

## Interview Series

# Kent, Optimum: RLNC, Memory Layer, and Scalability

Apr 16, 2025

**Min Jung** | Research Analyst

[miniung@prestolabs.io](mailto:miniung@prestolabs.io)

**Jaehyun Ha** | Research Analyst

[jaehyunha@prestolabs.io](mailto:jaehyunha@prestolabs.io)

## Contents

1. Let's start from the top — who are you, and what's the origin story behind Optimum?
2. What was the driving force behind building Optimum — and what problem are you trying to solve?
3. Could you explain what RLNC is in simple terms, how it works, and how it improves the blockchain ecosystem?
4. There have been multiple attempts in the past to improve block propagation speed and state storage issues. Why did you feel those approaches were insufficient?
5. What does it take to integrate Optimum into an existing blockchain?
6. Recently, there's been growing interest in scaling blockchains not through changes to the consensus layer, but by optimizing the infrastructure layer — especially the network layer. In your view, what are the strategic advantages of this approach?
7. Blockchain and crypto aren't just about technology. How are you working to build a strong user base and community around it?
8. What are some of the best Twitter accounts or articles to follow on related topics?
9. Looking ahead: what's the long-term vision for Optimum?
10. To wrap it up — favorite CT account, and any closing alpha for the users out there?



As blockchains evolve into a global decentralized computer, one major bottleneck remains: memory. Today's networks struggle with slow data propagation, inefficient storage, and latency that holds back everything from validator performance to real-time dApps.

Enter Optimum, co-founded by MIT professor and RLNC inventor Muriel Médard. Optimum is the first decentralized high-performance memory layer for the world computer—designed to supercharge any L1 or L2 with faster block propagation, more efficient storage, and real-time access to data. Whether you're a validator, protocol team, or dApp developer, Optimum unlocks speed, scalability, and better economics across the stack.

To understand how Optimum is redefining infrastructure for Web3, Presto Research sat down with Kent, Co-founder and COO of Optimum, to talk about what RLNC is and what building "memory for the internet of blockchains" really means.

---

## **1. Let's start from the top — who are you, and what's the origin story behind Optimum?**

I'm Kent Lin, Co-founder and COO of Optimum. I was deeply involved in the early days of blockchain through venture capital and startup incubation across US & Asia. Over time, I witnessed a recurring challenge: as blockchains (the "world computer") evolved, a critical component was missing—a dedicated memory layer to store, access, and update data efficiently.

Working in VC and advising numerous blockchain projects, I saw firsthand how the current infrastructure relied on makeshift solutions like gossip protocols and full-node replication. These approaches, while initially sufficient, quickly revealed their limitations in terms of speed, scalability, and cost-efficiency. I was in the process of getting my MBA from Harvard business school when I met Muriel Médard, her groundbreaking work on Random Linear Network Coding (RLNC) resonated with me. Combining her decades-long MIT research in high-performance data encoding alongside Kishori Konwar with my industry experience, we recognized that we could build a decentralized, high-performance memory infrastructure capable of powering blockchains at the level of traditional high-performance computers.

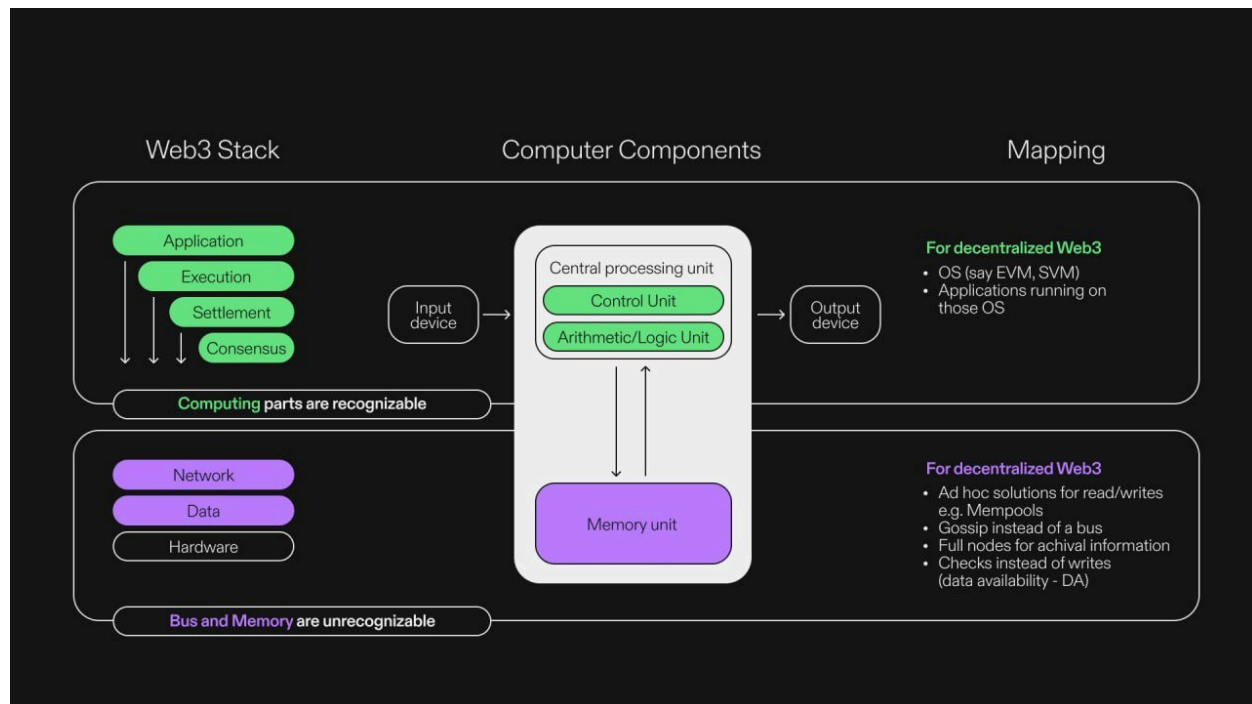
This understanding presented an incredible opportunity to improve Web3 as a whole. As a team, Muriel, Kishori, and I stopped what we were doing to go all in. I dropped out of Harvard and Muriel took a leave of absence from MIT. This is the background of how Optimum was born, a project founded on the vision of transforming how decentralized systems handle data, enabling real-time state updates, and paving the way for scalable, efficient Web3 applications.

## 2. What was the driving force behind building Optimum – and what problem are you trying to solve?

The driving force behind Optimum has stemmed from Muriel Médard's work in network communications. Her pioneering development of RLNC has optimized data propagation in cloud, industrial networks, aviation and wireless networks. When we turned our attention to Web3, we observed that its distributed network was burdened with inefficiencies: a convoluted, makeshift memory approach that leads to slow data propagation, redundant transmissions, and bloated state storage. Adoption of blockchain technology is rapidly progressing, signaling to us the urgency and opportunity of solving its bottlenecks.

Rather than applying incremental fixes, we saw a transformative opportunity: to build a decentralized high-performance memory architecture for Web3—much like the Random Access Memory (RAM) and memory bus that power traditional computers. RLNC is the mathematically optimal solution to resolve the critical bottlenecks of Web3.

**Figure 1: Web3 Becoming the World Computer**



Source: Optimum

## 3. RLNC (Random Linear Network Coding) is the core innovation behind Optimum. Could you explain what RLNC is in simple terms, how it works?

RLNC, or Random Linear Network Coding, is a cutting-edge approach to data encoding that underpins all of our work at Optimum. In simple terms, instead of sending raw data packets, RLNC mixes them using random linear equations. This means each packet becomes a coded

representation of the original data, and any sufficient collection of these packets can be used to perfectly reconstruct that data—even if some packets are lost along the way.

*\*Editor's Note*

*RLNC is a smart way to send data across networks. It splits a file into small pieces, mixes them with math, and sends these mixed pieces through different paths. Even if some pieces get lost, the receiver can still rebuild the original file as long as enough pieces arrive. This makes RLNC especially useful for tricky networks like wireless streaming or cloud storage, and ideal for dynamic, distributed computing systems.*

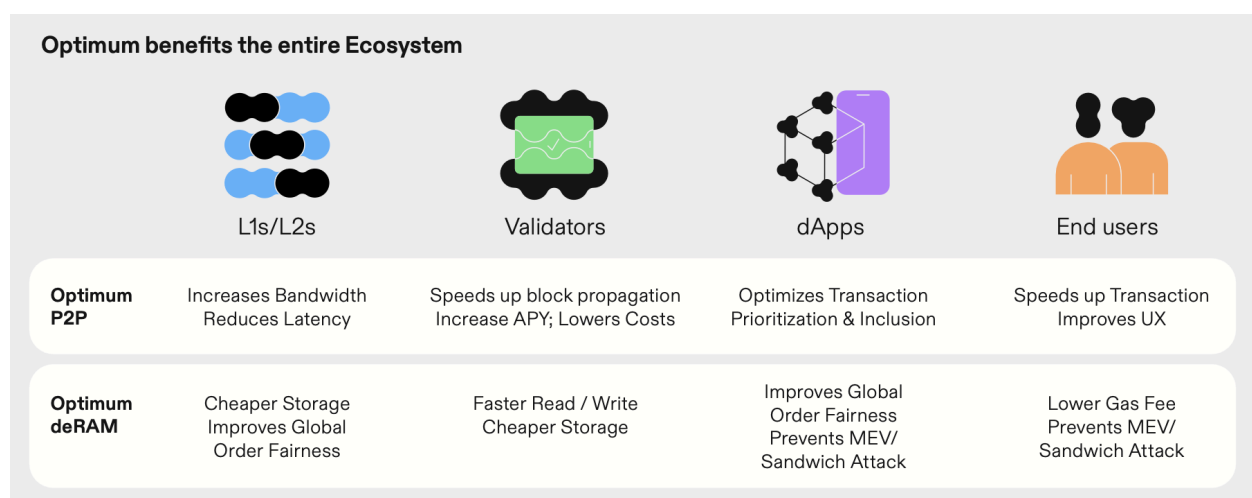
To learn more about:

[Introduction to Optimism with Muriel Médard](#)

### 3.1 How does it improve the blockchain ecosystem? Specifically, how do these improvements affect validators, dApps, and end users across L1s and L2s?

- *For validators:* Accelerated data propagation, lower operational costs, higher APY and MEV income
- *For L1 and L2 blockchains:* Faster block propagation, reduced bandwidth consumption, and optimized storage
- *For dApp developers:* Improved transaction relay and prioritization, enabling latency, throughput, and cost-sensitive apps
- *For end users:* Faster transactions and more responsive interfaces improve user experience. for trust and verifiability in decentralized systems.

**Figure 2: How Optimism Improves the Blockchain Ecosystem**



Source: Optimism

#### **4. There have been multiple attempts in the past to improve block propagation speed and state storage issues — for instance, through EIP-4844, state pruning, or optimized P2P protocols. Why did you feel those approaches were insufficient?**

While initiatives such as EIP-4844, state pruning, and various optimized P2P protocols have provided incremental improvements in block propagation and state storage, they only address symptoms rather than the root cause. These methods work around the existing issues—reducing blob costs or cutting redundant data—but they don't fundamentally resolve the inefficiencies inherent in the current memory architecture.

Our approach with Random Linear Network Coding (RLNC) is fundamentally different. RLNC doesn't just trim the edges—it rethinks data propagation and storage from the ground up. By mathematically encoding data into packets that can be reconstructed from any sufficient subset, we eliminate redundant transfers and ensure predictably fast, efficient updates. This creates a true memory layer for blockchains, overcoming the scalability challenges that prior incremental solutions have left untouched.

#### **5. What does it take to integrate Optimum into an existing blockchain? What are the potential friction points for adoption, and how did you design around them — particularly regarding your decision to offer a permissionless, API-based interface?**

Integrating Optimum is designed to be as seamless and low-friction as possible. Our permissionless, API-based interface allows node operators from any chain to plug our solution into their existing infrastructure without any disruption to the core consensus mechanism. This is how we do it:

##### *No Extra Hardware Required*

- Validators benefit simply by publishing to and subscribing from our OptimumP2P network using our API. For those who want to play a more active role in strengthening the network, running a lightweight contributor sidecar is optional, but it's not a requirement for accessing faster data propagation.

##### *Backward-compatible*

- Any given subset of validators can opt in to use Optimum; the underlying blockchain consensus remains unified. Validators using Optimum enjoy faster data delivery, while non-participants continue as usual, ensuring full network synchronization without compromise.

##### *Modular and Non-Intrusive Deployment*

- A validator is essentially a machine running a specific software (aka. client) for a given chain.

Optimum is designed as a side car, compatible with existing client softwares. In other words, Optimum is built to complement, not replace, existing protocols. You simply add our Flexnode software (or use our API to enhance performance).

Our approach makes it easy for any validator or node operator to leverage the benefits of high-performance, RLNC-powered memory.

**6. Recently, there's been growing interest in scaling blockchains not through changes to the execution layer, but by optimizing the infrastructure layer — especially the network layer. In your view, what are the strategic advantages of this approach?**

*\*Editor's Note*

*Lately, the blockchain industry has been buzzing with a shift toward the infrastructure layer. For years, the focus was on application-layer scalability solutions like Layer 2, sharding and consensus algorithms, but those areas have been thoroughly explored, leaving limited room for breakthroughs. Now, the trend is leaning toward optimizing the core infrastructure, like network layers and state access, to boost even more scalability. We'll dive deeper into this in our upcoming research piece!*

That's right, most projects up until now have been focused on improving the execution layers of blockchains. Now that increased volume is coming on-chain the space is beginning to hit a ceiling. We have identified that the most transformative infrastructure improvement is to be made at the network level. Blockchain is fundamentally a shared ledger maintained by a network of nodes. While computation or execution on individual nodes can be fast, network performance ultimately depends on how quickly nodes can sync data with one another.

Optimizing the network directly tackles the bottlenecks that slow down data propagation: reducing latency, cutting redundant transmissions, and lowering bandwidth usage. This is particularly impactful in decentralized systems, where data must travel across a globally distributed network under unpredictable conditions.

At Optimum, our RLNC-powered solution proves that network-level performance can dramatically accelerate the delivery of block data. We optimize the network layer and state read / write to unlock untapped performance gains. It's an underexplored, high-leverage opportunity to build a robust, scalable, and cost-effective Web3 environment that will support the demanding applications of tomorrow.

## **7. Blockchain and crypto aren't just about technology. How are you working to build a strong user base and community around it?**

You're completely right, Web3 is community-oriented and building a vibrant ecosystem of users, developers, validators, and enthusiasts is critical even for a B2B infra project.

At Optimum, we're actively nurturing a strong community through multiple channels. We engage with our audience on Discord and X through regular updates, AMAs, and interactive events such as live sessions at ETHDenver and Consensus. Our aim is to provide a transparent window into our technology and roadmap via explainer posts and technical deep dives, ensuring that everyone, from validators and node operators to dApp developers, can understand and leverage our innovations.

We work closely with our developer community by providing comprehensive documentation and onboarding support, and by actively encouraging feedback and collaboration to shape our roadmap. We've also created pathways for community members to contribute to the performance, reliability, and robustness of our network. By combining transparent communication with a culture of collaboration, we're building not just a user base, but a dynamic ecosystem dedicated to transforming decentralized computing for Web3.

## **8. What are some of the best Twitter accounts or articles to follow on related topics?**

I recommend following a mix of our official channels and thought leaders in the space. For instance, our Twitter account @get\_optimum shares regular updates on RLNC breakthroughs and product progress. You're also welcome to follow my twitter @kentlinyy and Muriel's twitter @MurielMedard. In terms of reading materials, sources like our index of technical papers and docs can be beneficial for those looking to dive deeper.


## **9. Looking ahead: what's the long-term vision for Optimum? Where do you want the project to be in 2–3 years?**

Our vision for Optimum is to become the foundational memory layer for a fully scalable, decentralized world computer. We're not just addressing immediate scalability bottlenecks, we're laying the groundwork for an ecosystem where every major blockchain, from leading Layer 1s to emerging Layer 2s, can seamlessly access a unified, high-performance memory infrastructure.

Our long-term goal is to see Optimum powering a multi-chain network, where products like OptimumP2P and deRAM work in tandem to provide fast, efficient data propagation and real-time state updates. This will dramatically improve validator yields through reduced latency and higher reliability, while also enabling dApp developers to build next-generation, latency-sensitive applications in trading, gaming, AI, and social platforms.

We envision a future where blockchains operate with the speed, efficiency, and resilience of traditional high-performance computers, unlocking new possibilities for mass adoption and innovative on-chain use cases.

**Figure 3: How Optimum Improves the Blockchain Ecosystem**




## OptimumP2P

OptimumP2P is a high-performance pub-sub protocol that dramatically reduces latency and optimizes the delivery of transactions, blocks, and data blobs.

This enhanced efficiency not only boosts network speed and scalability but also maximizes validator rewards—capturing every opportunity while using minimal bandwidth and CPU resources.

OptimumP2P seamlessly integrates with your existing node setup through a simple API call and provides a more robust, scalable, and efficient blockchain ecosystem.

[Validator Signup](#)



## Optimum deRAM

deRAM is a decentralized Random Access Memory (RAM) layer that ensures Atomicity, Consistency, and Durability (ACD). deRAM gives applications real-time read/write access to blockchain state, enabling fast, cheap storage and access.

deRAM unlocks the next wave of latency-sensitive, onchain use cases that today's blockchain infrastructure have struggled to support: trading, gaming, AI, social and more.

[Coming soon](#)

Source: Optimum

## 10. To wrap it up — favorite CT account, and any closing alpha for the users out there?

I follow and find great insights from technical voices like @\_weidai from 1kx, who dive deep into infrastructure and network optimization in decentralized systems.

Memecoins are hot this cycle—but which chain stays hot is anyone's guess. Today it's Solana, tomorrow Base, next week maybe BNB. What is clear: activity shifts fast. But imagine a project that benefits from any chain's activity. No matter which chain wins, its community wins. That project is called Optimum.

---

That's all! Thank you very much Kent. I really appreciate your time, a lot of exciting things are coming from Optimum.



## About Presto

Presto is an algorithmic trading firm where researchers and engineers solve challenging problems in global financial markets. Our core strength lies in combining engineering, mathematics, and science to navigate both digital asset and traditional finance markets with precision. Presto Research, our research unit, provides expert-driven insights to help navigate these markets effectively.

Find out more at <https://www.prestolabs.io>.

Follow Presto for more content: [X](#), [LinkedIn](#)

Follow Presto Research for latest research : [X](#), [Telegram](#)

---

## Authors

**Min Jung**, Research Analyst, [X](#), [Telegram](#), [LinkedIn](#)

**Jaehyun Ha**, Research Analyst, [X](#), [Telegram](#), [LinkedIn](#)

---

## Required Disclosures

*Any expression of opinion (which may be subject to change without notice) is personal to the author and the author makes no guarantee of any sort regarding accuracy or completeness of any information or analysis supplied. The views and opinions expressed herein are those of the author(s) and do not necessarily reflect the views of Presto or its affiliates. This material by itself, is not and should not be construed as an offer or a solicitation to deal in any investment product or to enter into any legal relations. This material is for informational purposes only and is only intended for sophisticated investors, and is not intended to provide accounting, legal, or tax advice, or investment recommendations, or an official statement of Presto or its affiliates. Presto, its affiliates and its employees make no representation and assume no liability to the accuracy or completeness of the information provided. Presto, its affiliates and its employees also do not warrant that such information and publications are accurate, up to date or applicable to the circumstances of any particular case. Certain statements in this document provide predictions and there is no guarantee that such predictions are currently accurate or will ultimately be realized. Prior results that are presented here are not guaranteed and prior results do not guarantee future performance. Recipients should consult their advisors before making any investment decision. Presto or its affiliates may have financial interests in, or relationships with, some of the assets, entities and/or publications discussed or otherwise referenced in the materials. Certain links that may be provided in the materials are provided for convenience and do not imply Presto's endorsement, or approval of any third-party websites or their content. Any use, review, retransmission, distribution, or reproduction of these materials, in whole or in part, is strictly prohibited in any form without the express written approval of Presto. Presto Research and related logos are trademarks of Presto, or its affiliates.*