Vulnerability Analysis of Cross-Layer Attacks on ROS 2

Hyunho Ryu[1], Sanghoon Lee[2], Kyung-Joon Park[3]
Department of Electrical Engineering & Computer Science (EECS)
Daegu Gyeongbuk Institute of Science and Technology (DGIST)
Daegu, Korea
{ryuhyunho, leesh2913, kjp}@dgist.ac.kr
[1]0000-0002-2860-1612, [2]0000-0002-8160-8952, [3]0000-0003-4807-6461

***Abstract*—Robot Operating System 2 (ROS 2) architecture based on the DDS is de facto standard of network architecture for distributed robot systems. However, ROS 2's distributed discovery mechanism expands the attack surface, making it vulnerable to security threats. If attacker infiltrates the network, data confidentiality and integrity can be compromised through traffic eavesdropping and modification. This can lead to Denial of Service (DoS) conditions and system malfunctions, thereby undermining the safety and reliability of robotic systems. This paper analyzes these vulnerabilities and verifies the feasibility of real attacks based on the expanded attack surface.**

***Index Terms*—Robot Operating System, DDS Security**

I. Introduction

Robot Operating System (ROS) is a widely adopted standard platform in robotics, providing software libraries and development tools that integrate sensors, actuators, and networks. It adopts a node-based structure, where each node is an independent process that executes a specific function [1]. Unlike the centralized master architecture of ROS 1, ROS 2 uses the Data Distribution Service (DDS) for communication, implementing efficient topic-based publish-subscribe messaging between distributed nodes. DDS is a based on a data centric design and implemented using the Real-Time-Publish-Subscrib (RTPS) standard [2]. This architecture provides high efficiency and low overhead, while meeting the requirements of various robot tasks through Quality of Service (QoS) policies, even in multi robot environment. Robot environments have very high demands on the safety and reliability of communication between nodes, and each node must always operate stably and be highly sensitive to communication delays or interruptions [3], [4].

ROS 2 supports dynamic discovery for distributed communication using multicast Simple Participant Discovery Protocol(SPDP). However during this process, SPDP packets are transmitted in plain text, exposing participant identifiers and unicast locators. Exposure creates attack surface that enables cross-layer exploits, including ARP-based attack. This study systematically analyzes these vulnerabilities and demonstrate the feasibility of cross-layer attacks in ROS 2.

This paper is supported by the Institute of Information and Communications Technology Planning and Evaluation (IITP) Grant funded by Korean Government [Ministry of Science and ICT (MSIT)] under Grant RS-2024-00442085.

II. Vulnerabilities in DDS process

A. ROS 2 DDS Discovery process

The discovery protocol is a necessary procedure for each participant to discover and interact with each other in a distributed network environment. It consists of two steps, SPDP and SEDP.

Step 1. Simple Participant Discovery Protocol (SPDP)

SPDP is an initial discovery protocol that enables all participants in a DDS network to announce their existence. Each DomainParticipant entity regularly transmits basic discovery information, including its GUID, unicast locator, and QoS flag, via UDP multicast. This allows other participants within the same domain to discover the existence of nodes and their fundamental communication parameters.

Step 2. Simple Endpoint Discovery Protocol (SEDP)

After the SPDP process, nodes share information about their endpoints, including the existence of data senders and receivers (DataWriter/DataReader), topic types and QoS policies.

During this process, DDS discovery information is transmitted in plaintext. This means that attackers can easily detect the existence of robot nodes using capturing network traffic. This demonstrates that the discovery process has potential security vulnerabilities.

B. Attack Surface

During the SPDP process, plaintext UDP packets are exposed to the network containing critical metadata such as identifying information like unicast IP/Port. This information is not encrypted during the DDS discovery phase, and the UDP/IP headers also unprotected. Attackers can use the exposed participant information to identify node existence and participant identifiers, which makes it easy to target specific attacks. Using the collected IP and MAC addresses, attackers can start malicious activities such as ARP sniffing, ARP spoofing, packet modification and hijacking.

III. Threat Model

This section presents a step by step attack chain, a scenario that exploits vulnerabilities in the discovery process of ROS 2 DDS based robot system.

A. Node Profile in Multicast Traffic

In the RTPS standard, multicast is used by default for the SPDP process, with 239.255.0.1 designated as the default multicast address. Each participant joins this multicast and regularly transmit SPDP packets. Even when a non-standard multicast address is configured in a specific robotic environment, the UDP/IP headers remain unencrypted and thus exposed in plaintext, allowing unicast locator information to be collected without restriction. Furthermore, RTPS port allocation rules are formally defined, and the port numbers increase sequentially, enabling active multicast traffic to be identified through simple network scanning. As a result, attackers can obtain information such as participant IP addresses and port numbers without requiring any keys or certificates.

B. ARP-based Man-In-The-Middle (MITM) Attacks

Exposure of SPDP packets allows attackers to obtain the IP addresses of robot nodes participating in the network. This enables ARP attacks without requiring additional information, such as the internal structure of ROS 2 DDS, encryption keys, or other secure data. For operational convenience, the ARP protocol does not require additional authentication procedures during the MAC address resolution process. It follows a last-writer-wins logic, meaning that even without sending an ARP request, the most recently received ARP reply is prioritized and applied. By exploiting this behavior, an attacker can poison the ARP cache of a target node by associating the attacker's MAC address with the robot's IP address. This feature can be coducted within the same network without administrator privileges and bypass ROS 2 layer protections (e.g., SROS 2), enabling cross-layer manipulation of communication.

By exploiting vulnerabilities in the discovery process, an attacker can gathering IP addresses of participating nodes and leverage them to conduct subsequent attacks across different layer. This highlights that information exposure during the discovery process not only creates an attack surface but also serves as a prerequisite for cross-layer attacks.

IV. Experiment

After obtaining information through the attack surface, and once MITM attack is successful, the attacker can attempt various follow-up attacks. This paper presents the possibility of attacks and their scope of impact depending on whether SROS 2 is enabled.

- **MITM attacks** involve intercepting and modifying RTPS payloads to inject robot control commands or modify sensor data. In configurations where SROS 2 is disabled, attackers can monitor and modify the contents of all packets. However, when SROS 2 is enabled, payloads are encrypted, making content based modification impossible. Even though techniques such as packet timing analysis can still be used to infer traffic patterns.
- **DoS attacks** can be carried out through ARP cache poisoning, and attackers can interrupt communication between nodes by controlling the traffic flow of the target node and intentionally blocking all packet transmission. In ROS 2, even if a MITM attack occurs through ARP cache poisoning, it cannot be detected at the ROS 2 layer, and nodes continue to believe that communication is normal. This attack method causes serious availability issues regardless of whether SROS 2 is enabled, as it blocks the communication path itself.

The experiments were carried out using ROS 2 Humble with Fast-DDS version 2.6.9, while the attacker operated from a standard computer environment outside of the ROS 2 framework.

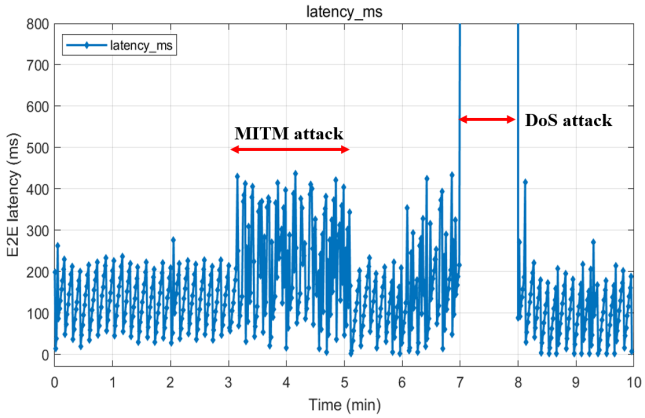

Fig. 1. Impact of vulnerability exploitation on system latency.

Fig 1 shows that network delay increases when a MITM attack is started, and traffic control enables intentional packet drops, leading to a DoS condition. In particular, a temporary delay increases occurs when packets are dropped through traffic control due to the ROS 2 retransmission mechanism.

V. Conclusion

In this paper, we analyzed vulnerabilities in the discovery process of ROS 2 DDS and validated the attack surface through practical experiments. We demonstrated that the SPDP process expands the attack surface and that a real MITM attack can cause network delays and even lead to a DoS condition. These vulnerabilities present significant challenges for robotic systems where real time performance and security are critical. As future work, we plan to propose strategies for detecting and mitigating cross-layer attacks within the ROS 2 layer.

References

[1] H.-S. Park, S. Lee, D. Um, H. Ryu, and K.-J. Park, "An analytical latency model of the data distribution service in ROS 2," in *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*. IEEE, 2025, pp. 1–10.

[2] S. Lee, H.-S. Park, J. Chae, and K.-J. Park, "Probabilistic latency analysis of the data distribution service in ROS 2," 2025. [Online]. Available: https://arxiv.org/abs/2508.10413

[3] G. Deng, G. Xu, Y. Zhou, T. Zhang, and Y. Liu, "On the (in) security of secure ROS 2," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 739–753.

[4] S. Lee, T. Kim, J. Chae, and K.-J. Park, "Optimizing ROS 2 communication for wireless robotic systems," 2025. [Online]. Available: https://arxiv.org/abs/2508.11366