BITF Earnings Call Transcript

Date: 2025-11-13

Quarter: 3

Operator: Thank you for standing by, and welcome to Bitfarms' Third Quarter 2025 Earnings Conference Call. I would now like to hand the call over to Jennifer Drew-Bear from Bitfarms' Investor Relations. Please go ahead.

Jennifer Drew-Bear: Thank you, and welcome to Bitfarms' Third Quarter 2025 Conference Call. With me on the call today are Ben Gagnon, Chief Executive Officer and Director; and Jonathan Mir, Chief Financial Officer. Before we begin, please note, this call is being webcast with an accompanying slide presentation. Today's press release and our presentation can be accessed on our website, bitfarms.com, under the Investors section. Turning to Slide 2. I'd like to remind everyone that certain forward-looking statements will be made during the call, and that future results could differ from those implied in this statement. The forward-looking information is based on certain assumptions and is subject to risks and uncertainties, and I invite you to consult Bitfarms' MD&A; for a complete list. Please note that references will be made to certain measures not recognized under IFRS and therefore, may not be comparable to similar measures presented by other companies. We invite listeners to refer to today's press release and our MD&A; for definition of the aforementioned non-IFRS measures and their reconciliations to IFRS measures. Please note that all financial references are denominated in U.S. dollars, unless otherwise noted. And now turning to Slide 3. It is my pleasure to turn the call over to Ben Gagnon, Chief Executive Officer and Director. Ben, please go ahead.

Ben Gagnon: Good morning, everyone, and welcome to Bitfarms' Third Quarter 2025 Earnings Call. We made strong, steady progress in Q3, building on the momentum from the first half of the year as we advance our transformation into a leading North American HPC and AI infrastructure company. Today, I'll walk you through our investment thesis, value proposition and key developments, including updates on our energy portfolio and site-specific advancements, all of which gives Bitfarms a competitive advantage to capitalize on the surge in demand for HPC and AI infrastructure. Turning to Slide 4. I would like to kick off today's call by outlining our market thesis, one that we believe differentiates us from our peers and best aligns Bitfarms with long-term investors in our transition to HPC and AI. Infrastructure is not a bubble. Since the invention of modern compute, the supply of compute has increased exponentially. As compute grows, so too does the data center industry that powers it. This is a trend that has a trajectory of over 20 years of exponential growth and an annualized growth rate of 8.8% behind it. This isn't a bubble. It's a reflection of a new paradigm that showed no signs of slowing down before AI and now as AI rewrites the rules of how humans interact with computers, the demand for data center capacity is accelerating. But the demand for compute and infrastructure has reached an impasse. Data centers that used to be measured in kilowatts are now being measured in megawatts and gigawatts. Racks that used to support 10 kilowatts are now being designed to support 370 kilowatts. The exponential increase in demand for power can no longer be met at the pace of the market demands. And as a result, the lease rates for data center infrastructure, which have grown at an average rate of 3% over the last 20 years, are now growing at an average rate of 12% since 2022, and we expect this trend to continue. Turning to Slide 5. Infrastructure is a bottleneck. As manufacturers continually introduce newer, more efficient chips and increase production every year, this trend continues to accelerate. Next year, NVIDIA alone is expected to be shipping somewhere between 10

and 15 gigawatts of GPUs. And that doesn't include, of course, AMD, Intel, Qualcomm and others who are also producing their own hardware with over 100 gigawatts of chips expected to be produced by 2030. While the supply of compute chips continues to increase, the growth in data center infrastructure is happening at a much slower pace. It is not silicon nor capital that will be the real bottleneck for continued growth in HPC and AI, but power and infrastructure. Over the next few years, the gap between the amount of chips that are being produced and the megawatts and the racks available to plug them in and operate them will continue to widen significantly. We strongly believe that as this dynamic continues to play out, the value and the economics will continue to move in favor of those who own the energy and data center infrastructure. We've watched this play out in the market with the contracts that have been announced in the industry to date. When Core Scientific and CoreWeave announced their landmark transaction in April of last year, the rates were contracted around $120 per kilowatt per month. As we've moved further along this curve that's shown on the slide, those rates have continued to trend upward. Most of the contracts over the past few months have been around $150 per kilowatt per month. As time goes on, this trend is expected to continue with analysts predicting a massive shortfall of nearly 45 gigawatts of power for data centers by 2030. Just within the last 2 weeks, Satya Nadella, the CEO of Microsoft, confirmed the shortfall when he publicly stated on a recent podcast that they have GPUs they cannot deploy. We believe that over time, the companies who've allocated and will continue to allocate billions of dollars into compute will be increasingly economically incentivized to pay rising prices in order to deploy their compute faster and with greater certainty, because every day they do not deploy is a day of revenue they will never recover and because their customers will simply move on to a competitor. With direct operating margins for new GPUs typically in the 80s or 90% range, this infrastructure expense is a modest cost driver for those who own the compute, equivalent to a low single-digit percentage of OpEx. If this cost were to double, it would not impact direct OpEx for the customer by more than a low single-digit percentage. These rates, which are largely inconsequential for the customer are very significant for Bitfarms as the developer. With OpEx costs that are largely fixed, every additional dollar earned in a lease goes to the bottom line. This is what Bitfarms is aiming to optimize for, not the fastest contract, but the highest value per megawatt and the greatest margins for the longest period of time with great customers. We believe this will be the primary driver of our multiple expansion and what drives shareholder value creation long term. Our investment thesis is clear and backed by decades of data. Our conviction is high, backed by consistent incoming demand. We don't want to cap our upside by signing leases prematurely. Instead, Bitfarms plans to optimize and achieve higher lease rates and margins through the following 3 strategic actions: one, prioritize infrastructure development first by minimizing the time between signing a lease and generating revenue for a customer, we will minimize the discount that would otherwise be applied to the lease rates and locked into multiyear contracts; two, take advantage of the increasing gap between supply of data center infrastructure and data center demand to lock in higher rates and greater margins under multiyear agreements; and three, while the industry is focused on NVIDIA GB200 and GB300, Bitfarms plans to leapfrog NVIDIA's Blackwell architecture and lead the industry in developing infrastructure for NVIDIA's next-generation Vera Rubin GPUs across 99% of our 2026 and 2027 development portfolio. With Vera Rubin GPUs expected to begin shipping in Q4 of 2026, and the infrastructure requirements to support them largely incompatible with facilities designed for Blackwell GPUs, we believe Vera Rubin infrastructure will be in the greatest demand and shortest supply in 2027 and will command significantly greater economics. Turning to Slide 6. We are able to take this approach because we have a robust balance sheet to fund development and know the value of what we own. While we don't have the largest portfolio of power among the public miners who are transitioning to HPC and AI, we do have the largest portfolios of power in each of the regions in which we operate, none of which are in Texas and all of which are either existing or emerging data center hubs. With consistent inbound demand for our sites, we have high conviction in the value of our unique energy portfolio, the demand for our power and our ability to develop next-generation HPC and AI infrastructure. We believe that not all megawatts are created equal. Our megawatts are strategically located in high-value areas that have multiyear waitlist to secure the power we have today. Our campuses are close to major metros and existing data center clusters, have ample access to major fiber trunk lines and undersea fiber optic cables and benefit from temperate climate compared to places

like Texas. While Texas is undisputably a great energy market and arguably the easiest market to grow and develop megawatts in the U.S., there are, of course, trade-offs. The trade-off to short-term development efficiencies is long-term operating inefficiencies. It is no secret that besides power, the primary challenge with data centers is cooling and cooling is becoming an increasingly more difficult problem to solve as energy density continues to increase with every generation of new hardware. Building and operating data centers in a hot, arid desert climate like Texas as opposed to cooler northern climates like Pennsylvania, Washington and Quebec means more CapEx and OpEx for cooling. This isn't an opinion. It's math and engineering. If we built our exact same data center for Panther Creek with the same design, equipment and materials in Texas, it would have a PuE of about 1.4 to about 1.5. Whereas in Pennsylvania, Quebec or Washington, it would be about 1.2 to 1.3. That means for every megawatt we are converting, more of those electrons are going to compute, which is the revenue-generating activity for customers as opposed to supporting revenue generation through cooling. Simply put, our megawatts are harder to get in higher demand areas, produce more value for customers and are worth more per megawatt. In Pennsylvania, we have the strategic foresight to acquire our 3 campuses and submit our energy applications in 2024 before the HPC and AI demand really came into play in the state earlier this year. This has positioned us with secured power at Panther Creek and Sharon and at the front of the queue with very well-advanced power applications at Scrubgrass. In Quebec, new power allocations are almost impossible to get with numerous data center applications denied by the province in the past year. Bitfarms has 170 megawatts operating with some of the cheapest power rates for data centers in North America and 100% renewable. 100% of these megawatts are currently being utilized for Bitcoin mining. And just in the last month, we confirmed that we will be able to convert our Bitcoin megawatts for HPC and AI. This means our Quebec portfolio represents a unique and strategic opportunity to increase total data center megawatts in the province by 25% from about 700 megawatts today, while fulfilling 2 strategic national and provincial objectives, the scaling back of Bitcoin mining megawatts while increasing HPC and AI infrastructure and data sovereignty. In Washington, we have 18 megawatts of secured power in the largest data center cluster on the West Coast with the cheapest power in the U.S. for data centers and 100% renewable. Because of this, the area has a 10-year wait list for power. Everybody is looking to grow here, and it is nearly impossible to do so outside of secured megawatts like ours. This means that despite the relatively smaller scale of Washington, sites in the area are in high demand by both enterprise and hyperscalers alike. I'd now like to spend a few minutes discussing Washington and the news we issued this morning in more detail. Turning to Slide 7. Earlier this morning, we announced plans for the conversion of our 18-megawatt Washington site to HPC and AI workloads. We signed a fully funded binding agreement for $128 million for all the critical IT infrastructure and building materials to develop the full 18 megawatts of gross capacity with anticipated industry-leading energy efficiency between 1.2 and 1.3 PuE. The state-of-the-art facility will feature: one, validated reference designs, ensuring compatibility and performance with NVIDIA GB300s; two, modular infrastructure, enabling phased deployment and scalability, reducing the downtime of Bitcoin mining revenues and ramping up our time to HPC and AI revenues; and three, proven thermal and power management systems critical for HPC and AI operations. The construction team is in Washington today with the general contractor and are kicking off the conversion of the Washington site, which is targeted for completion in December 2026. Turning to Slide 8. I would now like to discuss monetization strategy at Washington. With decade-long wait times for new power and the cheapest power in the U.S. for data centers, we are actively pursuing colocation for both hyperscaler and enterprise, where we can capitalize on the long wait times as previously discussed. This morning, for the first time, we announced we are also pursuing GPU as a service or cloud. While our focus is on developing next-generation Vera Rubin infrastructure across most of our portfolio, we believe there are some compelling reasons to potentially go with cloud as a monetization strategy at Moses Lake specifically. One, GPU as a service would enable us to capture the benefit of the lowest cost power for data centers in the U.S. for ourselves and generate what we expect to be above-market margins and returns for cloud. Two, the relatively smaller scale makes cloud at this site easier to execute and finance. We have more than enough liquidity to consider the site and strategy fully funded today and are in active discussions with leading GPU manufacturers on GPU sourcing and financing, which we believe could be done on very attractive terms. GPU financing could

materially reduce CapEx requirements and enhance expected returns. Three, we expect that by demonstrating our ability to execute across the entire stack, we will also be able to better understand customer needs, provide better quality service and negotiate better leases at our other facilities. Lastly, but most importantly, despite being less than 1% of our total development portfolio, we believe that the conversion of just our Moses Lake site to GPU as a service could produce more net operating income per year than we have ever generated with Bitcoin mining, providing the company with a strong cash flow foundation that would fund OpEx, G&A;, debt service and contribute to CapEx as we wind down our Bitcoin mining business. I will now walk through the rest of our sites in a bit more detail, starting with Panther Creek. Turning to Slide 9. Panther Creek is our flagship HPC and AI campus in Eastern Pennsylvania. As we've discussed previously, we have 350 megawatts of secured power with PPL. This power is contractually obligated to be delivered with 50 megawatts at the end of 2026 and 300 megawatts at the end of 2027. The site has sufficient acreage for the development of the entire 350 megawatts with capacity to go beyond that. Additionally, we have $200 million remaining on our project facility with Macquarie that is intended to finance Phase 1 of the project as well as a few long lead time expenses for Phase 2. We also have some exciting news around potential further capacity expansion at Panther Creek. Lately, there have been a number of developments, including the recent 403 letter from the Department of Energy and commitments to deploy more natural gas energy generation in Pennsylvania that have given us line of sight to expand beyond the existing 350 megawatts of secured power capacity. We have received positive indication on converting our existing interconnection service agreement, or ISA 60 megawatts to a firm energy service agreement, or ESA, of 60 megawatts to expand power to 410 megawatts and on a recent load study to expand power capacity to over 500 megawatts of growth capacity. With these positive developments that could meaningfully expand capacity at this campus and in line with our investment thesis, we are modifying our original Phase 1 designed for Blackwell GPUs and planning a new Phase 3 and Phase 4. The entire campus will now be developed for NVIDIA's Vera Rubin GPUs and their greater energy density to accommodate our new expectations on future expanded power capacity. This is expected to delay the energization of Phase 1 marginally from December 2026 into the first half of 2027, with no anticipated impacts to Phase 2 time lines. We believe this will enable the company to achieve significantly higher economics in line with our long-term thesis and strategy. Turning to Slide 10. Moving on to Sharon, where we have 110 megawatts of power secured by an ESA with FirstEnergy and PJM under development. We are currently operating 30 megawatts of Bitcoin mining on site, but have started development on an additional 80-megawatt substation, bringing the total available for HPC and AI uses to 110 megawatts. We expect to have the full 110-megawatt substation online by year-end 2026. We recently closed on the purchase of the land for the site, effectively ending our lease and enabling us to move forward with our planned development of HPC and AI infrastructure. Similarly to Panther Creek, we will be working to develop the campus for Vera Rubin GPUs, targeting site completion and revenue in the first half of 2027 for the full 110 megawatts of gross capacity. Turning to Slide 11. In Quebec, we have 170 megawatts of low-cost hydropower currently operating across multiple Bitcoin mining sites, almost all of which are within a roughly 90-minute drive from Montreal. This is an incredibly attractive opportunity for hyperscalers who are following what's called a regional campus strategy. This is something that was pioneered by Amazon, where smaller sites can be directly connected with direct fiber infrastructure in order to reduce the latency between sites below 2 milliseconds, enabling many sites to be connected together to function as one larger site. As I mentioned, it's almost impossible to grow organically in the province. And in October, we confirmed the ability to convert over our Bitcoin mining infrastructure to HPC and AI with regulators and utilities in the region. With that pathway clear, we are accelerating our plans in Quebec. We will focus our development efforts on the city of Sherbrooke, where we have 96 megawatts, robust fiber connectivity, a strong and developed local labor force and ample support from the local energy utility and municipality. We will be applying some of the standardized engineering and design plans completed for our Washington site to Sherbrooke in order to convert these facilities from Bitcoin mining into next-generation HPC and AI infrastructure adapted for Vera Rubin GPUs. Similar to Washington, Quebec has a cool climate and some of the lowest cost energy in North America for data centers. With strong unmet demand for GPU cloud in Montreal, Sherbrooke also represents a potential opportunity to scale up a cloud business in 2027 with VR200s, a strategy that we will evaluate as we

work through the engineering and development plans for Sherbrooke. The remaining 74 megawatts of Bitcoin mining in the province are earmarked for potential expansion in 2028, and we look forward to providing more detailed plans for Quebec in 2026. Turning to Slide 12. Last, but certainly not least, we have our Scrubgrass campus in Pennsylvania. This is about 30 minutes away from our Sharon, Pennsylvania campus on the western side of the state. With the exception of the new Panther Creek Phase 3 and Phase 4, which I spoke to a minute ago, this is the only power in our portfolio that is not 100% fully secured today. This is a very, very exciting development opportunity for Bitfarms. We believe this is the only campus outside of Texas for public miners converting to HPC and AI that has over 1 gigawatt of potential capacity. And while we have made great progress on developing the power story for this giga campus, there are still quite a few steps to be taken in order to contractually secure the power, which falls into 2 buckets. First, we have completed 3 conceptual load studies with FirstEnergy, starting with 250 megawatts, 500 and then 750 megawatts, thus moving over to what's called a detailed load study with FirstEnergy, which would eventually be converted over to firm service in an ESA. Second, we have made substantial progress on evaluating the potential to add additional generating capacity on site. This could be accomplished by building a 3- to 4-mile pipeline from our campus to the second largest natural gas pipeline in the U.S., the Tennessee Natural Gas Pipeline, which we have confirmed could supply up to 550 megawatts of natural gas, multiplying our generation capacity on site. We're still in the early stages of evaluating how we would expand the generating capacity, and we'll provide more details as we progress. Combined, the 2 buckets could potentially provide 1.3 gigawatts of gross capacity. And additionally, there is very good fiber infrastructure in the area with our 8 fiber infrastructure networks nearby and is in close proximity to Pittsburgh and Cleveland as well as the other data centers, which are starting to pop up throughout the state. The earliest time that we anticipate we could have additional power at this kind of scale implemented at Scrubgrass is around 2028. Though this is a longer lead time campus for us, we believe that with the forecast on power and demand for HPC and AI infrastructure, the timing for our giga campus will play-in well with the cycle, our investment thesis and our other development plans. Turning to Slide 13. To sum up, we believe that we are incredibly well positioned to execute against our investment thesis in 2026 and 2027 and maximize long-term shareholder value. One, we have a very unique portfolio of energy assets that we aim to fully convert to HPC and AI infrastructure. Two, we have announced our plans to convert our Washington site to HPC and AI workloads and lead the industry in the development of next-generation data centers for NVIDIA's Vera Rubin GPUs. Three, we are actively evaluating a potential cloud monetization strategy for our Washington site, which we believe would be a meaningful driver of cash flows and could eclipse any Bitcoin mining cash flows we have ever generated. Four, we are well capitalized to make our currently planned investments with a financial flexibility that exceeds $1 billion across cash, Bitcoin and our Panther Creek project facility with Macquarie, all of which are going to fund CapEx. As we continue to produce strong free cash flows from our Bitcoin mining operations that fund OpEx, G&A;, debt service and contribute to CapEx with no further planned minor CapEx. And lastly, we continue to execute on our U.S. pivot with the anticipated sale of our Paso Pe facility and our full LATAM exit. Our transition to U.S. GAAP for Q4, the establishment of our New York City office and working towards a U.S. redomicile in 2026. We believe this would give us significantly greater index inclusion and meaningfully improve the institutional composition of our cap table. I now have the pleasure to hand the call over to our new CFO, Jonathan Mir. Turning to Slide 14. Jonathan, over to you.

Jonathan Mir: Thank you, Ben, for the warm introduction. I'm excited to join Bitfarms at this pivotal moment in the company's transformation. My principal objectives as the new CFO are centered around capital allocation, capital sourcing and capital structure. I'm working hand-in-hand with the operations and development teams on the ground to ensure we implement financing plans that are appropriate for the company and its assets, efficient and support long-term shareholder value creation and that we are also allocating capital to its best possible risk-adjusted returns. With an extensive background in energy infrastructure strategy and financing, I believe there's an extraordinary opportunity to use our strong balance sheet, unique assets and the talents of our people to create value in the high-growth HPC/AI space. I look forward to working closely with the team to deliver on our strategy and capture the

exceptional long-term shareholder value that would accompany our successful execution. Turning to Slide 15. Today, Bitfarms has the strongest balance sheet and most available capital in the company's history. In Q3, we were able to execute across several initiatives. First and foremost, we recently completed a very successful convertible note offering, where we were able to upsize the offering to $588 million while improving on pricing, preserving upside and minimizing potential equity dilution through a 125% capped call. Bitfarms chose to issue convertible notes because they allow us to access capital at a lower coupon than straight debt and with less dilution than straight equity. The cash settled capped calls we purchased allow us to offset economic dilution up until $11.88 per share, representing a significant premium to the share price today. It is also important to highlight that investor commitment to Bitfarms is strong. 100% of institutional investors that management met with during the marketing process participated in the transaction and invested their capital in Bitfarms. We're thrilled with the outcome of this raise, and it will allow us to advance our pipeline in tangible ways. Second, we converted our previously announced $300 million debt facility with Macquarie to a project-specific financing facility dedicated to the development of our Panther Creek data center. Moving the debt facility from a corporate level to the asset level materially enhances financial flexibility for the entire company. In October, we drew an additional $50 million from the facility in order to accelerate development of the site for a total of $100 million drawn to date. Finally, we maintained steady and efficient mining operations throughout the quarter, achieving approximately $8 million in monthly free cash flow after G&A.; We expect to use this cash flow to support our HPC/AI development projects. Looking ahead, we anticipate continuing to use a mix of both corporate level and project level debt and equity financing as we advance our project milestones. On an ongoing basis, we will evaluate a wide range of opportunities and choose those that we believe support both a strong, stable balance sheet and realize the full potential shareholder value creation that would accompany the successful execution of our plans and fund milestone objectives. Turning to Slide 16. Let's focus now on our third quarter financial performance. In Q3, we achieved a total revenue of $84 million from continuing and discontinued operations. With the intention to sell the Paso Pe site in order to complete our Latin American exit, all revenue from that asset is classified as discontinuing operations. From continuing operations, we earned 520 Bitcoin and achieved revenue of $69 million, representing a year-over-year increase of 156% in revenue. For our continuing operations, our gross mining profit was $21 million, representing a gross mining margin of 35% and an average direct cost of $48,200 per Bitcoin mined. During the third quarter, we introduced a new program for digital asset management, Bitcoin 2.1, which is designed to offset Bitcoin production costs and achieve higher value per Bitcoin sold as a low-cost and low-risk funding mechanism for the energy infrastructure investments that define Bitfarms going forward. It is important to highlight that we are not a Bitcoin treasury company. The goal of this program is not to accumulate Bitcoin, but rather to offset the production cost of Bitcoin and by doing so, contribute to cost effectively funding our HPC/AI initiatives. This is a multi-strategy program that primarily sells both short and long-dated out-of-the-money covered calls on the Bitcoin and treasury as well as for Bitcoin production. During Q3, we incurred an all-in cost per Bitcoin of $82,400 from continuing operations. When considering our net gain of $13.3 million from derivatives against our all-in production costs, it would bring the effective all-in cost down to $55,200. Cash G&A; for Q3 was $14 million compared to $20 million in Q3 2024. The improvement was largely driven by lower professional services costs. Operating loss from continuing operations was $29 million for the quarter, including impairment charge of $9 million of nonfinancial assets. As a result, net loss from continuing operations for Q3 was $46 million or $0.08 per share. For the third quarter, our adjusted EBITDA from continuing operations was $20 million or 28% of revenue, up from $2 million or 8% of revenue year-over-year in Q3 2024 and up from $9 million or 15% of revenue in Q2 2025. Turning to Slide 17. Before we begin Q&A;, I'd like to reiterate our strong financial position and review our expected capital investment plans for the next 12 months. We are extremely well capitalized to fund our HPC/AI growth initiatives. We have a war chest of over $1 billion, comprised of roughly $820 million in cash and Bitcoin and the remaining $200 million available to draw from our Macquarie facility. With these funds, we expect to be able to fully finance the build-out of our Washington site and the initial phases of construction at our Sharon, Sherbrooke and Panther Creek sites. As we advance our development, the actual investment in our projects will be dependent on a number of factors. We are currently focused on executing on the

initial phases of our projects, beginning construction and securing long lead time items to ensure our project time lines. We will continuously evaluate a wide range of financing alternatives at both the corporate and project level, maximizing shareholder value with accretive financing will determine our choices as well as the need for a healthy balance sheet. In closing, I'll underscore that Bitfarms is in the strongest financial position in the company's history, and we have a clear vision of how we are going to best utilize this capital to advance our HPC/AI build-outs in North America. The entire Bitfarms team is incredibly enthusiastic and engaged about the opportunities ahead. With that, I'll now turn the call over to the operator for Q&A.;

Operator: Our first question comes from the line of Mike Colonnese of H.C. Wainwright.

Michael Colonnese: Appreciate all the color on the HPC strategy this morning. First for me, Ben, you mentioned that infrastructure for the Vera Rubin GPU should command a premium to the Blackwell infrastructure. Can you share more on how you guys are thinking about economics there and the CapEx differences?

Ben Gagnon: Thanks, Mike. Yes, happy to speak to that a little bit. There's kind of 2 driving forces there with our expectations on Vera Rubin economics. The first is that as the dynamic continues to play out where the infrastructure is going to be an increasingly greater and greater shortage, there's going to be a driver there that will drive the economics. And the second part of this is that the economics around supply and demand imbalance are really specific to GPU models. So if you look at H100s, H200s, the GBs, the 200s and 300s and then what's going to be the next series, the VR, there's a lot more infrastructure available to support those older GPUs, which have less specific requirements. And when you look at what's going to happen with the VR series, the energy density is going up from 190 kilowatts per rack with the GB300s to upwards of 370 kilowatts per rack with the VR200s. And so a lot of the infrastructure that's being built right now is not going to be compatible with the next generation. And as companies allocate all this money into those Vera Rubin GPUs, they're going to be very economically incentivized to deploy them. And what I spoke to with regards to our investment thesis earlier today, is that as this dynamic continues to play out, would you rather sit on your GPUs and not deploy them? Or would you rather pay a higher infrastructure expense in order to deploy them and start monetizing the asset. And really, the margins are so high on these GPUs, especially when the GPU is the newest, most cutting-edge state-of-the-art GPUs as the Vera Rubins will be in 2027, that the economic incentive to deploy those faster with very few options available should drive higher economics. We don't have a firm price point of exactly where that's going to lie, but we think the trend is abundantly clear that the economics next year and in 2027, they're just going to continue to get better and better, especially as the shortfall continues to get exacerbated.

Michael Colonnese: Really helpful color there, Ben. And how should we think about the wind down of your mining operations in the coming years, specifically as it relates to the pace and timing of hash rate coming offline as you start to convert and make further progress in converting your data centers over to HPC/AI?

Ben Gagnon: Yes, happy to speak to that. I mean the first area is the LATAM export that we've been working on. We obviously shut down our Argentina facility earlier this year. And I think one of the big areas here is the Paso Pe facility, which is an asset that's being held for sale. That represents around a little bit under 20% of our hash rate. And so that will impact the hash rate for the company rolling forward. But when we look at transactions like this, just like how we looked at the economics around shutting down the Argentina facility, we expect to pull forward a significant amount of expected free cash flow from those operations today so that we can reinvest them more immediately in the U.S., in North American HPC and AI infrastructure to greater effect. So while it should have an impact on the free cash flow from operations, really the impact is very mitigated by the fact that we're taking 1 to 2 years' worth of free cash flow from operations and bringing it forward for reinvestment now. And then we also have the derisking factor with regards to having less and less Bitcoin exposure or Bitcoin mining exposure, I should say. So as we move forward through 2026, the next sites that would be

coming offline, would be coming offline as we develop the HPC and AI infrastructure and they would get replaced. Washington would probably happen sometime in the -- probably middle of the year, and that would be about 1 exahash and everything else will kind of come off slowly as we convert over the facilities to HPC and AI. So it would be a bit of an orderly transformation, and we'll continue to update the market as we announce those plans.

Operator: Our next question comes from the line of Brett Knoblauch of Cantor Fitzgerald.

Brett Knoblauch: Thanks for a lot of the color on the different sites throughout the call. I guess when it comes to maybe your PA sites and getting additional power, I feel like that's kind of like the biggest catalyst maybe over the near term. I believe Stronghold was kind of in queue before you guys went out and acquired it, which was probably, I don't know, over a year ago now. Do you have any idea on an update of when you expect to maybe expand the power capacity at both Panther Creek and Scrubgrass. Is that a couple of months thing? Within 6 months thing? How should we think about the timing there?

Ben Gagnon: Thanks, Brett. Yes, it's a pretty exciting development there at Panther Creek because just over the last couple of weeks, we've received positive indications on the conversion of the ISA to an ESA as well as the expansion with an additional load study. It's a little too early to say exactly when that would come on to site. What we're planning here is an additional Phase 3 and Phase 4, which would come likely after Phase 2. But it's possible that the conversion of the ISA to an ESA could happen very quickly because all of the infrastructure is in place. There is no investments that need to be made. It's really just subject to the regulatory approval and signings and paperwork for all of that to be converted over. So I would think within the Phase 3, it's not really clear exactly when that's going to take place, but it could happen quickly. It could take several months. When it comes to a Phase 4, that's likely going to be a 2028 deal.

Brett Knoblauch: Awesome. And then on the GPU cloud as a service, the CapEx figure that you've noted on, I guess, maybe converting that Bitcoin mining to host GPUs, that was not including the GPUs, correct?

Ben Gagnon: Correct. That's not including GPUs and some of the construction costs associated with converting over the facility. So there will be additional expenses at the Washington site. We've had several conversations now with some of the leading GPU manufacturers, and we think that there's very attractive financing options on the GPUs as well that would really keep the CapEx requirement down to basically the infrastructure expense, and we'd be able to fund potentially up to 100% of the compute through these GPU manufacturers, which could be done on what we believe to be really attractive terms. And we also think that it would provide a significantly greater return profile on doing GPU as a service or cloud.

Brett Knoblauch: And from a capital allocation, I guess, standpoint, what is your guys' preference? Obviously, the PA sites appear to be leaning more towards colocation, Washington site cloud. Do you guys expect to kind of grow both businesses at the same time? Is there a preference for one to kind of get online sooner than the other?

Ben Gagnon: The expectation is that the Washington site will be the first site that's fully online. The Sharon site will probably be the second site that's fully online because Scrubgrass -- sorry, Panther Creek is split out into those 2 phases in 2027 with additional Phases 3 and Phase 4, which still needing to be confirmed. Our priority is managing the critical path and all the project management time lines that we have across our various facilities. But when we're looking at capital and how we'd allocate it across, it's managing the critical path, and it's also making sure that when it comes to looking at the opportunities around cloud, we're doing so in a way that makes sense and is affordable. And one of the benefits of doing it at Washington is a relatively smaller scale does make it very cost effective to do it. It's something that we could consider fully funded today. It's something that we could get financing for it at scale, whereas when you're looking at the really large campuses that we have in Pennsylvania, a

colocation strategy is going to be a lot easier to finance.

Operator: Our next question comes from the line of Stephen Glagola of JonesTrading.

Stephen Glagola: On the $128 million critical IT supply agreement for Washington, can you clarify the counterparty to that agreement? Is that T5? Or is that another firm? And then additionally, just a follow-up to the last one on the GPU cloud model potentially at Washington and Sherbrooke. Can you maybe elaborate on what factors make GPU as a service compelling relative to standard colocation in these markets? And sort of how are you evaluating both potential GPU risk and your, let's say, return on invested capital IRR hurdle for the cloud opportunity?

Ben Gagnon: Yes. Thanks, Stephen. When it comes to the supply agreement that we have for the Washington site, it's not with T5., it is with a large publicly traded American national company who serves and supplies data center equipment and data center services. The facility is really an attractive facility for both colocation and both cloud. But when you look at the opportunities that we have here to go fully up the stack and what that might mean for the company, both in terms of a free cash flow perspective as well as our ability to really demonstrate ourselves not only as a developer, but as an operator, I think there's a lot of tangible benefits there that will pay dividends in the long run. The conversion of the site according to our modeling and similar transactions that have happened in the market over the last couple of months, indicate that this one site could be worth significantly more than the entire Bitcoin mining business that the company has been operating for multiple years. And so that would provide us with a really strong free cash flow foundation as the Bitcoin mining business winds down. It will also enable us to better understand and better learn these facilities as we're looking to provide service and work with hyperscale and enterprise customers and neocloud customers on really large campuses. And so the benefit of doing it at the smaller facility is that we should be able to extract a lot of knowledge and value that we can apply to a lot of our other facilities as well.

Operator: Our next question comes from the line of Mike Grondahl of Northland.

Mike Grondahl: Ben, just curious, what would you describe as the 2 biggest challenges to maybe meeting your time lines for Washington, Sharon and Panther Creek? Like what's going to be the potential bottlenecks and how are you dealing with them?

Ben Gagnon: Mike, I mean the potential bottlenecks in construction are a little hard to forecast. I mean construction is something that is changing every single day on the ground. I think that the key way that you mitigate potential risk in construction is having great partners with your owners rep, your general contractors, having a great team of project managers internally who are making sure they're on track of everything, every step of the way, and they're trying to think forward on all the potential problems in managing that -- those critical paths. It's not possible, I think, to identify what would be the key bottleneck or the key risk. But I think with the team that we have in place, the strategic partners that we have in place and the kind of groups that we're working with on the contracting side or on the owner's rep side, we're in a really strong position to execute.

Mike Grondahl: Great. And then any rough guidelines or framework you can give us for sort of like 2026 CapEx?

Ben Gagnon: So when we're looking at 2026 CapEx, we've outlined some of the numbers for Washington. We're still working on clear path forward as we're revising for Vera Rubin. The real challenge with providing full CapEx figures for 2026 is that the Vera Rubin infrastructure is so new that even NVIDIA hasn't completed their validated reference designs to support that equipment and that infrastructure. So that's something that's adjusting in real time and still moving forward. We should expect to have a better indication of what CapEx looks like in 2026 in Q1. From our conversations that we're having with the various different engineering firms and suppliers and partners of NVIDIA, NVIDIA is going to be producing the first Vera Rubin GPUs and taking them for their own purposes in probably Q2 of next year. And so sometime in Q1, the reference design should be relatively final, and we should be clear in terms of what the CapEx implications are for 2027 and 2026.

Operator: Our next question comes from the line of Nick Giles of B.Riley.

Nick Giles: Appreciate all the detail here. Ben, you mentioned the higher rack density of the Vera Rubin gen and that it could make the rack density suited for Blackwells obsolete. And it wasn't that long ago that 100 kilowatts per rack was the high end of the rack density. So how are you thinking about future proofing as this trend continues? And are there any contract structures that could protect you from the need to upgrade later down the road?

Ben Gagnon: Thanks, Nick. It's a great question. The evolution of hardware is happening at a rapid pace, right? The GB200s were 150, the GB300s were 190 kilowatts per rack. And now the Vera Rubins are going to be over 370. And what that means is that your cooling needs to provide a lot more capacity in a very small footprint. It also means that your electrical distribution is very different. Most of the networking is more or less the same. But on the cooling and the electrical, it's a really big challenge. And one of the things that NVIDIA is looking at doing is increasing the voltage and even going to direct DC systems for the Vera Rubin technology. So they're looking at switching over to 800-volt DC. That doesn't mean that you necessarily have to go upwards of 800 volts or switch over to DC, but it does mean that as the increasing energy density continues to accelerate, you need to be rethinking your energy infrastructure and how you're actually building out these facilities. I think one of the ways that you try and do this is you try and build for the hardware at the time and then you try and lock that in with multiyear agreements, which help you to recover your investment and capitalize those investments over a long period of time. When you're signing an agreement for 5, 10, 15 years, most of the time, those agreements don't anticipate material upgrades to the infrastructure or any upgrades to the infrastructure. And so you're locking yourself in, the customer is locking themselves in with the infrastructure that they have in hand. And so I think the best way to mitigate those risks is to spread out your facilities, make sure you have a pipeline that exists over multiple years and make sure that you're building to the technology that's coming, not to the technology that already exists today because if you're building for today's technology by the time the facility is done, it's obsolete.

Nick Giles: I really appreciate that perspective. That takes me to my next question. You mentioned the pipeline. Obviously, you have a lot of growth in front of you, but how much time are you spending on M&A; opportunities? And where does that ultimately rank in terms of capital allocation?

Ben Gagnon: Virtually none, Nick. Our focus as a management team is execution, execution, execution. We don't believe that there is a tremendous value that comes for our shareholders for looking at opportunities that are 2029, 2030 and these kind of long lead time items. We believe the value comes from executing against our existing portfolio. And we continue to get inbounds in terms of new opportunities and growth opportunities, but none of them seem to compare at all with what we already have in hand. And so I think the best opportunity for us is to continue to execute against our existing pipeline. There will be a time in the future where we're going to want to continue to expand that pipeline. But that's probably an easy year or 1.5 years out from today.

Nick Giles: Got it. That's good to hear. Maybe one more, if I could, just for Jonathan. Sorry if I missed any commentary around this earlier, but how are you ultimately thinking about the Bitcoin treasury? Would you look to liquidate these holdings around the time that mining operations wind down? Or would those be separate time lines?

Jonathan Mir: So to be -- first, it's nice to meet you. So we are definitely not operating as a Bitcoin treasury company, and we don't want to be one. What we're doing right now through programs like Bitcoin 2.1 is offset Bitcoin production costs and achieve higher value per Bitcoin sold in a low-risk, low-cost funding mechanism for the energy infrastructure investments that define Bitcoin going forward. The program primarily sells short and long-dated out-of-the-money calls on the Bitcoin and the treasury as well as for Bitcoin production. So our efforts are focused around maximizing yield and minimizing costs. And we expect the Bitcoin treasury to wind down into strength as we allocate it to CapEx.

Operator: Our next question comes from the line of Martin Toner of ATB Capital Markets.

Martin Toner: Congrats on all this progress, guys. My question is around the GPUs. What's your confidence in being able to acquire them on a timely basis? And would you go through a distributor that comes with the financing or who might finance them?

Ben Gagnon: Thanks, Martin. Yes, happy to speak to that. We've had quite a few conversations with leading GPU manufacturers. As you probably know, NVIDIA produces GPUs themselves, but they also sell chips to a lot of OEM manufacturers. When you speak with those manufacturers, they often have finance programs in place, and those finance programs are -- can be pretty attractive, especially if you have the right infrastructure to ensure the quality and the lifespan of those GPUs. So going with an OEM manufacturer has a lot of benefits. They'll provide a full turnkey solution with regards to the servers themselves, and they can often come with financing. With our time line for end of next year on Washington, we're highly confident in sourcing our GPUs, and we believe that there's a lot of financing options out there that we are evaluating and could really juice up those return profiles.

Martin Toner: That's great. Is there a good exahash number to use for Q4?

Ben Gagnon: Our exahash should stay relatively consistent in Q4 when you're looking at our continuing operations. It's not possible right now to really forecast the impact or when the impact from the Paso Pe sale is going to happen. But the site continues to run today. It continues to hash. It continues to generate free cash flow. It's just not classified there under normal revenue according to IFRS standards, we have to hold that under discontinuing operations. But I think if you just look at the hash rate associated with our -- the rest of our portfolio, that will stay relatively constant -- it will stay constant throughout Q4, and then we'll make adjustments to it throughout 2026 as we execute on the HPC and AI development.

Martin Toner: Fantastic. Can you give us a sense for initial conversations with customers of the GPU as a service product, reaction and confidence in being able to like contract them on a timely basis?

Ben Gagnon: So conversations on the GPU front are really new for us because we've only started evaluating this in the last month or 2 as we've seen the market dynamic really take hold. I think the inbound demands that we've had across Washington and specifically Panther Creek is a lot. And when we're looking at what's the best way to service those customers, what's the best way to lock in long-term value under those agreements, there's a variety of different customers who are coming to us, and some of them want the GPUs included in there, and there's an associated premium that could be potentially extracted from that. So it's a little too early to indicate exactly what we would expect with economics, but we do believe the economics from our conversations and from the internal modeling that we've done and from the transactions that a lot of the companies in the space have announced in the last couple of months is very compelling, especially when we can execute it at a smaller site like Washington, which we can consider fully funded today.

Operator: Our next question comes from the line of Brian Dobson of Clear Street.

Brian Dobson: I guess more broadly speaking about Bitcoin mining, as more and more miners transition megawatts to HPC? How do you see the global hash rate evolving over the next few years?

Ben Gagnon: No, interesting question, Brian. Personally, I think the hash rate is going to continue to evolve at the same rate that it has been evolving. But if Bitcoin price is not moving up meaningfully, that would be a major headwind to further growth. I think what you'll see more likely is that Bitcoin miners will continue to rotate out to lower and lower cost jurisdictions. And I think one of the big dynamics that is taking place is that the public miners represented almost 1/3 of the entire network, and they all seem very keen on moving over to the higher economics associated with HPC and AI. So that removes a lot of the available and current existing infrastructure for Bitcoin mining. So there could be some potential headwinds in exahash growth for the network. But I think what you'll see is it's just going to rotate off to different jurisdictions. We've seen huge growth in the Middle East, in Africa. I think Russia is a very large booming market for Bitcoin mining right now. And I think the best opportunity for most miners in

the United States really is this transition to HPC and AI. And the economics are really going to drive that forward because the U.S. is the best market to invest in for HPC and AI, whereas Bitcoin mining is largely location agnostic. And it's happy to go to cheaper locations, higher-risk locations, more remote locations than HPC and AI is.

Brian Dobson: Yes, excellent. And then just a quick follow-up. So as you're reviewing your portfolio, do you see an opportunity to engage in this type of megawatt redeployment in a broader sense?

Ben Gagnon: When we're looking at whether or not we could redeploy our Bitcoin mining assets somewhere else, I think the opportunities are really few. And really, I don't think that's a great use of management's resources or time. I think the best opportunity is to basically bring forward what should be estimated free cash flow for mining operations today into cash and reinvest those into HPC and AI.

Operator: Our next question comes from the line of Michael Donovan of Compass Point.

Michael Donovan: Ben, you mentioned dollar per kilowatt trends. Can you quantify a premium on dollar per kilowatt that you're seeing for power secured in Pennsylvania or Washington versus Texas?

Ben Gagnon: Yes, it's a good question. There's a few variables that go into dollar per kilowatt on these leases. One is obviously time line, one is location. Another one is risk factors that go into the development time line. And so it's not really possible to pinpoint an exact price per location because there's multiple factors which come into play when you're looking at what the total lease rates can accumulate to. I think if you look around at the transactions that are here and you look around at kind of what Bitfarms could secure today at Pennsylvania before it's even really broken ground at our Panther Creek site, which we plan to do next month, we could probably lock in $140 to $150 per kilowatt per month. But I think when you look at that rate, that rate takes into consideration the location. It also takes into consideration the shovel has not been put in the ground yet. And what we don't want to do is we don't want to lock in a lot of discounts that would be associated with the build time line and the uncertainties around the build time line into a 10-, 15-year agreement. What we'd rather do is we'd rather execute against our construction milestones utilizing the substantial war chest that we have today. And the closer we can bring that window down from signing a lease to actually generating revenue from a customer, the more that we should expect to get. It's hard to put an exact price, but I would think that if that window was shorter, we could probably get upwards of $180 per kilowatt per month if we didn't have the risk and uncertainty priced into the time line that would bring it down to $140 to $150 per kilowatt today. That's internal estimates and modeling. So there's a lot of factors that go into that. And we also think that as you execute against 2026 and as the gap between data center supply and data center demand continues to exacerbate, those numbers could get even better. And when we look at how does the margins work out for these contracts, you're largely looking at pretty fixed OpEx. And so the difference for the company between getting $140 per kilowatt hour, $140 per kilowatt per month versus $150 or $180 is not only a huge increase in terms of the top line revenue, but it's an even larger increase in terms of the profit margin, in terms of what your adjusted EBITDA is going to be. And then not that all translates out into that multiple expansion that we're targeting with this transformation, right? So if you're getting a significantly higher free cash flow out of that operation, that's what the multiple expansion is going to be based on. So we really want to make sure that -- we're not pricing in those discounts. We're trying to maximize the dollar per kilowatt per month in the lease, and that's going to be the way that we achieve the highest multiple expansion for shareholders in the long term.

Michael Donovan: That's helpful, Ben. And you talked about connecting data centers to be one campus, and I was hoping you can unpack this a bit more. How can we think about distance between hauls or pods versus theoretical loss and performance for compute?

Ben Gagnon: Yes. There's a strategy that Amazon pioneered. It's called the regional campus strategy, and they've effectively determined that somewhere around 300 miles is the cost-effective range to build direct fiber infrastructure. But the real thing is the latency that you could get between your sites. Now obviously, when you're looking at these facilities, you're even concerned about the latency in rack and

in between racks or inside the facility to go from one rack to another rack on the other side of the facility. So that latency is becoming an increasingly bigger bottleneck as you're looking at performance on the high, high end of GPUs. But what we've seen is that most of our facilities in Montreal, where we'd be looking at this regional campus strategy, they're much closer than 300 miles. They're all within 90 minutes of Montreal. Many of them are 15-, 20-minute drive apart from each other. And so it would be possible to reduce the latency below 2 milliseconds with direct fiber. It would be pretty cost effective to do so. And you'd get a lot of benefits from doing that in terms of the scalability, given it's just so difficult to scale up new megawatts in the province.

Operator: I would now like to turn the conference back to Ben Gagnon for closing remarks. Sir?

Ben Gagnon: Thank you very much. I would like to thank everyone for attending our earnings call this morning. The management team is very excited. Our long-term investment strategy, we believe, is fully aligned with long-term investors. And we are really, really excited about the future of this company and what we're building at Bitfarms, and we appreciate your continued support. Thank you.

Operator: This concludes today's conference call. Thank you for participating. You may now disconnect.