

Speaker 1 ([00:00](#)):

This is episode 617 of the AWS Podcast, released on August 31st, 2023.

Hawn Nguyen-Loughren ([00:09](#)):

Hello, everyone, and welcome back to the Official AWS Podcast. We got some exciting updates. Amazon Managed Blockchain announces Access and Query. I'm Hawn Nguyen-Loughren, also known as Hawn Solo, your friendly neighborhood co-host of the Official AWS Podcast, and I'm joined by John Liu. Thank you for joining us.

John Liu ([00:27](#)):

Yeah, thank you for having me, Hawn.

Hawn Nguyen-Loughren ([00:29](#)):

Awesome. So tell us a little bit about yourself and what do you do at Amazon Web Services.

John Liu ([00:35](#)):

Yeah, I'm the head of product for Amazon and managed blockchain, or Web3 and blockchain services more generally. I've been at AWS for two years now, really looking after product strategy for public and private blockchains. Prior to this, I spent five years looking after product strategy in the public blockchain space, and before that, 18 years in traditional finance, a real good mix of enterprise use cases and innovative blockchain use cases.

Hawn Nguyen-Loughren ([01:01](#)):

Nice. So we're here to discuss a really cool service update called Amazon Managed Blockchain, AMB, which announces Access and Query. Now, this expands AMB's capability to help developers seamlessly interact with public blockchains so that they can build scalable applications quickly and securely, which is quite fascinating. So let's just start with the basic. So what is Amazon Managed Blockchain, AMB, all about?

John Liu ([01:28](#)):

You hit the nail on the head. Working with technology always happens to have a lot of undifferentiated lifting, especially when it comes to something new where you have these decentralized nodes, you have different type of data that you have to consume. It's almost like a new transaction layer that people have to interact with. So what we're trying to do is give developers a comprehensive service so that they can build much more easily and just focus on the application layer, for example, as opposed to worrying about, "Hey, is my server up and running?"

Hawn Nguyen-Loughren ([02:00](#)):

Yes, I truly love the underappreciated heavy lifting that we provide here, especially around blockchain. So what has this changed lately in the realm of customer experience?

John Liu ([02:10](#)):

So AMB has been running for three years when you include our private blockchain feature, and so we heard a lot from customer signals, "Can you make our job easier by introducing more blockchains and also move up the stack?" So what we've done really is to address those exact two needs. We announced

Access, which is going to be a serverless access to blockchain networks, very imaginative with our names. And we also announced Query, which actually helps developers work with data much more easily.

Hawn Nguyen-Loughren ([02:44](#)):

So let's just double click on that a little bit more. Can you tell us a little bit more about that? How do we make it easier for them?

John Liu ([02:51](#)):

Absolutely. So with AMB's Access serverless offerings, now what happens is, you can instantly call the standard remote procedure calls. Think of these as standard functions that exist on a node, and they let people send transactions to the blockchain or read the latest state of the blockchain. Maybe it also helps to paint the picture. Without these nodes, you really can't do anything with the public blockchains out there. All that information is there, all these transactions are there, but if you're not plugged into that network, you can't do anything. So this is what AMB Access is all about. We give the scalable servers, we give that consistent RPC that you can call so that you know it's there and it's instantly available to you. You don't have to go worry about downloading the data, syncing the node to the blockchain before you can use it. This is all very low level undifferentiated lift. That's what AMB Access helps solve.

([03:49](#)):

Now, AMB Query, on the other hand, addresses the data needs. We'll dive probably a little deeper into this, but I think a lot of people's conception is, "Wow, blockchain has this immutable dataset, so why do I need help getting this data?" It turns out, the way that blockchain stores data is optimized for decentralization, and that means that it has to be access from everywhere and it has to be stored in the optimal size. It gets very expensive to store this data everywhere at all times. That makes it really good for blockchain decentralization, but very difficult if you're trying to build an application and trying to get a query view. "Get me my balances from the beginning of time until now," that makes it actually a much more challenging task.

Hawn Nguyen-Loughren ([04:33](#)):

Yeah, interesting about the decentralization, because it makes it a little bit more flexible, but also in terms of customer experience and trying to get all that data, I just want to understand a little bit more about that. So what are some of the key consideration for customers in the overall experience with this architecture?

John Liu ([04:49](#)):

When it comes to what customers want, they just want to get this data, they just want to connect to the blockchain. Simple, right? No nodes, no special code that they have to configure, no upgrades or patches that they have to maintain just to stay connected to the chain. They want to make sure that this connection's reliable as well. Just like any type of enterprise application, you want to make sure that the nodes you're connecting to are going to be there. So typically, what happens is they set up clusters of nodes and they could be in different regions, especially as we're starting to get into these enterprise applications from, let's say the big banks, they usually have to have some type of failover. One region goes down, they have to make sure that they can still send and read transactions by connecting to another region.

([05:33](#)):

Finally, you got to deal with a scalable solution. Now, in the early days of blockchain, a lot of developers could run this themselves because they're innovating, they're pushing the boundaries. But when you start getting users, 20,000, 30,000 users, all of a sudden now you have to make sure that the nodes that you're maintaining and supporting these transactions can actually process all this traffic. And because of some of the low level nuances of taking a node and having to sync that node to the blockchain in order to actually make it usable, this is actually a much more difficult task.

Hawn Nguyen-Loughren (06:08):

Gotcha. I'm a little interested on the developer's experience. So say I'm the developer, I would like to plug into the AMB, with the query and et cetera. So how would that work? How would that look from the developer's perspective?

John Liu (06:21):

From developer's perspective, it's now super simple. You basically choose your endpoint, you hit that endpoint, and you're done. So for example, if I want to just connect to the Bitcoin network, what I would do is I would say, "Great, let me go to your Virginia region, connect to this, select the Bitcoin endpoint, call the remote procedure, call that I actually need." And that's it. I don't need to go and configure the instance of the server. I don't need to say, "Hey, is this supposed to be a more powerful server or less powerful server?" It's done.

(06:52):

It's the same thing with Query. If I want to go and get the latest balance of a particular wallet or a particular transaction, I just call an API, I'd give the chain ID, which tells me what blockchain do I want to pick the data from, and then submit, for example, the address that I'm interested in. And the address is the account. You can think of that as the equivalent of a bank account number where you want to pull a bunch of information from. It really is very simple compared to the traditional ways of having to, we discussed earlier, all that extra action that you need to take.

Hawn Nguyen-Loughren (07:24):

Gotcha. So it's a pretty simplification of the view of API, right? You're just enabling an API and hiding all of that extra burden of having to run the server and just leveraging the service. Is that right?

John Liu (07:35):

Absolutely.

Hawn Nguyen-Loughren (07:36):

That's awesome. So why did you choose to start with Bitcoin and Ethereum?

John Liu (07:41):

When we look at the way cryptocurrencies have evolved, I mean you hear the news, you look at it, it's almost like you see Bitcoin and Ethereum very often. This is what our customers have been asking for when it comes to the cryptocurrencies they want to support. And we hear from customers, well, Bitcoin, it's the largest cryptocurrency by market cap. It's the oldest cryptocurrency, and it's almost, as Larry Fink said recently, Larry Fink, the CEO of BlackRock, recently said, "Bitcoin is like an international asset. It's an asset that a lot of traditional finance customers are comfortable with dealing." Likewise, we chose Ethereum because it is the programmable network. All the smart contract development that you've

seen in the last three years. Whether it was decentralized finance, whether it's the non-fungible tokens, NFTs, that's all possible because you have this programmable layer on top of the decentralized network. And the biggest one of them, we'll call it the grandparent of them all, is going to be Ethereum. So these are the two chains we chose to support.

Hawn Nguyen-Loughren ([08:44](#)):

Gotcha. Based on what we're seeing in the market and the customer between those Bitcoin and Ethereum. So that's super cool. That's awesome. So in terms of what we're trying to solve, what are some frustration and challenges that organization are looking to resolve with this?

John Liu ([09:01](#)):

Yeah, I think we've covered a little bit of this, but it is worth repeating the challenges that they have. A lot of their challenges is around trying to get that reliable and scalable connection point so that they can be able to transact, basically submit their transactions to the blockchain. Anytime you want to do a transfer of a cryptocurrency, for example, or if you want to deploy a smart contract, you can think of that as a transaction that you're submitting into the blockchain network.

([09:28](#)):

Now, these go through these nodes. The Access is the one that abstracts away these servers or nodes. And the nodes, you have to make sure that they're connected to their blockchain and they have the right software updates. If they're not on the right patch, they might end up connecting to the wrong chain. And that's very high risk, because when you have a connection to the wrong chain, you might be submitting transactions to a chain that's no longer recognized by the main network and you might think that your transaction's gone through, or you might think that your processing deposit and transfers for someone else that's legitimate, but it's actually not being recorded on the main chain. That's got financial implications to it. So this is a pretty big challenge that a lot of developers have to work with.

([10:17](#)):

Secondly, when you're trying to scale this infrastructure, you can't just have your servers there and spin them all up and they're ready to go. What typically has to happen is, you have to sync these nodes to the blockchain. So you take a server, you listen to the blockchain from the beginning until the latest transaction, so it's got the full history on that node and it can connect and transact with the latest state of the blockchain. And that can take, depending on the blockchain, could be a few hours, it could be a few days, and sometimes that syncing process might actually fail. Just like any server, any type of software, the connection breaks or something, and you're forced to restart again. So you can imagine when you're trying to scale that up and down, whether it's you trying to submit transactions or read data, there's this almost big lag that you have to deal with as a developer.

([11:06](#)):

I think the final one, the one that we were really excited about helping our customer with with AMB Query, was the difficulty in how developers have to face when they're actually trying to get this data. When they're actually trying to get the data, it's not just about downloading the information from the blockchain itself, they also have to call smart contracts to get additional information. Sometimes they have to parse the smart contract logs to get what we know as internal transactions. They have to stitch all that together in order to then index the data and serve it in a fast way so the applications can get the data they need in a reasonable amount of time.

Hawn Nguyen-Loughren ([11:46](#)):

I really love the way that you're taking away that labor-intensive processing and reducing losing transaction and having to re-sync anything. So definitely gives that lift and helping developers focus on the innovation and what they're going to do. But I just like to understand a little bit more about how does AMB Access and AMB Query solve for this?

John Liu ([12:07](#)):

So let's start with AMB Access. So you mentioned earlier, it's all about this serverless, scalable, and secure access to the multiple blockchain networks. And we're starting this serverless offering with Bitcoin, which compliments AMB's existing dedicated server offering for Ethereum. And perhaps we can explain why there's a serverless offering and then there's going to be a dedicated server offering as well. Now, serverless offerings are really good for general application usage because that endpoint is there for you to instantly access. You don't have to wait, you don't have to configure any nodes. However, for certain type of enterprise customers, and now usually it falls into the finance ones, they may want to have some custom requirements for security and compliance. For example, they're the only ones who can end up being sending traffic to a particular node, or they might want to configure certain type of resiliency based on their regional exposures. Now, in such a case, a dedicated offering could be more useful. So we're still keeping the dedicated offering for Ethereum. This is kind of the choose your own chain, choose your own adventure based on what your use case is going to be.

([13:17](#)):

Now, the second thing that we're helping people with is, you don't have to worry about all that resiliency and scalability, as I mentioned earlier. You don't have to maintain multiple nodes, you don't have to patch them at all. We take care of all the patches for you. And the patches could be timed patches, scheduled patches, like the big upgrade that took place over the last September with Ethereum being upgraded from proof-of-work to proof-of-stake, and even some of other chains updates like Polygon is about to undergo a pretty big update for their blockchain technology as well. It can also go towards unplanned patches. There've been a few unplanned patches in Ethereum. Usually that has to do with some bug that has been found. It needs to be fixed. And if you're not on the right patch, as I mentioned earlier, you might end up being connected to the wrong chain, and now you're processing and dealing transactions which aren't "real."

([14:15](#)):

And finally, moving on, wrapping up with Query, the data side, we earlier discussed the challenges on getting that data already. What we do is we abstract all of that, downloading that, indexing, deparsing of the smart contracts, checking for my standard events, and we serve it to you through these APIs that instantly deliver you the data that you're asking for. In this case, it's very much like a traditional application data, API. If I'm building a banking application today, I'm almost used to just say, "Great, let me connect to the open banking API, let me connect to this bank account." A hidden API returns to me a particular account address, and it returns to me the balance that I have, and now you get the same experience using AMB Query.

Hawn Nguyen-Loughren ([15:02](#)):

Yeah. I can see how that really helps, especially if you're in a decentralized working framework, if you will. Definitely can see that giving that undifferentiated heavy lifting. So that's super cool. So going a little bit more about what you're describing, can you give me some more examples of that?

John Liu ([15:20](#)):

Sure. I like wrapping it up in a banking application, because I think a lot of people are familiar these days with the mobile banking application, the experience feels quite familiar. You open up your phone, you log in, and what's the first thing that you do? You end up seeing the current amount across your checking account, across your savings account. How much dollars do you have? Maybe you've got an investment account over there, you click in and you say, "Oh, I've got different types of stock." Maybe it's Amazon stock, maybe it's something else. But this is an experience that's very familiar to people.

[\(15:55\)](#):

And then the next thing they want to do is say, "Well, what have I done in the past? Have I been doing some trades? Has there been some deposits and transfers coming into this account?" Then they like to dig in, or at least I like to do a lot, is say, "Well, how has my balance been changing over time?" This just kind of gives me a track of my purchase habits, for example. Have I had a pretty big spending amount back in March or something and my dollar amount has dropped to almost zero? I like to see that type of information. The general sense there is, people expect to have a good view of their current information broken down by the assets that they have, and they also expect to have that view across the history.

[\(16:34\)](#):

So let's draw the analogy now to cryptocurrencies as the equivalent of that. What do you want to do when you come into a cryptocurrency wallet? First thing you want to do, you want to see the different assets that you currently hold. And these assets could be Bitcoin, it could be Ether, it could be stablecoins like USDC, and it could also be non-fungible tokens. Non-fungible tokens definitely got its time in the sun in 2022 and late 2021, but it continues to be used by a lot of big enterprises to create these new utilities. So I think non-fungible tokens, we're going to see more and more of that application, by the way.

[\(17:13\)](#):

So anyways, you see a balance of different cryptocurrencies and NFTs and you want all that to load within just a few seconds. You don't want to sit there and wait. And then next thing you do, you want to go ahead and start doing deposits and withdrawals, because it's not just some static view of cryptocurrency. These are assets that you can actually transfer or pay with. So you have your deposit withdrawal and transfer actions from the cryptocurrency application. In order for you to do that, you actually need to connect to the node, to access, because you now call these standard RPCs, the standard APIs in order to create transactions, sign the transactions, and then send the transaction. So you can say process a withdrawal or a transfer from your account.

[\(17:58\)](#):

And in order for you to see all this information that I was just talking about, the current balance, the transaction history, instead of you as a developer going in there and downloading all the data and indexing all of it yourselves, you now just call an endpoint, the query endpoints, submit the chain ID, submit the address, and then all the information returns to you in sub-seconds provided, of course, it's finalized information on the chain already.

Hawn Nguyen-Loughren [\(18:24\)](#):

Just double down a little bit more on the non-fungible tokens because I've seen that in reference to authentication or security and et cetera. Just like to hear from you, what are you seeing from the customers of how they're using that?

John Liu [\(18:36\)](#):

Yeah, non-fungible tokens, like blockchain in general, there's a wide variety of applications that you can get to, but the most, we'll say near term successful use cases actually goes around the big retail or maybe media enterprises unlocking utility for their consumers by holding these non-fungible tokens. I'll give you some good examples from public companies like Starbucks or Nike. They're doing a very good job of engaging with their consumers in this new digital world, this new digital medium, and unlocking the benefits for the holders of their non-fungible tokens. For example, Nike, you can go ahead and get yourself a digital sneaker. Once you have that digital sneaker, you might be able to go to an in-store event. And because you have this access pass, you're granted rights to actually go to a limited time opening, or you can map it to a physical sneaker that you can get by holding this digital twin.

[\(19:44\)](#):

Likewise with Starbucks, they use it as a consumer engagement token. So you're familiar with your Starbucks points, your Starbucks loyalty points. You go and buy yourself a bunch of lattes, you've earned a bunch of points and you might get your next one for free. Now it's being linked onto the blockchain. That transaction, the customization is actually going to be tracked within the non-fungible token. And the interesting part is, one could potentially take this non-fungible token as a consumer and sell it. Here's where the whole decentralized new form of value exchange start getting interesting from a traditional centralized solution, consumer can now participate in the value more and can also transfer the NFT to other people if they choose the options there for them to choose.

Hawn Nguyen-Loughren [\(20:33\)](#):

I'm never going to look at a Starbucks the same anymore. I'm going to think about non-fungible tokens every time I drink my latte. That's awesome. So last question that I have, the most important, is how do customers get started?

John Liu [\(20:45\)](#):

Yeah. The best place to get started is just hit on over to the Amazon Managed Blockchain website at AWS and click on getting started. And from there, there're going to be a few examples for you to walk through, something as simple as what we discussed, like go ahead and create a Bitcoin transaction, send that transaction, and read it by using AMB Query to something that's more full on. A lab would actually take you to creating a non-fungible token marketplace and be able to send, create, mint these NFTs using AMB Access.

Hawn Nguyen-Loughren [\(21:18\)](#):

Well, that's super cool. I'm going to have to check that out for sure. So thanks, and we'll definitely have those links in the show notes as well. So John, thank you so much for coming on the podcast today.

John Liu [\(21:29\)](#):

Thank you for having me.

Hawn Nguyen-Loughren [\(21:30\)](#):

And as always, we love to get your feedback. There is a link in the show notes to submit feedback. And until next time, keep on building.