Episode 9

Fusion Founders

David Kirtley

In the world, there's over 4,000 gigawatts of installed fossil fuel capacity, and our goal is to replace it all. I don't think one type of fusion is enough to do all that. We're going to try, we're going to move as fast as we can. But I think you're going to have different kinds of power in different locations, whether it's remote, military bases, giant factories, or data centers.

Our plan is to demonstrate electrons on the grid in 2028, then scale manufacturing as fast as possible to deploy these systems. Our goal is to get to making generators per day by 2030, rather than generators every few years. That's a big scale, a big lift.

As a human, I want other fusion and advanced, carbon-free power out there in the world, too, because we have that big of a need, and we need to move that fast.

Packy

The person you just heard from is David Kirtley, the founder and CEO of Helion Energy. The dates David mentioned, 2028 and 2030, are five and seven years away. He believes that Helion Energy will deliver fusion-generated electrons to its first customer in five years and to the grid in seven.

On the last episode, we talked about the joke that fusion is always 30 years away. But today that joke is far too pessimistic.

Julia

Helion is one of about 80 startups working to make fusion happen within the next decade in the bizarro relay marathon that is the fusion race. The baton is firmly in the hands of these startups. The outcome is still uncertain.

But the question isn't whether humanity will achieve commercial fusion, but which companies will, with what approaches, and when.

Packy

This is truly an extraordinary time to be alive. Between solar, fission, and fusion, we're entering a new era in human history. For the first time, we won't primarily produce energy by burning things. We'll be able to manufacture energy by capturing the sun's rays, splitting atoms apart, and fusing them together.

Fusion, many believe, will be one of humanity's greatest triumphs. We'll be able to generate energy in the same way that the stars do, right on earth. And we have one of the coolest jobs in the world because we get to talk to some of the people who are most likely to make it happen.

Julia

Today we'll be talking to five fusion founders and operators, each taking a different approach to generating fusion energy. We'll have David Kirtley at Helion Energy, JC Btaiche at Fuse Energy, Francesco Sciortino at Proxima Fusion, and Ryan Umstattd and Derek Sutherland at Zap Energy.

We'll also hear from a few of the investors we met on the last episode, Clay and Clea at Lower Carbon Capital, who tell us about a few of their portfolio companies like Commonwealth Fusion Systems and Avalanche Energy, and Ian Hogarth at Plural Platform, who will explain why he's making a concentrated bet on stellarators.

Packy

These interviews were one of the coolest parts of doing this series for me personally. Last week I was reading my son Dev a book on quantum physics because I'm trying to correct my lack of physics knowledge. One of the pages was about fission and fusion. The next night I had to miss story time for our conversation with David, and when I got off, it was surreal telling Dev, "Remember fusion, the thing where you smash two atoms together to make energy? I just got to talk to one of the people who's actually going to make that happen."

Dev only kind of cared, but it was one of those moments where it hit me that many of the people we've spoken to this season, including you, have a shot at making the world legitimately better for a lot of people, including our kids. It's why we started Age of Miracles with fission and fusion. We think this stuff really matters. Anyway, we have a lot to cover in this one.

Julia

Let's do it. First we'll answer the question of why now for fusion, and for startups in particular, by talking to the companies taking advantage of the moment. Next, we'll dive into the business models and the economics of a fusion plant. Finally, we'll discuss what a rollout of fusion across the globe might look like. That feels crazy to say. It makes fusion seem almost mundane, like an ordinary business.

But it's something these operators are already thinking about and planning for. If commercial fusion is going to change the world, it needs to be commercially viable. Let's start with one of venture capital's favorite questions: Why now? Why, after 80 years in government and academic labs, has fusion finally broken into the commercial sector?

And why might these companies, which have much smaller teams and smaller budgets than these international government projects, actually make fusion happen?

Packy

"Why now?" is maybe the most important question in fusion. There are four main categories of answers. One, government-funded breakthroughs. These startups really are standing on the shoulders of giant research programs. Two, the funding landscape. With climate change a looming threat, both governments and venture capitalists are willing to back companies that have a shot at contributing to a solution.

Three, technological advances. Better materials, software, and components have been game changers for the startups we've spoken to. And four, startup speed and iteration. Inspired by companies like SpaceX, startups are taking approaches that governments can't or won't to get to market more quickly. Clean Air Task Force's Cecilia Gonzalez has been in fusion for two decades, and she told us she's never seen a time as promising as right now.

Sehila Gonzalez

I have been in fusion for almost 20 years, and I had never seen the excitement and hope for the future that I have seen in the last two, three, five years. Before, fusion was something done in academia and some national labs. Now it's something that is in the Financial Times and The Economist. I think we are really in a good moment. We have private capital and the private sector providing flexibility and a more agile approach to fusion, which is really convenient.

We have a lot of knowledge generated in the public sector, in the traditional sector, which is very important because without knowledge you cannot progress. We have tools that we didn't have before. All the AI and new software tools for example. There's also all these developments in superconductors and these new elements that maybe have been created out of fusion but help fusion to progress.

We are in a time where tools that were not existing ten years ago are available now and more will come in the next year or so. Additionally, there's the need for new sources of energy, which have to be clean. So it's a really good combination of having new tools and having the need to have these new sources of energy which are making fusion go faster than ever.

Packy

Clay Dimas at Lower Carbon Capital said he came into his first fusion pitch skeptical but left sold because of the same trends.

Clay Dumas

Very early on, we were offered an introduction to Bob Mumgaard, CEO and co-founder of Commonwealth Fusion Systems. To be candid, we went in thinking fusion is always 30 years away. But we got this introduction from someone we respect, so we kept an open mind. We came out of a 1-hour conversation with Bob completely sold. We committed to invest in CFS's Series A, their first round of outside capital, and became curious about different paths towards commercializing fusion technology.

One big takeaway was that the trends benefiting CFS weren't exclusive to them and would give rise to other pathways to commercializing fusion. First, major advances in material science had big implications for magnets and superconductors, with broad applications beyond fusion like transmission and cancer research. Second, cheap compute and a generation of engineers steeped in machine learning were advancing towards AI. This had big implications for simulating super weird plasma conditions at 150 million degrees Celsius, where you're trying to figure out how these tiny particles interact with one another. Having greater fidelity of what was happening in those conditions by running computer models really sped up the rate of learning and physical world testing. CFS was on the cutting edge of this.

Julia

Each startup we spoke with takes advantage of all four trends in some way. We hope that by hearing from the founders directly, you'll come away with a greater appreciation for how close we might actually be and why. You might even leave this episode with thoughts on which company will get there first. That's one of the fun parts about watching a race.

So let's meet the fusion startups and hear why they think now is the right time to build fusion in their way. While fusion seems like a sci-fi technology, these are serious people with practical plans to bring it online. Their backgrounds range from years in top fusion labs to simple high school backgrounds.

Packy

Let's start with the 800-pound gorilla in fusion, Commonwealth Fusion Systems, or CFS. While we didn't speak with CEO Bob Mumgaard or chairman Dennis Whyte, we recommend listening to Dennis's conversation with Lex Friedman. We'll link to it in the notes and resources guide.

For now, here's Clay to explain what CFS does.

Clay Dumas

The first and probably best understood fusion reactor design is a tokamak. This is being commercialized now by Commonwealth Fusion Systems outside of Boston. It's received the most attention and dollars from researchers over the past 4-5 decades. Consequently, it's where the physics are most de-risked, which is why CFS has raised so much capital.

One downside for tokamaks historically has been their size, needing extremely powerful magnets to confine plasma at over 150 million degrees Celsius - hotter than the sun's center. For a long time, the leading concept was ITER, a multinational effort to develop a Q>10 reactor in France, now billions over budget and decades behind schedule.

CFS has turned the tokamak concept on its head, shrinking the reactor by making the magnets much more powerful using superconductors. That said, Tokamaks aren't fully understood; work remains to keep reactions continuous and contain the plasma. But there's a reason why CFS is often referred to as the safest and surest way to commercial fusion.

Julia

CFS is a giant in the space, spun up out of Whyte's lab at MIT where Mumgaard was a student. The company has raised $2 billion from investors, including Lower Carbon, Bill Gates, Breakthrough Energy, Alphabet, and Khosla Ventures. Like Clay said, it's the safest bet. Tokamaks are well understood, and thanks to advances in magnets, CFS can scale down its reactor, speed up development, and make fusion commercially viable.

Whyte is a legend in the field, having worked on ITER before heading to MIT. CFS expects to bring its first 200-megawatt electric plant online in the early 2030s, and the smart money is betting that they'll do it. But CFS certainly isn't the only startup in the fusion race. ITER to MIT to startup isn't the only route.

JC Btaiche took the most direct and least conventional route into fusion. He skipped college and started building.

JC Btaiche

Growing up, my father was a nuclear physicist, so I was fascinated about the universe and how things work. I really wanted to go see it. I was very disappointed when my father told me that we've never had a human physically see the universe and it's not quite possible. So I started googling how we can go to space and how a human can go to space and come back to tell the world about what they've seen.

As part of my googling, I found the only reasonable possible way to do that is if you build fusion drives, like fusion-powered rockets. So I started aggressively reading about fusion drives, fusion-powered rockets, and really wanted to build one. This led me to be motivated to sit in the classroom a little longer in high school.

I did some research in plasma physics when I was still in high school, which was my first formal exposure to fusion. From there, I realized that to learn the fastest, make the most progress, and have the biggest impact on the field, it's much better to build a company rather than sit in a classroom. So I decided to build a company essentially in lieu of going to college.

Julia

Perhaps because he hadn't been colored by years of research or experience in a lab, JC approached the fusion space with a fresh sheet of paper when starting Fuse Energy out of Canada. He asked the fundamental questions: starting today, what is the best approach to fusion and one that customers will be willing to buy?

JC Btaiche

There were a lot of academics at the time who were very convinced that their research was impressive and great, and it was time to spin it out. But I started looking at where governments are spending billions of dollars, because governments are the most incentivized to make this work. This led me to find out about the Z machine, which is one of the most successful nuclear experiments in the United States and in the world.

It's the highest source of X-rays and has the Guinness World Record for the highest temperature achieved on Earth. I was like, okay, why is no one building it? This was refurbished in 2007. It was pretty old technology, and it had reached very impressive results. It's ten times more efficient than lasers, and it's 10% the cost and size of NIF, which is the experiment that achieved ignition for the first time in history, but no one's paying attention to it.

There's a very clear roadmap that people within the field have laid out for the next generation of the Z machine, but no one was building it. Every time you asked, they'd say, "Oh, we know we need to do this, but it's taking time, we're trying to get the approval." So we just went and built it. Today I think we've built the world's first and highest energy pulse power driver ever built, and so we're working to essentially build the next generation of the Z machine, or MagLIF.

Packy

I'm going to break in here to explain. On the last episode, Andrew Cote talked about the Z pinch generator design. The Z machine is related, but slightly different. It's a specific facility at Sandia National Labs that uses Z Pinch principles to achieve high energy density conditions for research and potential fusion energy production.

It's the world's most powerful and efficient laboratory radiation source, using high magnetic fields associated with high electrical currents to produce high temperatures, pressures, and powerful x-rays. While Sandia uses the Z machine for research on nuclear weapons and validating physics models, it can also be used to incinerate things and generate fusion energy. Back to JC.

JC Btaiche

There are three main reasons we chose it. First, controlled implosion methods objectively lead the race for fusion. NIF has the highest result, Z is right after, and then there's the tokamaks. It's a very practical path with billions of dollars and decades of research behind it, albeit much of it being behind closed doors.

Second, this technology is critical for building a long-term mission company. It's important to pick a technology that could be commercializable. The technology to build the next generation of Z is immediately useful today for multiple national security needs.

Third, this approach has an intermediate step towards providing power, using fusion neutrons to bombard radioactive waste in a hybrid fusion-fission concept. It's a polarizing idea, but it's an option. These were the main reasons we chose to work on this.

Packy

In JC's answer, you hear a couple of the "why now" themes. One, Fuse is building the next generation of a technology, Z machines, initially developed by the government. Two, it's building something that might make sense for government funding and that customers might be able to buy sooner rather than later.

Julia

His idea about the intermediate step is really smart. It's a bit of a fission-fusion hybrid and a perfect bridge between the first half of the season and this one.

Before Fuse gets to fusion power, it can use neutrons from a fusion reaction to create a fission reaction from nuclear waste, kind of like a fast breeder reactor.

JC Btaiche

Traditionally, there's radioactive waste stored after traditional reactors work. This waste decays for hundreds or thousands of years. We can take a fast neutron from a fusion reaction, like deuterium-deuterium, and surround the fusion chamber with actinides (radioactive waste).

The neutron will accelerate the waste's decay rate, reducing the half-life from hundreds of thousands of years to tens of years. Because the decay happens faster, it releases more energy. This process addresses waste recycling, which is a $40 billion industry in the US alone.

The original concept was called the "incinerator" because of the Z machine. If done efficiently, we can produce power, becoming a power generating company separate from waste recycling. This dual-use approach could generate revenue from multiple customers.

Packy

Thanks for listening so far. Hang on, we'll be right back after a quick word from our sponsors.

Julia

Fuse's path to market involves generating revenue from near-term feasible activities, like disposing of nuclear waste and producing power from it, on the path to generating pure fusion energy. Most fusion companies need to build multiple generations of generators on their way to commercial fusion, and JC is betting he can start generating revenue earlier in that journey.

Others are taking different approaches, like Germany's Proxima Fusion, which is leveraging better simulation software to run as much of its design in silico as possible before touching metal. You heard the phrase "in silico" on the last episode when we discussed the world's largest stellarator, the Wendelstein 7-X at the Max Planck Institute in Munich. It's no coincidence you're hearing it again here. Proxima founder and CEO Francesco Sciortino worked on W7-X before launching Proxima and set up the company's operations nearby.

Francesco Sciortino

We created the company in January 2023, with a team mostly from Max Planck and one co-founder from MIT. I did my PhD at MIT with him, and Martin, another co-founder, joined us from Google X in California. The company aims to take the visionary Wendelstein 7-X stellarator project in northern Germany to the next level using simulation-enabled concepts and leveraging today's high-field superconducting magnets.

We can now design solutions to problems that have historically been complicated to deal with in magnetic confinement fusion experiments. The idea is to translate some of the complexity of tokamaks into a more predictable device, one that works like a microwave oven - you turn it on, it runs steady state with continuous operation and no surprises, and you turn it off when you choose. That's what we are chasing as a company.

Packy

Francesco listed a few big "why nows" for Proxima: building off the Wendelstein 7-X and partnering with Max Planck; new materials, specifically high-field superconducting magnets. Stellarators are a form of magnetic confinement fusion, like Tokamaks, so good magnets are key; and most importantly, software simulations.

Proxima plans to design and simulate the reactor, making trade-offs between physics and engineering in software before touching any metal. Recall that on the last episode, Ian Hogarth, whose Plural Platform led Proxima's seed round, told us that stellarators were the platonic ideal of fusion generators. But before these advancements, they were just too hard to build.

When we asked her to explain stellarators, Lower Carbon's Dr. Clea Koolster made a similar point. She said that stellarators have gone from impossible to imagine building, to possible.

Clea Kolster

Stellarators are like the ugly duckling of the tokamak. It's magnetic fusion energy, where you're confining plasma using a very strong magnetic field, but in a tokamak, you have the toroidal magnetic field, a poloidal magnetic field in the middle, and a current running through the plasma. The goal with fusion is to minimize instabilities in your plasma. The optimization between all those different magnetic fields and moving pieces is what either drives the instabilities or keeps them down.

Stellarators have a very twisty, crazy configuration that historically was impossible to imagine building or simulating. Now, with better computing, we can understand how that complicated, twisted magnetic field operates. That's what happened at the German Max Planck Institute with the W7-X reactor, the first to show plasma stability within a stellarator.

What we found exciting in two companies we've invested in is their approach to making stellarators simpler to build and easier to maintain, potentially through controls or laser patterning the magnetic field directly onto the material. The perceived benefit of this reactor configuration is that without the additional poloidal magnetic field, stellarators could be much smaller, potentially driving down costs significantly and requiring fewer materials to make them.

Packy

Francesco agrees that powerful software makes building stellarators more feasible, resulting in a power source as predictable and easy to operate as a microwave oven.

Not being a plasma physicist or fusion engineer myself, I asked him to explain what makes stellarators ideal.

Francesco Sciortino

By the concept itself, you don't have pulsed behavior. Every fusion concept involves some sort of up and down behavior, some form of either implosion or a sudden large amount of energy, or pulses that go over hours. But stellarators are the only concept that is truly steady state. You can build a stellarator that just runs, like a microwave oven. W7-X has demonstrated this, running for minutes last year with nothing happening after 20 seconds or so. That's one key advantage.

The other is that you are fully controlling your hot ionized matter. This plasma that we have to confine at 150 million degrees, you can confine it completely externally with some big coils. The challenge is, can you design coils that can go to high enough magnetic fields? Fusion power scales strongly with magnetic field intensity. If you can get this magnetic cage done well, you start addressing other aspects of the design. You have to support huge forces, deal with humongous heat fluxes, and lots of other things.

You need capability in designing and assessing trade-offs. That's the nature of Proxima Fusion, a group with these tools and understanding of where to put effort on physics and engineering questions. We believe we're transitioning from a physics focus to an engineering focus with a mindset on commercial viability. Stellarators, in our opinion, have a much better market fit. If you can deliver continuous baseload with a simpler device, you have a better future.

The question is, can you design and manufacture it? If W7-X hadn't been manufactured with incredible technical tolerances, it would be a bad idea to pursue this. But we've done it.

Julia

This physics and engineering trade-off has come up a few times, not just with stellarators. It's core to Zap Energy's approach. Zap is betting differently than many fusion companies. Instead of pushing to the outer limits of magnets or lasers, it's focused on a "sheared flow Z-pinch" approach.

They're betting that by building something less capital intensive and easier to engineer, they can iterate faster and get to market sooner. Ryan Umstattd, Zap VP of Product and Partnerships, explains.

Ryan Umstattd

Zap Energy: no magnets required. Traditional approaches to fusion require really big magnets or lasers, and Zap needs neither. Building something with less upfront capital cost is important to what we're doing at Zap. But equally important, if not more so, is iteration speed. Time is money. They're often interchangeable. If you can build something cheaper, you can also build it faster.

Fusion is hard. Decades of research have shown that. We're going to have to learn a lot, and we want to learn it as fast as possible. If we have an approach where we can design, build, and commission a device within a year, we have an opportunity to make very rapid progress. And I think that's what we're going to need to see to commercialize fusion energy.

Julia

Time is money. Zap is leaning into the fourth "why now?" - startup speed and iteration - and the company thinks that the fast lane is right down the middle of the other approaches. Remember in the last episode when we talked about the triple product: density, temperature, and confinement time?

Inertial confinement optimizes for density at the expense of time, and magnetic confinement optimizes for time at the expense of density. Derek Sutherland, a plasma physicist and fusioneer at Zap, explains how Zap plans to increase its triple product by splitting the difference.

Derek Sutherland

Zap sits between those extremes. We're a pulsed approach to fusion, but we're not getting to quite as high densities as inertial confinement, nor as long confinement times as magnetic confinement. We're splitting the difference on the triple product. The benefit is that you don't have to go extreme in any technology direction.

You don't need super intense high-tech repco magnets or really awesome lasers that tend to be expensive and need constant improvement. We have a very simple approach that's between those two, using largely off-the-shelf technology in a specific application that gives rise to a commercially attractive approach to fusion.

Packy

That doesn't mean that what Zap Energy is doing is easy. They're bringing a fresh approach to one of the oldest ideas in fusion, the Z-pinch, that generated false positives in the UK way back in 1952.

Derek Sutherland

The Z-pinch is the OG fusion concept. The principle is simple. You're mainly flowing a current. In a cylindrical coordinate system, if you flow an electric current in the plasma in the Z direction, it produces an azimuthal or circular magnetic field around that cylinder that compresses it to very high densities and temperatures. That's where the Z-pinch gets its name.

The problem is that without intervention, the Z-pinch plasma is unstable. Instabilities would crop up that would make the plasma terminate before you make enough energy to pay for everything you put in. In other words, it's hard to hit net gain without intervention.

ZAP's value add is using a new way of stabilizing the Z-pinch called shear flow stabilization. Think of it like a busy highway. You're stuck in the exit lane but wanted to go through, and there are fast cars going past you. You can't merge because of the shear flow between your exit lane and the highway. Similarly, we see experimentally and from theory that when we have enough shear, it stabilizes the Z-pinch for very long durations.

This means you can hold that plasma around long enough to make enough energy to pay for the energy that went into it. It reopens the Z-pinch as a path to net gain. That's what makes ZAP unique: the shear flow stabilization.

Packy

And how does zap measure progress on the path to net gain?

Derek Sutherland

Technically, the main thing you're changing in the Z-pinch as you scale up performance is the amount of electric current flowing in the plasma, and how hard you're pinching the plasma with that current. If you raise the current, you produce a larger magnetic field, compressing it to a higher density. It's like a piston being compressed harder and harder.

We can measure the temperatures, densities, and confinement times. That tells you what the triple product is. From that, you can derive the Q. The Q is the scientific gain, power out versus power in. That's how we internally measure our progress.

Packy

My role on this show is one I was born to play: asking the dumb questions. When Derek told us that progress was a function of current, I asked him why they couldn't just turn the current all the way up and achieve Q>1 today.

Derek Sutherland

Yeah, it's a physics and engineering application thing. Our scaling laws clearly predict the current required to hit scientific break-even. That's our guide, but realizing it in practice is more involved. We're mainly working on raising the current in our pinch. We need the correct pulse power system to do that. We're learning how to do it in the most efficient manner possible.

You take energy from a big capacitor bank and couple it efficiently to the plasma to make fusion. Learning how to optimize that efficiency and actually realize the current where you want it to flow isn't trivial. But we're making a lot of progress, and we see a path forward.

Julia

Coming into these conversations, I didn't realize fusion was far enough along that many companies now view it as more of an engineering challenge than a physics challenge. Engineering is hard obviously. Going from models and simulations to a working generator that produces more energy than it consumes is hard.

But talking to Ryan and Derek at Zap, you really feel they're on a path to pulling this off. Iteration speed builds momentum, and it's cool to see companies optimizing for speed.

Packy

We have to point out here that one of the big reasons these companies are able to iterate so quickly and so often is that they're not burdened by the same regulatory regime as fission.

Derek explains.

Derek Sutherland

From a regulatory standpoint, it's much faster to iterate with fusion, primarily because you don't have the main concerns that come with fission. We're not using any special nuclear materials. There's no uranium, thorium, or plutonium.

Criticality doesn't apply to fusion. It's not a nuclear chain reaction. So you don't have as much concern when trying to do experiments and prototypes, because it's just very safe.

Julia

Must be nice. David Kirtley at Helion, which also prioritizes fast iteration speeds, explained that regulation was the biggest risk to the business a few years ago, especially since their model is predicated on speed. But the regulatory situation for fusion has landed in a good place.

David Kirtley

Yeah, I would say regulation around fusion three or four years ago was the biggest risk for this technology. It wasn't clear who would regulate it or how, or if there was a default answer. So it was possible we'd build working fusion generators and then couldn't deploy them because there wasn't regulation. Not to say that the regulation would be too hard or too easy, it just didn't exist. That was a big risk for the company, for all the companies, but also for the technology in general.

We started working with the NRC, the Nuclear Regulatory Commission, a number of years ago. I've given talks at public meetings. We've been working with the technical staff and commissioners over the last three years or more to figure out where fusion fits in the regulatory landscape. Our goal was that it fits somewhere. It mattered more that we fit in the existing regulations than exactly where we fit in those regulations.

The NRC commissioners voted unanimously earlier this year that fusion be regulated under Part 30, which covers particle accelerators and hospital parts of the regulatory code for nuclear. That means we're regulated by the state. In Washington, the Department of Health regulates us rather than a federal body. That's really good for Helion because we've been working with them since 2018.

Our previous systems have all been regulated, licensed, and inspected because we want to make sure that fusion gets to the world quickly, but safely. That's an a priori requirement. Now we have the job of teaching not just Washington but all the other states about fusion. We're working with state regulators on how to actually regulate fusion, what's easy and hard about it, and how to do it safely.

Packy

Move fast and safely. We talked a lot about how regulation slows down construction projects when we discussed nuclear fission, and that the regulatory burden is a big reason nuclear plants end up being so expensive. But we also discussed the negative impact regulation has on iteration speed, and ultimately safety. By making it harder to test and iterate, regulators impede the development of safer nuclear reactors. That needs to change.

Fortunately, it seems fusion won't hit the same roadblocks, and Helion is taking advantage of the opportunity to test and iterate quickly. David told us that the secret to the company's speed is that it engineers systems that are easy to make in order to get on the grid as soon as possible.

David Kirtley

A lot of us at Helion came out of scientific and academic programs where we focused on discovering physics and doing new diagnostics, but not delivering a product. When we spun off Helion, our goal was to make electricity on the grid as soon as possible, even if sometimes it's not as fun or elegant.

What shortcuts can we make to move faster? That's been the mantra of Helion: how do we iterate really quickly? We're now building our 7th prototype. How do we actually get electrons demonstrated and on the grid as fast as possible and engineer systems that are easy to make? That has been the mantra as we've built all of these prototypes over the years.

Packy

Helion's approach is not without its detractors. Its 2028 target date with Microsoft is wildly aggressive. Some in the industry see Helion as a manifestation of Silicon Valley bravado, believing that in fusion you can't just move fast and iterate your way to success.

So we asked him what he thinks the company's doubters miss.

David Kirtley

A lot of it comes back to looking at how modern hardware technology companies operate. It's less about the physics and more about how you're building a company. Think about SpaceX and Tesla and many others. Modern aerospace is a good example of building and testing as fast as possible and iterating. In January, we were running our 6th generation system while physically building the 7th generation system and engineering the 8th generation system, all at the same time. That's how you speed up the process.

Our first peer-reviewed published thermonuclear fusion happened in 2011 on a small-scale system funded by the Department of Energy. Since then, we've built four more systems, iterating on that, increasing the yield, neutron output, and fusion reaction rate. About a year ago, we published that we were the first to do deuterium and helium-3 fusion in bulk. We were also the first private company to hit 100 million degrees, the key operating temperature for fusion. We've set those milestones and metrics all along.

But a lot of that comes back to the philosophy of how to build fast. If a diagnostic is going to take four years to build, it's too long. We're not going to build that diagnostic, even if it's the best one. Is there a cheaper diagnostic that's faster that I could build in six months? That's the one we're going to pick. We keep that at every stage, even though sometimes it's hard and a bit chaotic to have all of those parallel things happening at once.

Packy

Julia, you've got to be proud as a SpaceX alumn. The company came up a lot when we talked to fission companies. Here it is again as part of a 'why now' and 'how so fast' for fusion.

Julia

I think the reason it keeps coming up again and again is that while fission reactors and fusion generators are among the hardest machines in the world to build, so are reusable rockets. What SpaceX showed is that rapid experimentation and iteration times don't just apply to software or simple hardware.

A machine is a machine is a machine, granted, perhaps with more complexity and more pieces to it, but the same tight feedback loops should benefit all of them.

Packy

It makes sense when you take a step back: go fast to go fast, make mistakes, learn, iterate. That doesn't mean that what Helion is doing is easy, though.

It's doing a few things that are really hard, all rolled into one.

Julia

First, it's using magneto-inertial confinement fusion, a hybrid of magnetic confinement and inertial confinement.

David Kirtley

The goal is to take the best of magnetic confinement, which keeps that 100 million degree fuel from touching the wall, because you don't want that hot fuel to ever touch the wall. And the best of inertial confinement, which is don't hold on to it forever. Nature is unstable and doesn't like that. Instead, squeeze it and get to fusion as fast as possible.

And then add that third one that most people aren't doing, which is directly extract that electricity. The big trade-off is that all this has to happen fast. So it's all pulsed. That's the inertial part. The beauty is you get to do it fast. The trade-off is you have to, which means you need triggering systems that respond in nanoseconds, in billionths of a second.

Technology didn't exist ten or 20 years ago, and you need massive pulse power systems, big high voltage electronics. So we have to design and run these biggest power electronics in some way. Helion is more an electronics company than it is a fusion company. That's where a lot of our technology and our team focuses on, is those big power electronics.

Packy

Then there's the fuel Helion is using: helium-3, an element that I learned is common on the moon by watching "For All Mankind," but one that is not common at all here on Earth.

David Kirtley

I love "For All Mankind". If you need to start your business by going to the moon, you probably have a rough business model and road ahead, as maybe seen in the TV show. But for Helion, in fact, what we named the company after is the nucleus of a helium-3, which is called a helion. Helium-3 fuel is actually one of the older fuels.

The brilliant scientists who did a lot of the early work in fusion recognized helium-3 would be a great fuel because when you fuse deuterium and helium-3, it forms helium-4 (regular old balloon helium) and a proton - all charged particles, all electricity, all trapped in the magnetic field. There are two challenges. One, as you pointed out, there isn't a lot of helium-3 on Earth. So how do you get helium-3 to test it, and then how do you generate it in your system? The other is that it requires higher temperatures to operate. Both are negatives that you have to overcome.

We solved the first one with a patented high-efficiency energy recovery system. This lets you do fusion a lot cheaper, which is what got us really excited. You take two deuterium atoms (deuterons), fuse them together, and they make helium-3 in gas form. So if you have fusion already, you can make helium-3 to do helium-3 fusion. But the key to unlocking that is having really efficient fusion and efficiently putting electricity in and recovering it out. If you're not doing that process, it's hard to make helium-3, and then it's hard to burn the helium-3 to make electricity. It was a chicken-and-egg problem that required modern high-speed transistors and fiber optics to unlock.

Julia

And finally, it's doing fusion in a way that can directly harness the electricity. Remember in episode seven when Casey Handmer told us that nuclear reactors were stuck in the Stone Age?

We'll play it again because it's a great clip.

Casey Handmer

At the end of the day, if you're boiling water to make energy, you can make heat however you want. You can make it with nuclear power, coal, or gas. But you're still boiling water.

It's like Stone Age. "Ooh, make water hot, boil water. Turn turbine."

Julia

Turns out most fusion generators do something similar. David and Helion are trying to skip that step and just capture the electricity from the source, which is only now possible because of all sorts of technological advances.

David Kirtley

We've done fusion for a long time, but our goal is to do it differently. Our goal is to take lightweight isotopes of hydrogen and helium, fuse them together under intense pressure to form heavier atoms, and release a tremendous amount of energy. But we don't want to release heat energy. We want electricity. Our goal is to directly harness the electricity from that fusion reaction as electrons and get it out on the grid as soon as possible. Our systems are pulsed and electromagnetic, but the focus is always on how we can get electricity out of fusion as fast and efficiently as possible.

It takes a certain level of technology before this can actually happen. The first cars in the 1800s were electric, but the batteries, motors, and transistors didn't exist to make that niche product into a widespread one. Then gasoline engines took over, and we had 100 years of gasoline engines. We're only now at the place where we have efficient power electronics. With regenerative braking and lithium batteries, electric cars finally make sense.

If I was doing fusion in the 1950s, I'd be doing thermal fusion too, using the energy conversion methods available then. It's taken modern high voltage power electronics, fiber optics, gigahertz speed computing before we can really do fusion in a way that harnesses electricity directly.

Packy

I know I'm getting a little cheesy on this episode, and I've written about this idea before, but there's something magical to me about how these disparate branches on the tech tree - high voltage power electronics, fiber optics, gigahertz speed computing, machine learning, magnets - all developed for completely non-fusion related reasons, turn out to be critical to making fusion happen, potentially just in time to be a serious weapon in the fight against climate change.

There's this physical phenomenon, two light atoms fusing together to produce a heavier atom and a release of energy, that occurs naturally in the sun. Physicists figured this out about 80 years ago, and researchers have been working on it since. Now startups may finally be able to do it in an energy-profitable way because of all these seemingly random developments. I'm not a religious man, but at the very least, capitalism works in mysterious ways.

Julia

Amen. As we've discussed throughout this season, for an energy source to work in the capitalist system, it must compete by doing something no other energy source can, or by being cheaper than the alternatives. Fusion generators will live inside power plants. Companies need to convert the energy from the reaction into electricity, which they can sell directly to customers or into the grid.

While it's early and none of these companies have yet achieved Q>1, these aren't just research projects. These startups have had to design their companies with unit economics in mind. So we asked them to describe what the unit economics of a fusion plant might look like. JC at Fuse told us that he looks at three different variables.

JC Btaiche

I'll preface this by saying different fusion concepts may have slightly different ways to think about it, but in my mind, there are three main variables. The first is how much it costs to create the fusion conditions, which is like the dollars per joule delivered on the target. For example, NIF has a 400 megajoule laser, depositing roughly two megajoules to the target. The laser costs a few billion dollars, so their cost per joule delivered on target is roughly $2000. We think Fuse can eventually get to $40 per joule and even lower.

The second variable is how much energy you can produce from what you've delivered to the plasma. This is more of an engineering function. The more efficient the reaction becomes, the more energy you get, which will decrease your cost because you have the same input cost but get more power. NIF, for instance, hasn't upgraded their laser. But over ten years, they've increased the target efficiency by 1000x just by improving the target design, engineering, and physics.

The third variable is how much it costs to convert the fusion output into electricity. This is more predictable because it's essentially looking at fusion as a heat source and converting that into electricity, unless you're doing direct electricity conversion, which has less precedent. So we think about these three functions: dollar per joule delivered to the target, fusion gain, and cost to convert fusion gain into electricity. Fuse is focusing intensely on decreasing the dollar per joule delivered to the target and improving efficiency.

Packy

JC lists three factors: One, how expensive is it to deliver energy to the target? Two, what's the fusion gain or Q once you do? And three, how efficiently can you convert that fusion gain into electricity?

It's kind of a bottom-up approach. Ryan at Zap Energy explained how they think about unit economics from the top down, using overnight capital costs. Overnight capital costs are the hypothetical costs if projects were completed overnight and are often used in large infrastructure projects to normalize for the impact of time, including removing things like inflation and interest payments.

Ryan Umstattd

Yeah, and I'll caveat these with no one's yet built a commercial fusion plant. We're all doing our best to estimate the costs, using class four or five estimates. That's terminology from the Association for Cost Estimating for different ways to price out what you think your nth of a kind commercial unit might be versus your pilot plant. As we've gone through that with Zap, we started to see overnight capital costs in the range of $3,000 to $4,000 per kilowatt. That's about half the overnight capital cost of advanced nuclear today or solar thermal. But it's still significantly more than a natural gas turbine, which might be closer to $1,000 to $1,500 per kilowatt of overnight capital cost.

This leads to a power plant that you've spent money upfront to build, but now your operations and maintenance are quite affordable. Our fuel costs are practically negligible. For a fusion power plant, fusion is such an energy-dense fuel that you measure the annual input in kilograms, not train cars. I can ignore the fuel costs when it comes to operations and maintenance.

This leads to levelized costs of electricity in the range of $30 to $60 per megawatt hour in our estimates today, based on different input assumptions. So that makes us competitive. It's not the cheapest electricity source today, but I think the market's going to be in drastic need of an on-demand carbon-free source. Renewables just don't get us there by themselves. I'm comfortable that this is a competitive LCOE for what the market will need in that timeframe.

Julia

Francesco at Proxima points out that while being cost competitive is important, and the company's model suggests that it can be, getting it to be "cheap enough" works out in the short term because fusion energy is so compelling and versatile.

Francesco Sciortino

The cost of energy from fusion is, of course, uncertain. Our system analysis suggests it could be cheap. That's not to say it has to be cheap. Fusion is so compelling that if you get it done as a safe, abundant, clean source of electricity, process heat, possibly making hydrogen, then it doesn't really have to be the cheapest thing. We don't have to make it cheaper than photovoltaics today. We're not competing for that market.

So if you get to a power plant that makes an order of a gigawatt for $3 billion, you're in business. What you really want to look at is not the overnight cost of a new power plant, but rather the levelized cost of electricity (LCOE). The models tell us that we can achieve five cents per kilowatt hour electric, which will be extremely compelling if we manage to be anywhere in that order of magnitude. That's great.

Julia

Ian Hogarth of Plural, Proxima's co-lead investor, added that thinking of fusion in a vacuum undersells its strategic importance and the role that governments will play in supporting its initial growth.

Ian Hogarth

It's such a deeply strategic technology. There's going to be a race for fusion power globally in the 2030s, where countries will try to connect fusion to the grid faster than others, because it's going to be a massive new industry. It's going to underpin a lot of progress and opportunities. I think you're going to have very significant subsidies emerge in the first chapter of getting fusion on the grid. If you're thinking like a state, you're considering how much you'd pay to get the world's fusion industry based in your country.

You'd probably pay quite a bit, because it's going to be pretty strategic in lots of ways. So I think there's going to be some very significant state involvement in fusion in the early days, as there has been up until now, with the likes of 7X. There may also be a role that some of the largest technology companies in the world will play, like the offtake agreement that Microsoft has agreed with Helion.

This question of AI requiring more energy and fusion as a kind of baseload source of energy that doesn't have some of the downsides of fission is going to move up the agenda of large corporations as well.

Julia

That point Ian just made is a good one. We're going to need a lot more energy, thanks in large part to companies like Microsoft that are building power-hungry data centers to support the growing demand for AI.

Clea at Lower Carbon said that estimates of 5x electricity demand in the US by 2050 are probably conservative.

Clea Kolster

Ultimately, the demand for electricity in the US alone is going to grow at least 5x by 2050. And I'm pretty sure that's insanely conservative because when you look at computing demand alone, you need about ten years for that to be effectively the same demand of the entire US from an electricity standpoint. So we're set up for a lot of electricity demand.

Enter fusion. We don't really know exactly where that's going to be in 2050. But right now we're really excited about a lot of technologies that are going down the learning curve and could represent a significant piece of that energy pie in 2050. A lot of those are projected to be sub $70 per megawatt hour. And that's really valuable in terms of a firm source of electricity. Some might project that somewhere between 10 to 30% of the overall energy source pie.

Julia

Make no mistake, though, ultimately the goal of fusion is to come down the learning curve to the point where fusion can replace fossil fuels globally. Whether that happens by 2050 or sometime later, that's going to mean rolling out fusion plants across the globe.

We asked David at Helion what he thinks the rollout will look like once they demonstrate Q>1, and he was clear about the goal.

David Kirtley

In the world, there's over 4000 gigawatts of installed fossil fuel capacity, and our goal is to replace it all. I don't think one type of fusion is enough to do all that. We're going to try and move as fast as we can. But I think you're going to have different kinds of power in different locations, whether it's remote, military bases, giant factories, or data centers, and they're going to require different kinds of power.

Our plan is to demonstrate electrons on the grid in 2028, then scale and manufacture these systems as fast as possible to deploy them. Our goal is to get to making generators per day by 2030, rather than generators every few years. That's a big scale, a big lift.

As a human, I want other fusion and other types of advanced, carbon-free power out there in the world too, because we have that big of a need and we need to move that fast. That's my view. We're going to move as fast as we can, and we engineer the mass manufacturing into the systems right now, here in Everett, Washington.

Packy

And because Helion plans to directly convert the fusion reaction into energy, cutting out that third leg of the cost structure that JC mentioned, it believes it can make fusion really cheap - so cheap that it can compete directly with all other energy sources and win. We asked him what the world looks like when that happens, when we have abundant fusion energy. This is what he told us:

David Kirtley

The whole team thinks about this a lot. We believe we have an approach to fusion that can be low cost and generate electricity at a cent per kilowatt hour eventually. That's radically low cost. It means we can replace fossil fuels and stop climate change, but it also opens up new possibilities. Many parts of the world don't have the amount of low-cost electricity we do. We want to address markets in India, Africa, and Asia to improve their standard of living.

Our first customer is data centers. We're seeing AI growing at an enormous rate, and it's going to need power. Our data center and computer infrastructure will require massive amounts of power, and we want to support that. We want to support a world with massive computing available for everyone in their pocket and at home.

Julia

I think that's why we're all in this and why using atoms to generate clean energy is worth doing, even though it's so hard. We can stop climate change, open up new use cases, and bring energy to parts of the world that don't have access to it.

I think it's one of the most important projects humanity can undertake.

Packy

It's not going to be easy. Every person in fusion we spoke to pointed to challenges and risks their companies face, and they're non-trivial. Achieving Q>1 and eventually Q Infinity is one of the biggest challenges humanity has ever undertaken. We don't mean to gloss over those challenges, and we'll release full episodes with some of the founders so you can hear more details from them.

But I think the world's assumption is that fusion energy is practically impossible and impossibly far away. We wanted to do this episode to show you that it's simply not true anymore. I kind of got fusion-pilled when I wrote "The Fusion Race" in May. But Julia, I'd love to hear how your thoughts on fusion have changed after having these conversations.

Julia

I thought fusion was essentially a science project before we started working on this together. Fission was the area I knew and understood - we do these today, but have a go-to-market problem expanding the footprint. It's cool to see incremental progress being made with fusion. We've achieved Q>1, albeit for short periods, and we're making progress with different approaches.

I like the analogy of fission reactors to fusion reactors. It helps you understand that we're splitting atoms in one, merging atoms in the other, with different ways to set up your reactor. Each approach has its pros, cons, risks, and benefits. As you put it, Packy, we're in this race with different approaches and companies. Now it's about who can iterate, learn, and improve fastest to achieve consistent Q>1, and ultimately something commercially viable.

Packy

Well said. That was one of the things that blew my mind and made me want to write about fusion and include it in this season. With so many smart people and companies approaching this, each with a real shot, somebody's going to figure it out. They all see this "why now?" moment for fusion.

After writing that piece, I was surprised by how much of an engineering challenge it's become, as opposed to discovering new physics. We've known about this for 80 years, and now we have the software to design it, make the right engineering trade-offs, or we've seen SpaceX and know it's possible.

When I wrote about Helion's work with Microsoft, many thought they couldn't pull it off. After talking to David, I'd bet my own money on Helion achieving commercial fusion before 2030. It sounds aggressive, but they've been doing this since 2013 and are on their 8th generation of generator.

Julia

I love how fast they're moving. They're always working on three generators at once: one being tested, one coming out of design into final specs, and one they're designing for the future. They're all impacting each other. What they learn from version six influences tweaks on version seven before testing, which informs how they design the 8th prototype.

It's got me thinking: fission and fusion are different and at different lifecycle stages, but they're not often grouped together. Why isn't there just an "atomic energy" category? We discovered renewables, which are awesome, but here's the future - the potential of the atom, whether splitting or fusing it. Maybe there's a branding opportunity here?

Packy

Yeah, we talked to people about this a little bit. But to me, it seems like the fusion people are saying, "Give us a clean sheet of paper. We're being regulated by the health people, and it's way easier over here. Don't lump us in with fission. We actually love fission and think it's great and we want more of it.

But for our regulatory path, we really want to have our own kind.

Julia

I don't want to be lumped in with the baggage of fission, it's true. I get it. I'm actually happy that fusion seems to have found a less burdensome path than fission, which is really bogged down.

Packy

One of the things that's come up as we've had these conversations is fission's momentum right now. I feel like maybe one of the biggest risks to getting change to actually happen, because it does take so long, is that if Helion hits the grid in 2030, this momentum we have with fission might fade.

People might say, "Whatever, we'll just go with fusion." Does that worry you at all? I can see a lot of this momentum that fission has potentially being redirected.

Julia

If we truly can get there that quickly and then scale up, great. The question is, how quickly can you roll this out commercially? It's all atomic. It has these wonderful benefits: extremely energy dense, extremely low carbon footprint, and reliable. Both fission and fusion have those appealing qualities of energy that we don't have with the intermittency of renewables. Either one, honestly, whatever's going to work best, let's do it. I still believe in having a diversity of energy sources in case your supply chain sucks. You can lean on another, and you get strength through diversity. So it's probably good to have both.

We know how good these fission reactors are. They've been around 40-60 years now. They're now starting to license them for 80 years. No question they're going to stick around, especially now that we've seen public attitudes changing. I don't think fission is really going anywhere, but it is a good question of how much we're really going to see it expand and what that race looks like in terms of timing between fission and fusion.

Packy

I am fully nuke-pilled after this season. Coming into the season, I was already pretty nuke-pilled and now fully nuke-pilled.

In reality, we should be building fission reactors for the next 50-100 years, whatever that number ends up being. The thing that worries me is maybe we get caught in this in-between spot where fusion hasn't proven it can scale up, but it seems so close that people lose support for fission.

Maybe we don't change the regulations on the fission side, but fusion hasn't figured out how to scale up yet. That puts us in a weird spot. That would be the biggest risk that I can see.

Julia

I think it's a big risk, and I think it's actually important that you still have the two cohorts, the fission and fusion camp, really championing what needs to happen for their specific field to push the ball forward. We just need to be pushing for it on all fronts.

Packy

We'll talk about what the future looks like in the next episode, but I came out of this really thinking that there is a clear path on both sides of fission and fusion together putting a shitload of clean, essentially limitless energy on the map. If we can do this, if we can not get in our own way, the world is going to be a much different and better place if we pull it off.

Julia

Yeah, I think we know what we need to do in both categories. Fortunately, there's just a bunch of awesome startups and others who are working on pushing the ball down the field. Super exciting times.

Packy

Amen. On the next episode, sadly the last episode of season one of Age of Miracles, we'll look to the future and discuss what the world might look like if everything goes just right.

If we have cheap, abundant energy for all.

Julia

I'm so excited for our final episode. It's gonna be sad to wrap up the season with you, Packy, but I think this last one's gonna be a great one.

It's been so fun to have taken the time to dive into fusion today.

Packy

See you next week.

Julia

See you next week.

Packy

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. See you next week.