2020 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
CYBER TECHNICAL SESSION
AUGUST 11-13, 2020 - NOVI, MICHIGAN

Cyberattack Detection and Bus Segmentation in Ground Vehicles

Ryan Elder[1], Courtney Westrick[2], Peter Moldenhauer[1]

[1]Southwest Research Institute, San Antonio, TX
[2]Ground Vehicle Systems Center, Warren, MI

ABSTRACT

This paper describes an approach to secure previously deployed vehicles by using bus monitoring and segmentation to remove malicious messages from the CAN bus. Modern automotive buses were designed for reliability rather than security. This lack of security means that any node on the bus can transmit a message to any other node and the receiver cannot verify the sender or that the message is unaltered. The intrusion detection and prevention system seeks to solve that issue by actively monitoring traffic on all connected busses, alerting an operator when an error is detected and removing flagged messages from the bus. The system will eventually be installed on an Interim Armored Vehicle (IAV) Stryker.

Citation: 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.

1. INTRODUCTION

Modern automotive buses such as SAE J1939 and Controller Area Network (CAN) were designed for reliability rather than security. As a result, the protocols allow any element on a bus to transmit to all of the other elements without verification of sender identity or message integrity. With the increase in connectivity seen in the ground vehicle domain, this lack of security means that military vehicle bus networks may be vulnerable to malicious or enemy cyber threats.

This security weakness has been addressed by some commercial automotive manufacturers through a combination of bus segmentation and a central gateway. This method typically moves critical safety nodes (e.g., steering, brakes, acceleration) away from easily accessible ports and only allows whitelisted or pre-approved messages from the open port to the safety critical nodes. This approach is excellent for securing newly manufactured vehicle buses but would be difficult to implement on vehicles already on the road as their networks are not designed for permanent segmentation.

In order to secure previously deployed vehicles, GVSC and SwRI have developed an approach that uses bus monitoring and segmentation to identify and then remove malicious messages. This pairs bus segmentation with an intrusion detection system (IDS). An IDS is a security tool deployed on a network or system to monitor network traffic and flag anomalous behavior, and within this system, it is used to actively monitors traffic on all

DISTRIBUTION A. Approved for public release; distribution unlimited. OPSEC #4184 connected busses, alert an operator when an error is detected and remove flagged messages from the bus.

2. BACKGROUND

Automotive technology has historically focused on reliability rather than security, as demonstrated in the development of the CAN bus. CAN was developed in 1986 with the purpose of allowing the different electronic control units (ECU) on a vehicle to communicate with each other [1]. One of the key elements of the protocol is the arbitration ID which not only serves to identify the message but also indicates the priority of the message, allowing for arbitration between messages simultaneously coming from the different nodes on a vehicle.

CAN is implemented using a variety of protocols depending on the manufacturer and vehicle. SAE J1939 is one such protocol that defines implementation in vehicles with diesel engines, such as many military ground vehicles, and is used in this project.

In order to ensure the real-time nature of the bus, messages are broadcast from one node to every other node on the bus. This action is done without any verification that the message is actually originating from the specified node, meaning that a single compromised node on a vehicle can then freely impersonate any other node.

These vulnerabilities were exploited in 2014 when researchers were able to compromise a Jeep Cherokee remotely through its internet connected telematics system [2]. Because of the lack of security measures present on the bus, it was possible to pivot from a low-impact target such as an entertainment system to higher value targets such as the transmission.

Attacks such as these had been possible before this research; however, the scope had been limited as it required physical access to the vehicle. Remote access greatly increased the risk to automotive manufactures as attacks could theoretically now be carried out against multiple vehicles at once, leading Fiat Chrysler to issue a massive recall of millions of affected vehicles [3].

The publication of this attack has brought to attention the need for increased security in automobiles. Many different technologies are being developed as well as adopted from traditional networks.

2.1. IDS Basics

There are two primary methods of detection that can be implemented by intrusion defense systems: signature-based detection and anomaly-based detection. Signature-based detection uses the characteristics of previously identified malicious packets 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. Anomaly-based detection is trained using normal traffic to describe what typical operation of the system looks like. Packets are then flagged if they vary from that training data behavior.

Flagging that data can result in several different outcomes. Many systems limit their response to the detection and logging of malicious events, typically used in situations where false positives would have significant impact on the functionality of the host system. However, other IDSs are linked to prevention systems. Known as intrusion detection and prevention systems (IDPSs), these prevention methods depend on where the IDS is placed on the network. In traditional networks, it is possible to drop packets before they are sent to other nodes. The broadcast nature of CAN means that mitigation methods must be deployed to each node on the vehicle. Typically, these measures would need to be installed in the initial design of the vehicle or a full gateway would need to be installed at each ECU.

2.2. Bus Segmentation Basics

One common method for addressing security issues on the CAN bus is through the segmentation of the bus. By segmenting the network, it is possible to ensure that even if one node on a vehicle is compromised, the entire vehicle will not be affected. Typically, segmentation is used to separate ECUs with remote connectivity, such as the telematics and entertainment units, from critical systems such as the drive train.

The segmentation in the CAN bus is accomplished through the use of a gateway that routes the traffic on the bus. Those gateways can further examine the packets that need to be passed among the different segments of the network and block unexpected messages.

3. SYSTEM ARCHITECTURE

In order to solve the security issues present on the CAN bus, SwRI and GVSC developed an IDPS that combined a custom IDS with a bus segmentation solution that allows for the removal of malicious packets from the CAN bus.

The design is flexible and can be quickly adapted to other CAN-based protocols. During normal operation, all bus nodes are connected and communicating directly with each other while the IDPS only monitors traffic. When monitoring, all messages are scanned by the detection algorithm, which uses anomaly-based and signature-based techniques to detect anomalies. Once an anomaly is detected, the IDPS identifies the responsible node and then segments the bus to move the malicious node to an isolated, secondary bus. A central gateway is inserted between the primary and secondary buses, and the malicious messages are removed from the primary bus. Any single node can be switched from the primary onto the secondary bus.

The result is an IDPS that can quickly remove malicious messages from the bus and avoid introducing a single point of failure onto the vehicle (e.g. run all nodes through the central gateway). The IDPS is transparent to the vehicle during normal operation and is designed to allow all traffic through in the event of an IDPS issue. This system improves integrity and availability of the automobile by flagging malicious messages and removing them from the bus.

3.1. Intrusion Detection

The anomaly detection module has been tested against many different types of attacks, including random injection, denial of service, and varied timing. In addition, the system employs a signature-based whitelisting system that checks the arbitration ID of each packet against the Database Container (DBC) file for that vehicle. If the ID is not present in the DBC file, the packet is flagged as malicious. Figure 1 shows a screenshot of the IDPS in action.

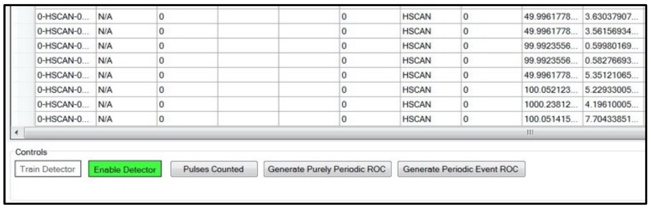

0-HSCAN-0...	N/A	0			0	HSCAN	0	49.9961778...	3.63037907...
0-HSCAN-0...	N/A	0			0	HSCAN	0	49.9961778...	3.56156934...
0-HSCAN-0...	N/A	0			0	HSCAN	0	99.9923556...	0.59980169...
0-HSCAN-0...	N/A	0			0	HSCAN	0	99.9923556...	0.58276693...
0-HSCAN-0...	N/A	0			0	HSCAN	0	49.9961778...	5.35121065...
0-HSCAN-0...	N/A	0			0	HSCAN	0	100.052123...	5.22933005...
0-HSCAN-0...	N/A	0			0	HSCAN	0	1000.23812...	4.19610005...
0-HSCAN-0...	N/A	0			0	HSCAN	0	100.051415...	7.70433851...

Figure 1. Anomaly Detection Software

In addition to whitelisting (allowing only specified messages) of CAN messages, the system is trained to understand what is typical of the environment. Malicious activity is detected by identifying messages that are statistical outliers when compared to the mean inter-message interval timing. Variations of the inter-message interval spacing are accounted for by applying the algorithm using different multipliers of the standard deviation.

Several statistical outlier cases are considered. One case entails a higher priority CAN message that causes a purely periodic (PP) message to miss its scheduled timeslot, and thus its inter-message interval is larger or smaller than the mean. Equation (1) summarizes one of the outlier detection algorithms used, where x is the current inter-message interval, N is a multiplier used to tune the detection algorithm, α is the standard deviation and *mean* is the mean inter-message interval.

$$(x + N * \alpha) < mean \qquad (1)$$

The training determines the timing of each PP arbitration ID that is then used to calculate the *mean* time. The quality and size of the data set directly impacts detection accuracy. Vehicle specific DBC files are also necessary to determine valid arbitration IDs for the whitelist filtering. Traffic on the busses can vary greatly from one vehicle to the next, but the system can be easily adapted to new configurations.

3.2. Bus Segmentation

Under normal operation, the IDPS's primary function is to monitor traffic. If a malicious message is detected, the IDPS segments the bus by moving the agitating node to the secondary bus. Once the bus is segmented, the IDPS acts as a gateway between the primary and secondary bus and filters messages from the agitator, removing them from the primary bus.

A typical bus is configured with several nodes (e.g. Engine, Transmission) connected, as shown in Figure 2.

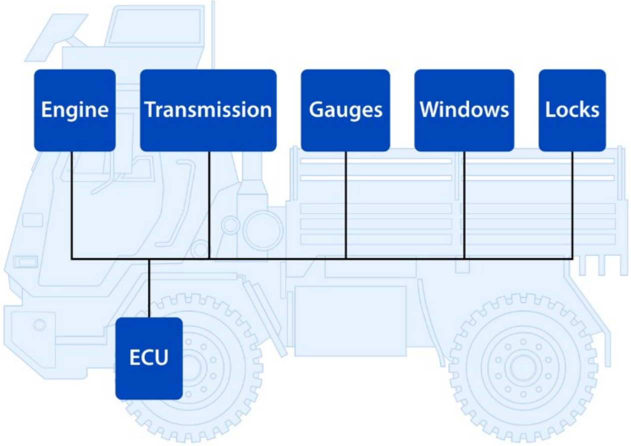

Figure 2: Typical Bus Configuration

Commercial automotive manufacturers have secured their buses through a combination of bus segmentation and a central gateway, an approach which protects critical nodes and minimizes impact on bus reliability. The method used in this paper employs a variation of bus segmentation that passively monitors bus traffic until malicious messages are detected. The bus outline with the IDPS installed in a passive monitoring state is shown in Figure 3.

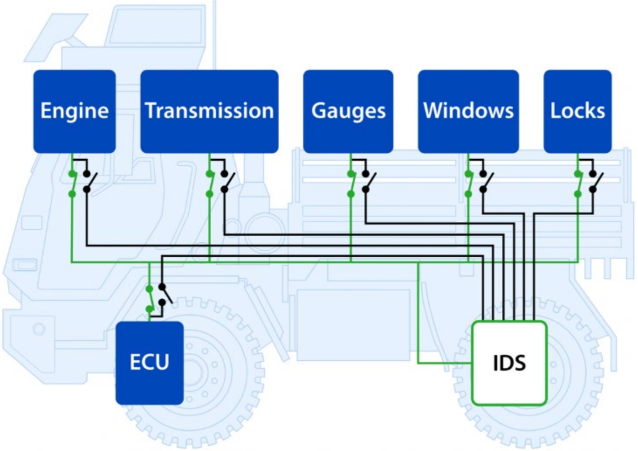

Figure 3: IDPS Passive State

When a malicious message is detected, the suspect node is segmented to the secondary bus (red lines) as shown in Figure 4. When segmented, the IDPS is a gateway between the two busses, filtering out malicious messages.

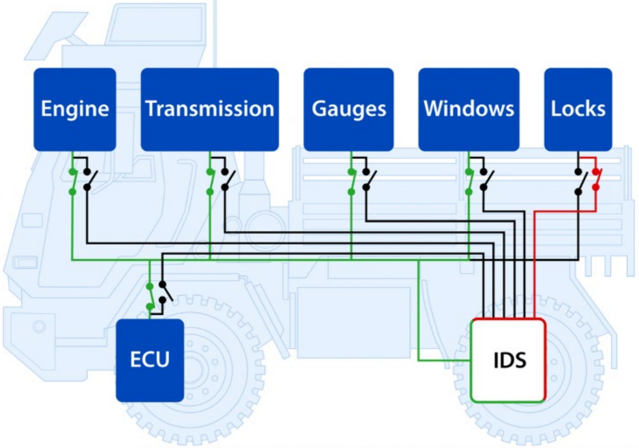

Figure 4: IDPS Active Removal of Malicious Messages

Because of the open-broadcast nature of J1939, the IDPS will not always know which node sent the malicious message. In order to determine the node that the malicious messages originated from, it may be necessary to iterate through the nodes and isolate them until the malicious performance only occurs on the isolation bus. Training data will be used to prioritize nodes for isolation (e.g. diagnostic ports may have a high priority).

4. PERFORMANCE

A series of tests were performed on the IDPS in order to determine its effectiveness. The tests were run using a simulated vehicle environment to emulate the different ECUs on the vehicle. They were connected to the IDPS, which operates on an embedded computer running an SELinux kernel.

Traffic was then simulated based on previously recorded traffic from vehicles using the J1939 protocol. Attacks were inserted into the bus traffic based on the attack vectors described below. Those generated attacks were then compared with the log files from the IDPS in order to determine the accuracy of the system.

Bus switching was tested using a relay switch connected to the IDPS through a serial connection. Upon detection of an attack, the attacking node is switched off of the main bus to an isolated bus using relays. Success of the bus switching was verified by reviewing log files of the main bus and the isolation bus.

4.1. Attack Vectors

To test the IDPS, determining what attack vectors it should prevent is necessary. The following test cases were created in order to ensure that the system accomplished its stated goal of increasing security on the vehicle.

1. Normal – While not an attack, this test case is necessary to ensure that the system does not detect attacks when normal traffic is present on the bus. This test case uses real vehicle data.
2. Arbitrary Injection – Fuzzing the system is a common way to find vulnerabilities. This test case looks for that behavior by adding arbitrarily created packets to prerecorded traffic at arbitrary timing.
3. Bus Flood – Injecting packets on the CAN bus at the maximum speed of the bus can lead to packets being overwritten. This test injects high speed packets, both replayed and fuzzed, to overwrite legitimate traffic.
4. Throttling – A compromised ECU can manipulate packets on the bus by increasing the frequency with which packets are sent. This test case ensures that the IDPS can detect the increase of legitimate packets.
5. Whitelist – Each vehicle has a DBC file that determines what each message on the vehicle means, and the system should screen packets that do not follow those formats. This test case adds packets that are generated randomly to the prerecorded file.
6. Diagnostic – It is important to determine that legitimate diagnostic messages that would not normally be present in training data are not flagged. Diagnostic messages are inserted into the prerecorded data in this test case.

4.2. Results

In order to determine the success of the system development, target thresholds for the detection rates of the system were set. The true positive rate is to exceed 95%, while the false positive and false negative rates are to remain below 5%.

Overall, the detection rates on the system have been excellent. A packet is designated a true positive if it was injected as an attack and the IDPS successfully flags it as so. A false positive occurs when a packet from the original recorded data is flagged as an attack. A true negative is a packet from the original recorded data that is not flagged by the IDPS. A false negative occurs when a packet that was injected as an attack is not flagged as malicious. The overall statistics have been created by averaging the results of each of the packets. In order to ensure that vehicle operators do not need to continually handle alerts that are not relevant, minimizing the false positive rate was a significant goal for the project. This goal has been completed so far while maintaining good true positive detection.

5. FUTURE IDPS DEVELOPMENT

While the results of the project have been largely positive, the system can still be improved through a variety of steps. The first method for improving IDPS performance would be performing deeper packet inspection. This step could involve a more in-depth approach to anomaly-based detection. Currently, the algorithm only examines the arbitration ID rather than the contents of the packets. Deeper inspections of the data could lead to analyzing the timing of recurring data payloads and detecting if they are disrupted.

Another element that could improve accuracy is detection of missing period packets on the bus. Currently, the IDPS does well in detecting packets that have been inserted onto the bus; however, message dropping is also an attack vector able to compromise an ECU. Dropping detection would help mitigate that issue.

Additional work should also be done to characterize the diagnostic messages that can be sent with J1939 protocol and determine if they are malicious in nature. The IDPS currently assumes that diagnostic messages sent at normal rates are benign. Determining what diagnostic messages can be implemented and when they should be sent would limit an attacker's ability to investigate and alter ECU software. For example, specific diagnostic messages could be flagged when they occur while the vehicle is in use, as they should only be used while the vehicle is out of use and under repair.

Further work can also be done to refine the physical mitigations triggered by the IDPS. The largest opportunity for improvement is increasing the number of nodes that can be segmented at once. This step would involve the creation of a quarantined zone rather than the segmentation of one bus at a time. The process of filtering out malicious packets before forwarding them to the clean zone would still be implemented, making it easier to contain a large-scale compromise of the vehicle.

There are also discussions in place about possibly adapting this project for other busses. The modular nature of the system should allow for adoption to other protocols and would allow other vehicle systems to benefit from the security increase that an IDPS can provide.

6. CONCLUSION

As reliance on automotive busses increases to support a wide variety of remote access technologies, the lack of security measures built into those busses increasingly exposes them to the external world. An intrusion detection system is an excellent way to introduce security measures into a bus without requiring redesign of the vehicle. Bus segmentation further enhances those benefits by adding prevention to the IDPS. Taking such measures will ensure that ground vehicles can continue to take advantage of the wide variety of measures available in the automotive space.

7. REFERENCES

[1] W. Lawrenz, *CAN System Engineering*, 2nd ed.London: Springer, 2013, p. 3.

[2] C. Miller and C. Valasek, “Remote Exploitation of an Unaltered Passenger Vehicle”, BlackHat USA, 2015.

[3] A. Greenberg, ‘Hackers Remotely Kill a Jeep on the Highway-With Me in It’, 2015. [Online]. Available: https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/ [Accessed: 7- May-2020].