DoubleZero Protocol

December 2, 2024

Accelerating communication in high-performance distributed systems to
Increase Bandwidth and Reduce Latency

Austin Federa, Andrew McConnell, Mateo Ward
austin@doublezero.xyz, andrew@malbeclabs.com, mateo@malbeclabs.com

Special Thanks to Dr. Nihar Shah, Chris Remus, Malbec Labs, Dr. Kevin Bowers, and the Firedancer team

Abstract

The DoubleZero protocol is a decentralized framework for creating and managing high-performance permissionless networks, optimized for distributed systems like blockchain. It enables permissionless contributions of underutilized private fiber links to create a dynamic and expansive network. The DoubleZero network has a low-level permissionless architecture capable of seamlessly bridging heterogeneous networks. Together, it creates a meshed transport layer that filters and serves traffic across low-latency and high-bandwidth routes, allowing distributed systems to send and receive information efficiently. Through a permissionless controller that lives on a public blockchain, the protocol attracts new contributions with incentives and manages the network configuration and routing in response to demand spikes, outages, and other disruptions. This network infrastructure layer enables systems like Layer 1 blockchains to be freed of communication bottlenecks and approach the maximum performance that physics allows. In short, the DoubleZero network is a new internet optimized for distributed systems.

Introduction

The current pace of improvement in throughput and latency of decentralized systems is unable to keep up with the needs of developers and users. This is most noticeable in Layer 1 blockchains, but also true for Layer 2 networks and non-blockchain distributed systems. This is despite substantial improvements in the computing capabilities of individual validators.[1] However, these systems are increasingly bottlenecked by bandwidth limitations and variable latency in communication between validators, rather than compute capabilities within validators. The key to unlocking higher performance for these systems is to optimize data flows. The DoubleZero protocol is designed to do this by combining individually-contributed fiber links into a synchronized network that filters spam, increases bandwidth, lowers latency, and removes jitter from communication. In short, it is designed to improve the performance of blockchains – and all decentralized distributed systems – by advancing reliability and scalability.

The physical network layer that blockchains use has not changed since the advent of Bitcoin. This underdevelopment limits the scalability and decentralization of performant blockchains, making them both more centralized and more inefficient. The hardware requirements needed for validators to consume the firehose of inbound transactions escalate quickly. Consensus is reached slowly, as validators propose blocks and vote over public internet paths prone to high jitter and inconsistent routing.

Compared to the public internet, the DoubleZero network offers two improvements to blockchains. First, inbound transactions can be edge-filtered (i.e. removed of spam and duplicates) by specialized hardware prior to being sent over the DoubleZero network. This allows blockchains to benefit from shared system-wide filtration resources rather than needing each individual validator to provision sufficient resources. Second, outbound messages can be explicitly routed, tracked, and prioritized to improve efficiency. Validators can control latency and reduce jitter (i.e. random variation) in that latency in reaching consensus. Taken together, the DoubleZero network can achieve levels of performance that are otherwise unobtainable without the centralized co-location of systems.

The DoubleZero network architecture maps conceptually to two concentric rings: an outer ingress/egress ring, and an inner data flow ring. The outer ring of the network interfaces with the public internet, and DoubleZero's network devices here use performant hardware (such as FPGAs) to mitigate distributed denial-of-service attacks, verify signatures, and filter duplicate transactions. Servers on the inner ring build consensus with this filtered traffic over optimally-routed dedicated bandwidth lines. Together, these two rings provide a scalable and permissionless information service.

Indeed, the DoubleZero network can be used for optimizing any distributed system. Blockchain is not unique in this bottleneck, as nearly every application of time-sensitive, high-performance compute is limited by ingress and egress of data flow rather than the act of doing the work. For instance, within the blockchain world, RPC nodes, MEV systems, and even Layer 2 chains can

[1] For instance, the blockchain space has seen a new generation of high-performance validator clients (e.g. Firedancer and Reth) and a new generation of high-performance chains (e.g. Sui and Monad) that improve on old systems by a literal order of magnitude.

benefit from the filtered and low-latency links provided by the network. Outside of the blockchain world, content delivery networks, gaming applications, large language models, and enterprise users can benefit from the network for similar reasons. For many of these applications, their high-priority traffic otherwise flows over the public internet, indiscriminately co-mingled with non-time sensitive traffic. The DoubleZero network offers an alternate and improved model.

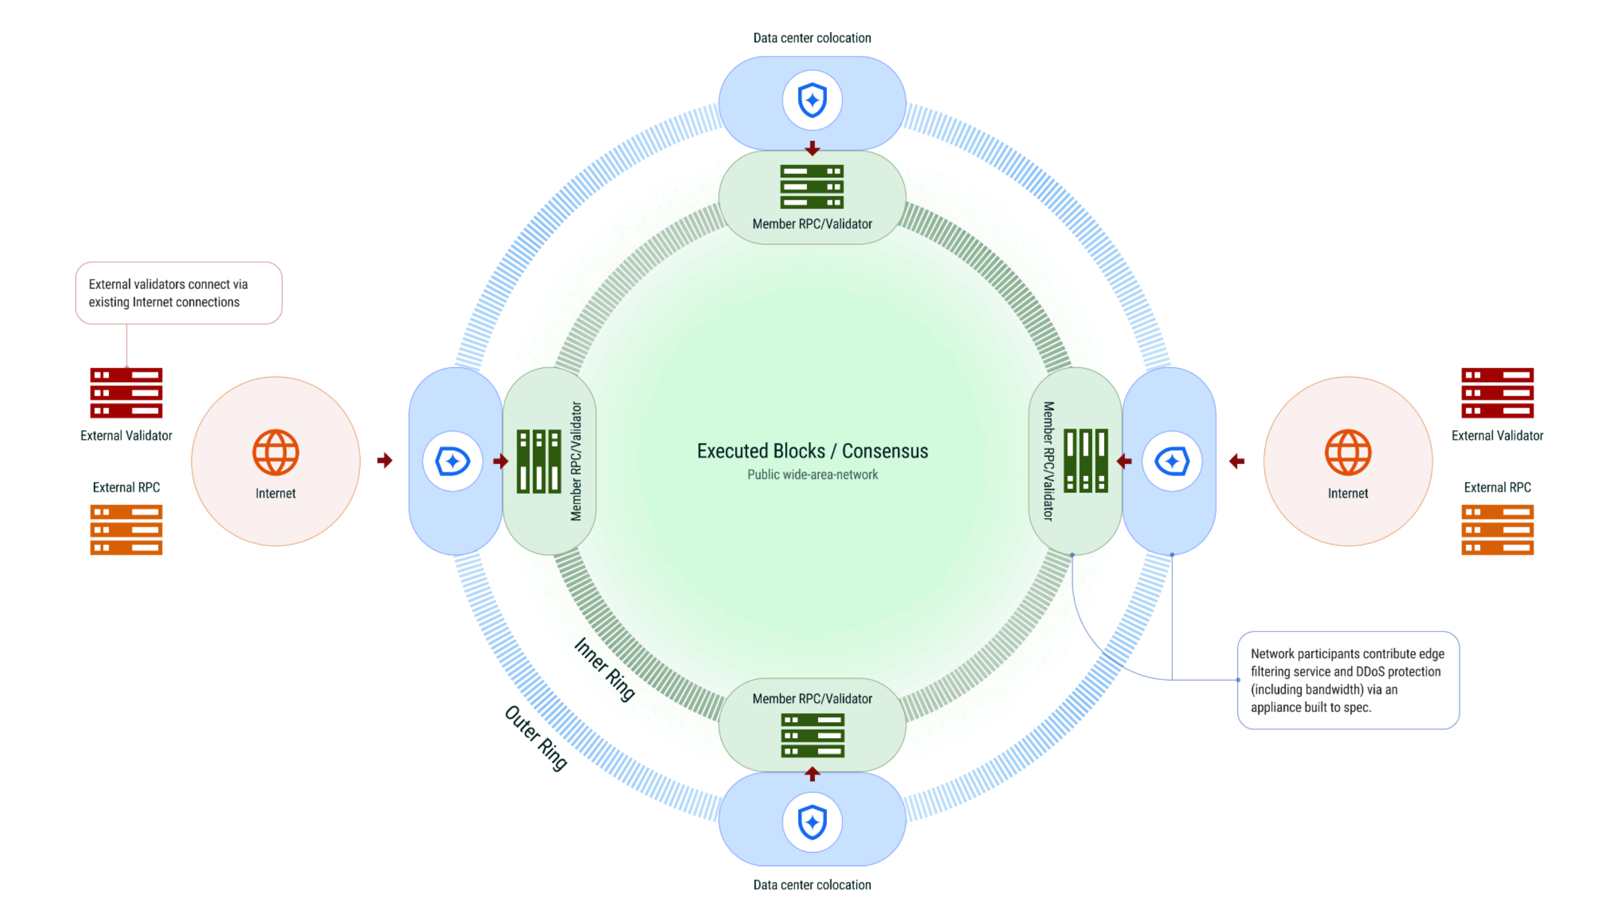

Figure 1. The DoubleZero Conceptual Network Diagram

The DoubleZero protocol is not just a technical achievement – it is a new economic model for bandwidth and communication. On the provider side, it allows private businesses to monetize underutilized private links they may have purchased or leased from a telco or network service provider. Rather than sitting idle, entities can contribute their links into the DoubleZero system and open new revenue streams for themselves. On the user and operator side, it allows distributed systems to enjoy the benefits of private networking without relying on centralized systems and long-term contracts. Instead, these systems can be flexible and continuously fine-tune their throughput requirements, network topology, or communication prioritization. Most importantly, the DoubleZero protocol not only matches suppliers and users, but also turns individual contributions into a unified global network. Individual links from suppliers – which are fixed in their routes, limited in their reach, and offered without ubiquitous global redundancy – may line up poorly with demand needs from modern systems that demand flexibility, reach, and resiliency. The DoubleZero protocol stitches these individual links into a robust and expansive system, leveraging the same network effects that made the internet what it is today.

Like the public internet, the distributed ethos of the DoubleZero network is a strength rather than a weakness, with multiple independent contributors creating something more valuable than the sum of its parts. Ultimately, the protocol is the first of its kind as an “N1” – a base layer of neutral and performant physical infrastructure. It is upon this N1 that distributed systems and applications (network-related, such as N2s or otherwise) can be built.

A core belief of the DoubleZero project is that abundance breeds innovation. Every step-function change in availability or corresponding reduction in latency opens up entirely new types of applications and uses for distributed systems. From this first principle, DoubleZero sets a north star of increasing bandwidth and reducing latency for distributed systems worldwide to support a more fair and open future for software.

System Architecture

DoubleZero begins with a simple but profound idea – what if each node operator (validator, RPC, etc.) was not solely responsible for data ingest and egress? We propose that by decoupling filtration and verification from transaction inclusion, block production, and execution, we can create a parallel and protected transaction flow that enables node operators to optimize their network operations.

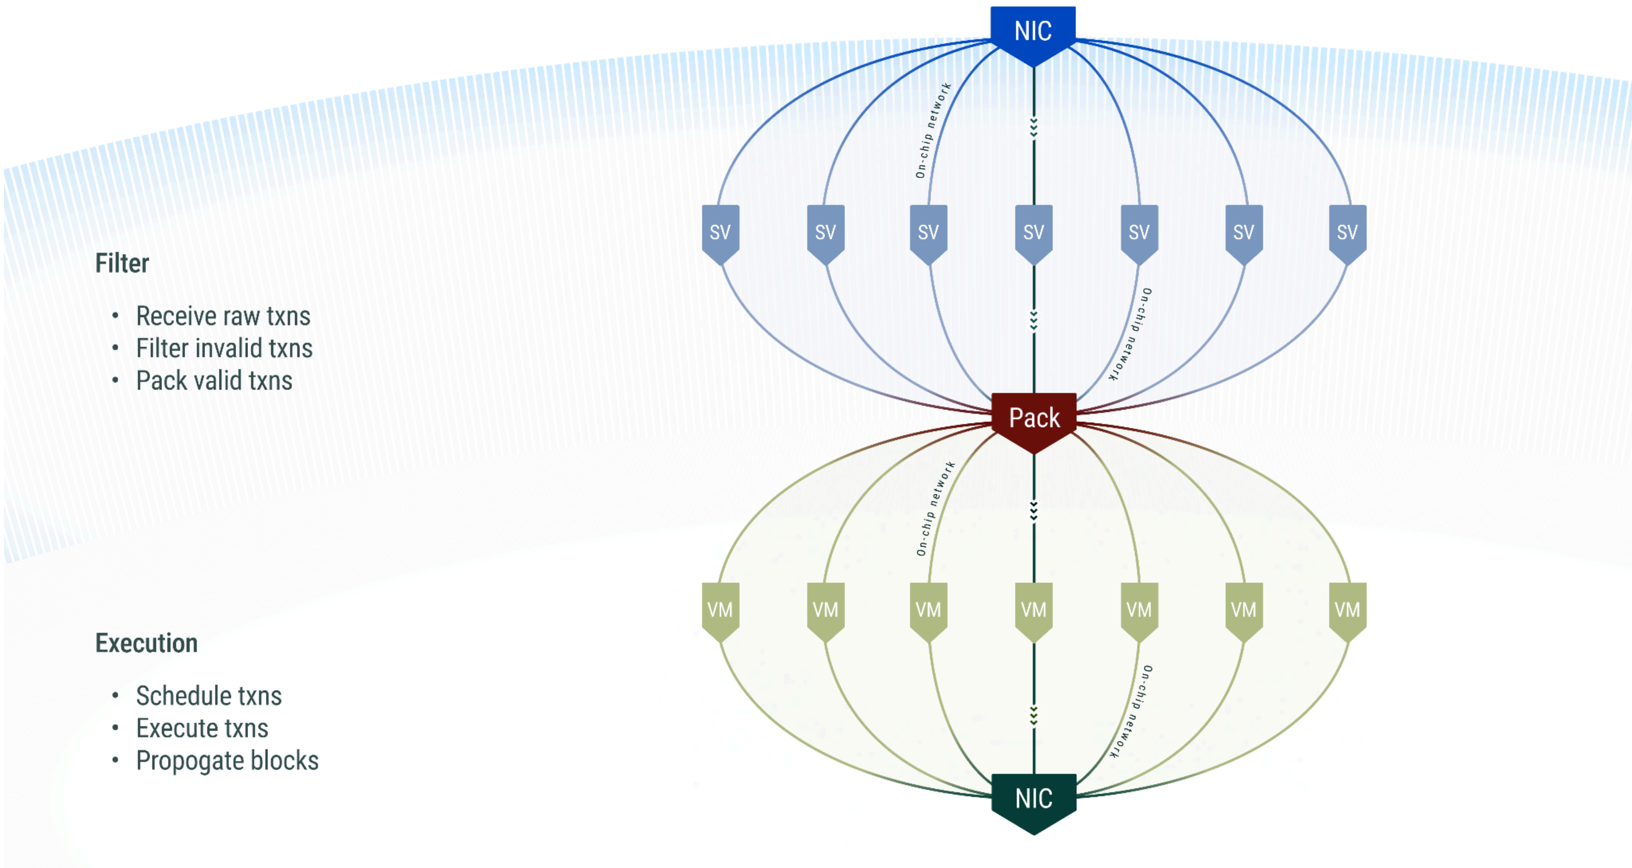

Figure 2. Existing Proof-of-Stake Blockchain Validation Process.

While a validator's role is commonly understood as a single operation, there are two underlying steps: producing a block and finding consensus for a block. Block production begins when a flood of unfiltered inbound data arrives at the validator. The validator performs deduplication, filtration, and signature verification operations, reducing the possible pool. The validator then selects a subset of viable transactions to include, constructs the block, and propagates its proposed block to the other validators for consensus. The other validators either approve the block and append it to the chain, or they reject the block and continue onwards.[2]

[2] While there is a difference between the mechanics of leader and leaderless blockchains, particularly with respect to ordering of blocks, this subtlety can be safely ignored here.

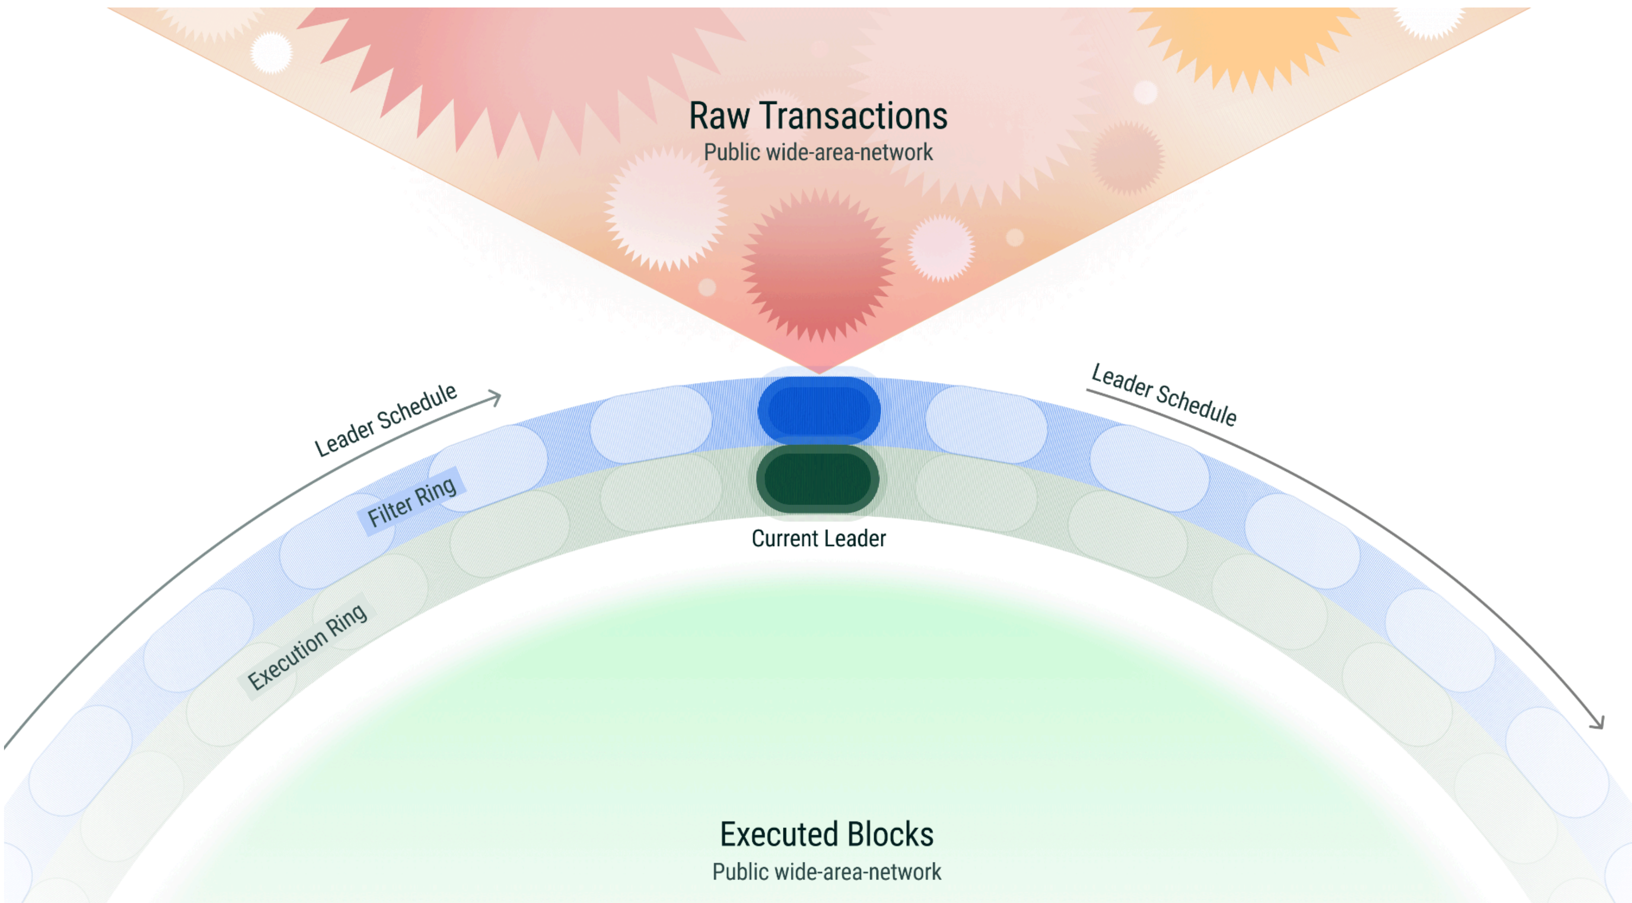

Figure 3. Data Flow as Rings.

On most networks today, these processes are handled by one x86 machine's worth of resources.[3] That CPU's cores are linked by an on-chip network within the machine. We propose to expand this to a global network: the DoubleZero network. By placing FPGA appliances to filter inbound traffic at key ingress points at the network edge, the DoubleZero network acts as a common shield for all local validators (i.e. validators that are downstream of the devices connected to the inner ring of the DoubleZero network). Compared to validators, these FPGA appliances can handle orders of magnitude more traffic as they perform the specific tasks of removing spam, de-duplicating a transaction set, and verifying signatures – and this is all done via open-source software that can be inspected to ensure fair treatment of inbound transactions. Downstream validators receive a

[3] Even PBS uses only one x86 machine's resources at a time.

substantially smaller transaction set and only have to re-verify signatures on the transactions included in the final block, to maintain the blockchain's security. Their resources – typically a single x86 machine – are significantly freed to perform tasks related to block production, transaction execution, indexing, and so on.

If we disintermediate the filter and execution rings to provide more flexibility, we can distribute ingress resources more efficiently and more resiliently. Through resource sharing, we no longer need to provision each individual validator with enough resources to meet global demand; rather, global demand can be met through infrastructure sharing which requires far fewer resources in aggregate.

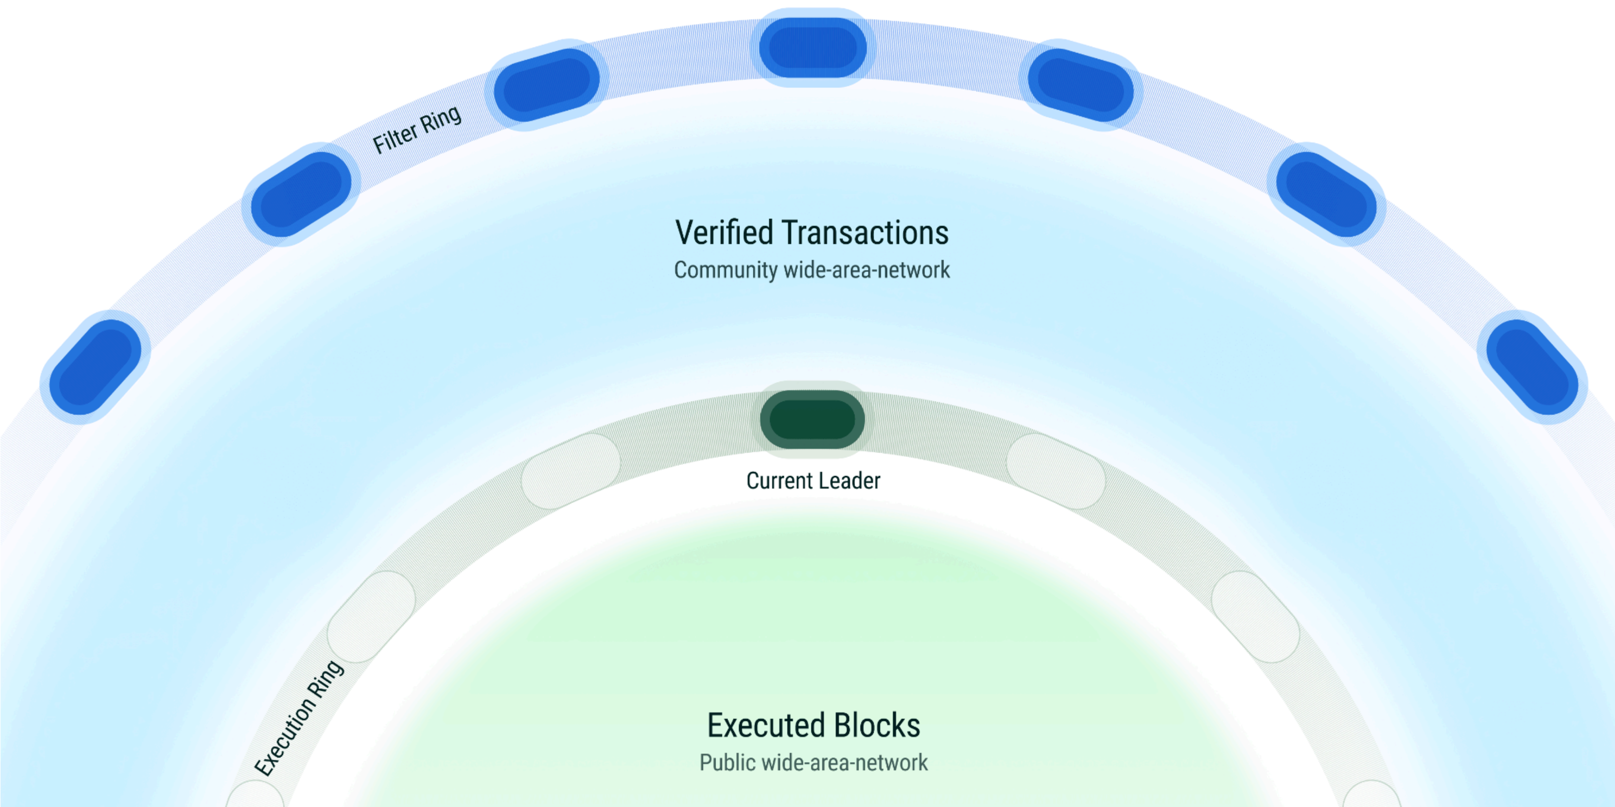

Figure 4. Separating the Filter and Execution Ring

Further optimization of traffic flow is achieved by support for multicast traffic on the inner ring of the DoubleZero network, which functions as a single network domain under common control of protocol participants. Value-added functions like block or shred propagation can be included via multicast.[4]

[4] Some networks make use of shreds, or pieces of a block, which are transmitted before the full block is built.

Having low-latency access to the newest state can provide significant advantages for high-performance applications. A denial-of-service attack, for example, would be much more difficult to successfully pull off under this model – it would require many terabits per second of traffic hitting many geographically-distributed sites simultaneously. It would no longer be an attack against an individual validator or even a single blockchain; it would be an attack against hundreds of data centers and ISPs around the world. A successful attack would be several orders of magnitude more difficult and resource-intensive – and underlines that using infrastructure sharing creates a more robust and resilient network architecture for all.

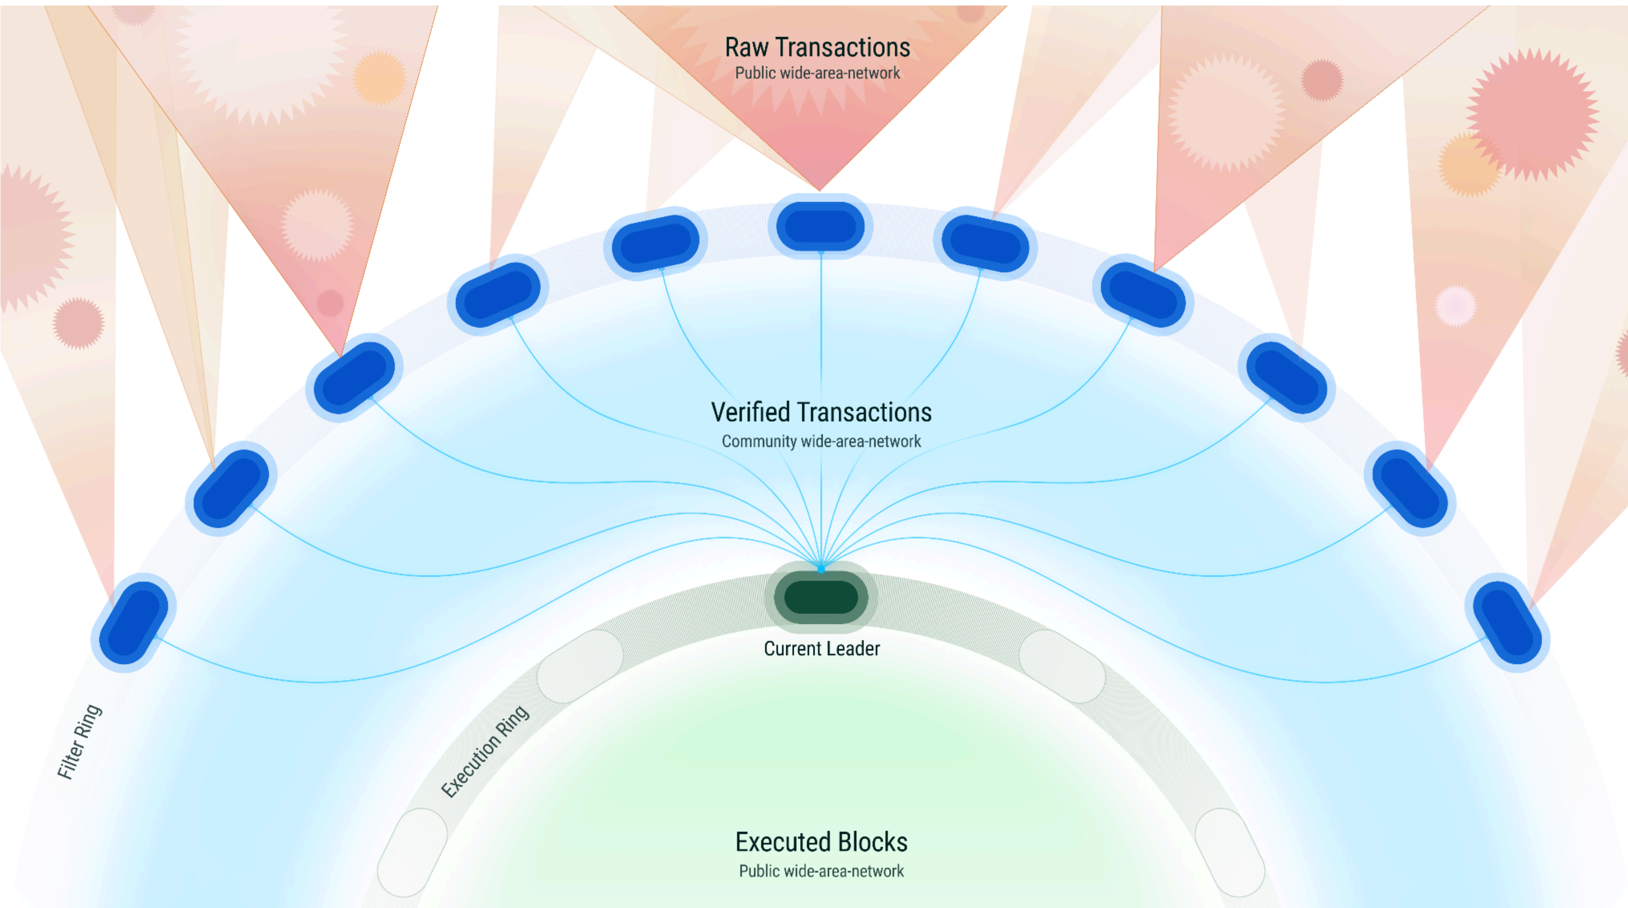

Figure 5. A DDoS Attack on the DoubleZero Network

The Physical DoubleZero Network

Permissionless Global Connectivity for Blockchain Traffic

The physical layer of the DoubleZero network is its foundation. It consists of two key components: network devices at key ingress/egress points of the network and provisioned bandwidth across the network. Together, they form a meshed network that comprises the DoubleZero network. While individual contributions power the DoubleZero network, the network makes use of orchestration technology to aggregate those heterogeneous contributions. This allows the network to leverage the strengths of private networks (which are limited in reach but operate on a unified standard) and public networks (which are expansive in reach but negotiate multiple standards). DoubleZero combines the best of both worlds.

Permissionless global *software* abstracted networks are nothing new – look no further than BitTorrent and the Tor browser. Software abstraction allows these protocols to operate over the internet, but therein lies their greatest limitations – they are software networks running on top of the public internet and are still governed by the same fundamental limitations of all public internet traffic.

The architecture of the DoubleZero network is designed to connect validators, RPCs, and other systems at the network edge to a core underlay connected at DoubleZero Exchange points (DZXs). DZXs enable efficient connectivity between multiple data centers within a metro area in a similar way as public internet exchanges connect different networks together. These underlay links are provided by network contributors, and are the core transit layer of the protocol. The inner ring exchange points allow ubiquitous interconnectivity to happen on a global scale, regardless of how many contributors participate in DoubleZero.

Network Devices

The network devices on the DoubleZero network serve two functions. First, they enable data links contributed by individuals and organizations to operate as one cohesive network with prioritization. Second, they implement filtration, verification, and spam protection.

Network contributors on the DoubleZero network are free to utilize any suitable hardware for these tasks, but the network holds itself to high standards. For instance, at launch, the DoubleZero network uses commercially available FPGAs (to balance performance, speed, and cost), and to run integrated appliances that can handle both filtering and routing. One sample deployment alone can handle deduplication, filtering, and signature verification for multiple Gbps of inbound data. These devices also manage connectivity to downstream validators and other client devices.

While the deployment of advanced filtration and signature verification technology to data centers and other validator facilities would improve the state of blockchain networks on its own, DoubleZero proposes to go a step further and create a new permissionless transport layer for blockchains and distributed systems that is fully integrated into the filtration service. This can run across a variety of hardware and will expand as the network grows.

Fiber Links

The fiber links on the DoubleZero network provide the low-latency high-bandwidth connections between different points on the globe. Network contributors provision the connections and they encode a service level agreement for the quality of that connection.

The fiber links are either directly owned by the network contributor, or leased by them from the major network service providers. Network contributors add them to the network and commit to a service-level agreement for each link. This commitment establishes the main characteristics of the link, including endpoint locations, bandwidth, latency, and compliant MTU size. (As an illustration, a network contributor for a link from Los Angeles to Singapore may commit to 10 GBps at 85ms latency with 1600 MTU.) Performance of the links is monitored and tested by the DoubleZero network to ensure compliant links are eligible to earn rewards, while non-compliant links may be disqualified for rewards and potentially removed from the network.

The DoubleZero project participants believe that there is substantial spare capacity in the fiber market. First, a modern fiber link, whether terrestrial or subsea, can support hundreds of terabits of data per second. Second, there is substantial "dark fiber," or fiber that has been installed but does not operate. For instance, the Federal Communications Commission [2007] estimates in its last publication that only 35% of the fiber installed in the United States is being utilized, and 65% remains unutilized. And since then, fiber capacity has only grown dramatically; rising from reaching 15% to 46% of US locations from 2015 to 2023 (per the Fiber Deployment Annual Report of 2023) and reaching 70% of EU39 households (per the FTTH Council Europe in 2024). Third, many enterprises, in buying or leasing fiber, tend to over-provision capacity, building it for worst-case rather than routine outcomes. This suggests that the DoubleZero network could benefit from contributors repurposing already-installed links.

Contrary to popular belief, data centers within a city or region are not isolated from one another. A vast network of short-haul fiber and other routing services connect data centers within a metropolitan area, allowing for direct connections between them and thus allowing the expansion of the DoubleZero network. This helps reduce incentives for co-location and expands the geographic reach of DoubleZero.

The DoubleZero network's core focus is on fiber. However, network contributors are free to provision other communications services, such as microwave or satellite technology, and attach it to the network as an application (an "N2").

Community Governed and Operated Networks

Even the most performant physical infrastructure networks see the occasional hiccup. Network-wide issues are where the final component of the DoubleZero system comes into play — a streamlined operational process for bringing together network contributors to diagnose and repair the problem, while users fall back to the public internet. (Failures with a single network contributor's infrastructure are the responsibility of that network contributor and do not need community-level operations.)

Centralized systems generally enjoy advantages over decentralized systems in times of network failure or issues. While centralized systems have more single points of failure, they can quickly muster technical resources to resolve issues (By contrast, decentralized networks such as blockchains are more resilient, but require coordination among many unaffiliated stakeholders to troubleshoot issues.) In its mature state, DoubleZero contributors will be able to commit similar resources towards routing around disruptions and maintaining network cohesion. As the network grows, the DoubleZero network aims to minimize the gap by implementing best practices, developing comparable operational procedures, and simulating failure situations from the start. There are many existing and successful models of multiple network operators working together to operate and maintain international networks – for example, consortium sub-sea cable systems.

Network Creation

The DoubleZero project is a new economic model for filtration and routing services over global networks. In the DoubleZero protocol, they can be both provisioned and utilized permissionlessly through a blockchain-based smart contract that verifies the integrity of the interactions. Existing solutions today operate on behalf of a specific customer in B2B relationships and do not act to increase the aggregate performance of the distributed system.

This flexible model is key to unlocking network effects. The DoubleZero protocol automatically turns individual contributions into a holistic network that is more valuable than the sum of its parts. A provider can contribute a single and even redundant link – it might have low value to a distributed system in isolation, but is highly valuable when operating alongside links from others. Indeed, the federated nature of the DoubleZero project is its core strength.

Smart Contract-Defined Network

The DoubleZero protocol organizes itself with verified data stored in onchain smart contracts. The distributed control plane uses this information to reach consensus on routing and prioritization decisions for the network. This provides a transparent and permissionless experience for both contributors and users.

Contributors connect to the network and set a service level agreement enshrined in a smart contract, which includes endpoint locations, bandwidth, latency, and compliant MTU size. They send those details, along with the initial duration of the contribution, to the smart contract. Validators and other users can opt-in to predetermined services and routes, such as a specific L1's traffic, state propagation, and validator set, or they can author their own contract to define a routing request.

As inputs come in, the distributed controller uses smart contract data to identify the optimal configuration for the network. One key component of this configuration is the correct price, so that users pay for the cost of using the network in proportion to the request they impose on the system and providers earn incentives for contributing to the network in proportion to the value they generate in providing links. Ultimately, the project aspires to offer tiered pricing, where users can pay for priority access or base access to a given route to maximize network efficiency. The

DoubleZero network aims to offer a high degree of customization and control to its services and pass the surplus back to providers – a major advantage over the public internet, where traffic gets commingled regardless of priority

Incentives and Alignment

The open ethos of the DoubleZero project hinges on internal alignment and reciprocity by its network contributors. For the network to work, they must meet their service level agreements and have incentives to catch defection on the part of others.

A system of checks and balances, based on the same cryptoeconomic values that secure proof-of-stake networks today, will underpin the DoubleZero system and ensure incentives are aligned for cooperation rather than defection. For instance, links that perform below their stated service level agreements will be penalized and ineligible for rewards (and eventually may be excluded altogether), while links that meet those obligations will be eligible for rewards. Unlike blockchain systems, physical networks lack the abstract verifiability of a block – where any node anywhere in the world can validate its correctness. Physical networks need physical connections (peers) to verify integrity, which today is only possible with redundant physical connections. DoubleZero will require new innovation for single point verification.

Network Effects

The DoubleZero project goes beyond a standard marketplace, where providers and users meet and negotiate bilaterally. It is a true network – each link is more valuable as part of the whole than it would be in isolation.

Consider a link provider that has a few disconnected links, which represent underutilized capacity on their internal networks. While the entity could sell services over those routes piecemeal, each route would require substantial overhead in finding buyers and negotiating contracts. More importantly, each link has limited value, as it serves only one particular route.

But the provider who offers those links to DoubleZero benefits in two ways. First, the process of adding contributions to the platform is easy and standardized (all one needs to do is install and configure a few switches), requiring little overhead. Second and more critically, the link grows in value by its association with neighboring links. A disconnected link can potentially become a highly connected link on the DoubleZero network, making it substantially more valuable than it would be on its own.

Network effects are present in many other markets. One example: airlines. A single airline route between two cities is rarely offered in isolation, as there is not likely much customer demand. Instead, they are bundled into a carrier’s larger network and interact with other routes through hubs. These bundles are substantially more valuable; they become even more valuable when the carrier joins an airline alliance and implicitly ties their routes with routes in other parts of the world. DoubleZero is the platform that takes individual routes and turns them into a thriving cross-connected network.

The DoubleZero network can even benefit from redundant links, short links, or low-performance links. While the network is high-performance and expansive holistically, these types of contributions add redundancy, capacity, and interconnectedness. Thus, the DoubleZero project can allow providers to monetize capacity that otherwise would have almost no value in isolation.

Taken together, providers who make unused capacity on their network available to DoubleZero face lower transaction frictions and generate more user surplus from the network economics, especially when compared to selling that capacity in isolation. In turn, the DoubleZero network passes that surplus onto the providers.

Blockchain Use Cases

Blockchains, as the highest-profile distributed system archetype, holistically represent the strongest use case for the DoubleZero network. A blockchain's dual objectives of being public and performant often work against one another in the status quo. Across contexts as diverse as Layer 1 blockchains, RPC nodes, MEV systems, and Layer 2 chains, the DoubleZero network can push towards attaining both.

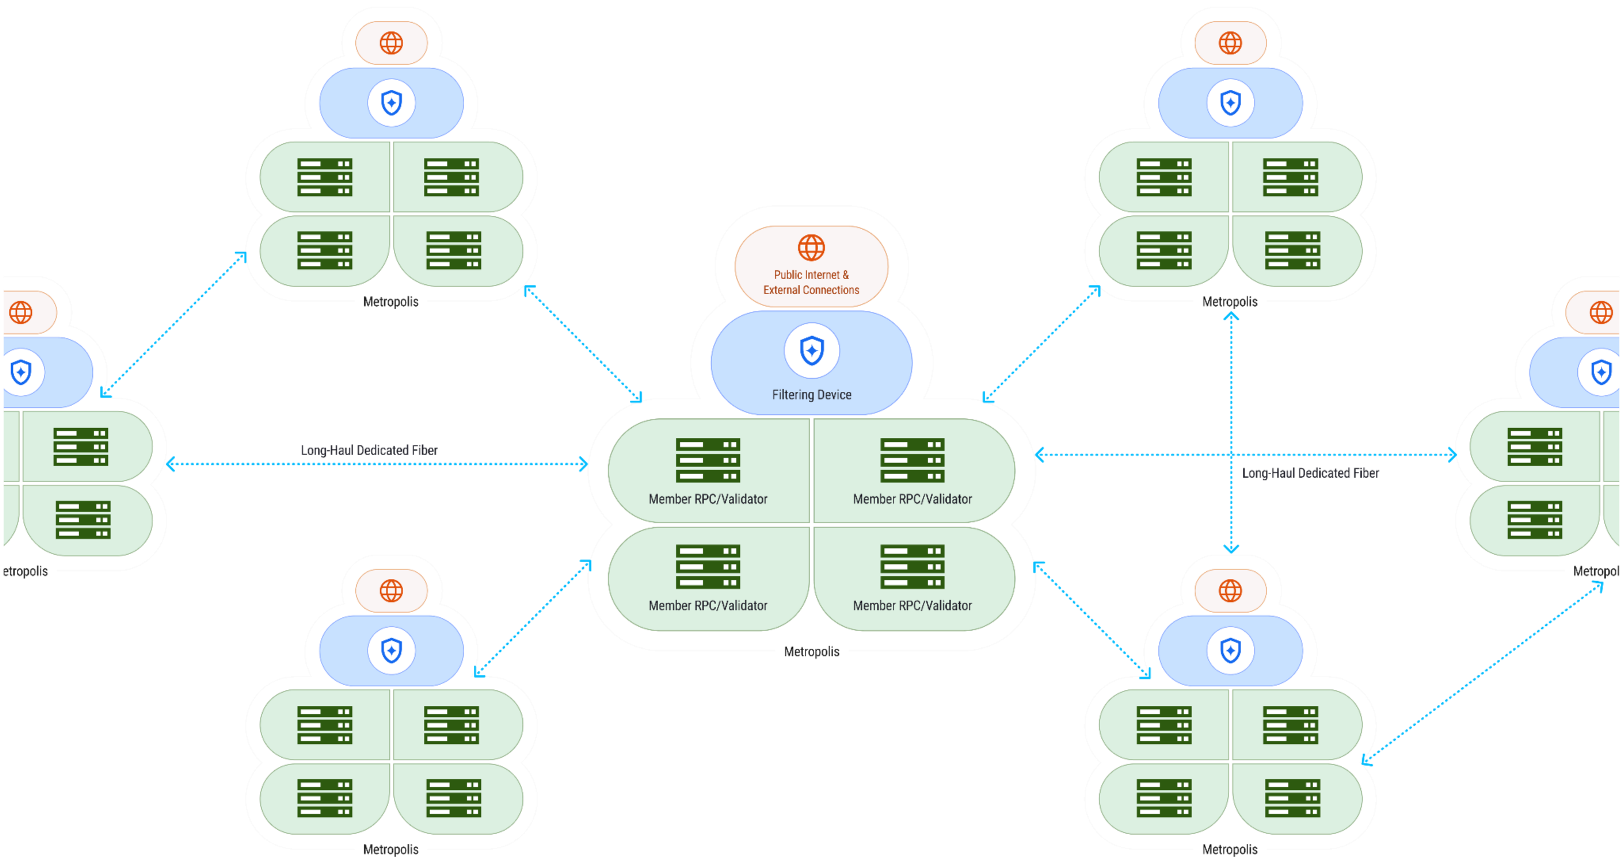

Figure 6. A conceptual map of the DoubleZero Network

Layer 1 Blockchains

A Layer 1 blockchain is the distributed system that stands to gain the most from the DoubleZero network. In the status quo, each validator faces a firehose of inbound transactions (including spam and duplicates). Consensus is reached slowly as validators propose blocks and vote over public internet paths prone to high jitter and inconsistent routing. In tandem, users suffer from poor experiences, whether it be competing with a flood of spam or waiting for delayed finality. Even today's most performant blockchains fall short of providing both the web2-style experiences users expect and the professional central venue trader experience.

The industry has developed some solutions, but each comes with disadvantages. Mempools serve as a central holding pen for transactions (to lower the ingestion burden on individual validators), but this adds latency to the system and broadcasts transaction intentions well before their inclusions. Quality-of-service solutions have been tested on blockchains like Solana wherein validators accept transactions predominantly from those who have staked capital, but this excludes entities who do not have sufficient stake from full participation, does not guarantee economic efficiency, and adds additional latency from the proxy step. Validators can invest more in hardware to keep pace, but this of course inhibits the development of a diverse and decentralized validator set and does not address aggregate network performance.

The DoubleZero network offers a more robust solution in two ways. First, it provides the filtering technology to filter inbound traffic that comes over the public internet. This step utilizes scalable hardware running open-source code, mitigates distributed denial-of-service attacks, verifies signatures, and removes duplicate transactions. This separation of filtration, deduplication, and signature verification from block production, execution, and consensus frees up substantial resources to focus on block production. This is cooperative infrastructure sharing, where the system in the aggregate can share resources rather than every validator individually needing them.[5]

Second, it routes the filtered traffic – whether user transactions, blocks, or consensus votes – to the validator set over its dedicated bandwidth links with minimal latency and jitter. This ensures that blocks are built, shared, and ratified quickly. Under this networking paradigm, high-performance validator clients or consensus algorithms can come much closer to their theoretical maximums. Indeed, this can be optimized further as the DoubleZero network is expected to support multicast functionality, which allows for better propagation of state transitions, blocks, and shreds.

There is one final advantage of using the DoubleZero network for Layer 1 blockchain communication. Packets traversing over the network are expected to be traceable so that users can verify their inclusion or exclusion from intra-validator communication. For systems that are sensitive to concerns of censorship, this is a beneficial by-product of transparency in traffic routing.

[5] There are many other successful examples of cooperative infrastructure sharing, including mobile networking, TowerCos, or subsea fiber optic cable system consortia.

RPC Nodes

Remote Procedure Call (RPC) nodes can similarly benefit from the DoubleZero protocol. As the primary interface between users and the chain, they face especially severe filtration and communication challenges.

RPC nodes comprise the infrastructure that mediate user interactions on a blockchain. The primary role of an RPC node is to receive inbound transactions from users and to forward those transactions to either the leader or the chain's mempool. The secondary role is to provide an updated view of the chain to downstream applications, such as DeFi applications, block explorers, or gas estimators that help users form efficient transactions. RPC nodes are effectively non-voting validators – they track state transitions on a blockchain, but do not participate in consensus.

There are three challenges that RPC nodes face that could be mitigated by the DoubleZero protocol. First, RPC nodes are directly exposed to traffic surges from users, particularly around broad community events like airdrops or NFT mints. Because they specifically receive the transaction flow via the public internet, they are especially vulnerable to distributed denial-of-service attacks.

Second, the role of the RPC node is to forward user transactions to the block leader or the mempool, depending on the blockchain. Transaction deliverability matters immensely for an RPC, especially since the highest-value transactions (e.g. arbitrages) tend to be both competitive and time-sensitive. Resilient and low-latency connections from the RPC node to the block leader or mempool can assist RPC nodes in these responsibilities.

Third, RPC nodes return information from the blockchain to downstream users, wallets, and dApps. They are primary providers of data about the state of a network, providing everything from wallet-balance information to available liquidity for a given trading pair on a DEX. In order to do this, they need to ingest, validate, and index changes as close to real time as possible. Resilient and low-latency connections back from the validators to the RPC nodes can ensure that RPCs have up-to-date views of the blockchain.

The DoubleZero protocol can make the RPC ecosystem more competitive. RPCs already vie aggressively amongst themselves on timeliness of information both to and from the chain. The DoubleZero network opens an avenue for them to improve service yet further.

Maximal Extractable Value

The DoubleZero network can support a more fair and open maximal extractable value (MEV) ecosystem. These systems are extremely latency-sensitive, such that even small improvements in latency and deliverability can deliver additional value.

MEV refers to the value that block producers can take by ordering transactions in a block – by strategically adding a transaction to close a price dislocation caused by a large trade. In practice, mature MEV systems typically situate themselves as adjacent to the validators. They receive proposals from external traders (known as “searchers”), identify a profitable ordering of transactions, and forward the ordering to the block leader along with some additional payment.

Reliability and latency in communication matter on two dimensions. First, MEV systems benefit from a more real-time view of the state of a network. With fresher data, they can construct more profitable transaction ordering, allowing them to craft more efficient blocks. Second, even holding the opportunity set fixed, MEV systems that can relay proposals to validators quickly can use the extra time to find better orderings. Block construction is a complex combinatorial problem that requires approximation and simulation algorithms to find local optima – any extra time allows those algorithms to run longer.

Put another way, MEV systems benefit from faster retrieval of data (read operations) and faster delivery of blocks or bundles to the current leader (write operations). Increasing bandwidth and reducing latency improve both of these factors.

Layer 2 Blockchains

Layer 2 blockchains that use central sequencers or a distributed set can utilize the DoubleZero network too. In a typical Layer 2, a single sequencer posts inbound transactions to a Layer 1 chain; a set of Layer 2 validators then use those to update the chain state. There are standard ways that a Layer 2 chain can utilize the DoubleZero protocol, such as filtering inbound transactions to the sequencer, posting to or retrieving from the Layer 1 without latency, or verifying updated state among the validators quickly.

There are two more critical ways that Layer 2 chains can benefit from the DoubleZero network. First, most Layer 2 chains aspire to support multiple sequencers (posting transactions either in rotation or in tandem) to reduce centralization, censorship, and liveness risks. This architecture shift requires substantial coordination between the sequencers to prevent wasted computation for both the Layer 1 chain and the Layer 2 validator set. The DoubleZero network can smooth that communication between sequencers.

Second, many Layer 2 chains aspire to scale using data availability layers, which store transaction data, smart contract data, and even off-chain data on separate systems that can deliver information on demand. But the efficacy of this modular architecture is limited by the bandwidth connections between the modules, and low-bandwidth connections can lead to new bottlenecks in the underlying chain. Again, the DoubleZero network can ensure that the data availability layers keep pace with Layer 2 chains.

Other Blockchain Contexts

The DoubleZero network offers advantages in other blockchain contexts: state synchronization, indexers, multiple leaders, and network extensions, among others.

One challenge with high-throughput blockchain systems is onboarding new validator nodes. These nodes must synchronize to the current state by replaying transactions from the chain start, as outlined by MegaETH [2024]. For high-throughput blockchains, this can impose too much pressure on existing resources like bandwidth. The new validator node cannot simply ingest transactions at the current speed of the blockchain but must ingest *substantially faster* in order to catch up to the current state. The high-bandwidth connections of the DoubleZero network can be useful in overcoming these constraints and making it easier for new validators to join. By the same principle, it can also help validators that fall behind come back online.

As blockchains scale, supporting infrastructure like indexers will become more important. Like RPC nodes, indexers need to receive updated state quickly so that their end users can form quick and reactive decisions. State propagation through DoubleZero reduces delays to indexing fresh data, and can help keep distributed indexer instances current.

Many Layer 1 blockchains aspire to have multiple leaders simultaneously produce blocks in different parts of the world. Without tight coordination between leaders, though, this can quickly lead to redundant or outdated transactions and wasted compute. Like multiple sequencers on Layer 2 chains, multiple leaders on Layer 1 chains need rapid global state synchronization; and the DoubleZero network can facilitate this better than the public internet.

Lastly, there are a whole class of blockchain extension services that do not fit nicely into the “Layer 1” or “Layer 2” buckets. These uncategorized services are sometimes referred to as “network extensions.” Nearly every network today has the concept of a compute operation that can be run on only a subset of a network’s validators, because it requires specialized hardware or software. ZK proovers, GPU workloads, and asynchronous execution platforms all can benefit from the increased bandwidth and reduced latency links DoubleZero provides.

Other Use Cases

The DoubleZero network has uses beyond blockchain technology. Any setting in which distributed systems have to communicate quickly, precisely, and voluminously can benefit from its structure. This section lays out four additional examples.

Content Delivery Networks

At a high level, content delivery networks (CDNs) look similar to the DoubleZero network. CDNs make content delivery more efficient for users by filtering public requests and routing them efficiently through their network. This, coupled with local caches that hold content, allows for users to enjoy reliable and fast experiences for a given web application.

A new CDN can rely upon the DoubleZero network to bootstrap itself more quickly and focus on its core competencies. CDNs excel at forecasting demand and provisioning content accordingly; but in the status quo, they must also filter traffic and establish networks with low latencies and high bandwidth to offer an end product.[6] A new CDN built atop the DoubleZero network could return to those core competencies only, and leave the filtering and routing to the protocol.

[6] The incumbents, like Akamai and Cloudflare, mostly establish long-term leases from network service providers to structure their network.

Moreover, the DoubleZero network can boost the operations of existing CDNs. One major advantage of the DoubleZero protocol is its flexibility: Users do not need to contract for expansive and long-term leases, but can customize their usage for specific routes and time intervals. Thus, CDNs that maintain their own network but lack coverage in certain regions or need more bandwidth across central corridors can use DoubleZero accordingly, offering better experiences to users and relieving strain on their networks alike. This could be especially powerful as multicast functionality comes to the network and allows for more efficient delivery of content.

Online Gaming

Latency makes and breaks online games. Even mild lags ruin the user experience and restrict players to simple turn-based games. Traditional game developers have used two models: peer-to-peer models or central servers. The DoubleZero network can augment both.

Some games use peer-to-peer models, whereby the end players connect directly over the public internet.[7] These are relatively efficient, particularly when there are few players, but these models are directly constrained by the limitations of public internet. Latency, constrained bandwidth, and jitter can lead to frustrated users. The DoubleZero network, with its dedicated links and its flexible usage model, can deliver stability to the connections and make the games more enjoyable.

Other games use the model of a central server to which all players connect. These models are particularly used when there are many players. Historically, these games have been built atop cloud providers, which have their own set of advantages and disadvantages. Cloud providers offer stable performance for players close to the edges of their network, but players far from those edges once again face the limitations of public internet. Moreover, this generates substantial counterparty risk with respect to the cloud provider for the game developer.[8]

The DoubleZero network offers improvements on both fronts. As a network whose links come from a diverse set of contributors, it aims to have wider reach than the networks built by single cloud providers. And as a permissionless and decentralized network, it limits counterparty risk to the underlying user.

Training Large Language Models

As large language models have exploded in popularity, new architectures have been proposed to optimize the training of those models. In particular, one class of proposals is to distribute training amongst many data centers. It is difficult to co-locate large numbers of GPUs in a single facility and align their spare capacity, and so large tech firms have designed systems to coordinate asynchronous and distributed resources such that they can train unified models, e.g. Gangidi et al [2024].

But one limiting factor for distributed training is often bandwidth, as Patel et al [2024] explain, particularly when training is deeply distributed or even mildly concurrent. The DoubleZero network

[7] OatmealDome [2024] offers a particularly detailed description of Splatoon as an example of this.
[8] Epic Games, whose game Fortnite runs on Amazon Web Services and who is also fighting a bitter lawsuit against Apple for exclusion from the App Store, offers a useful example of the delicate balance that game developers must strike with the hyperscalers.

can ease this restriction. Models can be trained more seamlessly across a range of facilities and GPU availabilities, with the DoubleZero network acting as the high-bandwidth connection between those resources to transfer or duplicate snapshots of the model. This ensures that spare compute does not sit idle while snapshots slowly transfer over the public internet. Given its open ethos, the DoubleZero system can help both an entity improve performance across its own distributed fleet and an entity renting machines from independent fleets, and thus allow for more rapid innovation in the field.

Enterprise Usage

Private internet has existed for enterprises for a long time. These private networks, where network service providers lease bandwidth to companies for their internal operations, allow companies to bypass the limitations of the public internet. Indeed, many of DoubleZero's anticipated connections will have this underlying dynamic.

However, this business model is inflexible. Bandwidth is sold in fixed capacities and for long durations. Leases must be negotiated, and new customers must follow time-consuming onboarding procedures. This squares poorly with enterprise usage, where bandwidth needs are often flexible (e.g. an unanticipated surge of traffic). Cloud providers fill this business need better, but they come with their own downsides – high cost, limited network reach, ingress/egress congestion at major data center locations, and ecosystem lock-in.

The DoubleZero protocol can offer a better way. By giving users flexibility to shift latency, throughput, and duration – and by allowing them to purchase and utilize that bandwidth permissionlessly, without locked-in leases or contracts – DoubleZero's alternative solution can better fit the business needs of the fast-moving marketplace. In fact, DoubleZero can even act as a buffer for existing leases, allowing enterprises to respond nimbly to changing conditions.

Conclusion

At their foundations, performant centralized systems and performant decentralized systems look remarkably similar. Consider the modern x86 processor, with its twenty-four cores, integrated GPU, and specialized execution silicon. While it acts as a single entity, its hardware architecture is deeply distributed: twenty-four independent cores working on separate tasks, connected by a high bandwidth interconnect on the chip that moves data and instructions between the cores at defined ingress and egress points.

But the x86 processor also offers a cautionary tale. The limiting factor for the chip's performance is almost never the cores, but the movement of data between cores. Doubling the cores does not yield double the performance unless the on-chip interconnect, memory controllers, and other data flow systems can make similar gains in efficiency. Without these optimizations, each core wastes more cycles waiting for instructions than actually processing instructions.

Decentralized applications today are increasingly hitting bottlenecks in data flow, limiting their performance. DoubleZero, with its robust filtering, configurable routing, and generous bandwidth, is the solution to unshackle them. This is the new internet for modern distributed systems.

What DoubleZero does is simple: Increase Bandwidth, Reduce Latency.

References

Federal Communications Commission (2007). *Statistics of communications common carriers (2006-2007)*.

Fiber Broadband Association and Cartesian (2023). *Fiber Deployment Annual Report 2023*.

FTTH Council Europe (April 5, 2024). “EU39 reaches 70% FTTH/B coverage.” *Shaping Europe's digital future*. European Commission.

Gangidi, A., Miao, R., Zheng, S., Bondu, S., Goes, G., Morsy, H., Puri, R., Riftadi, M., Shetty, A., Yang, J., Zhang, S., Fernandez, M., Gandham, S., and Zeng, H. (2024). *RDMA over Ethernet for Distributed Training at Meta Scale*. ACM SIGCOMM 2024 Conference.

MegaLabs (2024). “Megaeth: Unveiling the First real-time blockchain.” MegaETH Research blog.

OatmealDome (2024). “Splatoon 2's netcode and matchmaking: An in-depth look.” Personal blog.

Patel, D., Nishball, D., and Ontiveros, J.E. (September 4, 2024). “Multi-Datacenter Training: OpenAI’s Ambitious Plan To Beat Google’s Infrastructure.” *Semi-Analysis*.

Schwarz-Schilling, C., Saleh, F., Thiery, T., Pan, J., Shah, N. and Monnot, B. (2023). *Time is money: Strategic timing games in proof-of-stake protocols*. arXiv:2305.09032.