in order to do that. You'd have to basically rip out all the walls, maybe the ceilings, to run 10x the amount of cabling. So that's one important thing to understand.

Then the land that that data center sits on may also not be supportive. One of the really interesting learnings was that every megawatt of power capacity has to have backup generation because these are mission-critical facilities. You expect to be able to access the data or use the data or whatever you have there at any point of time. What you'll see is every single data center has backup generators, and these things are quite large. If you have a plot of land that maybe supports 50 megawatts, you could literally have 50 container-size backup generators on site. You can't really 10x that without changing the land footprint. Essentially, all of these are land constraints. So that's the first key constraint.



The second constraint is energy. In a lot of these locations, there just isn't the energy available. The interesting thing to realize about the energy market is typically we think about the grid and how large it is, but energy is very local. If you think about the topology of a network, you literally have to have supply and demand balanced at every point. There's nodes that all of the wires connect to, so it's really many local markets with local supply and demand characteristics as opposed to one homogeneous amorphous market across the U.S. So that's the second constraint—the actual energy—and there's two pieces to the energy constraint.

One is actually generating the energy, and the second is getting it to the right spot. On the energy generation, I don't think this is going to be a surprise to anyone, but there's been a push to shift away from fossil-fuel-generating sources to renewables. What's interesting about that shift is that the average size of a renewable generating plant is much smaller. To put this in perspective, if you have a gas-fired power plant, it may produce something like 300 megawatts on average. But the average solar farm or the average wind farm may only produce 25 megawatts or 30 megawatts. What that means is you need 10 times the number of plants to produce the same energy, and you actually may need much more than that because renewables are intermittent. They're not dispatchable. The nameplate capacity is a little bit different than the actual output.

So generation is a big problem because previously you had generation assets close to these demand centers. Well, now that's no longer the case. If you're going to build 10 solar farms or 20 solar farms to satisfy the same amount of demand, you can't really do that in a major population center. You need to have that further out. So that's the one major challenge is actually just generating the power.

The second challenge is getting the power to the right spot. If you can imagine those 20 solar farms or wind farms located in some more remote area, you need to build transmission lines to get it to the demand center. So those are the two key challenges. One is just you need a lot more land to build new data centers. You can't retrofit existing ones. And then the second one is you need to power those. In order to power those, you need to build generation – and a lot of it – and then you need to build transmission to get it to the right spot.

[00:16:07] Andrew Johnson: Let's go back to what kicked off this topic. You recently visited Northern Virginia, which is also known as the data center capital of the world. I don't know if you called it that, but certainly, the stats that you laid out allow it to lay claim to that. And you came away with some interesting observations specific to Northern Virginia as a case study. What were your learnings from that trip specifically?

[00:16:27] Chris Silvestre: Some context on Northern Virginia: it is called Data Center Alley. The reason it's called that is because it's still the largest data center market in the world, as I mentioned. For some numbers, it is about 2500 megawatts of capacity. That's two times greater than the second major data center, which is Beijing, and it's four times greater than the number two U.S. city, which is Dallas. So this is a really, really big data center market.

Again, it's happened that way because it makes sense to be close to everything else Internet-related, close to the storage, and close to the hyperscalers. By doing that, you minimize the data transfer costs. So that's why it's still the major data center capital of the world and projected to be. The learning that I came away with is that the demand is very real. I visited five different data center companies when I was out there, and they are completely sold out. They have no capacity left. Any data centers that they had on the books, in progress of being built, or slated to come online soon are completely sold out.



The reason for that is, as I mentioned, the power capacity of these AI workloads is so much higher than previous workloads. If a data center company – let's say Equinix or a Digital Realty – if they had something on the books they were planning two years ago, maybe it was a 100-megawatt facility somewhere, maybe in Virginia. And all of a sudden, a customer recognizes that the compute required is 10x what they previously thought. Well, you may now have two customers purchasing all of that capacity, whereas previously you may have had 10. So that's essentially what's happened is anything that was available was very quickly spoken for and very quickly purchased.

The second interesting thing is that they're running up against very real constraints. One of the companies that I visited was a private company called Aligned, and they had planned a new data center in Virginia and Loudoun County. This was in 2022. At the end of 2022, the local utility, Dominion Energy, actually called them and said, "Hey, you need to put this project on hold because we cannot supply you with power." I think they signed an LOI for the land, they started pre-selling customers, and they actually had to mothball the project because they were informed by their utility that they could literally physically not get the power. That's what I saw.

The demand is very real, everyone is sold out, and it's because the constraints are also very real. There was a lot of optimism about being able to solve those challenges. One of the ways that they can do that is transmission. I mentioned that power is actually very local. It's many thousands of nodes where the supply and demand has to balance. So what's interesting is the problem can be solved with either generation or transmission to redistribute the load throughout the grid. It's not necessarily that you have to build 20 solar farms far out. I mean, we need that as well. But in the short term, you can do things like connect one center that's oversupplied with the center that's undersupplied. Those are part of the plans to alleviate some of the near-term challenges and some of the more constrained areas.

[00:19:34] Andrew Johnson: You mentioned optimism, and certainly, humans and businesses, if anything, adapt to these types of constraints over time. My last question is how are businesses and governments responding to all of this? And to bring it back to our clients' portfolios, what, if any, are the investment implications around this? In other words, what are the risks and opportunities that may come from this transition?

[00:19:58] Chris Silvestre: I'll just start with the demand side. I mentioned the demand is very real. If we think about it from the perspective of these corporations, mainly the hyperscalers that are buying all of this demand, it's a risk-reward calculation. I think they're probably sitting there thinking, "Well, if AI really is going to deliver on the promise, then we need to be at the table. We need to be a part of this." Maybe the worst-case scenario is that we bought too much data center capacity, and that's okay. It will be absorbed over time. Instead of three years, it will be absorbed over five years or six years or whatnot. So that's really what's creating the impetus for signing deals and mopping up all of the capacity.

In terms of what businesses and governments are doing, I mean, that's what businesses are doing. They're being proactive about it. They're buying everything in sight, and they're actually also structuring some very interesting deals. One that most individuals have probably heard of is the deal Amazon did with Talen. So Talen runs a number of nuclear facilities, and Amazon actually purchased a data center campus from Talen Energy onsite at one of these nuclear facilities, which is fascinating. The nuclear facility has two generation units. And the idea is that the data center campus would be supplied by one of those units. So businesses are responding in creative ways, such as Amazon did with Talen. That's a very unique deal. The likelihood of there being many of those is



limited, but there was one, and there may be another.

So I think businesses are doing what they can and making sure that they have a seat at the table to balance that risk-reward. Governments are also acting, and it's interesting because most people have a perception that governments can't get their act together. I think the U.S. especially has a great ability to do that when it's necessary. With respect to AI or high-performance computing, I think there is a sense that this is a priority, and it may be a matter of national security, certainly a matter of national interest. As a result of that, we are seeing governments at various levels act. I think the most prominent example of that is the Nuclear Production Tax Credit that was passed and the IRA that was passed. I think that was in recognition of the need to ensure that energy and power are not a roadblock to the U.S. being a leader in this space. They are mobilizing – through legislation, through incentives – to encourage the private sector to actually make power available to support this trend.

In addition to that, there are other levels of government or organizations that are acting. I mentioned transmission. One of the biggest backlogs or barriers has been actually getting projects connected to the grid. And this predates AI actually. I'll connect it to a comment I made before, which is as you move from fossil-based generation sources to renewables, you just need a lot more of those types of projects. I mentioned kind of a 10 to 1 ratio between solar farms to replace one fossil-based generating asset.

So what that's done is that's created a big strain on the organizations that are responsible for approving those projects and getting them connected to the grid. These are regional transmission organizations that actually determine or provide approval to a generation asset to connect. What you had happen is as renewables were gaining prominence, the number of projects essentially requesting permission to connect to the grid ballooned, and these organizations didn't have the ability to actually process all of those applications, which came at the absolute worst time when we have a new incremental demand driver in AI.

What we're also seeing is those organizations implementing reforms to speed the process and essentially review more applications and grant approval to more applications to actually connect to the grid. At different levels of government and organizations, I think there's a recognition that this is something that has to be solved, something that has to be solved quickly, and we are seeing action as a result of that.

To give a very specific example, PJM is one of the largest regional transmission organizations in the U.S. Their queue to connect to the grid was essentially at a standstill for the past few years. The reasons for that are varied and probably too specific to get into on this particular call. Suffice to say, they made some reforms, and just in these past couple of months, they announced that they had approved 300 projects to connect to the grid. Those take time to build, to build the asset, to build the transmission, but the point is, mobilization is underway.

In terms of the opportunities, they are varied. I think in the whole utilities ecosystem, there are potential opportunities. The most obvious is utilities themselves, and that can happen in two ways. There are fully regulated utilities, and there are unregulated power markets. There are potential opportunities with independent power producers. If anyone's familiar with those companies, the stock prices are all up 5x; the market has caught on to that. But then there are all sorts of opportunities in the ecosystem that are less direct.

For example, I mentioned that data centers need backup generators, so that's a big source of potential incremental demand for companies like Eaton or Cummins or ABB, and that's literally because for every megawatt



of power capacity for a data center, you need a backup generation. So if you think roughly 10x, well, that's incremental demand. The backlogs are multiple years for transformers, so those companies are benefiting. Services, companies like Quanta Services or MYR Group or Mastek. These companies actually build transmission lines, and they help the data center to connect to the grid.

So those companies are experiencing incremental demand. Battery companies. Tesla has been a big beneficiary of this. Or Enersource. So there's a whole bunch of companies in the ecosystem that could potentially benefit, and I think they already are benefiting. The question is, will they continue to benefit? We've already seen the first phase of it play out, and the question is now shifting to well, how sustainable is this demand? Of course, everyone knows that NVIDIA stock is at all-time highs. It's up 10x. The key question is, well, can they keep it up?

And then there are maybe some less obvious ones, like natural gas. The whole natural gas ecosystem seems to be benefiting because there's a question about how that incremental generation will actually be produced. So far, in the last couple of years, the thought has been that we are heading full steam ahead with renewables. But, of course, there are problems with renewables. It's an intermittent resource. When the sun is not shining, the solar farms aren't producing. When the wind isn't blowing, the wind farms aren't producing electricity.

So we need some solution to that. We need either battery to store that power and deliver it to the grid at the right time. Or if we can't build enough of those renewables, maybe there is more fossil fuel generation in the near future. I think the whole natural gas ecosystem is optimistic about that and that could lead to the need for more pipeline capacity in midstream. It could lead to the need for more E&P (exploration and production). It could also potentially result in more gas-fired power plants.

Many, many opportunities, different ecosystems are impacted, and it's just a question of, as always, what's not priced in yet.

[00:27:15] Andrew Johnson: Alright, Chris, we've covered a lot of ground here today. I always learn a lot in my conversations with you, both on the podcast as well as off. I appreciate you taking the time to share some of those learnings with our listeners today.

[00:27:27] Chris Silvestre: Thanks a lot. Appreciate it and look forward to being back.

[00:27:31] Andrew Johnson: Hey everyone. Andrew here again. To subscribe to the Art of Boring Podcast, go to mawer.com. That's M A W E R dot com forward slash podcast or wherever you download your podcasts. If you enjoyed this episode, leave a review on iTunes, which will help more people discover the "Be Boring, Make Money" philosophy. Thanks for listening.

