Episode 4

Fission Founders, Part 1: Manufacturing Nuclear

Matt Slotkin

It's worth noting what we build a lot of today in the United States. We build a lot of wind, solar, and combustion. Those are all things that are centrally manufactured. We basically print them in a factory, and then we staple them to the ground.

That's a lot easier than building a giant concrete civil works project with thousands of on-site people that need to be found, trained, hired, relocated, and housed. It's really complicated.

Julia

That's Matt Slotkin, the founder of nuclear startup Blue Energy. On today's episode and next week's, we're talking to founders like Matt who have seen the big problems with nuclear and are building companies designed to address them. We'll also hear from the investors who are backing them.

We wanted to kick off this episode with Matt because he's building almost exactly what Nick Touran said was one of his favorite ideas in episode three.

Nick Touran

My other favorite idea is to build a factory of large reactors - a shipyard factory. That's a really interesting concept. If we wanted to get serious about decarbonizing quickly, that's the kind of thing we should be looking into.

It's like bringing Henry Ford's factory production at shipyard scale to floating reactors.

Packy

Many people do want to get really serious about decarbonizing quickly.

So what's Matt building to do it?

Matt Slotkin

Our approach is to centrally manufacture them in shipyards. The entire power plant is manufactured in the shipyard, where you already have a pre-existing labor force and equipment.

This ends up being a near-shore, offshore nuclear power plant. There are many benefits to having a nuclear power plant in the water instead of on land.

Packy

Matt is bringing Nick's dream idea to life: build reactors in shipyards and float them offshore. We'll dive into more detail with Matt later, but how cool is that?

Julia, I have to admit, this season has been a rollercoaster for me. After our conversations for episode three, I was nearly convinced that we just needed big companies to start building more large reactors with government support, maybe even some central planning, and that nuclear might not make sense for startups, except in some edge cases.

There's no doubt that we need to get better at building large-scale nuclear plants, putting gigawatts on the grid, but I was beginning to think that was the only viable path forward. And then we talked to the founders and they brought me right back.

Julia

The founders building nuclear startups are betting that by starting with a fresh sheet of paper, they can radically improve nuclear energy production by improving the reactors themselves or the size and form factor. Some founders are rethinking key components, such as the cooling methodology or fuel type, to make them more passively safe or efficient in fuel use. Others aim to rework the reactor design slightly to make them more manufacturable.

Some are sizing down reactors significantly to serve off-grid locations where diesel generators are currently the sole power source. While there's excitement about these new models, not everyone believes this is the answer to building more nuclear. They're probably wise not to deem these our silver bullet, as much is still unproven.

Some founders, like Isaiah Taylor of Valor Atomics, whose great-grandfather worked on the Manhattan Project, are less concerned about the exact design or size we're building. We just need to bring the cost of building way down. We shouldn't simply aim to match China on the cost of building large-scale nuclear plants. We should tap into the entrepreneurial spirit that moves America forward and build better, cheaper, and faster.

Isaiah Taylor

There are a couple ways we can sanity check ourselves. One is that China is doing this about seven or eight times cheaper than we are. That's a really obvious one. How are other countries with completely separate regulatory frameworks doing this?

China's doing it about seven or eight times cheaper. So there's a sanity check that we're doing it very wrong. But even China is doing this based on old technology that, frankly, they've stolen and is 15-20, even 30 years out of date from how we know to build nuclear reactors very cheaply.

So even old technology and old build patterns are seven or eight times cheaper when you just do them in a different regulatory framework.

Julia

The bet nuclear fission startups are making is that even though nuclear fission is a 70-year-old technology, we can build it better today. The term "first principles thinking" has become a cliché in tech, but it's a useful way to approach building things. It means breaking down complex problems into their most basic fundamental elements and then reassembling from the ground up.

Packy

Talking to the founders got me fired up, and that's why. We heard variations of "the laws of physics say it's possible to do this much better, so let's use them as our starting point." The founders we spoke with are working on different approaches, as Julia mentioned, from the ground up: different designs, fuel types, sizes, and more. Broadly, we can put them into two categories. One, manufacturing: take existing approved designs, but manufacture them instead of constructing them on site to make them cheaper over time. Two, advanced reactors: create new reactor designs that come with built-in advantages and open up new use cases for nuclear.

Originally we were just going to do one episode for founders, but after talking to them we realized there was too much to fit into just one episode. So we're splitting it in two along those two lines. Today we're talking to the founders focused on taking existing designs, scaling them down maybe with some tweaks, and manufacturing them. Next episode we'll talk to people building advanced reactors with new designs beyond pressurized water reactors.

There's some overlap, to be sure. Some of the companies working on advanced reactors or serving off-grid uses are also rethinking how to build nuclear more efficiently, and we'll include some of their thoughts today. We'll also talk to the investors backing some of these ambitious, if off the beaten venture path, startups to get a sense for how they think about investing in such atom and capital-intensive companies. For today's episode, it's time to manufacture.

Julia

We've talked about this idea of experience curves or learning curves a few times, and it's particularly relevant here. Experience curves describe how unit costs decline as cumulative production increases, and it's often captured by something called Wright's law, which is like Moore's law for manufacturing based on production.

In 1936, Theodore Paul Wright, an aerospace engineer, assistant MIT professor, and administrator, noticed something while working as the director of the Curtiss-Wright Corporation. Every time the company doubled its airplane production, the cost of each plane decreased by a fixed percentage. That observation, which would be named Wright's law, held across industries from semiconductors to cars to consumer electronics to batteries, and most pertinently, to renewable energy like solar and wind.

The chart you see that shows the precipitous decline in the price of solar and wind is the result of those technologies' experience curves. Simply put, the more units of a product you manufacture, the cheaper each unit becomes. But nuclear, as we've discussed, has largely been constructed, not manufactured, and certainly not mass manufactured. The thesis behind SMRs is that since they are smaller and built more modularly, they can be mass manufactured and thus move away from the tyranny of bespoke construction projects.

Packy

When we spoke to Mark Nelson, he made fun of me because I used the joke term "large modular reactor." But there's a point there, right? Aren't the cores of AP 1000s manufactured, and it's just that you have to take it on site and build things around it?

So it turns, on balance, into a construction project. Are there differences besides size, or is it really like, if it's just smaller, you can do more of it in the manufacturing facility, and do less construction?

Julia

A lot of it is size. Some people give the whole SMR movement a hard time for being a marketing ploy. Yes, largely they are just much smaller. But manufacturing smaller things is more scalable. Think about having to manufacture something massive and the type of machinery you'd need for it. The bigger it is, the more bespoke it feels, even in a factory setting. If you're thinking small, you're really printing them out. There's something to be said for smaller being more conducive to manufacturability overall.

It's smart that people are saying we have these large but more modular reactors than the initial ones. To your point, though, when you need to bring a very large reactor core into what is basically a construction project, on net it becomes a construction project. The real challenge for SMRs is going to be, can you really have them come off the production line and just put them on a truck, drop them somewhere and have them turn on? Or will you end up needing to construct shielding or other safety features around them? Time will tell how much these SMRs do look, on net, to be much less like a construction project than these large reactors.

Packy

When it comes to manufacturing nuclear capacity, we went deep with three founders doing just that. Matt Slotkin of Blue Energy, a nuclear integrator, licensing other companies' reactor designs and utilizing

shipbuilding factories to build floating nuclear power plants. Bret Kugelmass of Last Energy, who's building 20 megawatt small modular reactors using tried and true pressurized water reactor designs. And Isaiah Taylor of Valor Atomics, who wants to use nuclear to cheaply produce hydrocarbons.

Matt and Bret are working with the existing design of pressurized water reactors, and working to scale their production. Isaiah is taking a new approach and selling into a different market. But his thoughts on manufacturing were so helpful to our understanding that we decided to include him in both episodes. Each of them is taking a wildly different approach to manufacturing nuclear.

But one thing they have in common is that they're looking outside of the nuclear industry for inspiration. Instead of viewing nuclear reactors as some special class of particularly difficult to build thing, they view them as just another pretty hard to build thing.

On the last episode, we basically tried to answer the question: given how the nuclear industry works, what can we change to speed up and cost down what we currently do in order to build more nuclear? But Matt, Brett and Isaiah are asking a different question: Given that we want to build a big, complex machine, which industries do that best that we can steal ideas from? In other words, what can we take from the industries that benefit from experience curves in order to put nuclear on that same trajectory?

Julia

So, first things first. How far off are we on cost compared to those other industries?

What kind of opportunity does manufacturing nuclear plants in factories present? Isaiah breaks this down as simply as possible.

Isaiah Taylor

When I was evaluating whether or not I should start Valor and forecasting into the future, I wanted to know, if we could actually operate this thing at scale, if we could be bringing nuclear reactors off the assembly line and installing them into this hydrocarbon production grid, what's the actual cost of that? What could we get that down to? I looked at a bunch of other industries that create large, complex pieces of machinery, and I tried to give them a general t-shirt size of how much this thing costs.

I used a funny metric that you might laugh at, but I found quite useful. I said, let's just take the volume of an industrial piece of equipment - a large diesel generator, a combine, a bus, a semi truck - and find out what the volume and sale price was. The dumbest metric we can think of, but you could apply it across an extremely large range of industrial machinery. Well, it turns out no matter how complex and complicated and small scale of a thing that you're building, in the industrial manufacturing world, things generally fall between $500 to $10,000 per cubic meter.

If you want a cubic meter of industrial machine, if that thing is super dense and complicated, it's going to be like $10,000 per cubic meter. If it's more sparse and less complicated, something like a combine tractor, it's like $500. The highest I could possibly find outside the world of nuclear is the Falcon Nine rocket, which is about $50,000 per cubic meter.

If you try to do that with something like a nuclear power plant, you're looking into the $500,000 to $2 million per cubic meter range. So it's almost not even worth looking at how we can make this process more efficient, or how we can bring the costs down. You actually need to clear the slate and say, how do you actually manufacture these, and how do you manufacture them for the right customer with the right target, in a way that allows you to use more traditional, normal manufacturing processes to get more inside that $500 to $10,000 per cubic meter envelope.

Packy

Did you catch that? A nuclear power plant, which uses a practically free physical reaction to heat up water to spin a turbine, costs 10 to 40 times more per cubic meter than a Falcon Nine rocket.

Isaiah suggested that the difference between construction (building things bespoke on site without repeated processes or automation) and manufacturing (building things repeatedly in a factory, often with the help of

automation) are so huge that efficiencies and cost cutting don't really matter, as long as you can move from construction to manufacturing in the first place. That's one of the big bets that this batch of nuclear startups is making, that by building the same thing over and over again - by printing out nuclear reactors, as Matt puts it - you can dramatically lower their costs. You can bring them in line with other things that we manufacture, and draw lessons from hundreds of years of manufacturing experience in those industries.

Julia

To be clear, the concept of manufacturing reactors isn't new. Isaiah told us that the first thing a lot of people realize when they get into nuclear is that we should probably be manufacturing reactors.

Nick Touran highlighted that the NRC actually granted a license for shipyard nuclear manufacturing back in 1982.

Packy

SMRs are not a new idea, but the concept itself is only part of the story. Last Energy and Blue Energy get to start with a fresh sheet of paper and craft a strategy that incorporates lessons from previous attempts. They're a perfect lens through which to look at the opportunity in manufacturing nuclear reactors today.

Like Matt, Bret at Last Energy is looking less to earlier SMR companies and more to other industries for inspiration.

Bret Kugelmass

In our early rollout, we're going to use best practices from offshore oil and gas platforms, where they put an extremely high premium on field labor because their field is out on an ocean platform. Given the historical record of nuclear companies, most of the money gets spent on construction mess-ups, delays in schedule, and high interest rates on construction loans.

We go after that very intensely. We treat it as if it's an offshore platform and say we will do anything and pay almost anything to minimize skilled trade labor in the field. That comes from offshore oil and gas first and foremost.

As we begin to level up our throughput, we'll transition more to airplane manufacturing techniques, and then beyond that, into automotive manufacturing techniques.

Julia

Bret says he will pay almost anything to minimize skilled trade labor in the field. The cost difference between construction (skilled trade labor in the field) and manufacturing (non-skilled labor that is efficiently used) is so high that it's worth investing more in manufacturing to do as little construction on site as possible.

What's amazing is that we heard versions of the same idea from Bret, Isaiah, and Matt. When all three noticed the same thing, we paid attention. Here's Isaiah.

Isaiah Taylor

The trade-offs that generally happen in manufacturing processes, like autonomous processes, or even in nuclear, tend to focus a lot on heat efficiency at the expense of manufacturability. If you can squeeze a couple percent of heat efficiency out of a process, but there's another human in that process, or that tank is a weird shape that's hard to manufacture, that's a trade-off that generally happens.

I'm on the complete opposite side. I'm willing to lose points in heat efficiency if this thing is ten times easier to manufacture.

Packy

So there's the opportunity: if you can manufacture reactors and minimize on-site construction as much as possible. These entrepreneurs believe that you can make nuclear so cheap that you don't even need to sweat little efficiency gains until you're further down the experience curve. But how do you build a company around that idea? Nuclear is hard. Manufacturing is hard. Venture capital is more typically used to start software companies.

The best way to understand how smart founders are thinking about building hard nuclear manufacturing startups is to hear how they've thought through the problems themselves. To be clear, again, there's no guarantee that they'll succeed, and we'll discuss some of the challenges later in this episode and in the next one. But for now, let's have Matt and Bret walk us through how they decided to enter the fray and how they planned to win.

Julia

The path Blue Energy's Matt Slotkin took to nuclear is straight out of Trey Stevens and Markie Wagner's "Choose Good Quests." After spending time at Bridgewater and founding a successful software startup, he asked himself, "What are the most important problems? What are the problems I'm working on, and how do they compare?" In that journey, he realized he'd accumulated the skills and resources to do something really big.

After exploring a few industries, he landed on nuclear. In the intro to the episode, Matt said the question was, "What are the big problems, and what's the best way to solve them?" The answer he came to was one you're probably familiar with at this point in the season.

Matt Slotkin

The biggest problem with nuclear today is it's a terrible business. So to me, the big challenge is how to make it more attractive? Part of that is standardization. Every time we build a nuclear power plant in the United States, it's a different group of people in a different place, with different requirements, constraints, environmental factors, and permitting. It becomes very difficult to finance.

If we made it easy to finance, easy to regulate, and straightforward for both of those things to happen, I think we'd actually bring it down the curve. It's worth noticing what we build a lot of today in the United States. We build a lot of wind and solar and combustion. Those are all things that are essentially manufactured. We basically print them in a factory, and then we staple them to the ground.

That's a lot easier than building a giant concrete civil works project with thousands of on-site people that need to be found, trained, hired, relocated, housed. It's really complicated.

Julia

So Matt and his co-founder, Jake Jurewicz, analyzed the nuclear industry and spent hundreds of hours talking to people and visiting plants. They came up with two core challenges that need to be balanced.

Matt Slotkin

One challenge is creating something that someone will buy - that's the construction and manufacturing challenge. The other is creating something that can be regulated. I think every nuclear project is a bet on balancing these two things. Advanced reactor groups bet that if they build new reactors, they'll convince people that while it's unproven regulatorily, it'll have such a good business case because it'll be much cheaper and safer.

Inversely, incumbents say the regulatory process is so burdensome and difficult that they'll build what they know how to build. Yes, it has challenges, but it's the only thing that's worked so far. Right now, there isn't a big success story in the nuclear startup world or the fission world because no bet has been proven correct on that balance. We tried to look at that balance and take the fewest risks on both sides to build something that strikes the right balance.

From the regulatory side, we don't want to design or invent a reactor. We don't want to use anything the NRC doesn't already understand. Many nuclear startups have looked at the NRC and decided to go elsewhere because it's less burdensome. That's a different problem where they don't necessarily have the infrastructure to do the regulation, and that hasn't really panned out in many cases.

We want to be in the United States, use a regulator that actually has a track record, and use something they really understand. We want to use pre-certified reactors, like the SMR-size light water reactors that power submarines, aircraft carriers, and almost every commercial power plant in the world. We want to use this reactor technology and then figure out how to build it very cheaply.

Our approach is to centrally manufacture them in shipyards. The entire power plant is manufactured in the shipyard, where you already have a pre-existing labor force and equipment. You don't have to hire, train, or house people. Maybe most importantly, once a shipyard builds the first one, the same people can build the second one and actually get along the learning curve, which nuclear hasn't really embarked upon.

It's more expensive to build nuclear now than it was 20-40 years ago, inflation-adjusted. We think that by using technology people understand and that has been regulated, using centralized manufacturing in a process that can produce multi-billion dollar cruise ships and oil rigs in under two years, you have an opportunity to build something that can be regulated and be much more cheaply and efficiently built, which can radically reduce the cost.

Packy

I love the idea of printing out nuclear reactors, putting them on ships, and then sending them around the world where they sit offshore to power progress on land.

Blue Energy is treating reactors as just another hard thing to be built and tapping into the expertise, skill and talent in the shipbuilding industry to start lower down on the experience curve than they would if they started from scratch. And the thesis goes, as these experienced shipbuilders build more and more nuclear reactors, the cost of each reactor will come down, too.

Matt Slotkin

One of the beautiful things about shipyard manufacturing is that it's not a hypothetical technology. These shipyards are extremely efficient, well-run, and know their learning curves so well that they'll price them in. If you want to buy ten ships, they'll tell you how much the first one will cost and how much the 10th one will cost. By the 10th of a kind, for a ship with similar complexity to what we're talking about, it's 35% cheaper. That's a big deal in the energy world where you're basically selling electricity, a commodity. But I think that's just the beginning.

Many nuclear power plants never get to that 10th-of-a-kind point. They build one and then move on, maybe building another in a totally different place with a different group of people, making it even more expensive than the first. But getting to a place where you're printing these out and within building ten of them, you're 35% cheaper, is significant. You continue that curve down, and by the 50th one, you're 50% cheaper. It gets much cheaper, much more quickly.

It's easy to be cynical and think that nuclear power is just complicated and expensive, but historically, that's not the case. In the '50s and '60s, we were building nuclear power plants for $800 per kilowatt, radically cheaper than coal and cheaper than almost any other energy source today. There are reasons why it got very expensive, many well-intended, but it's not inherent to the technology. Tapping into any kind of learning curve would radically help nuclear.

Julia

We heard Matt's point loud and clear. It's not inevitable that nuclear will get cheaper if you simply start manufacturing it. But there's a lot of precedent that says this should be possible. Bret Kugelmass also came to nuclear after startup success, and he identified many of the same challenges and opportunities that Matt did.

Nuclear isn't inherently expensive to build, and there's no reason it shouldn't be much cheaper than it is today. Last Energy is building SMRs based on standard pressurized water reactor designs, innovating not on the technology itself, but on the business side, the reactor size, who they sell to, and how they go to market.

Bret Kugelmass

What does a gigawatt coal plant cost? Let's say a $1 billion dollars. What does a gigawatt nuclear plant cost? $10 billion. What's the functional difference between a coal plant and a nuclear plant? A nuclear plant should be cheaper because you don't need all the fuel handling equipment, coal yards, or trains. You can drive the fuel up on a truck. If anything, a nuclear plant should be cheaper than a coal plant on a per megawatt basis, yet it's ten times as expensive. So where did all that money get spent? Regulatory requirements.

After studying this for two years, we realized this was the problem, not technology. It would have been fine if we just kept building my favorite two plants, Point Beach 1 and 2. There's 1100 megawatts for you at less than $1,000 per kilowatt, built in under three years, 1968 to 1971. Ideally, we'd be building 10,000 of those everywhere we need power plants across the world and the grid. I have no problem with 1968 technology.

But how do we get from here to there? You still need to put together a lot of money to do that, and people with money aren't necessarily going to fund that idea or take the risk. So our thought was to go small and go to different regulatory environments where their paradigm is fundamentally different.

Julia

Bret founded Last Energy to build small and build in the UK, where differences of opinion after the Manhattan Project pushed the government to create its own independent nuclear industry and regulatory framework. Bret believes they'll be able to move more quickly through the regulatory process in the UK than the US. The trade-off is using standard reactor design to minimize time spent introducing regulators to a novel design, which would likely take years and add uncertainty to the process.

Plus, Bret believes the current reactor designs work great. We have many of them efficiently and safely operating in the US and globally. So why change anything?

Bret Kugelmass

We pride ourselves on being the least innovative nuclear company out there. We focus on innovation around manufacturability, not nuclear technology. We won't touch the core physics or chemistry.

As much as people like to say these reactors are proven because national labs built them once, you've got to operate something for 30 years before you begin to see what random corrosion issues happen, which bolts wobble loose on your pumps due to the frequency of how you arranged your pipe hangers. You need decades of expertise. This affects your economics today because it affects the rate at which you can borrow construction capital and the operational performance into the future.

This is why we decided to settle on the standard PWR, shrink it down, and design it for manufacturability. We think about how to build 10,000 of something a year, even on our first power plant. Granted, there will be design improvements and iterations, but I direct our engineering team to think about the big picture, high-throughput manufacturing. That's how we need to make design trade-off decisions.

Packy

Thanks for listening so far. We'll be right back after a quick word from our sponsors. One thing I've noticed, talking to nuclear startups and hard tech startups more broadly, is how important it is to have a strong strategy from the beginning. The need to spend a lot on capex early really makes you think twice about where you're heading.

Plus, as a16z American Dynamism partner Catherine Boyle explained, the dynamics in these hard tech categories are just different.

Katherine Boyle

One of the things unique about this category, especially with American Dynamism companies, is that these are often terrible categories with one extraordinary winner, or a handful of extraordinary winners. The power laws are so much more pronounced. Part of that is because of the regulatory aspect, where once you can get through the regulatory hurdle, you have a much deeper moat.

If you compound that with manufacturing hardware, physical moats that are built with these companies. That's why you see something like a SpaceX, which is the most valuable private company in the US.

But when you look at where the next ten aerospace companies are, there's such a larger delta between the winners and the runners up, which is not true of cybersecurity or typical SaaS investing. There's many more winners where the gains are distributed across many. So this looks like a very different category.

When we talk about nuclear, defense, or aerospace as a whole, these are categories where you really do have to be in the top handful of companies and work with those companies, and they become extraordinary outcomes.

Packy

Getting the strategy right is critical. My favorite strategy book is Richard Rumelt's "Good Strategy, Bad Strategy." In it he writes that a good strategy has three parts. One, diagnosis: it names or classifies a situation, linking facts into patterns and suggesting more attention be paid to some issues and less to others.

Two, the guiding policy, which outlines an overall approach for overcoming the obstacles highlighted by the diagnosis. And three, coherent actions: the resource deployments, policies and maneuvers that are undertaken should be consistent and coordinated. The coordination of action provides the most basic source of leverage or advantage available in strategy. Done right, strategy is based on a realistic assessment of the challenges and opportunities, a general direction about what's important. From there, companies can make better decisions about which specific actions to take.

A good strategy doesn't guarantee success, but it gives a company the best shot. Talking to Matt and Bret and to the founders you'll hear from in the next episode made my strategy nerd soul happy.

Julia

Let's dive into Blue Energy and Last Energy, and where they look similar in what they're trying to do with building out their SMR facilities. Both companies believe that manufacturing and doing things repeatably in high volumes is going to bring them down the experience curve, but they each have a different approach. Blue Energy is looking at using shipyards as the place for their manufacturing.

These are places where you have an experienced workforce that has been building ships for many years, repeatedly, and has a lot of similar processes. They are doing more of an integrator partnership approach, whereas Last Energy is going to stand up their own manufacturing facilities to manufacture their SMRs.

Blue Energy operates as an integrator, needing to be in partnership with other organizations like these shipyard builders. Last Energy, since they're standing up something themselves and vertically integrating, is able to have more control.

That said, Blue Energy gets to move faster since they're integrating with someone else and also benefit from the experience of those shipbuilders. It's interesting to see the pros and cons of the different approaches towards how they want to set up their manufacturing.

Packy

It's interesting to pick apart the interlinked decisions that companies make. Blue Energy is trying to build 100 megawatt reactors and Last Energy is building 20 megawatt reactors. It's amazing to me that Last Energy is even standing up their own manufacturing facilities to do this. But I can't imagine Blue Energy taking some VC money and going out to stand up shipyards that can produce these five times larger reactors.

Everything is linked. It'd be nice to have access to a shipyard that's all yours with the workforce as part of your team. Maybe down the line they grow and can do acquisitions or vertically integrate in another way, but you almost have to give up some of that control if you want to go larger. It's just fun to see how all those different pieces fit together in the companies' strategies.

Julia

Blue Energy and Last Energy are thinking differently about how they'll ultimately sell their power to end users. Blue Energy, especially because they're building bigger SMRs, is considering selling directly into the grid for states like California or countries like Poland, setting up long-term power purchase agreements about what price is paid for that electricity over a long period of time.

Last Energy, on the other hand, is targeting the industrial market - because 20mw is a great size for that industry, and then sell them 'behind the meter' directly to those businesses instead of going through the grid.

Packy

This is why the strategy stuff is so fun. Once you start looking at these companies, even though they're manufacturing nuclear reactors, you see how all these decisions influence each other. If you're going big, you should sell into the grid, which affects your business model, sales team, and hiring decisions.

Last Energy is selling to industrials, and interestingly, Microsoft just announced they're exploring using nuclear energy to power some of their data centers. A 20-megawatt facility feels like a good fit for that, as we discussed with Bret.

Bret Kugelmass

Yeah, Microsoft's the big news right now. Actually, one of our early investors is the former chairman of the board of Microsoft, Dave Marquardt. So I'll just leave it at that.

Packy

It's all integrated from diagnosis through guiding policy, through coherent actions. Microsoft recently announced hiring someone to implement global small modular reactor and microreactor strategies to power data centers. They and others will need a lot of power to run energy-intensive data centers to keep up with growing AI demand. Those data centers need a lot of power, but not a gigawatt. In addition to being easier to manufacture, Last Energy's 20-megawatt reactors are well-sized to meet the needs of large data centers and industrial sites. It's a match made in heaven.

Bret told us the company has already signed eight contracts representing 51 20-megawatt power plants to date, also under a PPA structure. Given the capital intensity of manufacturing nuclear reactors, the strategy extends to financing. One of my favorite things to discuss with nuclear entrepreneurs is how they think about financing their businesses. Venture capital is great for many things, but can be an awkward fit for nuclear. Josh Wolfe told us that Lux looked at SMRs, but they didn't make sense for his venture fund at the time.

Josh Wolfe

What about these small modular reactors? Could you build 100-megawatt reactors that are much cheaper and just add them as you need electricity? That seemed like a good idea. We probably saw about six or seven different schemes, some of which are being commercialized today.

We did not fund them because we thought it was too long, too much bureaucracy, too much regulatory framework. You had this weird chicken-and-egg problem with the Nuclear Regulatory Commission basically saying, "We're only going to help sponsor something if there's a customer."

And a customer would say, "Well, we're only going to put money in if you're going to get the approval." So there was this chicken and egg problem that created this paradox and this inertia, and you never saw anything really fun there.

Julia

The beauty of entrepreneurship is that founders can study the issues that prior attempts faced and design around them. Plus, as we talked about earlier, both Bret and Matt had founding experience under their belts by the time they came to nuclear. So they understood what VCs would and wouldn't finance.

Financing a nuclear reactor startup is a challenge, but it's solvable. Matt told us about the different funding alternatives available.

Matt Slotkin

There's the literal funding of the company, which is done through equity, like any startup might, to fund the team, the design, and the development of this platform. And then there's the literal plants themselves. There's a couple different ways to do that. One of the SpaceX approaches that I think was amazing is they built the rockets on their own balance sheet. If you can raise that much money and do that, it lets you move so much more quickly.

If we wanted to get to a point where we were doing regular project financing, it's up to us to make it as cookie-cutter for those folks as possible. I think that's how most energy products get built, with project financing, which is primarily debt. There's good reason to believe that would be done on a project-by-project basis.

That's also part of what makes energy more complicated than software - you've got something great, but you've got to align the capital each time you want to deploy it. And that can be a challenge.

Packy

Bret thinks about funding Last Energy in a similar way.

Bret Kugelmass

Let's differentiate between two types of equity: corporate equity and project equity. Corporate equity is what most people think of from the VC community. This is where you sell part of your top-level company, including all your IP, to an investor for a percentage and amount of money. You then use those proceeds to grow value for all shareholders.

Project equity, sometimes called project finance, involves setting up a special purpose vehicle (SPV). You put certain assets in this vehicle, which is wholly owned, a subsidiary, or joint venture of your corporate entity. You can sell shares of this SPV to different people. This is still equity, but a different type, and it's non-dilutive to the top company.

So yes, you can sell equity, but we're selling equity in the power plant corporate entity, in the physical asset and real asset itself.

Packy

To Josh's point, funding a nuclear startup purely with venture capital doesn't make sense. Both Bret and Matt plan to use a mix of VC dollars to fund the company and project financing to fund the projects. Each is taking different approaches to the challenges Josh laid out.

Last Energy, for example, is focusing its efforts in the UK, where regulation is somewhat less burdensome, and pre-signing customers to smooth the regulatory process and make projects financeable. It's much easier to ask the capital markets to fund a project with signed contracts and long-term guaranteed cash flows on the other side.

Blue Energy made trade-offs explicitly designed to ease the burden of working with the NRC, using a design they're already comfortable with and a manufacturing approach they believe will make all approvals after the first go much more smoothly. He explained this by walking us through his thought process on how long it should take to manufacture each reactor.

Matt Slotkin

Let's go backwards. To start building, you need a construction license from the NRC, which will optimistically take a minimum of two years. We're not showing them new technology; we're taking a reactor they've already looked at and putting it in new packaging. We could be blissfully surprised or disappointed by that process, but two years would be optimistically realistic. To get to that point, you need about two years minimum to find a site and do the necessary work.

You can't just submit an application with an idea; you need a concrete proposal. The license is given to a particular site, not a design. In a great case, you could be looking at something in six years. These are not technically constrained problems, but people and bureaucracy problems. I don't mean that cynically; it's literally about working with people to collaborate, get approvals, and be allowed to do things.

The exciting thing is, whether the first one is in six, seven, or eight years, you now have a fully baked supply chain to start printing these out. This isn't about building one and going through the whole process again for a second. Now we have a shipyard with staff who know how to make these things. They can make a second one, and the regulator has seen this exact design with all the questions they've already asked.

Julia

Venture capitalists typically operate funds with a ten-year lifecycle, meaning they must return money to their limited partners within ten years. For a startup to be successful in a VC's portfolio, it needs to go public or get acquired roughly within ten years. So the bet on Blue Energy from a VC perspective is that the company will get its first reactor generating power, have PPAs for future reactors, and a clear path to much faster rollouts within those ten years.

All the hardest work will have been done, making it a company the public markets would be willing to invest in. Oklo, a nuclear startup founded in 2013, recently announced it's going public via SPAC co-led by Sam Altman, exactly a decade after its founding. Its co-founder and CEO, Jake DeWitte, told us there are real advantages to being public, chiefly that it helps with project financing the company's plans through the PPAs it's signed. He's been pleasantly surprised by public market investors' appetite for and familiarity with nuclear.

Jake DeWitte

I've found an incredible amount of familiarity, savviness, depth, and nuanced understanding in the public investor markets. It's been awesome, and I think that's a really good sign for the industry as a whole, because it's supporting a lot of demand. I would argue there's probably more demand than there is actual peer exposure in the public markets, which is good. This also helps pull in support in the private markets, which is great throughout the whole pipeline. I think that's really a positive here.

It's also important because investors are increasingly seeing that there's differentiation, that not all nuclear things are going to look the same. They can take different approaches, and investors are getting smart across the board on those things. Think about what's in the public domain now. You've got a small light water reactor, a small gas-cooled reactor, and a small liquid metal fast reactor. It's kind of cool to see that exposure. Investors are getting pretty smart on this stuff. I think that's going to carry through.

It also creates a good signal on the venture side of things, which is really important for sustaining the pipelines. I have to say the reception and uptake has exceeded our expectations, and we were already pretty optimistic. One of the hard things here is that all of us are meaningfully pre-revenue for a while, and most nuclear companies are going to be like that for a bit. But investors understand what that means and they know how to look at the risk from that.

They're seeing that there are several ways to understand how milestones come together and how progress gets made, which is awesome. This supports what that roadmap looks like, and then allows you to focus, iterate, and communicate accordingly.

Packy

What Jake said about public market awareness creating more support in private markets is crucial. The same day we spoke to Jake, I'd talked to a pre-seed nuclear founder who cited Oklo as an example of a company that would generate most of its revenue and impact after its first decade, but could get its investors liquidity within that ten-year window. More examples like Oklo should make it easier for new nuclear startups to raise venture capital.

Ten years isn't a hard rule. Many VC agreements with their LPs give options to extend fund life by a couple years, often multiple times. No one will complain about a 100x company taking a couple more years to get liquid. But it's helpful for investors to see big outcomes are possible within about a decade.

Julia

Investors are tempted by the sheer size of the opportunity. Global energy markets are in the $4-6 trillion range annually, practically limitless. Energy has the useful feature that the cheaper and the more we produce, the more people use it.

Operational reactors can be a great business. Blue Energy plans to manufacture 100 megawatt reactors. Let's do some back-of-the-envelope math: Assuming a 90% capacity factor (slightly lower than current US levels), plants would run 90% of available hours in a year, or 7,844 hours.

That's 784.4 gigawatt hours of electricity, or 784,400 megawatt hours. At a rough wholesale price of $40/MWh, you're looking at $31 million in annual revenue per reactor. These can run for decades, and Matt isn't just thinking about making one reactor; he wants to print hundreds.

Matt Slotkin

We're talking to shipyards in South Korea that make about 50 big ships a year, four a month. In terms of complexity and steel tonnage, what we're describing is similar to one of those ships. A single shipyard, even at 25% capacity, could produce one reactor per month. These are small modular reactors, around 100 MW. That's over a gig a year.

By 2038, you're at six gigs, which is a lot of ARR for us, but not hugely meaningful for global energy. If there's high demand for secure, clean, reliable power on demand, and you took over the whole shipyard at 100% capacity, by 2038 you're at around 35 gigawatts, about $25 billion of energy demand.

If you're adding a new shipyard at 100% capacity every year, you're at well over 150 gigawatts in the next 50 years. We could end up anywhere in that range. The whole point of the design and the reason why this is exciting is that there's a world where you can ramp up, and there are 200+ shipyards globally of this size.

Packy

A lot has to go right for Blue Energy to reach 150 gigawatts of energy. If they get there, that's $100 billion per year in essentially annual recurring revenue. That seems crazy, and it's not guaranteed, but energy markets really are that large. Of the 10 largest companies in the world by revenue, 7 are energy companies. Of the top 25, 13 are energy companies.

We spent time talking how these businesses are financed because it's important, perhaps as important as reactor design or manufacturing process. One main reason nuclear has struggled is because, as Matt put it, it's a terrible business, but there's nothing inherently terrible about it. Uranium is an extremely energy-dense fuel source. Regulation is a real issue, but entrepreneurs view it as part of the design challenge, not a non-starter.

Matt, Bret, and Isaiah believe that nuclear reactors can be manufactured as cheaply as any other big, complex machine. If we start treating them as such, financing remains an enormous risk. Will these companies get the funding, both venture and project financing, needed to stay alive long enough to ride down that experience curve?

Julia

Josh Wolfe told us that to be successful, nuclear needs to tap into the two emotions that motivate: fear and greed.

Josh Wolfe

The greed aspect typically needs one story for people to rally around, one narrative of success. It could be a public company or somebody who made an investment in an existing utility, like Exelon Chicago, or someone seeing runaway success that people suddenly galvanize around. That's the beautiful thing about markets.

Somebody, whether they want to leave a legacy, truly believe in it, or think they're going to make a fortune, even if it's a 1% or one basis point probability of working. But the magnitude if it does is so enormous that it appeals to that vector of greed.

Julia

When it comes to manufacturing more nuclear capacity, in the immortal words of Gordon Gekko, greed is good. Trillion dollar markets are up for grabs in the energy transition.

Packy

I think we need to rebrand the energy transition. It sounds a little corporate and boring to the electronaissance, by the way.

Julia

All right, I dig it. Trillion dollar markets now up for grabs in the electronaissance, then. The capitalist bet here is that with a big enough prize at stake, entrepreneurs will figure out ways to work through all of the issues we discussed in episodes two and three.

That the best way to decarbonize isn't through degrowth or even government intervention, but through good old-fashioned capitalism.

Packy

That was one of the most interesting things to me in these conversations. I've been so worked up about regulation as an outsider, someone who doesn't have to actually deal with the NRC. The founders' response has essentially been, "It is what it is, we'll play the game on the field." Obviously, they'd all be happier if the regulatory process was smoother. They're signing themselves up for years of paperwork, meetings, and delays.

There's no question that we could decarbonize more quickly and get more energy on the grid faster if the NRC adopted a more pragmatic stance. But the incentives here, both to make an impact and financially, are so strong that the right kind of people - often second-time founders with resources and experience to take big swings - are willing to go through the pain for even that 1% shot at making it out the other side.

There's this great quote from Thoreau that "Trade and commerce seem to be made of rubber, because they always manage to bounce over the obstacles which legislators are continually putting in their way." That one seems apt here.

Julia

I'm optimistic about this category of startups, these Blue Energies and Last Energies of the world that are saying, "We're not going to innovate on design, we're going to focus on manufacturability, and then we're going to go to market with this new power purchase agreement model." I think there's a lot of potential there. Solar going small and improving manufacturing might not feel like it's going to have this big impact the way that gigascale reactors do.

But I think we'll see that this will pay off. They will come down the experience curve and the productivity learning curve. In the end, we'll wake up in 10-15 years to dozens of gigawatts of new nuclear capacity being produced

every month.

Packy

On the next episode, we have five more startups for you. We're going to go even more sci-fi and talk to founders who are building advanced reactors using different fuels and coolants, serving end markets like the military, the global hydrocarbon market, and even the moon.

Thanks for listening and see you next week on the next episode of Age of Miracles. Thank you for listening and watching to this episode of Age of Miracles. If you like what you hear, please rate, subscribe, and share.

And if you're feeling really generous, tell us what you think in the comments. Plus, we have a ton of resources and references in our resource hub if you want to go deeper, and we've linked them all in the show notes below.