Summary:

During a monthly community call, the Nosana team provided updates on various topics. Laurens discussed technical updates, including changes to the Nosana Node software and the introduction of semi-fungible tokens for market access. Jesse addressed the recent changes to staking and APY rates, explaining the reasoning behind them. Sean provided updates on Test Grid Phase 2, including the number of inference jobs being run and the reopening of registration for new nodes. The team also discussed plans for onboarding clients and the potential for future revenue streams. They emphasized the importance of client growth and adoption in driving the success of the network.

Transcript:

Rochelle:
Welcome to our first monthly community call. We're so excited to start having these on the last Thursday of every month. Today is the 30th of May, and it's 5:00 p.m. CET, so that means it's time to go. Today we have Sean, Laurens, and Jesse speaking on this community call. We're going to cover dev and technical updates, the staking and APY change (I know that's a hot one for you guys), and Test Grid, and we'll go over the use cases to really hammer in why we are doing what we're doing. So let's start it off. You guys want to say hi?

Sean:
Afternoon, everybody.

Laurens:
Hello, hello. It's good to see you all here.

Jesse:
Hey, good afternoon.

Rochelle:
So yeah, let's just get right into it. Laurens, can you give us some updates? Everyone's wanting to know what everyone's been up to and the team behind the scenes, so take it away.

Laurens:
Sure. Sure. I'll dive a bit into the technical things that we've been working on and that we've released in the last month or so. Most of you probably already know this, but of course, we made big changes to the Nosana Node software. We basically rewrote it from scratch, giving us lots of cool new possibilities to build features quicker and better with a really solid structure there. We made a system where we could now automatically match a GPU to the correct markets. So every time you start up your node, you don't have to specify your market anymore, but there's that it detects your GPU, does a benchmark, and based on your GPU, you will get access to the right markets. And these markets are basically protected with a token—a non-fungible token at first. That's all protected in the smart contract on Solana. So only if you have this non-fungible token that belongs to the collection of the market can you, as a node, get into that market.
So everybody that got accepted into Test Grid received this non-fungible token for his or her correct markets, and we wanted to also make sure that you can automatically switch markets if we detect that you have a new GPU or a different GPU. In order to do this, what's happened is that if you want to change GPUs, you basically have to send your non-fungible token back and it would detect that you would send it back. That's one

transaction, and then you would get a new transaction receiving your new non-fungible token for the new markets.
This gave some problems, especially with the Solana congestion that's happened in the past where people did send back their token, but it wasn't properly detected because of a transaction timeout or something else going wrong. So they basically lost their token, but they didn't get a new one. So we decided to switch up the system and move to semi-fungible tokens with a couple of cool features in it.
So a benefit is that we don't have to create basically a new token for every node. So there's basically one token per market, and we can give an amount, like one token of these semi-fungible tokens, to each node, and they are looked at. But most importantly, we have Atomic Swap now. So if you want to change markets, the sending back of your token and receiving the new token happens in one transaction instead of two separate transactions, which means that you will only basically lose your old token if you also get the new token. So the scenario where you basically send back your old token but don't get a new one, that problem will be gone.
That being said, there are still quite a lot of nodes right now on the old token, and throughout the next week we will work on switching everyone to this new semi-fungible token system to make it more foolproof to switch markets in the future.
And then yeah, of course there's Test Grid Phase 2 launched, but Sean will tell a bit more about that later on. So I won't go into detail about that.
We are also looking into multi-GPU nodes. Right now, there's one special market for an enterprise client of us that has multi-GPU nodes, so we're testing with that as well. It's not there yet for the auto markets. We are testing and experimenting with that as well. We have an upgraded benchmark job to basically test your GPU capacity, and in the future, this will also be used to detect cheaters basically. And we're slowly rolling out this benchmark job across all markets. There were some problems with the Docker registry, and we're tackling those problems as well. So that's also going to be very exciting. We're going to see more and more different and diverse jobs coming in besides the well-known Whisper and Stable Diffusion jobs that we have been doing.
Another really important thing that's coming up and that we're working really hard on is the improved Explorer. Right now, Explorer is kind of slow. It has to fetch all the jobs from the blockchain, which is just slow. And the more jobs we're posting and with all the new notes on board, there's quite a lot of jobs being done at the moment on our Test Grid Phase 2. So Explorer is already getting a bit slower and a bit slugly. So we're working on a new, improved Explorer and a new dashboard that's going to be way snappier, and have more filter options. And we're also working on that same dashboard on a node dashboard, where you can have a bit more information about your node, see what job you are currently running, and see how much you have already earned. Basically, your own node dashboard in a nice UI on our website.
For me, the most exciting new feature that we've been working very hard on is that it's maybe a bit technical, but it will open up a lot of possibilities for clients that want to use the Nosana Network. And I will also give a little demo about it. It's basically the possibility to first expose your node to the public when you are running a job and also forwarding local port that's running in a container when you are running a job through basically that proxy. So we're exposing web service possibilities in a job. And this sounds a bit fake, so maybe it's good for me if I just show a little demo of this. Let me try to see if I can share my screen. And Discord, I never did it before, but hopefully it works. It's a bit technical, maybe, but this is the text section, so bear with me. Can you all see my screen?

Rochelle:
Yes, we can see it very well.

Laurens:
Cool. All right, so here I have my terminal and I'll basically show how a simple job definition file will look like. So when jobs are being posted to the Nosana Network, the job is described in a job definition file. And it looks like this. It's a bit technical, but yeah, it's really flexible. It gives lots of possibilities for people to post different jobs.

This is like a container job. So you take a Docker container from the Docker registry, for example, in this case just a plain Ubuntu, and you can just send the commands to it.
So normally, somebody would send this to a market and a node would pick it up and get paid for his compute time. For the purpose of the demo, I'll just run it locally right away in my node, so without posting it to the blockchain. So I installed the Nosana Node CLI, so I can do Nosana Node run, and then I run, for example, this job definition file that I just showed, which basically just counts to 10 and then post the results back to the blockchain. So these are the raw results that are being posted back to the blockchain.
But then what really sets us apart is, of course, that we have GPU support. So you can add a little flag here, GPU is true. And that's all that is needed for clients that are posting jobs on the Nosana Network to be able to get GPUs. You can all see my screen still, right? I see there are some people who have problems with my screen, but most people can see my screen.

Jesse:
I cannot see it. It's loading for me as well, but I'm on a mobile phone, so I don't know if that's on my end. You can see it, right, Rochelle?

Rochelle:
Yes, I can see it. And I'm also recording this display, so afterwards everyone else will be able to see it too if they can't see it right now at the moment.

Laurens:
Okay, that's good. Then I'll just continue. Yeah, this one is not that great for sharing screens, apparently. So now I'll just run this GPU example, which shows that you can actually see the GPU that I'm running on my load. It's a pretty poor GPU, it's just on my laptop, but at least you can see that you can post a job and retrieve the GPU.
But now the new cool feature is that people can now expose ports from within their job. So it's for clients who want to use this. It's basically as simple as adding this line to your job definition file. Basically expose a port that's running in this Docker container. So this job is being posted and then yeah, the magic is all happening on the node side, where the node opens up a proxy to be able to expose himself through a domain, and then he forwards this local port that's running in a sandbox container in your environment on the nodes and is exposing that surface to the one that is posting the job. So this is a Jupyter Notebook Docker and job definition file. It basically starts like a web surface with a Jupyter Notebook that you can use to get access to the GPU. So I'll run that script right now. Nosana nodes run job samples. Jupyter Notebook.
So this is familiar to most of you. It's pulling the necessary Docker containers that are needed. But now, yeah, the magic has happened. So now there is an endpoint exposed for my specific note and for this specific job. So it's basically my node URL... Let me open up a Chrome window. So basically, it basically opens up a proxy for your specific node. So even if you're behind a router and not a firewall, you don't need to expose your ports or exposure nodes or have a public server or something. This works through a fast reverse proxy. And of course, this is testing for me locally, so that's why I have my own domain name. But this will be something like node.nosana.io and then your own node, which is your own node endpoint.
And then if I go here, because I'm still running that job, I see this Jupyter Notebook, so this is available as like a web server. So I exposed a local or that's running in my container because I'm running this job as a node and I'm exposing it to the public, so the one that posted this job can get access to a web service that's running on my node. So for example, I could even open up a terminal and I'm basically in my sandbox job environment, where I can even see my GPU here. So here's my GPU, right? And this is all exposed through a web service. So you can launch decentralized services, you could even use it to host a site on the Nosana Network, but mostly we built this for clients that want to do quicker AI inferences. So right now, every AI inference is like one job. One job is one AI inference, but this will change. So this means that one job will basically open up a session, like maybe an API to an AI model, and then while the job is running, the client can post multiple

prompts to this AI and have very quick AI inferences, like a peer to peer connection from the user to the Nosana Node. And yeah, that's because of this peer-to-peer connection, they can run AI inferences way quicker than what we can do right now. And at the end of the session, then it's settled on chain by basically finishing the job that ensures the session with the Nosana Nodes and then the payments are done on chain. Maybe it's a bit technical, but it's a really cool feature that I wanted to share with you all because it opens up lots of possibilities. And yeah, we're talking to quite some clients that are in need of these features, and I'm personally really excited to get this feature up and running.

Rochelle:
Thanks, Laurens, for the share. Thank you for the share. That was an awesome presentation. Again, anyone, if you weren't able to view it, I am recording this and you'll be able to see it afterwards when I post it. There's a lot of information and I can't wait to have it all condensed in text format for you all to go back over again as well. Moving on, I'd like to bring Jesse on to talk about the staking and APY changes.

Jesse:
Thanks, Rochelle. Thanks for the demo as well, Laurens. I enjoyed it. I couldn't see the screen, but I'll be watching the recording. So yeah, I'll go into the staking and the recent APY changes people have seen. It's been quite a hot topic, so I think it's worthwhile to cover the reasoning and the structure behind this so people understand better. And I will cover a few questions as well that have been asked frequently in the community. For some context, I think it's good if I give a bit of background on how the staking protocol started and where we came from so it's a bit more clear how things work and why certain things happen this way.
So the initial staking protocol was released, I think, nearly two years ago, so it's been around quite a while. I was one of the core engineers who designed the staking system. Nosana has quite a cool staking system where the stakers are able to receive rewards, like any fraction of rewards that can be distributed at any moment to all the stakers. So even if there is one NOS being distributed to stakers, it's possible to do this in real-time across all the stakers, even if there are thousands of them, which was a really innovative system. I think we were the first to do this on Solana, and it's very cool, and it's been performing extremely well for us over time.
But of course, Nosana has seen a lot of changes over the years. And most obviously, we've done a big pivot at the at the end of last year, where we were focusing heavily on the AI and GPU markets. And this also meant that we had to revisit some of the staking incentives that were in place.
At the time, we were using the mining pool to reward stakers. So basically, we're using a constant reward emission from our mining pool to increase the stake rewards. And this is next to the network rewards that are already automatically distributed to stakers. So currently on the Nosana Network, any transaction that happens as being part of purchasing GPU power, a fraction of that gets automatically distributed to stakers, and that, of course, adds to the APY. And next to that, there is the mining pool, which is the x-ray incentive right now to bootstrap the staking pool, and it's, at the moment, still the majority of the rewards being distributed.
So when we pivoted to this new market, we needed to sort of revisit the total timeline and the ambition of the project, which has drastically increased. And we noted that we needed to lower the mining pool reward emission in order to make sure the project can go on for many more years, which is required.
So in January we did the first of two adjustments to the mining pool emission, resulting in a lower APY rate, and I think maybe around a week a bit more ago we did the second step of this process. So this was a two-step process that was announced as the first step in January, and at the moment we're at an emission rate that is very healthy for the network and allows us to go on for many, many more years as is required.
But of course, there are a lot of questions and people were, I think, surprised by this event, which, of course... Yeah, that's not nice. And I really understand that people were concerned and had questions about why this works like it does. But I think the most important thing to understand is that it's a really important adjustment to make sure the network is in a healthy state for growth in the coming years. And that's just something that's absolutely necessary in order to make sure we can get to the places where things are going right now.

So yeah, to cover maybe a few points that were mentioned in the community, and I think it's good to highlight them and share some thoughts on how it happened. First of all, the APY is, of course, not solely determined by this pool. There are rewards coming in, which at the moment aren't that high, but as more and more nodes join the network, and Test Grid Phase 2 grows, this will contribute more and more to the APY of the network and that will contribute to the APY that you will see over the long term of the project.
The mining pool is there just basically for this bootstrapping phase, which now is extended to be longer than initially planned. And another important part of the APY is that the more people stake, the lower the APY; the less people stake, the higher the APY. So every time adjustments, maybe through external factors or by changes in the emission rate of the mining pool come to the network, there will be an equilibrium that's disturbed, right? So right now, APY has decreased, so people will start unstaking and that will, again, impact APY, which will increase. So there is a dynamic going on, and at some point there will be an equilibrium where people are happy staking in the network. But right now, as the change was put through a little while back, there has been a pretty big decrease.
So one of the questions that was asked a few times and that I would like to touch upon is why we didn't announce this beforehand. Why didn't we put out a notice a couple of weeks before that this would happen? I think it's a very fair question. Of course, this did cross our minds, but yeah, for several reasons, we did not. And one of them is that we did not know exactly when this would happen. So it was actually postponed a few times. The timing was quite unclear. So announcing it would probably also lead to a lot of speculation, which, again, would lead to changes in the APYs because people would speculate on the staking pool. So we didn't want to cause that.
And also, we don't really believe in announcing news. So we don't announce news for big positive things coming up. We don't think that's a nice and healthy thing to do, but we also don't announce news or speculate on these negative ways. And also, on the first decrease in January of the emission rates, we also did not announce it beforehand because we didn't want to give anyone an advantage in that sense to speculate. And also, what if we had to postpone it or something? That would also create scenarios where people would not be happy. So we decided not to announce anything beforehand.
Of course, when we made the first adjustments in January, we said that there would be another part coming, the second step of this process, which has happened now. But yeah, that's the reasoning behind it.
And one other point that I wanted to touch upon is that some people remarked that they shouldn't be able to adjust their current staking duration that they have. So for example, someone suggested having their unstaking period alongside the APY. So there's many reasons for this too. First of all, it's quite complicated to adjust the staking protocol. Very complicated, actually. And we've only once ever adjusted the staking protocol in its history. And this was actually quite soon after we launched it, because back then, if you unstaked your tokens, you would have to wait for the complete unstake duration of your stake to get anything back. And it was quite a heavy measure, but that's how the staking system was designed and how it was documented and described. But unfortunately, I think a lot of people didn't understand this back then and people would stake for a year, and after unstaking, they would then realize they would not get any tokens back until 365 days later.
So what we did was do the linear unstake release. That's the only change we ever implemented. And we think it was fair because no one was harmed by that. There was no one who got worse from that upgrade. For everyone that staked, it was kind of better that they would get their tokens back gradually after unstaking. So it means if you unstake today, you will get your tokens back every second until you get everything back at the end of the unstake duration.
That was the only update we made in this case. It's a very complicated upgrade to do something like that. And second of all, it wouldn't be very fair to everyone, as, for example, people that unstaked maybe half a year or a year ago would've to wait a full year, whereas people that were unstaked would have to wait a little bit later after that update; they would have to wait half as long to unstake, which isn't very fair to everyone in the system. And there is a lot more nuance and complications there that we have to consider as well.

So updating staking protocols is just not an option in general. Until we get to a more decentralized governance system, I wouldn't expect any updates to ever come to the core staking protocol, which is also separate from the rewards protocol.
But those are just some more context, thoughts, and reasoning behind why these changes were put in place. And the most important message here is that this was an extremely important update for Nosana to be able to grow and get to the point where network rewards will become the majority of the APY in the network. As we're going to discuss in a little bit, we will see some adoption soon on Test Grid Phase 2 as transactions grow. And yeah, we're very confident that with the current staking configuration, the network can last and succeed in the biggest ways. I'll hand it back to you, Rochelle.

Rochelle:
Thank you for that. That was a big update. And now we're going to move on to another hot topic, which is Test Grid. So I will take this over to Sean and have Sean give us some updates and info about it.

Sean:
Thanks, Rochelle, and thanks to everyone for being here today. I see a few nodes in the audience, so a lot of what I'm going to say won't be news to you, but hopefully there'll be a couple of surprises for you. So we started the node component of Phase 2 at the end of April, April 29th, and it was a busy couple of weeks for the whole team at that point, dealing with a lot of troubleshooting, bugs, and onboarding issues, but overall, this process went a lot smoother than it did in Phase 1, so we're happy with the improvements that we made there. But there's always room to do better, and we've got quite a list of things to make better.
A lot of defects have already come out in the first four weeks. Some are minor, some are major. A couple of times, we've had to get all of our nodes to restart. At least once, we've had to abandon the ship for 12 hours until the fix was available. We've fought with Solana timeouts, and we've fought with RPC issues. And that's great. The whole purpose of doing this, the whole purpose of Test Grid is to flush this stuff out.
As Laurens mentioned earlier, we pushed a large benchmark job across all markets two days ago, and that really didn't go so well. But those issues have now been resolved and we started pushing those jobs back out again today into select markets, and they'll be pushed out across all markets in the coming two days.
Now everyone loves statistics, right? So here are a few for you. Since the start of Test Grid Phase 1, back in December, we have factually run almost half a million inference jobs on our markets. Now, our markets are made up entirely of nodes provided by individual members of the Test Grid. So there are no data centers involved, and there are no enterprises. We've got some testers with multiple nodes, but that's the extent of it. And our nodes are earning, every single one of them is earning every day, every minute, and every second. And we're going to eclipse those numbers by a large margin very soon because we're executing almost 70,000 inference jobs per week since the start of Phase 2.
From a node perspective, we've been maxing out with about 550 to 600 nodes online at any one time, concurrently running inference jobs. So this is not, some nodes are getting jobs and others are waiting. When we max out the job requester, which we do quite regularly, we will have jobs being pushed to 500 to 600 nodes concurrently. Now, we had over a thousand nodes register for Phase 2, but for reasons unknown to us, about 30% of them haven't shown up. We don't know why.
So we're not completely satisfied yet. We really want to see what this engine can do. We really want to hit the thousand nodes concurrently running jobs, and we really want to push our daily inference load over 100,000 inferences a day. So the only way we're going to do that is to get more nodes. And to accomplish that, we're going to reopen registration on Monday. So this is for new nodes that are not currently in Test Grid. All newly registered nodes will go on a wait list, and we will onboard nodes in markets where they're needed for nodes that are capable of running the advanced models that we are beginning to test with the benchmarks. We've got a lot of interest from nodes who missed the cutoff for Phase 2, so hopefully we can fill our ranks real quick. So for those of you who've been accumulating parts and GPUs and motherboards, get busy, get your rigs built, and you'll be able to register them on Monday.

We'll be moving forward with a lot of new things over the next couple of weeks. And during June, we'll start to inject external customer project inference workloads into the mix, which Jesse will give us an update on in just a minute. So I'll leave it at that. We've got a few things still to test to make this happen, some of which Laurens touched on earlier, so they're going to be busy to get ready, but we've got huge faith in our engineering team. They work very, very hard to make a lot of things happen for the project and for all of our nodes. So when we say June, we'll make it.
Finally, before I hand back to Rochelle, I would really like to thank all of the Test Grid community members who, like us, have been extremely busy the past weeks with troubleshooting, triage, and helping their fellow nodes when they run into trouble. We really, really cannot do this without you. Your support is critical to the success of Test Grid and to the project, and we're really proud of all of you. We've got a great community here and we're really, really happy that you're here with us. So thank you all for listening, and I'll hand it back to you, Rochelle.

Rochelle:
Thank you so much, Sean. Yeah, it's amazing. I'm sure everyone is stoked right now with the news that registration is opening up on Monday. Now, moving on from Test Grid, we know why we're testing and everything, but how does this fit into the bigger picture?

Jesse:
We've seen some incredible updates today, like Laurens showing us in detail what's been going on in the engineering side and the features that have been finalized or being tested and wrapped up now are all, of course, building up to the essential and really exciting part where we're going to be running production workloads, or at least real-world workloads on the Nosana Network. We're getting closer and closer to that point, Sean just mentioned it. We're even going to grow the number of nodes over the coming months, or at least the wait list so we can be ready to also support increasing numbers of client workloads.
But to go in a bit more detail on that side, there are a few clients that are lined up to onboard onto the Nosana Test Grid infrastructure, and we're very excited for that moment. We're confident at this point that the features are in place to start running these workloads. So yeah, as soon as the things that Laurens demoed now are all rolled out and we've been getting really good benchmark details and analyzed the market's performance and capacity of the grid at this point, we are confident that we can start running some of these workloads in June. So next month.
There will be some more details shared and some announcements made. We will do some AMAs as well with first projects that will be onboarded onto Test Grid. This will start next week, so that's when the first details will be shared. But it'll be great to see the real-world use cases and some real-world adaptation of our grid that we've been working on for such a long time. And that's now finally at the point where it can really contribute and be used.
So I'll leave it at that. So we'll be sharing over the next month more details about what kind of workloads will be run. And of course, nodes will notice that, yeah, workloads will be more diverse and there will be a lot of real world use cases among them.
Before ending this, I would like to make a call to action to everyone. As we've been growing a lot as a community and as a project, we're now truly a global community ecosystem, and we have a global network of GPUs. And at this point, we would love to get more community members involved in actively promoting Nosana, representing Nosana and being generally involved with the project. So we are looking for people to host community meetups around the world and to represent Nosana at different events where they're based. So if you're eager to contribute and would like to participate this way, please reach out. We're really going to roll out and increase these community participations and spread the Nosana world across the world. So here, I'll hand it back to you, Rochelle.

Rochelle:

Thank you so much. Yes. If you guys really want to start getting involved more, please reach out to me in a DM. We have a channel set aside for those of you who want to organize, have meetups, and really spread the word about how cool our project is. So give me a shout.
And, okay, drum roll. It's time to open this up for community Q&A. We're pretty good on time; remember, keep mindful; we've got about 20 minutes scheduled for this now. We don't want to run out of time. I will say this just to head up: we do not announce listings, partnerships, marketing plans, future marketing plans, or anything else we do not discuss and comment on. So please refrain from asking “wen listing," “wen moon," and all of that stuff. I appreciate it. We want some quality questions here, so if you would like to ask a question, raise your hand and request to speak, and I will pop you up here so you can ask.
Okay, we have a member here. We have; it's not showing. There we go. Okay, shoot. Ask away.

Community Member:
Hello, can you hear me?

Rochelle:
Yes, you sound great.

Community Member:
Okay. Thank you for this call. Basically, you covered almost everything I wanted to ask. Only one thing is left. I think it was information that you need to stake some tokens to run the node. And I don't think you covered this question. How much and maybe for different markets, different amount of tokens you need to stake or how will it work? Thank you.

Rochelle:
Thanks. And who would like to answer this question for us?

Laurens:
I can take this question. Yeah. So right now, the necessary stake that you need as a node to join Test Grid is still zero. So at the moment, you do need to have a stake account, and that's created when you start up your Nosana Node, but to not complicate things. Currently, we are requiring zero, no stake, but, of course, it's already built into our smart contracts on Solana that, yeah, a stake at some point is necessary for markets. And it's basically to make sure that the quality of the jobs is good so that there are no cheaters who try to spoof their GPU or just submit fake results back to the chain.
And we will have lots of quality control measures in place. And having stake is one of them, because that means that if we detect a malicious node, like a cheater, scammer, in the system, then we can, in a decentralized way, detect this. And because you need to have NOS staked to join certain markets, it'll mean we can slash bad nodes. The network can basically decide, Okay, this node is like a cheater, so we slash his nodes, so we slash his stake.
So yeah, in the future, it'll definitely be necessary, especially for better markets. You will need to have a stake to become a node, but at the moment, it's not necessary yet. And how exactly this works with how many nodes per market—that still has to be worked out yet, but the core structure is already there in the smart contracts.

Rochelle:
Thank you. Thank you very much. Does that cover it?

Community Member:
Thank you. Can I have one more question?

Rochelle:

Yes, really quick.

Jesse:
Sure, go ahead.

Community Member:
Yeah, this is regarding some GPUs on the market. You have no definition between the RTX A 6000 and just the 6000 other generation new one. It costs the same nodes per second. What's important for your grid to have? Is it performance or RAM capacity or what's the case here?

Sean:
Yeah, there's no scientific formula for this at the moment. This was simply a case where that market was not big enough. So when we have enough A6000s and enough 6000s to make markets of their own, they will be split.

Rochelle:
Thanks for that, guys. Let's open up more community questions. Anyone? Oh, I see one here.

Community Member:
Hi.

Rochelle:
Go ahead and ask away.

Community Member:
Okay, great. Just a quick one. Basically, I think on Solana there's two type of NFT. One is for the nodes, another one, I think, is much earlier, which is the Nosana Burner Phone. Could you share with us any plan for that Burner Phone in the future?

Jesse:
I can take that. You must be an OG, remembering the burner phones. It was a while back. At the moment, there is no concrete plan for the Burner Phones. They do not have a utility right now in the GPU markets or in our GPU network. So at the moment, they don't have a utility and we don't have plans for rolling them out further as of now. The initial idea came from before we went into GPUs and AI, and back then there was a utility for them, but it's no longer very applicable. So yeah, there's not much more we can say at this moment. So if you have a Burner Phone, yeah, that's something nice to have in your wallet and there are no concrete plans, as of now, for the future of Burner Phones.

Community Member:
That's great; thank you.

Rochelle:
Oh, here we go, JB.

Community Member:
Great. This is a bit of a good question. It's a 40,000-foot question, I guess, but the question is what is going to impede the hypergrowth of mainnet? So as a layperson, kind of an investor kind of question, but if you woke up tomorrow and were able to instantly tap into unlimited demand, what do you foresee would be the main

limitation to matching network growth to that demand? Would it be technical bottlenecks? Would it be staffing limitations? Would it be capital limitations? And how would you plan for that?

Jesse:
Great question. I think there's a whole lot that really depends on that type of demand. We're working now with a few clients that will start running workloads soon and we're matching clients with the network capacity we have right now. I think if we look at mainnet adoption and mainnet growth, what we would aim for, what would be like the KPI of network growth, would definitely be inferences done. So we want to just run more inferences. And if that's sort of simple inferences like Llama 3 inferences or different LLM model inferences, that's sort of really well suited now for the network. So if the demand would come solely from that, I think we could scale up quiet... Technically, we would be ready to scale that up just with the number of nodes scaling up. There's going to be some issues here and there.
Then there are a lot of things that still have to be added. There is also demand for running more fine-tuning jobs or small training jobs. There's demand coming from networks that need racks or prompt augmentation systems that have private data or data locality. These are systems we still have to develop as Nosana as we grow. So there are definitely pools of demand we can expand into by adding more features to the network that are planned.
So that would be a limiting factor if demand is more complicated. But when it comes to simple inference jobs, we should soon be at a point where scaling is more important than finding the demand. Of course, there is a whole lot of support needed. And yeah, it's a good question. I think we will find out where the bottleneck is, but it's probably not going to be on the node side. It's the current capacity we have now and the potential growth of nodes that we're seeing is very promising.
So for now, we don't foresee any massive blockers in the short term for growth on that side.

Community Member:
Great. Thank you.

Rochelle:
Thank you so much. Good question. Okay, moving on. Anyone else have any questions? Raise your hand. Here we go.

Community Member:
Hey guys. Can you hear me?

Sean:
How's it going?

Community Member:
Yeah, great to be here. So my question is, so where in the summer is going to get revenue because currently we are getting paid jobs in NOS, and I guess it's from teams treasury, right? And we're now having more than new clients. So how are they paying for the jobs we will be running? They're buying from the market or somehow playing it for the team.

Jesse:
I'll take this one. So yeah, good observation. So clients that when we get to Mainnet and clients start running real jobs and purchasing GPU power for their inferences, they will buy that from the market. So they would buy those NOS from the market to purchase that GPU power. So that's where a lot of utility and value come from. So buying NOS tokens gives you GPU power and gives you AI inferences. That's the value and the utility of the token is that you can basically run AI using the NOS token. So it would be from the market.

Community Member:
And how is Team Nosana getting benefits from it? In revenue meaning.

Jesse:
Yeah, there is not a direct revenue model there for the team. The team has done a funding round before, a while back, and has set out to build this network. And as we go forward, there might be professional services needed and revenue streams possible for commercial products on Nosana that's very foreseeable, and that might be potential future revenue streams, but there is not a revenue stream directly attached to the market. So yeah, that's basically how it works. And as it's an open, decentralized network, anyone will be able to build professional services around it. But if this network truly grows and institutional adoption and the inferences really grow out, then there are a lot of opportunities for monetizing services around the network, which is where Nosana as a team can also generate revenue streams from in the future.

Rochelle:
We have time for one more question. So here we go. Tesseract, you are on stage. Oh my gosh. Nice Avatar. Okay, go ahead and unmute yourself and ask.

Community Member:
All right, so JBSol touched on the important topic of demand and supply. I understand that, from the supply point of view, we are fine right now. So my question is about how does Nosana plan to drive adoption among developers and businesses? I am not asking about the specific examples here, of course.

Jesse:
Yeah, great topic. Great. These are the right topics to talk about. I think these are the important ones. Network success will mean more inferences and more inferences should come from more clients. And so yeah, the growth of clients is really important. We think it's definitely more important than the number of nodes, how many inferences, and whether we're running on the network.
Right now, we're in the Test Grid phase. We're aiming for smaller experimental projects to work with. We're not going to run mission-critical production workloads on a Test Grid, not a truly tested network. So automatically, we're looking at projects that are doing a lot of R&D, like innovation budgets and projects that are running newer, cutting-edge models. So in that whole sector of small businesses, startups, and medium-sized businesses with experimental budgets, as well as university groups and research groups that want to get access to specific arts and scale in different ways.
So those are the market areas that we're looking at because that gives us the opportunity to, first of all, showcase and run the latest models, and second of all, it gives us the opportunity to make sure that we have all the technical features in place to start serving more production, like more mission-critical workloads and at higher scales.
So I think to answer your question, we're still looking in that first corner, especially as we're going now into the Test Grid demand onboarding phase. We're looking more into working closely with these partners. We do have a really nice collection of GPUs already that are accessible in a convenient way for companies, so the value proposition is already pretty strong. It's just that we need to mature this technology and make sure it is ready to actually grow to more web-to-real-world workloads soon.

Community Member:
Yeah, that's a great answer. Thank you so much.

Rochelle:

Thank you so much. And thank you for that answer, Jesse. So that was our last question and it was a good one and a great answer. Positive. I just want to remind everybody that, of course, this is monthly, so our next call will be the last Thursday of next month. And I'm looking forward to having you all join us again. Stay tuned. We had that big news about registration opening back up for Monday, so pay attention to our socials and our Twitter. We are also hiring and we'll be posting more jobs, so follow us on LinkedIn for those things. And yeah, thank you everyone.

Laurens:
Thanks, everyone.

Sean:
Thanks, everyone. See you in the Test Grid chat.