

WAKU-RLN-RELAY: Privacy-Preserving Peer-to-Peer Economic Spam Protection

Sanaz Taheri-Boshrooyeh^{*†}, Oskar Thorén^{*†}, Barry Whitehat[‡], Wei Jie Koh[§], Onur Kilic[¶], and Kobi Gurkan^{||}

^{*}Vac Research and Development, [†]Status Research and Development, Singapore, [‡]Unaffiliated,

[§]Independent, [¶]Unaffiliated, ^{||}cLabs

sanaz@status.im, oskar@status.im, barrywhitehat@protonmail.com,

contact@kohweijie.com, onurkilic1004@gmail.com, me@kobi.one

Abstract—WAKU-RLN-RELAY is an anonymous peer-to-peer gossip-based routing protocol that features a privacy-preserving spam-protection with cryptographically guaranteed economic incentives. While being an anonymous routing protocol where routed messages are not attributable to their origin, it allows global identification and removal of spammers. It addresses the performance and privacy issues of its counterparts including proof-of-work and reputation-based schemes. Its light computational overhead makes it suitable for resource-limited environments. The spam protection works by limiting the messaging rate of each network participant where rate violation results in financial punishment. We deploy the novel construct of rate-limiting nullifier to enforce the message rate limit. We provide a proof-of-concept implementation of WAKU-RLN-RELAY to prove the efficiency and feasibility of our solution.

Index Terms—P2P, Spam Protection, Messaging, zkSNARKs, Zero-Knowledge, Anonymity, Routing, Pub/Sub, Gossipsub

I. INTRODUCTION

WAKU [1] is a family of peer-to-peer (p2p) protocols that enable anonymous and censorship-resistant communication over a network of heterogeneous peers including resource-restricted devices. It features WAKU-RELAY protocol [2] which is an anonymous gossip-based Pub/Sub protocol. Peers congregate around topics and can send messages to topics rather than individuals. Messages are routed using the gossip protocol. Receiver anonymity results from the gossip nature of the routing protocol. Sender anonymity is protected by anonymizing protocol messages i.e., removing personally identifiable information (PII) that binds a message to its owner e.g., IP address and digital signatures.

Being anonymous and an open messaging network, WAKU-RELAY is prone to spam activities where spammers send bulk messages, exhaust the resources of the entire p2p network and cause a denial of service attack. Global identification and removal of spammers are required to protect network resources and liveness. However, as messages are not attributable to their origin, conventional spam-prevention techniques like IP blocking are not effective.

None of the state-of-the-art p2p spam protection techniques i.e., *Proof of Work* (PoW) [3] deployed by Whisper¹ and *Peer scoring* [4] adopted by Libp2p GossipSub provide global spam-protection. Moreover, PoW is computationally expensive hence not suitable for resource-constrained devices. The peer scoring is also prone to censorship and inexpensive attacks

where millions of bots can be deployed to send bulk messages.

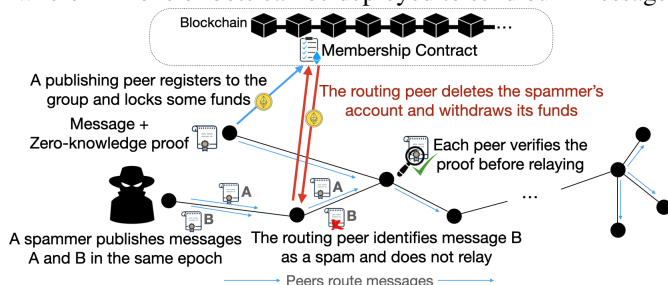


Figure 1. WAKU-RLN-RELAY Overview.

WAKU-RLN-RELAY deals with the aforementioned issues. It extends WAKU-RELAY with the novel construct of Rate Limiting Nullifiers (RLN) to enable an anonymous spam-protected p2p gossip-based Pub/Sub routing protocol which 1) is computationally efficient and suitable for resource-restricted devices 2) preserves user anonymity 3) controls spammers globally 4) mitigates Sybil attack 5) has built-in economic incentives where spammers are financially punished and those who find spammers are rewarded.

II. PRELIMINARIES: RATE LIMITING NULLIFIER

RLN [5] is a zero-knowledge and rate-limited signaling framework in which each user can only send 1 message for each *External Nullifier*. External nullifier is an application-dependent value and resembles a voting booth [6]. The group of authorized users is represented by a Merkle tree called membership tree whose leaves are members public keys pk . The corresponding secret keys sk are known to their owning users. pk s are cryptographic hash of sk s. The membership tree is stored in a smart contract deployed on the Ethereum blockchain. Users need to stake some Eth in the contract to join the group. User removal is done by passing a member's secret key to the contract. Subsequently, a portion of the staked fund of the deleted member is burnt and a portion is given to whoever does deletion.

Only group members can signal. Users use zero-knowledge proof (specifically Zero-Knowledge Succinct Non-Interactive Argument of Knowledge [7]) to prove their membership to the group without disclosing which members they are. In a nutshell, zero-knowledge proof systems enable one party to prove the possession of certain information to a verifier without revealing any other information. Each RLN signal consists of the following parts $(m, \emptyset, \phi, [sk], \pi)$. m is the message. \emptyset is

¹<https://eips.ethereum.org/EIPS/eip-627>

the external nullifier. $\phi = H(sk, \emptyset)$ is called *internal nullifier* which is member's unique fingerprint for that external nullifier. $[sk]$ is a Shamir [8] share of sk deterministically derived from m , sk and \emptyset . π is a zkSNARK proof that asserts 1) sk belongs to the membership tree 2) $[sk]$ is a valid share of sk 3) the internal nullifier ϕ is correctly calculated.

A signal is deemed valid if its π field gets verified by the verification algorithm supplied by the underlying zkSNARK proof system. When a user creates two different messages for the same external nullifier, its signals end up with identical internal nullifiers which signify double-signaling. The two signals will each reveal a distinct share of the user's secret key by which one can reconstruct the sk and remove the user.

III. WAKU-RLN-RELAY CONSTRUCTION

In this section, we explain the end-to-end integration of RLN in WAKU-RLN-RELAY and our design choices to achieve a cost-efficient and scalable solution that suites resource-limited devices.

In WAKU-RLN-RELAY, a smart contract is deployed on the public Ethereum blockchain which acts merely as a registry keeping an ordered list of users public keys, unlike the initially proposed RLN construction which maintains the membership tree on the smart contract. The Merkle tree is maintained off-chain by individual peers. This design choice enables constant complexity registration and deletion operations (as opposed to logarithmic complexity in on-chain tree storage) hence optimizing gas consumption by an order of magnitude.

In contrast to RLN where signals are kept on the contract, in WAKU-RLN-RELAY, messages are stored off-chain and distributed via p2p routing protocol of WAKU-RELAY. Thus, we achieve higher message propagation speed as opposed to the on-chain case where messages should be mined before being visible to the network. Moreover, we save our users the gas price that they have to otherwise pay to insert their messages to the contract.

An overview of WAKU-RLN-RELAY is provided in Figure 1. Peers that belong to the same GossipSub layer i.e., subscribed to the same topic form an RLN group. Peers route messages that belong to peers registered in the RLN group.

Registration: A peer submits a transaction to the membership contract embodying its public key pk as well as v amount of Eth to be staked on the contract.

External nullifier We use *epoch* as the external nullifier. *epoch* is defined as the number of T seconds that elapsed since the Unix epoch. Peers monitor the current *epoch* locally and are allowed to publish one message per *epoch*.

Publishing: Each peer is allowed to send one message m per *epoch*. Message publishing in the network is the same as the RLN framework. The message owner, attaches the nullifiers, together with a share of its secret key, and the zero-knowledge proof part to the message. The peer wraps this data inside a GossipSub message and submits it to the network as in the normal routing protocol.

Group Synchronization: Publishing peers must always stay in sync with the latest state of the group. Otherwise, by making proof of membership to an old version of the membership tree

they can risk exposing the index of their public key in the tree hence compromising their anonymity. Upon member update, the membership contract emits update events by listening to which peers can update their local trees.

Routing and Slashing: A routing peer follows the regular routing protocol of WAKU-RELAY i.e., GossipSub protocol [2] and additionally does the verification steps of the RLN framework. That is, it verifies the RLN proof and discards messages with invalid proof. The routing peer also validates the epoch of the incoming message against its local epoch to see if their difference exceeds a threshold Thr in which case the message is considered invalid and gets dropped ($Thr = \frac{D}{T}$ where D is the maximum network delay and T is the length of the *epoch*). Epoch validation prevents a newly registered peer from spamming the system by messaging for all the past epochs. Next, the internal nullifier is verified to spot spam activity i.e., double-signaling. To do this, each routing peer locally keeps a record of the secret key share $[sk]$ and the internal nullifier ϕ of all of its incoming messages for the past Thr epochs. This list is called a nullifier map. The routing peer checks every new message against this list to spot spam messages i.e., messages with identical internal nullifiers. Note that the nullifier map suffices to hold messages the belong to the last Thr epochs because older messages are considered invalid by default. If all the preceding checks pass, the message gets relayed.

IV. SECURITY AND PERFORMANCE ANALYSIS

Performance: The proof of concept implementation of WAKU-RLN-RELAY is available in [9], [10] which uses RLN library [11]. Generating membership proof to a group size of 2^{32} takes $\approx 0.5s$ on an iPhone 8 [11]. Proof verification run time is constant and takes $\approx 30ms$. Each peer persists a 32B public and secret keys and a $\approx 3.89MB$ prover key [11]. Also, storage of membership tree with depth 20 takes up 67MB from each peer (this can be optimized to 0.128KB using [12]).

Security: WAKU-RLN-RELAY achieves global spam protection while preserving user anonymity, namely peers 1) do not disclose any piece of PII in any phase 2) prove their compliance with the messaging rate without leaving any trace to their public keys. Sybil attack is also mitigated by making registration expensive. Economic incentives are guaranteed cryptographically via secret sharing.

REFERENCES

- [1] "<https://rfc.vac.dev/spec/10/>."
- [2] "<https://rfc.vac.dev/spec/11/>."
- [3] C. Dwork and M. Naor, "Pricing via processing or combatting junk mail," in *Annual international cryptology conference*. Springer, 1992.
- [4] "<https://github.com/libp2p/specs/blob/master/pubsub/gossipsub/gossipsub-v1.1.md>."
- [5] "<https://ethresear.ch/t/semaphore-rln-rate-limiting-nullifier-for-spam-prevention-in-anonymous-p2p-setting/5009>."
- [6] "<https://semaphore.appliedzkp.org/>."
- [7] J. Groth, "On the size of pairing-based non-interactive arguments," in *EUROCRYPT*. Springer, 2016.
- [8] A. Shamir, "How to share a secret," *CACM*, vol. 22, 1979.
- [9] "<https://github.com/status-im/nim-waku>."
- [10] "<https://rfc.vac.dev/spec/17/>."
- [11] "<https://github.com/kilic/rln>."
- [12] "<https://github.com/vacp2p/research/blob/master/rln-research/merkle-tree-update.md>."