Securing Distributed Disaster Early Warning Networks

Kumudu Senarathne
Dept. of Electronic & Telecommunication Engineering
Universiy of Moratuwa
Sri Lanka
senarathnegdkd.22@uom.lk

Dileeka Dias
Dept. of Electronic & Telecommunication Engineering
University of Moratuwa
Sri Lanka
dileeka@uom.lk

Raj Prasanna
Joint Centre for Disaster Research
Massey University
New Zealand
R.Prasanna@massey.ac.nz

***Abstract*—This paper proposes a secure and cost-effective communication and management approach for distributed disaster early warning networks to ensure reliable, low-latency alert dissemination. While decentralized architectures reduce latency by enabling direct device communication, they are highly vulnerable to attacks due to their exposure in public community settings and the challenges posed by Network Address Translation (NAT) and firewalls. Traditional NAT traversal techniques often lack the built-in security required for such sensitive infrastructure.**

We identify Nebula, an open-source overlay networking tool, as a viable solution because it is lightweight enough for low-resource devices such as the Raspberry Pi typically found in decentralized early warning systems, supports NAT traversal, and provides end-to-end encryption. Preliminary experiments conducted in an AWS environment—designed to simulate real-world Raspberry Shake sensor deployments—confirm Nebula's ability to establish direct peer-to-peer (P2P) connectivity with a single hop and effectively encrypt traffic.

***Index Terms*—Disaster early warning, Peer-to-Peer communication, Nebula, security, cost-effective**

I. Introduction

Natural disasters have become a significant threat, causing severe damage to people and infrastructure [1]. Disaster early warning networks provide timely alerts, helping people take protective action before hazards. However, ensuring the reliability of these systems is essential, as false alarms can mislead the community.

Traditional centralized networks introduce an additional delay in communication, as data must pass through a central server before reaching the end device. Furthermore, if the central server is interrupted during a disaster, the entire network gets interrupted. Moreover, the high load at the central point of the network affects the performance of each device. A decentralized architecture overcomes these limitations by distributing data processing to end devices. In this decentralized model, sensors are installed at community houses and connected to the network using its available household internet connection. Exposure to the public environment makes devices vulnerable to third-party attacks leading to loss of trust on the early warning network. To ensure security in such community-based networks, sensors work behind Network Address Translation (NATs) and firewalls, making it challenging to establish direct peer-to-peer (P2P) communication between sensors. Several NAT traversal techniques, such as hole punching [2], are used to connect devices directly. However, they do not have built-in security features to secure data communication from attacks.

Communication methods used in a community-based disaster early warning system should be able to penetrate firewalls and NAT to enable direct P2P communication, yet should have built-in security features. In some instances of real-world deployments of early warning systems, a mesh network of ground motion detection devices such as Raspberry Shake sensors are used as the end devices [1]. Hence, any security features should be implemented on such low-resource devices. Furthermore, this network should have self-healing ability during loss of connection and should be cost-effective to support large-scale community-based deployments.

In this paper, a cost-effective and secure communication method is proposed with an enhanced management approach for disaster early warning networks. Section II of this paper summarizes the existing solutions, and the proposed methodology and preliminary results are presented under Section III.

II. Literature Survey

Implementation and testing of a decentralized earthquake early warning (EEW) system using MEMS-based Raspberry Shake sensors across New Zealand is reported in [1]. A novel SD-WAN technology, Zerotier [3], was used to establish P2P communication, which is based on UDP (User Datagram Protocol) hole punching. Experiments are conducted on multiple earthquake scenarios to measure the system latency. The paper concludes that the community-based EEW networks are technically feasible using Zerotier, which saves precious seconds in detection despite being costly.

Quick UDP Internet Connections (QUIC) is another communication method tested for P2P communication in EEW systems in [4]. QUIC was introduced by Google, and it is the foundation for the HTTP/3 protocol [5]. This reduces the round-trip delay by combining the transport layer stream abstraction, the encryption handling, and parts of the application layer into a single layer. The authors tested QUIC against Zerotier and concluded that a lower latency can be achieved with QUIC over UDP hole punching for P2P communication.

https://doi.org/10.54389/VMWW6958

TABLE I
SUMMARY OF P2P COMMUNICATION METHODS.

Platform/Tool	Description	Open-source/Closed-source	Summary
Hamachi	Use STUN (Session Traversal Utilities for NAT) servers to find the public IP and port of peers, and if direct connectivity fails, relay traffic via a TURN (Traversal Using Relay NAT) server	Closed-source	Wolinsky et al. [2] documented that Hamachi supports P2P communication using its central management and security infrastructure. But it disables users from hosting their own infrastructure.
Netbird	Based on WireGuard protocol and establishes a full mesh encrypted tunnel between devices. A web interface is available for the network administration.	Open-source	Vojdan et al [6] evaluated Netbird for a secure edge cloud continuum and concluded that it is suitable for high-performance and low-latency environments when the traffic is tunneled using UDP protocol but not for TCP
Tailscale	Built on WireGuard [7]. Establishes end-to-end encrypted direct connectivity between peers using minimal configuration. This uses STUN for NAT traversal and relay the traffic via DERP servers if the direct connectivity fails.	Client code is open source, but the control server code is closed	Imam Mehrab [8] evaluated Zerotier against Tailscale and concluded that, Zerotier outperforms Tailscale in terms of direct connectivity success rate, faster connection establishment, resource utilization, throughput, and lower latency.
Nebula	A software-defined networking (SDN) solution introduced by Slack [9]. Designed for high performance and automatic peer discovery, making it ideal for large-scale deployments.	Open-source	Carpio et al [9] investigated the Nebula for edge computing environments and concluded that it is a light-weight and reliable solution.
Zoom	Establish direct P2P connectivity using STUN and TURN servers. The coordination of direct connectivity establishment is performed via ICE negotiations. Dedicated servers are required for STUN/TURN servers.	Open-source	Andersson et al [10] evaluated the impact of NAT traversal on P2P communication and concluded that ICE negotiations add significant delay in direct connectivity establishment. And it does not encrypt the traffic.

However, more work is required before applying this in real-world deployments.

Although there are several P2P communication methods available in the literature, they have not been directly used for disaster early warning networks. Table I shows a summary of these methods.

Among the communication methods described above, Nebula is selected for further study as it is an open-source software capable of establishing direct P2P communication, lightweight to run on Raspberry Pi devices, provides data security by encrypting the traffic, and supports NAT traversal. Nebula is freely available on GitHub. It has a Certificate Authority (CA) to issue certificates for each client for authentication. For the initial connection setup, clients contact the lighthouse server to get details about the other clients. Once the information has been obtained from the lighthouse, clients can connect with each other directly [11].

III. PRELIMINARY TESTS AND RESULT

Our expectation is to evaluate and enhance Nebula's performance to meet the requirements of decentralized disaster early warning systems.

A. Proposed Implementation Setup

In the real-world disaster early warning networks, Raspberry Shake sensors are used as the end devices. Our experimental setup should be as close as possible to how Nebula would behave when implemented on Raspberry-pi devices distributed across an area.

The experimental setup is shown in Figure 1. It deploys client VMs across different AWS Availability Zones (AZs)—physically isolated datacenters—for network separation. A t4g.micro VM (ARM-based, low-cost Graviton2) type is chosen to simulate a Raspberry Shake sensor behind unique public IPs and firewalls.

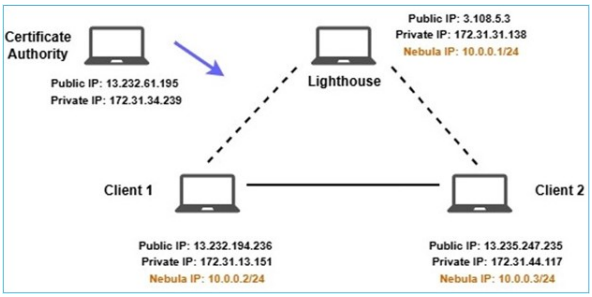

Fig. 1. Proposed Implementation Setup.

B. Preliminary Results

Initially, we tested Nebula's capability to establish direct P2P connectivity. After executing the Nebula on each client,

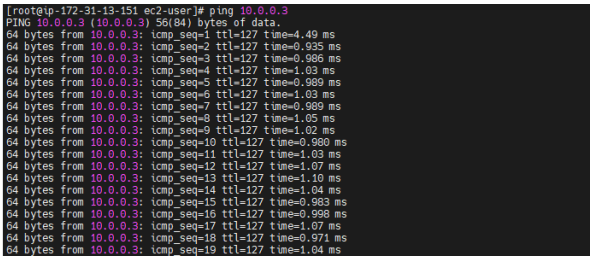

Fig. 2. Ping Test Result.

we performed a *ping* test, and the *traceroute* result was also analyzed. *Ping* test is used to verify the destination device is reachable and to measure the round-trip time (latency). In the *ping* test results, the first round-trip value is comparatively higher as shown in Figure 2, and this occurs because the client contacts the lighthouse first for the initial peer discovery.

Traceroute tests help map the path an IP packet takes from source to destination. As shown in the Figure 3, the *traceroute* result contains only one hop. This confirms that Nebula has established direct peer-to-peer connectivity.

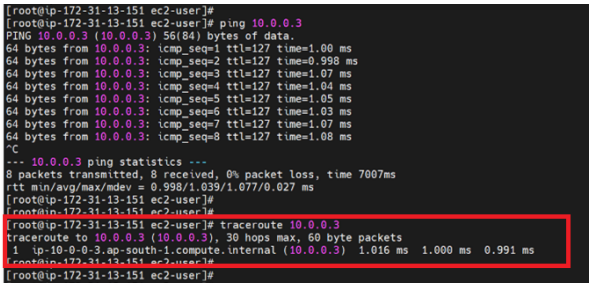

Fig. 3. Traceroute Result.

Further, clear ICMP packets were visible on the Nebula interface, and there was no readable content on the physical interface, as shown in Fig. 4. This confirms that Nebula encrypts the traffic effectively.

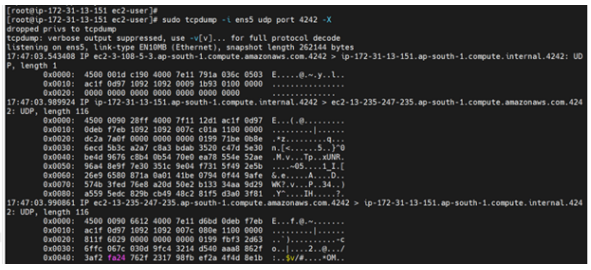

Fig. 4. Traffic Capture on Physical Interface.

IV. Future Work

Communication latency is a critical factor for a distributed disaster early warning network. We expect to perform latency tests multiple times at different times of the day over a specific period of time. Further, direct peer-to-peer connectivity, packet encryption, resource utilization, and the effect of encryption overhead will be tested. Client reconnection to the network after a network disruption will also be evaluated to ensure the self-healing ability of the network.

In the default Nebula setup, CA generates the client key and the client certificate files and sends them to the client over the internet. Due to the security vulnerability of sending the client key over the internet, we propose to generate the key file within the client and send the public key file to the CA to sign.

Furthermore, after installing the relevant Nebula version, a configuration file needs to be executed on each device. Necessary changes should be made in this file to update the certificate and key file path, Nebula IPs, and lighthouse details. Editing this configuration file manually for each device is not possible in large-scale deployments. Therefore, we propose to automate the configuration file generation. In order to reduce the management complexities, we expect to build a management portal that enhances the visibility of the network.

Finally, the proposed system performance will be compared with the Zerotier by implementing and testing Zerotier in the same experimental setup under similar conditions. With that applicability of Nebula disaster early warning systems to ensure secure peer-to-peer communication will be evaluated.

V. Conclusion

The objective of this research is to propose a secure and cost-effective communication and management approach for distributed disaster early warning networks. This work-in-progress paper proposes Nebula as a cost-effective and secure communication method, and key findings from initial tests are summarized. Nebula is a open-source, secure overlay communication method and the preliminary results prove its suitability for disaster early warning networks. We propose future enhancements to the proposed solution fulfilling the key requirements of distributed disaster early warning systems.

References

[1] Prasanna, Raj, Chanthujan Chandrakumar, Rasika Nandana, Caroline Holden, Amal Punchihewa, Julia S. Becker, Seokho Jeong et al. ”“Saving Precious Seconds”—A novel approach to implementing a low-cost earthquake early warning system with node-level detection and alert generation.” In Informatics, vol. 9, no. 1, p. 25. MDPI, 2022.

[2] Wolinsky, David Isaac, Kyungyong Lee, Tae Woong Choi, P. Oscar Boykin, and Renato Figueiredo. ”Virtual Private Overlays: Secure Group Commounication in NAT-Constrained Environments.” arXiv preprint arXiv:1001.2569 (2010).

[3] “ZeroTier - Global Area Networking,” Zerotier.com, 2023. https://www.zerotier.com

[4] Hong, Benjamin, Chanthujan Chandrakumar, Danuka Ravishan, and Raj Prasanna. ”A peer-to-peer communication method for distributed earthquake early warning networks: Preliminary findings.” (2024).

[5] Gbur, Konrad Yuri, and Florian Tschorsch. ”A QUIC (K) way through your firewall?.” arXiv preprint arXiv:2107.05939 (2021).

[6] Kjorveziroski, Vojdan, Cristina Bernad, Katja Gilly, and Sonja Filiposka. ”Full-mesh VPN performance evaluation for a secure edge-cloud continuum.” Software: Practice and Experience 54, no. 8 (2024): 1543-1564.

[7] “How Tailscale works,” tailscale.com. https://tailscale.com/blog/how-tailscale-works

[8] Mehrab, Abu Imam. ”A New Approach to Peer-to-Peer VPN Connectivity: Achieving Seamless Communication Without Firewalls.”

[9] Carpio, Francisco, Marc Michalke, and Admela Jukan. ”Engineering and experimentally benchmarking a serverless edge computing system.” In 2021 IEEE Global Communications Conference (GLOBECOM), pp. 1-6. IEEE, 2021.

[10] Maenpaa, Jouni, Veera Andersson, Gonzalo Camarillo, and Ari Keranen. ”Impact of network address translator traversal on delays in peer-to-peer session initiation protocol.” In 2010 IEEE Global Telecommunications Conference GLOBECOM 2010, pp. 1-6. IEEE, 2010.

[11] TheOrangeOne, “Nebula mesh network- an introduction,” TheOrangeOne, 2021.https://theorangeone.net/posts/nebula-intro/ (accessed Feb. 20, 2026).