2022 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
CYBER TECHNICAL SESSION
AUGUST 16-18, 2022 - NOVI, MICHIGAN

Automotive Ethernet Cyberattack Defense in Ground Vehicles

Peter Moldenhauer[1], Jonathan Esquivel[1]

[1]Southwest Research Institute, San Antonio, TX

ABSTRACT

This paper describes the strategies and challenges involved to secure vehicles which use automotive Ethernet-based networks. Since the early 1990's, the Controller Area Network (CAN) bus has been the standard in automotive networking systems. However, automotive Ethernet is becoming more common in recent years and is considered the future in automotive networking. This new technology has unique advantages over traditional CAN bus networks (e.g. higher bandwidth that can support hashing and encryption), and it still requires additional security measures such as monitoring and detection of anomalies to better secure the vehicle. Southwest Research Institute (SwRI) has previously developed a CAN-only intrusion detection system (IDS) which protects a vehicle's CAN bus by actively monitoring traffic and flagging messages that are identified as anomalies. SwRI successfully implemented the ability to read, train, and detect on automotive Ethernet data in the IDS. The integration of automotive Ethernet in the IDS unveiled numerous challenges and lessons learned throughout its development.

Citation: P. Moldenhauer, J. Esquivel "Automotive Ethernet Cyberattack Defense in Ground Vehicles", In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 16-18, 2022.

1 INTRODUCTION

Ground vehicles are rapidly becoming more complex. As a result, the advanced hardware and software capabilities demand high performance network architectures. Automotive Ethernet-based networks are better suited for supporting this new functionality compared to the traditional Controller Area Network (CAN) bus. An example of this is how automotive Ethernet supports significantly higher throughput rates compared to that of a CAN bus. This expanded bandwidth that is offered by automotive Ethernet allows for easy processing of large amounts of data generated by the multitude of sensors and actuators found in modern ground vehicles. Despite the benefits, there are numerous vulnerabilities of Ethernet, which is then inherited by automotive Ethernet [1]. In addition, many of the security concerns that are present in a CAN bus network persist in an automotive Ethernet-based network. These cybersecurity vulnerabilities present a serious problem that needs to be addressed. Commercial and military ground vehicles alike are at risk of enemy cyber threats.

A solution to this cybersecurity problem has previously been identified through the development of an automotive intrusion detection system (IDS).

This IDS has been presented at GVSETS 2020 [3] and 2021 [2]. Initially developed to only run on a CAN bus, this IDS uses application layer CAN data monitoring to identify anomalous messages in the network. More specifically, the IDS is a monitoring and security tool that watches network traffic in a vehicle and identifies anomalous behavior [2].

To respond to the technological shift from CAN to automotive Ethernet networks, the IDS was expanded to read in automotive Ethernet data in addition to CAN data. This newly added functionality in the IDS allows for the training and anomaly detection of Ethernet data by leveraging the previously CAN-only algorithms in the software. However, this new addition did not come without challenges.

2 BACKGROUND

Ground vehicle technology has historically focused on vehicle durability, reliability, and performance over vehicle security. Even though automotive Ethernet greatly improves performance in a vehicles network, it is still vulnerable to cyber-attacks. As the name suggests, automotive Ethernet is indeed Ethernet-based and thus susceptible to the standard Ethernet-based threats. Some of these threats include Man-in-the-Middle (MITM), protocol fuzzing, and stack-based buffer overflows. Ethernet is also a well-known technology which is extensively described in literature that many people have experience with through day-to-day use both at work and at home. Therefore, the number of capable potential attackers is considered greater than compared to other automotive bus technologies [4].

Despite the major differences between the automotive Ethernet and CAN protocols, the threat models are similar for both types of networks. Attacks on the communication network can occur in the form of deliberately inserted faulty messages or intentional interference with the transmission of correct messages (e.g., manipulation, delay, removal or, replay of messages) [5]. Some of these specific threats include sniffing, spoofing, and denial of service (DoS).

Between the typical Ethernet-based threats, the commonality of experienced Ethernet-based attackers and the standard CAN-based threats that are still applicable, automotive Ethernet is far from secure. As automotive cyberattacks become increasingly prevalent and complex, the need for increased security grows along with it.

2.1 IDS Basics

An automotive IDS is designed to identify anomalous packets in a vehicles network and then send an alert based on the flagged packets [2]. These anomalous packets are generated from unexpected or unusual events (e.g., intrusions designed to control the vehicle or reduce mission readiness) that occur in the vehicles network. There are various strategies an IDS can employ to identify these intrusions. Two of these techniques can be categorized as signature-based and anomaly-based. Signature-based detection uses the characteristics of previously identified attacks to uncover anomalies so packets that do not match any of the recorded signatures are not flagged. Anomaly-based detection examines the behavioral characteristics of the traffic rather than the contents. These characteristics can include timing and distinct sequencing patterns of traffic packets.

2.2 Ethernet Integration Strategy

In the previous CAN-only IDS that was presented at GVSETS 2020 [3] and 2021 [2], the IDS would identify, or key, packets by the arbitration id (ARB ID) of each CAN frame. Acceptance criteria would be built for each key from packets collected in a training set. Detection would then use these acceptance criterium to accept or reject each packet by using an incoming packet’s ARB ID to check if the packet should get flagged as an anomaly.

This identification strategy works for CAN data since the communications are carried out through a bus and can easily be identified by their ARB ID. In the case of automotive Ethernet, which is an IP-based communication, this method does not work as there is not a direct id to describe an ethernet packet in a meaningful way. Automotive Ethernet also communicates with packets of higher verbosity which then increases the needed complexity of detection schemes. To adapt the IDS to fit the automotive Ethernet use case, there were investigations to find a new method of keying automotive Ethernet packets and a push to develop a model for handling new Ethernet packet behaviors.

The team decided to use an identifier key composed of the source and destination's Ethernet (MAC address), IP address, and port for each packet (see Figure 1). This identifier key was able to act as a synonym to the ARB ID in the CAN-only detection. The more verbose key information contained in automotive Ethernet packets (compared to CAN) allowed the team to key off smaller sets of packets and create comprehensive acceptance criteria during training. Rather than describing the behavior of an individual key group, the verbose key allowed descriptions of hierarchical groups, e.g., all packets with the same destination IP, which enables more complex network modeling.

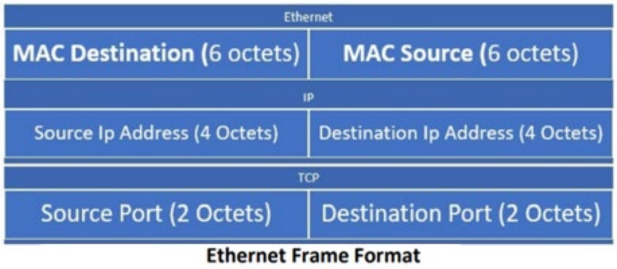
Ethernet

MAC Destination (6 octets) | MAC Source (6 octets)

IP

Source Ip Address (4 Octets) | Destination Ip Address (4 Octets)

TCP

Source Port (2 Octets) | Destination Port (2 Octets)

Ethernet Frame Format

Figure 1. Automotive Ethernet identifier key. Key was composed of both source device information and destination device

3 SYSTEM ARCHITECTURE

The research team leveraged the strategies of the application layer CAN detection algorithms to also accept automotive Ethernet data. By reworking these proven detection techniques, the security of automotive Ethernet can be greatly improved. The team aimed to solve security issues present in an automotive Ethernet-based network in addition to CAN bus networks.

The IDS monitors Ethernet traffic by dissecting the packet into a key composed of source and destination information. This yields the ability to keep track of timings for every key group and a whitelist of valid package signatures. Only certain devices are expected to communicate with other devices, as well as keep track of previous communications and this signature is an essential part to the acceptance criteria. Automotive Ethernet communications follow similar behavior to automotive CAN communications in that each device generates packets at a typical, observable, rate. Each incoming packet is used with the previous packet of a given key to create a time delta that is ran through an acceptance check to determine if the packet should be flagged as anomalous. The anomalies from this first layer, packet attribution layer, are then fed through a second layer, anomaly alerting layer, to make the final decision if an alert should be sent out that an attack is happening on the system.

The result is an IDS that can not only detect anomalous automotive Ethernet packets but also identify the critical nodes of the automotive Ethernet network. Figure 2 describes part of the network the research team used for testing where d1 and d6 are seen as critical nodes.

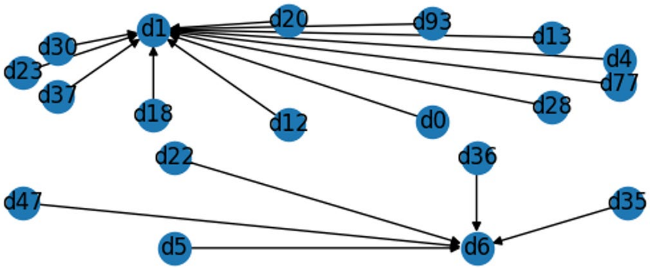

Figure 2. Key nodes of network. D1 and D6 are ciritcal nodes as they receive communications from many nodes

By utilizing the source and destination devices as keys, the team can identify the core source and destination nodes of the network which are the critical points to apply the IDS. Applying tailored monitoring tools to these key nodes, one can greatly increase the success of the IDS as well as gain a deeper understanding of the vehicle network.

3.1 Ethernet Detection Algorithms

The primary methods used for detection are to address timing (signature based) and whitelisting (anomaly based) anomalies. Timing anomalies are addressed by constructing a time delta composed of an incoming packet and the previous packet for a given key. During training, each packet is grouped by its key and the time deltas between all groups are analyzed. The team then models acceptance criterium by estimating a multimodal gaussian curve to fit to the histogram describing all observed packet deltas. This is shown in Figure 3 where the Y-axis represents the likelihood a packet would arrive at the X-axis time delta. The vertical red lines represent local maxima (some of which are very small) and the green windows represent local acceptance criteria. For example, one can see that the packets of this key signature often arrive with a time delta of 0.01 seconds but has significant variation in its tails. During detection, the time delta is compared to the expected time deltas seen in training to determine if the IDS should flag the packet as an anomaly.

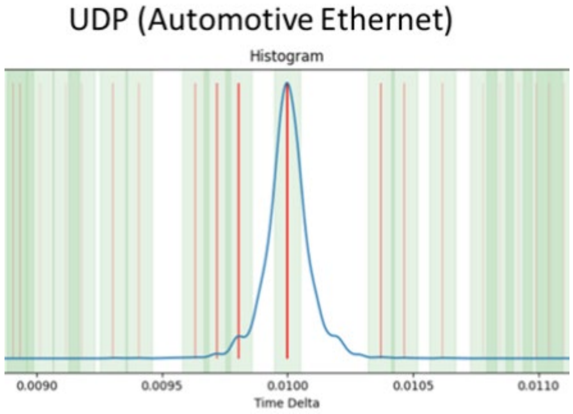

Figure 3. Packet delta probability distribution. The Y-axis represents the "likelihood a packet would arrive" at the -axis "time delta"

Whitelisting anomalies are addressed by flagging devices that communicate with addresses not seen in training. This is an advancement from the CAN-only detection as we can distinguish devices that primarily send packets from devices that primarily receive packets.

4 PERFORMANCE

To determine the success of the IDS, automotive Ethernet traffic was captured for both training and detection. These capture files were then used to train and validate the algorithm offline, in different scenarios, such as a set of normal driving, a set of the vehicle in idle, and a final set of the vehicle in idle under attack. The first dataset was a recording under normal driving conditions and was split into training and testing subsets. These subsets used different lengths of splits to estimate the minimum data required to train a useful model. The next dataset was a recording with the vehicle idling in the "ON" state. This set was to serve as a baseline for the attack set with similar conditions. The final dataset was a recording with the vehicle in idle while an anomalous (address spoofed) packet insertion attack was sent on the network. This final set was used to test the IDS with a simulated threat and evaluate its ability to correctly flag anomalous traffic.

4.1 Training

The first stage of the analysis was set to evaluate not only if the team could accurately model normal traffic but to determine how much data was required to minimize false flagging. To do this, the team would train on different sizes of datasets and test the performance on the last half of the normal driving recording. A summary of these results can be seen in Figure 4. The test results revealed that a longer training time yielded more accurate characterization along with a lower false positive rate. The team found that the IDS can yield a false positive rate of less than 5% with less than a minute of recording. With focus on the packet attribution layer, the IDS was able to model certain communications better than others with the smaller sets of data. This is likely due to some packets requiring a longer training set to properly model.

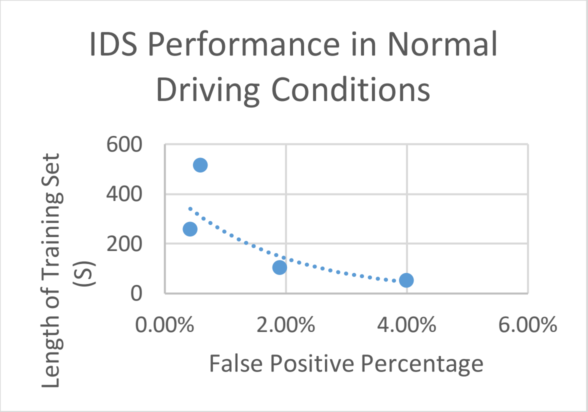

Figure 4. Test results summary. Note lower false positive rates with longer data recordings

4.2 Baseline and Attack sets

The next stage of analysis was to evaluate the IDS's capability to detect true anomalous traffic. The baseline was a recording of network traffic with the vehicle in idle. This baseline would then be used to compare performance against the attack set. Results for the baseline reveal positive results with less than 1% false positives that the model was able to successfully recall normal traffic. The IDS was also able to recognize and flag the attack set immediately following the first attacks resulting in a successful proof of concept.

4.2 True Positive Analysis

The final test was an attack set where the vehicle was in idle, and a throttling attack was sent out on the vehicle networking targeting a set of specific communications. In a throttling attack, an attacker attempts to insert an anomalous packet into the network to effectively "cancel out" the original good message by overwriting the message with similar message containing harmful data. In this test the model was able to recognize attacks but was unable to attribute the exact packet causing the anomaly.

It is important to note that the CAN-only IDS uses a two-layer detection scheme, the first layer focuses on flagging individual packets and the second focuses on interpreting the first layer to properly warn the user only when there is high confidence that there is an attack, not just a false positive. While the automotive Ethernet IDS is still able to maintain a low 2nd layer false positive and still properly catch and flag anomalies, there were unique challenges discovered in the 1st layer detection.

In CAN, when a similar attack is conducted only the attacked communication(s) are affected. In automotive Ethernet this attack caused unrelated communications to be affected resulting in a large number of falsely attributed packets being flagged as the cause of the anomaly. The second layer can mitigate poor packet attribution as shown by the results of low false positives in normal detection. In addition, the second layer is the most important as it is there to filter noise from the first layer and only alerts the operator when there is high confidence of an attack. The increased number of flagging's in the packet attribution layer (due to irrelevant packets being affected by an attack) will still trigger the anomalous decision layer to alert the operator. Packet attribution is still a highly sought-after capability, one that the CAN-only IDS achieves, and further digital forensics developments should take place to enable Automotive Ethernet IDS to improve this capability.

5 FUTURE WORK

Integrating automotive Ethernet detection capabilities into the IDS has been largely successful and there is additional work that can be done to improve the overall effectiveness of the IDS. Firstly, further investigation is recommended to develop by-packet anomaly attribution and an analysis of unrelated traffic during an attack. The IDS can recognize an attack but is unable to maintain the packet attribution. Another option is to log the events and investigate embedded digital forensic techniques to run near real-time analysis of the anomaly events to identify a group of potential packets responsible for the anomaly.

Secondly, the automotive Ethernet and CAN bus software modules of the IDS should be merged. CAN detection logic and the Ethernet detection logic run independent from each other and in the future these modules should run at the same time to enable deployment on more diverse system environments.

Finally, the IDS needs to be tested on more automotive Ethernet capable vehicles. For this project, the Ethernet integration was based off data captures from one (1) vehicle. Vehicles can operate in many ways: slow driving, fast driving, high maneuverability situations, or simply idle and the IDS should be able to model data in any of these states. Certain patterns and network traffic may also be due to the manufacturer and the IDS should be generalized enough to learn any system with minimal adaptations to the training pipeline. To improve the overall effectiveness of the IDS, the research team must do extensive testing with more Ethernet vehicle data. This includes testing with vehicle data in different environmental conditions and network throughputs to build a fieldable IDS.

6 CONCLUSION

Cyberattack defense in ground vehicles is a topic that needs to be addressed. As technology advances, so does the potential for security threats. With automotive Ethernet enabled vehicles growing in numbers each day, security solutions such as IDS's need to evolve to work with these vehicles. Preliminary effort to port to automotive Ethernet have shown promise, and more work is required to further increase the effectiveness of the ported IDS. Automotive Ethernet enabled ground vehicles *can* be secure in the future, but defense solutions must keep advancing with technology. The automotive Ethernet capabilities that were integrated into the IDS is a big step in that direction.

7 REFERENCES

[1] Jukka Manner Timo Kiravuo, Mikko Särelä. A Survey of Ethernet LAN Security.
IEEE Communications Surveys & Tutorials, 15(3):1477–1491, 2013.

[2] J. Wolford, C. Westrick, P. Moldenhauer, "Cyberattack Defense Through Digital Fingerprinting, Detection Algorithms, and Bus Segmentation in Ground Vehicles", In Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium (GVSETS), NDIA, Novi, MI, Aug. 10-12, 2021.

[3] R. Elder, C. Westrick, P. Moldenhauer, "Cyberattack Detection and Bus Segmentation in Ground Vehicles",
In Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium (GVSETS), NDIA, Novi,
MI, Aug. 11-13, 2020.

[4] Christopher Corbett, Elmar Schoch, Frank Kargl, Preussner Felix. Automotive Ethernet: Security opportunity or challenge? Lecture Notes in Informatics (LNI). 2016

[5] M. Ziehensack, R. Pallierer, "Secure Automotive Ethernet for automated driving – Multi-level security architecture", [Online]. Available:
https://www.elektrobit.com/trends/automotive-ethernet-automated-driving-multi-level-security/
[Accessed: 8-March-2022]