

WPRC Summary

Stablecoin session

Talk about stablecoins usually opens with market size. Market cap is now around \$300 billion, and onchain transfer volume gets compared with traditional payment networks. This session started somewhere else: what friction do stablecoins actually remove?

We began with USDT. Stablecoins first took hold in the crypto market, not in everyday payments. Traders could buy and sell assets in dollar terms inside an exchange and move funds between exchanges without a bank account in the loop. That alone was worth a lot.

So the question shifted to where stablecoins solve a problem that is hard to solve any other way. One example was trade with countries that have little access to the traditional financial system. When a bank transfer is hard to send at all, or has to pass through several correspondent banks, a stablecoin is more than a small fee saving. It can be a payment route that didn't exist before.

Korea looks different. Domestic transfers are already fast and cheap, and cards are everywhere. Companies also can't freely accept stablecoins and convert them to won yet. Crypto trading and cash-out for corporations are opening in stages, but an ordinary company is still far from being able to take stablecoins, book them, and convert them to won as a matter of routine. The technology works; there just aren't many places in Korea that need it yet.

Choosing a stablecoin

Different users judge this by completely different criteria. Companies care most about how well the peg has held and whether the reserves and redemption process can be trusted. Retail users care more about where they can spend it and how much they earn in rewards. For people moving large sums, liquidity matters most.

On- and off-ramps

However convenient stablecoins are, getting the money back into a bank account eventually means going through a licensed financial institution, because anti-money-laundering and know-your-customer checks have to happen somewhere.

Off-ramps also charge fees, which is why people want to skip them altogether. One way is to link a stablecoin balance to a card that works at ordinary merchants. The user never converts coins to won or dollars; settlement happens behind the scenes when they swipe. These stablecoin-linked cards are growing quickly. What's interesting is that they connect stablecoins to the card networks rather than compete with them.

Won stablecoins

This was the hardest topic. Korean consumers already have fast transfers and convenient cards, so it's hard to see why they would swap won for a stablecoin. That doesn't make a won stablecoin useless.

A company that earns revenue in dollar stablecoins still has to pay salaries and taxes, so at some point it needs won. If dollar stablecoins are everywhere but there's no onchain asset to convert them into won, sellers and intermediaries keep paying the FX costs. The first real demand for a won stablecoin may come from settling dollar stablecoins, not from convenience-store purchases.

Takeaway

The hard part of stablecoins may no longer be the technology for issuing them. It's figuring out where and why people would use them. Where there's no alternative, they're already in use. In Korea, where traditional finance is fast and cheap, that reason still has to be built. So the next question is less about building a better stablecoin and more about who will build an economy where people can keep spending the stablecoins they receive without routing them back to a bank account.

AI agent payments session

We broke an agent's purchase into stages and spent most of the time on what happens after it finds the product. Agents are already fairly good at searching and comparing. The hard questions come next: who gave the agent permission to spend money, how much it may spend, and who is responsible when something goes wrong.

Agent identity

What is an agent's identity? It could be the company that built the service, like ChatGPT. It could be a delegate acting for the user. Or each running agent instance could be treated as a party in its own right. The industry today mostly leans toward the second view.

Google's AP2 records what the user authorized as a Mandate. If the user is present at checkout, they can review the final cart and sign it. If they want the agent to buy on its own while they're away, they can sign conditions such as price or timing in advance. The point is less to establish who the agent is than to prove what it was authorized to do.

Visa's Trusted Agent Protocol (TAP) comes at a similar problem from the merchant side. A merchant wants to know whether a request comes from a properly registered agent or a malicious bot, and which user it's acting for and why. TAP carries those signals in signatures. That still doesn't make an agent instance an independent identity. When a session restarts or the model changes, there isn't yet a good basis for treating it as the same economic actor.

Disputes

Identity matters because of disputes. When a user later says "I never bought this," the merchant needs something to show the transaction was legitimate.

This is where authentication and authorization part ways. Asking a person to confirm every payment is closer to authentication. Letting the user set conditions up front and having the agent act within them is closer to authorization. For agent payments to scale, they'll have to move toward authorization.

Most card dispute rules, though, were written for people paying themselves. Under Visa's CE3.0, for example, a merchant answering a "this wasn't me" claim can present two earlier legitimate transactions and show that the IP address, device ID, user ID, and shipping address match.

Once an agent sits in the middle, those signals may not survive. The IP may belong to the agent's server, and the consumer's device data may never reach the merchant. Newer protocols are starting to carry this information explicitly; Visa TAP defines a way to pass the consumer's IP and device data in the Payment Instruction. So the issue isn't simply that taking the human out erases the evidence. It's what should replace the evidence that a person paying in their own browser used to leave behind.

What blockchain adds

Onchain payments can be conditional and can carry spending limits, but much of that can also be built on existing card and account systems. The session gave two reasons onchain still matters: programmability and interoperability.

If an agent has to pay while the person is asleep, the person can't decide each time, so a policy has to decide instead. Rules like "spend no more than \$100 at a time," "pay only this merchant," or "use these funds only within this pool" can sit close to the money itself onchain.

Crossing borders changes the picture too. Within one country, existing payment networks may work fine. Once you deal with foreign merchants and several currencies, you run into which countries a payment provider supports, overseas entities, settlement currencies, and FX spreads all at once. There, a stablecoin is less a payment method than a shared asset linking different countries' financial systems.

Credit and consumer protection

Onchain payments don't solve everything. Credit and consumer protection remain largely unsolved.

Most onchain payments so far are prepaid: you spend money you already hold. The card payments people are used to run on credit. In Korea, after the Asian financial crisis, the

government introduced income tax deductions and other policies to encourage card use. Besides boosting spending, the goal was to move cash transactions onto cards and make taxation more transparent. A prepaid wallet can't easily take over that role.

Refunds are similar. Many bank and card restrictions feel frustrating, but a lot of them exist so money can be returned when something goes wrong. Blockchain sells finality and non-custody: once a transaction is confirmed, nobody can reverse it at will. For a consumer, that same property can be a weakness. Which brings back the question: when an agent buys the wrong thing, how do you get your money back?

Takeaway

Agent payments aren't held back by missing technology. Agents can sign, take delegated authority, work within limits, and move money conditionally. The remaining problem is that payment systems were built on the assumption that a person pays. When an agent takes that person's place, the authentication signals and dispute evidence that used to appear on their own have to be produced some other way.

That's why AP2's Mandate and Visa TAP are interesting. They don't add another payment method; they try to keep signed evidence of what the user authorized and what the agent actually did. Agent payments may not really open up when agents become technically able to pay. They're more likely to open up when card issuers, banks, merchants, and payment networks start accepting those signed delegations and transaction records as evidence in real disputes. Until then, asking a person to confirm once before payment isn't an awkward stopgap. Under today's dispute rules, it's a reasonable default.