

Twitter Spaces Transcript - January 15, 2026 (Edited for clarification and readability)

Host: Greg Reyes (@joceyreyes209)

Guest: Josh Claman, CEO of Accelsius (@AccelsiusATX)

Greg Reyes: So, I think what I'd like to do, Josh, if it's okay, is just give folks a little bit of your background and kick things off. Does that work for you?

Josh Claman: Sure, that's fine. I can dive right into that if it works.

Greg Reyes: Yeah, okay, perfect. So, let me start out by introducing myself. I'm Greg Reyes. I often opine and share my thoughts about various companies that I'm an investor in, and the lexicon that I use, and it's actually worked quite well for me over the years, is the exact same lexicon that I used when I identified a company called Brocade Communications, that I became the CEO of, and turned into an industry powerhouse with over a \$24 billion market cap and 90% market share. And that lexicon basically is large TAM, unique novel technology, world class management team, a go to market strategy that will allow the company to scale, unit economics that will allow the company to become highly profitable. So, that's what I look at when I'm trying to assess and evaluate companies, and I became aware of one that happens to fit right into my wheelhouse, which happens to be data centers and compute. And I was introduced to Accelsius as a company that's part of a portfolio of a publicly traded venture capital company called Innventure. And I was introduced to Josh recently, and it was suggested that I invite Josh to participate in a spaces, to talk directly about Accelsius, and I'm super excited he's here.

And just by way of background, Josh is a very seasoned industry executive. As a matter of fact, when we were prepping for this call, it turns out that there are a number of folks that we both know from Josh's time at Dell. He has over 30 years of data center experience. He's a seasoned executive with global experience. He ran Dell's UK business. And since joining Accelsius in 2022, I think he's done a superb job at positioning Accelsius to emerge, as potentially, the market disruptor in one of the most important building blocks in the AI data center ecosystem, which is cooling.

So, with that, I'd like to turn it over to you, Josh, and just start the conversation by maybe having you talk a little bit about how do you think about data center cooling as it's evolved over the last 25 years? Me being a dinosaur, back in the day, data center cooling was just cold air blown over servers. So, things have changed quite a bit, and I

thought maybe you could walk people through the evolution of the technology, and where it's going in the future, and how Accelsius fits into that lexicon.

Josh Claman: Sure, very happy to. There are a couple sort of themes I'll weave into this as well in terms of the data center sector culture as well, which has implications on new tech adoption, where the sector is today, et cetera.

But essentially, for 30, 40 years, the data center sector has been air cooled. There have been very small incremental changes to what's called rack density, so that the kilowatts per rack, until very recently, I think the average kilowatt per rack was about 7.5 and it would go up based on the next generation of server or switch equipment by a kilowatt or so. And I think that sort of sets the stage for some of the challenges the data center sector has now.

It's a relatively conservative sector. It was air cooled, Greg, you said way back when, but actually, when we founded Accelsius 3.5 years ago, the data centers were air cooled. I remember having just three, even two years ago, pretty emphatic debates about air is fine. We've been talking about liquid cooling since the 1980s, we didn't need it, et cetera. Suddenly, everyone does need it.

And, there have been advances over the years, over the decades, in terms of slightly more efficient air handlers, slightly more efficient or sophisticated containment strategies on cold air and hot, hot aisles, et cetera, which have led to some incremental efficiency gains. But the sector, overall, this may be one of several provocative statements I make today, but was pretty inefficient. It got used to spending about 40% of the energy coming into the data center, on cooling, primarily cooling, and other overhead.

And so, it could be argued that the data center sector should have adopted liquid cooling 30 years ago and, and, and leave it behind, which it did with some IBM mainframes and Cray computers, et cetera, which were liquid cooled in the 1980s and 1990s. Air is a terrible transport of heat. It's probably one of the worst and this tolerance of inefficiency in the data center really stuck around until the advent of AI.

And this is when you see the processors just increased dramatically in wattage. So, the watts per socket go from 200— I mean, they're about 125 watts per socket to 150 to 175 and get very incremental increases over years. And suddenly, they're going from 300 to 600 to 1,200 to 2,200. All in a matter of 2 to 3 years.

And the other issue here is sort of the requirement in AI workloads to be proximate. You don't want a lot of latency, you want to pack as many servers and switches into a rack

as you can. So not only are the socket wattages, the wattage per chip going up, the wattage per server, and the wattage per rack is going up.

So, I think three, four years ago, the average density, again, was about 7.5 kilowatts. When you're talking about AI and many of you are probably following the industry or in the industry, but those densities are going to 100 kilowatts, 200, 300, and now some of the reference designs that are coming out in the next year or 2 are one megawatt. Not really racks, but pods of compute. So, the challenges are just enormous to power and cool these things.

Let me take a pause there, Greg, and you can redirect if you like.

Greg Reyes: Sure. So, if I take what you said, Josh, in addition to the topic of energy and powering these next generation AI data centers, the demand for cooling, it would seem, is one of the most critical building blocks to making the vision of AI data centers viable and scalable, which kind of causes me to want to pivot to a topic that you and I discussed, when we were prepping for this call, is that there was a lot of gnashing of teeth, and confusion, and loss of market cap, of companies in the cooling industry, when, at CES, Jensen Huang said during his keynote, that Rubin would be able to be cooled with 45 degrees Celsius water, and at the end of the day, didn't require chillers, and we saw JCI, Vertiv, and other companies see their stocks get hit as a result of that.

And there's a lot of confusion around it. I'm sure it's a question that you've been getting quite a bit since that comment was made. I know, Accelsius put a white paper out on the topic, but it seems to me there's some confusion, and for investors, to the extent that they understand the ground truth of the cooling issue, there's an opportunity for investors to profit from it. Can you help people understand that comment and maybe what Jensen implied and what the reality is?

Josh Claman: Yeah, maybe a little context first. So, a little, by the way, 80% of data centers in North America are still air cooled, by the way. So, the transition is happening. It seems like it's happening at sort of speed of a rocket. But there's still a lot of data centers around that are air cooled, and we can come back to that topic in a few minutes if you like.

But basically, if you have a liquid cooled data center, whatever the cooling architecture actually is in the white space or within the racks, the IT racks, you have, essentially, what's called a facility water loop or primary water loop, and that primary water loop is cooled using one of several mechanisms, but typically, it's cooled by using chillers. And chillers, within a chiller, there's a compressor, just like in your home air conditioning. And those are really power hungry devices.

So, the less you can rely on chillers or rely on chillers with their compressors turned on, rather than bypassed, the more you can save money. So, the math works something like this, that for every one degree C, you do not have to cool your primary water loop to reject that heat that you've captured through the IT equipment, you save about 4% energy per year.

So, when Accelsius was founded, one of our basic differentiators was that with our two phase cooling technology, we could cool these power dense IT racks with warmer, relatively warmer facility water, somewhere between 8 and 10 degrees C warmer than you would need for single phase water. And I'm throwing out some terms here that we probably want to revisit, and I'm aware of that, but just to talk about the water.

So, what we've gone to the market with that said, if you can not cool that water loop, if you don't have to turn compressors on, and if you initially design a data center around two phase cooling, you don't even have to spend the CAPEX on those chillers. Then you benefit from that. You get more free cooling because that water is warmer.

When Jensen came out in his keynote a few days ago, basically he said, we've now designed, we have a reference architecture with a power dense rack, based on Rubin chips that can utilize 45 degrees C water over the chip. Now frankly, from an environmental standpoint, that's fantastic.

We do have, we are a little bit cynical about how you would do that. That means you would have to rely on a delta T, a temperature change, from the inlet to the rack to the outlet of the rack, of very, very few degrees, and typically single phase water doesn't operate that efficiently.

So, what Jensen might have been intimating is that he's used one or more of a very few number of levers to increase to allow the increase of that water temperature. You can increase, for example, flow rate of the water. You can increase the surface area of the chips, because larger surface areas are easier to cool. You can increase the throttle point of those chips. Meaning you simply run those chips hotter. We believe that's really the main one he's hitting on here.

Now, we applaud that. It makes data centers more sustainable. It's really what Accelsius has been focused on from the beginning. But even if you're bringing in warmer water, we can even bring in warmer water, right, than single phase water, because of the efficiency of the two phase loop.

Greg Reyes: So, it sounds as though his comment was more outward facing and directed at trying to assuage concerns of municipalities about water consumption associated with AI data centers, and power versus a specific declarative comment as it

relates to specific technologies required in order to effectively cool hotter-running AI data center infrastructure. Is that a fair summary?

Josh Claman: I think it's fair. We've obviously been involved in some speculation about what exactly he was getting at. I do think, though, this comes at a time when there's a lot of mostly local regional protests around data center builds. They use water, they use a lot of power.

And so, I think the entire data center sector, from Jensen on down, needs to really focus on making, we call it making every watt count. So, you can imagine if you go and build a, in a community, 100 megawatt or one gigawatt data center, and you're utilizing 20, 30% of that power on cooling, you're gonna want to take a material chunk out of that with efficient design.

And I think that's what, essentially, Jensen was saying, we've designed our rack to run at 45C, kind of over to you guys now, over to the data center architects and operators rather than have NVIDIA, kind of the point of a lot of these protests, but I think his narrative is heading the right direction. And I think, and frankly, the data center sector, overall, really needs to take a lead on sustainability and efficiency, and including power and water use, and minimizing the impact on local communities, or else it's going to face more protests and more extreme protests.

Greg Reyes: Okay, so maybe Josh now is a good time to wade into the next level of discussion, which is, if I take everything that you just said and synthesize it, we've got a mega trend in AI data centers. The president has even come out and commented about the fact that those building these data centers are gonna have to participate in solving the power issue. We've got this startup that we can only own through Innventure because they own approximately probably a little bit less than 50% of Accelsius.

And here you are as one of the only companies offering a two phase solution that doesn't require chillers, doesn't require water, uses a non-conductive refrigerant, and you've got one other competitor, Zutacore, you've announced a relationship with DarkNX, which is— it was hard to find out details about them.

How do you fit into this? And if you take the Dell'Oro market research that says by 2029, there's going to be \$7 billion in cooling, liquid cooling spent, and that may grow to as much as \$25 billion over the next 5 years. What can you tell us? Like, are you on the launching pad? Are we close to an inflection point? How do you see this happening over the next few years? And, yeah, help me understand why I shouldn't be as excited as I am, but, how should we think about the evolution of the market and the company and the market opportunity as it relates to Accelsius?

Josh Claman: Well, I think you should be very excited. Obviously, we as a team are, and most of our investors are. I think we're really on the ground floor of something. Now, it doesn't come without some obstacles in our way.

So, when we founded Accelsius, there were about three basic architectures for cooling. All of them nascent in nature. One was immersion, putting the entire server in a pool of nonconductive fluid into kind of a bathtub. We saw some, we thought that's interesting, but we saw some real issues with cost, and sloppiness, and service protocols, and warranties, et cetera. But they had a lot of hype, three or four years ago. The reason I mention this, I often get questioned about it.

The other, the first mover in this market was water. We call it single phase water to distinguish what we do with that. But essentially, it's water mixed with a biocide called glycol. And it's pumped over through plates that are mounted on top of the chips to cool the chips and carry the heat off.

The issue with water is kind of multifaceted. One, you're pumping conductive fluid at incredibly high flow rates through very expensive server equipment. And if there's a leak, you're now destroying maybe 3 to \$5 million of equipment within a rack. And by the way, there are leaks all over. So, if anyone says there are no leaks on water, just ignore them. They lack credibility. There are leaks everywhere and it's very damaging.

Thermodynamically, there are more intriguing issues, though. One is that water and flow, I mean, heat, heat transport and flow rate are exactly related. So, it's a linear relationship. So, the rule of thumb is that for every kilowatt you want to cool, you need to pump about 1.5 liters per minute of fluid as a flow rate.

So, imagine now you have 100 kilowatt racks, or 500 kilowatt racks, or one meg racks. So, one meg rack, that's 1,500 liters per minute of flow rate, going through very small orifices. The reason you need to do that is that you're cooling through just convection, so the water is going over the heat source, it's carrying the heat off. There's no magic there. There's no exciting physics that happened there.

There's a limit. So ASHRAE and other associations like that believe the limit on effectively cooling with single-phase water may be about 1,500 or 1,600 watts per chip. The Rubin chips and chips like that go beyond that already. So, and that leads to another, probably, discussion that you might want to have.

So, there's a thermodynamic issue there, which we sort of translate as lack of thermal headroom. So, it will work for the H100, H200, Blackwell, probably the Rubin, because of a slight increase in surface area. It may not work for the chipset after that. And remember, in the old days, a new processor generation would come along every maybe

three or four years. Now they're coming along every 12 to 18 months. So, you want to preserve the investment in that infrastructure. It's a big issue for these data center operators and for their investors.

So, there are some other thermodynamic issues, which are a little bit more arcane and probably not appropriate for this audience, but they're important.

And so, so we, founded Accelsius to capture, kind of harness the magic of two phase physics, which basically means we pump a refrigerant into a plate that's mounted on the chip, just like water. But we orchestrate a phase change, we boil. We have a boiling regime. And it's that phase change to vapor that captures a huge amount of heat energy.

So that does a couple things. One, it solves the, it solves the headroom issue. So, it protects the investment in the cooling architecture through several generations of chips. We have yet to really find a limit on that, but it also does it with a nonconductive fluid.

So, we will leak too, by the way. There are a lot of hands on these things. They're third party service organizations, they're third party integrators. So, we cannot say we are gonna leak less than water. Although we will do everything we cannot to leak as do the water vendors, by the way, right? But when we leak, we leak a little vapor, it's nonconductive. You can actually drop a server into a pool of our fluid, and it won't damage any of the expensive GPUs, and, and so that's another benefit.

Our flow rates are actually much lower. Our flow rates actually increase with heat densities as well, but at a fraction of what water does. So, what that means is, we're not sort of self destructing, we're not damaging the microchannels and internal, sort of, micro finning structures within that loop, just with our flow rate. And so, it's just a much more practical, I think, fit for purpose architecture for cooling in the data center.

I'll say one more thing, Greg, that it sort of amuses us that there's pushback on this, and we can talk about some of the reasons for that, because I think they're important. But the pushback— because every other industrial sector that your audience would be aware of, moves from air to single phase to two phase. Your car has two phase in it, your home air conditioning is a two phase solution, nuclear power plants, what cools them is two phase. And every other sector, it moves from single phase to two phase. So, it should be no surprise to anyone that the concentration of heat in these processors are gonna hit a point where you have to go from water, pure convection cooling, to a two phase approach to cooling.

Greg Reyes: Okay, I really appreciate that explanation, and I'm not gonna ask you to say these words, but I'm going to say them, and you can tell me whether or not you think I'm in the ballpark.

So, we've got an industry that, for all practical purposes, will be moving to two phase cooling, because it is the, at this moment in time, only way to solve the data center cooling issue. It's a TAM that's going to go from a few hundred million to potentially 25 billion for two phase cooling. As best as I could tell, Accelsius is the only company that is good to go, deploying a solution that is available to the customer.

From a market share standpoint, my expectation is that your hope is to get more than your fair share. So, sounds like there's a big opportunity. You guys seem like you've got novel technology that solves the core problem. Is it reaching too far to say that in the round that you announced your B round on Monday, it kind of struck me that the two largest participants, Johnson Controls and Legrand are investors in Accelsius, and, obviously doing a little background work, that they seem to be powerhouses in the data center cooling market, Johnson Controls, big footprint in North America, Legrand, big footprint in Europe.

What can you, what's the inference that folks should take away from that? Because it's not your typical venture round, clearly. Am I getting that right?

Josh Claman: I think you are, and there are a number of dimensions there. One is, I think your first statement is correct. I think two phase cooling in the data center is going to be the dominant architecture over time.

Now, the narrative that we're fighting against is when— everyone, there are very few people in the sector that don't agree with that statement. The pushback we get is, when will it be needed? Which we find is sort of an interesting narrative to throw at us, because again, we're better now than water. For a lot of reasons, a few of which I've gone into. We're gonna be required later. So, in what other space do you say, I'm only gonna adopt that better technology when I absolutely need it? I don't think you see that. So, it's sort of a false narrative.

So, I agree with your optimism, and you might expect me to as CEO of Accelsius, in terms of the market adoption overall.

The B round, which was \$65 million on a healthy valuation. I think the reason for that is that people do understand this, that the really sophisticated thermal engineering teams in the data center sector understand that it's going to head in this direction.

Our relationship with JCI and Legrand is really interesting. We didn't really go after institutional money for this round. We went after potential commercial partnerships. Because the one really valid point I see to a slowness in adoption— I think the adoption curve is happening very quickly, but if you're living in it, it feels slow and frustrating— I think the reticence would come from, there's a lack of global supply, and alternative vendors for two phase cooling.

You mentioned Accelsius is an early stage company. We've just done a B round. It is if you're a hyperscaler, it would be hard to say, we're gonna move to two phase and we're gonna be completely reliant on this early stage company.

So, we struck these partnerships with Legrand and JCI very consciously to expand our reach, the credibility, potentially manufacturing partners, et cetera. But certainly more sources of two phase, more flavors of two phase technology than just Accelsius could produce. And I think this is gonna be a real accelerator to adoption out there, because as far as, in all of our conversations throughout the sector, from hyperscalers on down to colos to enterprise data centers, the valid pushback we're hearing is, we need a global supply chain for two phase technologies.

So, we are doing what a lot of early stage companies don't do, we're orchestrating that tipping point through these massive global partnerships. And we're really excited about them.

Greg Reyes: So, if— by the way, I think that that is a brilliant move, and it positions the company to scale. And in the past, I think that you've talked about and Bill Haskell from Innventure has talked about a number of proof of concepts. There has been an announcement with DarkNX.

When should people think about us starting to hit the knee of the curve in terms of scaling revenue? And, there's a comment that Bill Haskell had made back in the May Innventure earnings call, where he alluded to the fact that Accelsius could be profitable delivering 100 racks per month. From an investment standpoint, again, I shouldn't say everybody, but I suspect that 99.9% of the people on this call are here because they want to learn about Accelsius, and they're owning it through Innventure.

But when should folks be thinking about the revenue ramp, and is the comment about 100 racks per month being the key to getting to profitability? What are you comfortable sharing about that?

Josh Claman: Well, first, sort of a contextual point I'll make, is that I think what has, what makes it hard, and you went through this journey with Brocade, no doubt, but what makes it hard for the data center sector to adopt new technology is the, the, the

conservatism of the industry, right? And the cultural conservatism creates this sort of, no one got fired for buying IBM, right? I mean, you wanna, you wanna see other people using this technology.

So, we've done a lot of proof of concepts through the last year, et cetera. Now we're hitting production deals like the DarkNX deal, and you'll see more announcements coming forward. Once we get a few of those, the momentum's going to increase and we'll accelerate toward a tipping point. I think you'll see that towards the end of this year, maybe early next year, but it's not gonna be much farther than that because I think the other thing that we've seen is that, and sort of the misconception is that everyone's using water.

The fact is that I mentioned that about 80% of data centers in North America are still air cooled. So, so water deployments, actual production deployments, are still sort of in the, in the mid single digits, as a percent. So, what we saw was that when the first wave of water based deployments was sort of completed, we saw a real bump in our, the interest in our company, the leads we got from operators, et cetera, et cetera, other clients.

And I think the reason was this relatively poor experience with the water based architectures. And so, we think there's nothing ahead of us but acceleration toward the tipping point. I can't call the quarter, but I can say that we have a pipeline, an opportunity pipeline, of qualified opportunities that's about 8 times our revenue plan for 2026. And that will lead into higher growth in 2027. I think we are on track to hit cash flow positive by the end of this year. So, so I think a lot of what you said and the way you characterized it is accurate, but I wanted to explain some of the context of why we're so bullish about this. I mean, just the announcement with DarkNX has led to more opportunities, right? Because people are looking for a better architecture to cool these high density racks. And they're just waiting for something to gain enough ground to have credibility and support.

Greg Reyes: But it's so funny. I think you and I had a chuckle when we were prepping for this call about some of the similarities in terms of market evolution that you're experiencing that I experienced at Brocade. And I'm going to ask you a zinger right now, and, I don't know if you know the answer to this, but I can remember in the early days of Brocade, it was conventional wisdom that the storage area networking market would evolve identically to the data communications market that, you know, first you had Novell, with software that allowed file and print sharing, and then, we saw clusters, and then we saw switches, and then we saw wide area networks.

And I concluded that we were gonna completely cut out the idea of having very small storage clusters, and that customers would want to start with switches, and just so happened that I was right, and we really didn't have any competition. And, we grew very rapidly and got 90% plus market share.

That there seems to be a conventional wisdom that the market is going to go from driven air to one phase to two phase. And it kind of seems to me, like, the market may just completely leap over one phase and go directly to two phase for a whole bunch of reasons, not the least of which is not needing chillers and having a nonconductive refrigerant, et cetera, et cetera.

Am I completely on Mars with that statement? Is it possible if 90— 80% of the market is still using forced air, that they may just pivot to two phase as their first deployment?

Josh Claman: I think so. I do think so. I think that you have to realize that we talked about the conservative nature of the sector, but also the sort of incremental innovations for the last 30 or 40 years. Suddenly, these operators are really challenged. They have to figure out how to power these racks. They have to figure out how to cool these racks. They have to figure out how to generate energy behind the meter. I mean, it's not an enviable position to be in.

And I think in the initial wave of liquid cooling, they simply said, what is everyone else doing? I'm not gonna spend a lot of time climbing the learning curve on liquid cooling. I'm just going to do what everyone else is doing. At the very least, I won't lose my job doing that.

But water, in so many ways, is an impractical, just a fit, in terms of an architecture in the data center. And that is, by the way, you probably may not realize, Greg, how provocative that statement is. But the older and grumpier I get, the more provocative I become. But I do think that, because the narrative is that single phase water is the incumbent, but actually with kind of single, mid-single digit adoption, it's really not incumbent. The narrative is that it's incumbent.

So, I do believe with some of these sort of poor initial experiences with water— and by the way, an entire cottage industry has formed around this around water testing and filter changes and flushing of networks and all of this stuff— that there's gonna be a jump into two phase. I really do believe that. I think it's analogous to your experience.

Greg Reyes: Yeah, it's, it's, it just, if you have the ability to think critically and objectively, it just seems like a logical progression.

Um, so I guess at this point, Josh, before we open it up to questions, I'd like to give you the microphone to do a 60 second elevator pitch. If people are listening because they're excited about investing in the space, the data center ecosystem, and if they believe that Accelsius in cooling is a critical investment sector, and they can own Accelsius by purchasing Innventure stock, what's your elevator pitch to an investor as to why they should want to be involved?

Josh Claman: Yeah, I think we've talked about some of the big issues that I think we're a better solution now, we'll probably be a required solution. We're the leader, we have an incredible team, we have 5 PhDs on the team that are experts in thermodynamics and materials science, et cetera. We've really designed the team not just to release a product. We've designed a team to stay ahead, to release products along the way and stay ahead of the market and continue to be a leader.

I think when you think about any sector, any analogy, you would want what we have, which is an urgency in the market to solve a problem, and differentiators, whether it's lower flow rate, higher heat transfer coefficients, more future headroom as an architecture that doesn't damage the equipment it's designed to protect, a total cost of ownership that's lower than water. You're gonna, it's an enviable position to be in.

And so, I won't go into all the risks in front of us and the nervousness and paranoia that comes from being an early stage CEO. Many of you on the call probably experience that, but I think we are in a very, very solid position to continue to drive adoption and maintain leadership in this space.

Greg Reyes: Oh, Josh, I can't thank you enough for making the time. It's been great getting to know you, and hopefully the folks listening in have gotten some value out of the time. And I guess in the last 15 minutes or so, maybe, if you're okay with it, we'll open it up to some questions from the listeners.

Josh Claman: Sure, absolutely.

Greg Reyes: Okay, any questions? Okay. Well, Josh, I am not a pro at opening it up for questions. Usually, I see people raising their hands, but I'm not seeing any at the moment. Okay. Well, if somebody wants to DM me a question, I'll ask it for Josh. I think I'm making a mistake in terms of how I'm operating here. Bear with me. All right, well, I'm not getting any questions, and I am, I'm a little bit embarrassed to say that I am not a master of the Twitter Spaces hellscape. So...

Josh, I hope that you've gotten some value out of this experience. I know I can speak for people listening in and tell you that I think that virtually everybody listening is very much appreciative. And I'm hoping that maybe we can have one in another quarter or 2

and people can continue to track Accelsius' journey. So, with that, we'll call it a wrap, and I can't thank you enough for making the time.

Josh Claman: Okay, well, thank you. I appreciate it.

End of Transcript