National Marine Electronics Association

NMEA 2000® Vulnerability to Cyberattacks and Mitigation

Chris Quigley Warwick Control Technologies Ltd
Paul Sumpner Digital Yacht Ltd
Mark Oslund NMEA 2024 Director of Standards and Technical Activities

Abstract

NMEA 2000 is a plug-and-play communications CAN-based standard used for connecting marine sensors and display units within ships and boats. It sits amongst other NMEA marine communications protocols from NMEA 0183 at the lower-end through to the Ethernet-based NMEA ONENET standard. NMEA 2000 itself uses many of the features that are in common with SAEJ1939 and ISO11783. The standard has enabled the easy integration of electronic devices into a vessel. However, as with all CAN-based protocols, several vulnerabilities to cyberattacks have been identified. Many are at the CAN level, whilst others are in common with those protocols from the SAEJ1939 family of protocols.

Purpose

This paper will discuss the known vulnerabilities that have been identified with the NMEA 2000 protocol. These include weaknesses with the address claim and transport protocols, and covert communication channels using methods based on steganography. Techniques that can improve the robustness of NMEA 2000 to cyberattacks are described.

1 Introduction

NMEA 2000 is a CAN-based higher layer protocol used for the integration of marine electronics. It is now the de facto technology for integration of marine devices. The growth of NMEA 2000 and its Parameter Group Numbers (PGNs) has gone from navigation and sensors, now through to applications such as electric propulsion and entertainment. It sits amongst other protocols that can be used in marine applications such as CANopen, SAE J1939 and the two other National Marine Electronic Association (NMEA) specified protocols (0183 and OneNet). An example of a yacht using a variety of CAN high layer protocols is shown in [1], the vessel in this case using NMEA 2000, SAE J1939, CANopen and proprietary CAN.

NMEA 0183 provides one-way communications and as an older technology typically runs at 4.8Kbit/s. Devices are either “Talkers” or ”Listeners”. NMEA 0183 allows a single talker and several listeners on one circuit. All data is transmitted in the form of sentences that can contain ASCII characters. NMEA 0183 does not use any authentication or encryption.

NMEA OneNet is an emerging standard for marine electronic devices based on Internet Protocol, Version 6 (IPv6) and the IEEE 802.3 Ethernet Local Area Network. It provides a common network infrastructure for marine devices and/or services on IPv6. All OneNet application protocols, such as PGN Messages, are designed to use a standard IPv6 network protocol stack. OneNet can coexist with other protocols and services that operate parallel on the same network. The standard also specifies mechanisms for connecting OneNet networks, NMEA 2000 networks, and other networks via gateway devices.

2 NMEA 2000 Key Features

NMEA 2000 is now the main backbone for most marine vessels (recreational, workboat, small car ferries, coastguard vessels). Most installations have in the region of 25 to 50 devices on a network. Some larger installations have more than 50 devices spread across several NMEA 2000 networks. Typically, devices are connected via off-the-shelf connectors, cables, T-pieces and the network terminated at either end by off-the-shelf 120 Ohm terminators as shown in Figure 1.

A common misconception about NMEA 2000 is that it is simply SAE J1939 for marine applications. However, NMEA 2000's compared to SAE J1939 can be summarized as follows:

- NMEA 2000 is always 250 Kbit/s with a maximum of 50 physical devices on one network.
- Specifies a set of standardized messages called Parameter Group Number (PGN), each one has a unique number.
- Fast Packet Protocol is an additional transport protocol for rapid transmission of up to 223 bytes (31 CAN frames).
- A device must pass a certification test before it can be marketed as a NMEA 2000 device.
- Mandatory PGNs to be supported:
    - Product Info – includes part numbers and current drawn by the device.
    - Configuration Info – an ASCII description on how the device has been installed.
    - Tx and Rx list – provides a list of PGNs that the device sends and receives.
- Source Address claiming is dynamic addressing only, no fixed addresses.
- Commanded Address – is a mandatory service that can be used to address a specific device and change its Source Address.
- NAME Instance – allows System & Device Instance to be changed via a service over CAN.

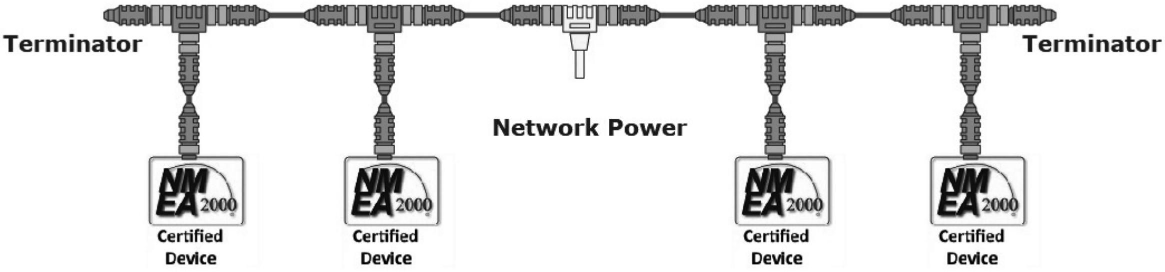

Courtesy of Maretron N2KBuilder®

Figure 1 Typical installation for NMEA 2000

3 NMEA 2000 Protocol Vulnerabilities and Mitigations

Recent papers have highlighted some of the issues of CAN bus and NMEA 2000 in terms of cybersecurity [2, 3]. These include issues such as spoofing, Denial of Service (DoS) and sniffing or eavesdropping the CAN bus information. The sniffing or eavesdropping of the NMEA 2000 network is in part mitigated by the fact that the NMEA 2000 specifications must be purchased from the NMEA. However, this does not make the information 100% confidential since the information on PGN encoding is often leaked or reverse engineered. These top-level issues have led to research into the vulnerability of NMEA 2000 in more detail.

The vulnerabilities of NMEA 2000 to cyberattacks has been broken down into the following three Protocol Groups as shown in Table 1. The table also shows example vulnerabilities that had been identified at the time of writing. This is by no means an exhaustive list and is an ongoing area of research.

Protocol Group	**Cyberattack / Vulnerability**	**Impact**
CAN	Janus	Message
	High Priority CAN ID DoS	Network
	Frame Spoofing	Message
	Relay / Man-In-The-Middle	Message
	Double Receive	Message
	Bus Off	Device
	Freeze Doom Loop	Network
SAE J1939/ISO11783	Address Claim Hunter	Device
	Transport Protocol	Message
	Commanded Address	Device
NMEA 2000	NAME Instance	Device
	Fast Packet Sequence	Message
	Data Instance Hopping	Message
	PGN Timing Attack	Network
	Steganography in Fields	Message
	Packet Sniffing	Message

Table 1 : Categories of cyberattacks for NMEA 2000

CAN Level Vulnerabilities are the same vulnerabilities that all CAN-based protocols are susceptible to. *Vulnerabilities of the SAE J1939 Family of Protocols* are those that are common to all protocols that are derived from SAE J1939. These include ISO11783 (ISO Bus), NMEA 2000 and Recreational Vehicle Communications (RV-C). *NMEA 2000 Specific Vulnerabilities* are those affecting features that have been added to create NMEA 2000 such as the Fast Packet Protocol. Each vulnerability is assessed in terms of its impact and whether it corrupts/destroys a message, device or entire network, e.g.:

- *Message* Level – vulnerability results in the corruption or destruction of a message (e.g. Single Frame, Fast Packet or BAM/CMDT).
- *Device* Level – vulnerability results in the shutting down or destruction of an entire device (e.g. sensor, actuator, MFD etc.).
- *Network* Level – vulnerability results in the denial of service or shutting down of a single network.

3.1 CAN-Level Vulnerabilities

3.2 Janus Attack

The Janus attack [4] is a low-level CAN protocol attack where a single CAN frame contains two different payload contents. This is named after the Roman God Janus, the god of two faces or of transitions. With the Janus Attack, a targeted device sees a different payload than other devices. This attack could be used to transmit a frame to evade an Intrusion Detection System

(IDS), or it could put two different actuators into inconsistent states (e.g. moving a pair of motors in different directions). It breaks the atomic multicast feature of CAN (where every device sees the same frame). The attack works by exploiting the CAN protocol synchronization rules and targets devices that have different sample points. One of the main and easily implemented detections against this attack, is devices should have sample points set as close to each other as possible. NMEA 2000 device certification provides mitigation by checking that a device's sample point is to the NMEA 2000 requirement. This does not eliminate the risk but it does significantly reduce the risk.

3.3 High-Priority CAN ID - Denial of Service

The process by which the CAN protocol ensures that the one CAN message will always win access to the network in the case when two devices try to transmit at the same time, results in the feature that the lowest value CAN ID always wins arbitration for network access. This can be misused if a malicious device transmits CAN identifier 0x00000000 as often as possible (which is the highest priority ID for a 29-bit CAN bus). This results in a Denial of Service (DoS) for other devices wanting to access the network. This is also referred to as the Bus Flood Attack in another publication [5]. A high priority CAN ID can cause this effect. However, any CAN message if sent at a fast enough rate can use too much CAN bus bandwidth, resulting in a DoS [2]. Mitigations include monitoring of bus load, allow/deny lists, monitoring of CAN message update rates and then raising an alert by some means (probably not over the CAN bus due to the DoS state). This could be carried out in software, by a 3rd party device or using a secured CAN transceiver (such as NXP TJA115x family).

3.4 Frame Spoofing (Simple/Adaptive/Error Passive)

This is a type of attack in which a receiver accepts a fake frame as if it came from a legitimate sender. There are numerous ways in which it can be achieved at the CAN level [5]. An example for NMEA 2000, could be vessel speed sent by a malicious device on the network whilst the actual vessel speed sensor has been disconnected from the network. From the point of view of the attacker, it is important that the original device is disconnected so that it does not send the same CAN ID as the malicious device. NMEA 2000 has been identified as being vulnerable to this type of attack in a previous study [2]. Mitigation strategies include the use of a secured CAN transceiver, authenticating/watermarking of messages or fingerprinting of the network so that message transmitters can be verified [12, 13, 14].

3.5 Relay / Man-In-The-Middle

A Relay or Man-In-The-Middle attack can be seen as a two-way spoof in which the communications between two devices in interrupted [2]. The mitigation for this is similar to that which can be used for spoofing.

3.6 Double Receive Attack

The Double Receive Attack has recently been highlighted by Tindell [5] and is an exploitation of a feature of the CAN protocol that is in ISO11898 and includes a warning for it. The protocol defines that a receiver accepts a frame as finished at the second-to-last bit of the EOF field and that the transmitter accepts it as finished at the last bit of the EOF field. There is a very small chance of a bit error in the last bit of the EOF field. This means it should be recessive, but the transmitter sees a dominant bit and then signals an error. The result of this error is that the frame is put into arbitration again. All receivers will have already accepted the frame and passed it up to the application software. However, because of this bit error, the transmitter will send the frame again and the receivers will receive the same frame again. To be able to make a Double Receive Attack, an attacker could use the general purpose I/O of a microcontroller to insert the dominant bit into the aforementioned EOF field.

Mitigation for the double frame reception can be achieved by including a sequence number or counter into the frame data field. Receiving devices then expect this to increase or decrease in each instance of the frame that is received. It should also be noted that this approach would also protect against failures of the communications between the main microcontroller and the CAN controller in which data field values are not being updated.

In NMEA 2000, a sequence counter is a part of some single frame PGNs and all transport protocols packets (Fast Packet, BAM, CMDT). Some single frame PGNs do not have a sequence counter. It is desirable that all newly specified single frame PGNs have a sequence counter to mitigate against this attack and other failures with similar symptoms. Legacy PGNs that do not have a sequence counter can be protected by a 3rd party Intrusion Detection System (IDS) that monitors the update rate for single frame PGNs.

3.7 Bus Off

The Bus-off Attack is another one highlighted by Tindell [5] and is where a specific ECU is targeted and driven offline whilst all the other ECUs continue to operate. This could be used as part of a wider attack (such as a spoofing attack or denial-of-service attack). The Bus-off Attack is a low-level protocol attack achieved by disturbing the CAN bus when the Device Under Attack is transmitting a message. Instead of targeting a specific frame, all frames from the same device are targeted. This forces the Transmit Error Counter (TEC) above 255 and the device's CAN controller automatically goes bus-off. Most devices will try to recover automatically, requiring the attack to be repeated. Mitigation strategies include automatic recovery from Bus-Off and monitoring of the network for this type of situation. NMEA 2000 requires a Heartbeat message to be sent periodically which includes a field with the CAN controller state. This could be useful in monitoring for this type of attack.

3.8 Freeze Doom Loop

The Freeze Doom Loop attack is another one highlighted by Tindell [5]. It is a low-level attack that effectively freezes bus traffic for an arbitrary time and could be used to delay a specific CAN frame or to generally reduce the bandwidth of the CAN bus. In the original study it is stated that it is difficult to detect. The symptom of this attack is a DoS of the CAN bus. Mitigation strategies include timing analysis and using a device with a CAN controller that can detect an overload condition.

3.9 SAE J1939-Level Vulnerabilities

ISO11783 – NAME Convention

Self-Configurable Address	Industry Group	Device Class Instance	Device Class	Reserved	Function	Function Instance	ECU Instance	Manufacturer Code	Unique Number
1-bit	3-bit	4-bit	7-bit	1 bit	8-bit	5-bit	3-bit	11-bit	21-bit

NMEA 2000 – NAME Convention

Reserved (set to 1)	Industry Group	System Instance	Device Class	Reserved	Device Function	Device Instance (Upper)	Device Instance (Lower)	Manufacturer Code	Unique Number
1-bit	3-bit	4-bit	7-bit	1-bit	8-bit	5-bit	3-bit	11-bit	21-bit

Figure 2 : ISO11783 and NMEA 2000 NAME Field Comparison

3.10 Address Claim Hunter

The Address Claim Hunter is an algorithm that hunts address claim messages and attempts to kill devices by forcing them into the state where they cannot claim a valid address. It does this by monitoring the bus for Address Claim messages (maybe from a particular manufacturer) and claiming any attempt by a NMEA 2000 device to claim a particular Source Address by claiming it with a higher priority NAME field. The first studies known to report a vulnerability in the SAE J1939 address claim functionality was in 2018 [6, 7]. These were particularly concerned with any protocol from the SAE J1939 "family" of protocols that uses the dynamic address claim such as NMEA 2000. This is the primary method that NMEA 2000 uses and therefore it is particularly susceptible to this. A more NMEA 2000 specific discussion of this problem is discussed in [9, 10].

As far as NMEA 2000 is concerned, the attacks split into two types:

- **Illegal NAME** – those which are illegal as per the protocol specifications. Therefore, it would not be expected to occur on a network.
- **Legal NAME** – those which are legal as per the protocol specifications. Therefore, it would be expected to occur on the network.

Illegal NAME	Legal NAME
CAN With No NAME e.g. data field all zeros. Other illegal values e.g. Industry Group not equal to 4 Manufacturer Code equal to 0	NAME plausible according to protocol. NAME plausible according to certified device list. NAME plausible according to network snapshot.

Table 2 Examples of Illegal NAME and Legal NAME – Address Claim Hunter Attacks

Examples of Illegal NAME and Legal NAME Address Claim Hunter attacks are compared in Table 2. Since there is a variety of attack approaches that are possible, it makes 100% protection from Address Claim Hunter attacks extremely difficult.

3.11 Illegal NAME Address Claim Hunter

Illegal NAME Address Claim Hunter algorithms use NAME field values that you really should never see on a NMEA 2000 network and therefore devices should be able to detect these easily and reject them. It should be noted that tests carried out by the authors of this paper on a random selection of NMEA 2000 devices suggest that most devices are susceptible to these types of attack. Types of Illegal NAMEs include the CAN with No NAME, so called since it involves a device NAME which is all zeroes (e.g. 00 00 00 00 00 00 00 00).
Other illegal NAME settings such as:

- Self-Configurable Address bit or first Reserved set to 1.
- Reserved should always be 0.
- Industry Group should always be 4 = Marine.

The first and easiest approach to mitigate against the occurrence of an Illegal NAME or an Illegal NAME Address Claim Hunter attack is to check the above fields for plausibility. If an Address Claim message is found to have an Illegal NAME, then it can be rejected. This approach is compliant to the NMEA 2000 certification tool tests. An attractive additional action may also to be to send an NMEA 2000 Alert.

3.12 Legal NAME Address Claim Hunter

The next step in checking the plausibility of a NAME field is to check whether it contains an implausible Class and Function combination. However, a more sophisticated device such an IDS could check whether a device is an actual NMEA 2000 certified device by cross-checking the NAME with some other information that should be available from the device.
NMEA 2000 devices could easily implement a device NAME plausibility check which will make the system more robust. There is however still the possibility that a malicious device could mimic a Certified Device to shut down the network and therefore other mitigations could include:

- Fixed Addressing – solves the problem but is against the plug and play nature of NMEA 2000.
- Snapshot of network during installation – e.g. by some kind of IDS.
- Fingerprinting of the network using its physical properties was a way to ensure that an Address Claim message is transmitted by the expected device. There are many examples of these. A study by Cho and Shin [12] used the tiny variations in bit timing characteristics (clock skew) between CAN devices to identify the correct sender of a message. Another study by Shin and Cho resulted in the filing of a patent using a fingerprint of the analogue levels of the CAN signals [14]. Another method by Avatefipour et al [13] used a time and frequency domain analysis of the physical signal as a way of fingerprinting messages from different CAN devices. It has been pointed out that these methods may be prone to false

positives [5]. This means that they are unlikely to be useful for identifying a single instance of a rogue message but would be useful for providing information on longer term trends.

- A technique proposed in [16] uses a bit-banged implementation of a CAN controller (rather than using an actual CAN controller) to detect the attack. Once detected, it uses the CAN error protocol to destroy the offending CAN frame. In practice, this is a complicated method to implement. It could be a useful technique to help protect legacy devices. Plausibility checks as a way to decide whether a device should accept an Address Claim message are much easier to implement for future devices.
- Design the vessel network architecture using multiple networks to ensure a high level of decoupling of functional systems.

3.13 Transport Protocol Attack

Broadcast Announcement Messages (BAM) and Connection Management Data Transfer (CMDT) are transport protocols used within the SAE J1939 family of protocols for messages greater than 8-bytes of data. Since these are formed from multiple CAN frames, the opportunity to disrupt the flow of frames is a possible attack and has been highlighted in previous studies [6, 8]. Mitigation strategies could include detection and alerting to the corruption of a Transport Protocol message.

3.14 Commanded Address Attack

Commanded Address is a standard feature of the SAE J1939 protocol (in part 81) to allow another device or diagnostic tool to change the Source Address of another device by sending it a message. The device is addressed directly by its NAME field and results in it claiming the new Source Address in the Commanded Address message. This can be used by a malicious device to constantly change the Source Address of a device under attack resulting in it at least being partially withdrawn from network activity and creating confusion for other devices on the network. Mitigation strategies include detection of when the Commanded Address is happening a lot to a particular device or even limiting access to this service by requiring a certain higher security/login level be reached before using the service.

3.15 NMEA 2000 Vulnerabilities

3.16 NAME Instance Attack

Sections 8.4.1 and 8.4.2 of the NMEA 2000 specification [15] describe two fields in the Address Claim message that are field programmable. There are two Instance fields in the NAME:

- System Instance
- Device Instance

There is a provision in the NMEA 2000 protocol to change these values using a Complex Command. A device can respond with a NACK if it does not allow the changing of these fields. If it does support the changing of these fields, it changes the values and then acknowledges by sending an Address Claim with the new Instance value. The problem is if your NMEA 2000 devices does support the changing of the Instance fields in the NAME, there is no limit to how often this can be done. Therefore, it could be changed continuously and cause a lot of disruption on the network. Mitigation strategies include detection of when the Complex Command service is happening a lot to a particular device or even limiting access to this service by requiring a certain higher security/login level be reached before using the service.

3.17 Fast Packet Sequence Attack

The fast packet protocol is unique to NMEA 2000 and allows a burst of data transfer up to 223 bytes over 31 CAN frames. Similar to other transport protocol attacks, it manifests itself as an interruption of the flow of packets. Mitigation strategies could include detection and alerting to the corruption of a Fast Packet message.

3.18 Data Instance Hopping

Many PGNs within the NMEA 2000 specification have an Instance field so that the protocol can support several different instances of the same data as described in section 8.4.3 of the NMEA 2000 specification [15]. An example of this includes fluid level which may have values to represent the level from various tanks around the vessel. Another example is battery module voltage, state of charge etc. The instance value can be used to represent the values from a number of different battery modules.

The data instance for these PGNs can be changed by the Complex Command service. However, the ability to change leaves devices open to an attack in which a malicious operator or device can continually address specific devices and change the data instance. The result of this is confusion of the control system and other devices will not know what the data actually represents. Mitigation strategies include detection of when the Complex Command service is happening a lot to a particular device or even limiting access to this service by requiring a certain higher security/login level be reached before using the service.

3.19 PGN Timing Attack

Many PGNs within the NMEA 2000 specification have the ability to have the Priority in the CAN identifier and/or the PGN Update Rate changed using a Complex Command as described in section D.4.10 of the NMEA 2000 specification [15]. If a PGN's Priority value is reduced, it in fact raises the priority of access to the network of the CAN message. If a PGN's Update Rate value is reduced, the PGN is transmitted more often onto the network. The impact from abuse of these features could be a serious corruption of network timing such that the network becomes overloaded and message delivery is delayed.

Mitigation strategies include detection of when the Complex Command service is happening a lot to a particular device or even limiting access to this service by requiring a certain higher security/login level be reached before using the service. Other strategies worth using include Message Update Rate Analysis and Bus Load Monitoring.

3.20 Steganography in NMEA 2000 Fields (Covert Communication Channels)

Encrypted communications can look immediately suspicious to defenders and detection tools. Conversely Steganography allows hackers to hide data in a way that would be difficult to easily catch. To even be able to catch steganography, you first have to know the technique, and then you have to know which file(s) to analyse. Steganography is different to Encryption. The key difference between encryption and steganography is that for the former, the message can be seen but no one can work out its meaning unless they can successfully decrypt it. With steganography, the fact that a message has been sent is a secret and therefore unknown.

One way that steganography can manifest itself in NMEA 2000 is by hiding information in the least-significant bits of the signals sent within a CAN message. The data field of a CAN message carries the signals that are used in the control system. If you examine the length of typical signals that are specified within various CAN standards, it is found that they usually have more than enough resolution for the task. It could be said that the signals are over-specified in that the resolution provided in greater than needed. This over-specification can lead to a reduction in the space available in a CAN frame which could have been used for other signals. The over-specification also leaves the signal vulnerable to abuse from steganography techniques using the LSBs of the signal. Consider PGN Engine Parameters, Rapid Update (1F200), field 2 is Engine Speed which is 16 bit and scaled at 0.25 RPM per bit. Therefore, the question to ask, would you notice if the least significant two bits were used for hidden data?

Steganography in NMEA 2000 creates weaknesses and opportunities.

- Use to initiate an attack upon certain conditions being met. E.g. via a gateway from OneNet, IoT, J1939 in a field smuggled in a PGN which reaches a malicious NMEA 2000 device installed a long time ago in the vessel. Upon reaching a certain set of circumstances, the trigger for the attack can be smuggled in
- Communicate information on which device or manufacturer to attack.
- Use to pass configuration data, even changes to a device's flash across the NMEA 2000 network.
- Watermarking, in which the hidden data is used to authenticate and protect communications between two devices. An example would be a transmitting device could use steganography to hide the authentication code. A receiving device which knows the steganography algorithm is able to receive and retrieve the hidden authentication code.
- Hiding images or stolen data.
- Download of malware
- Espionage

3.21 NMEA 2000 and Packet Sniffing

One of the strengths of NMEA 2000 is that it is easy to access the network and read the data. Easy access to the PGNs and their associated fields makes the diagnosis of issues with the appropriate diagnostic tools relatively easy. However, the ease of reading of NMEA 2000 PGN fields could be seen as a security risk [1]. Some CAN-based applications are considering various encryption methods for signals carried in the CAN data field. This has the disadvantages of increased processing for the encryption/decryption algorithms, and diagnostic tools would need to be privy to the encryption/decryption methods to be able to view the PGN fields in any meaningful way.

The NMEA keeps the NMEA 2000 PGNs secret to some degree since the specifications must be purchased. This only discourages access rather than preventing access to the network. Various diagnostic tools are available to view NMEA 2000 data and these could also be used to reverse engineer the values held within NMEA 2000 PGNs. There may be the need for certain new PGNs to be encrypted in the future.

4 Gaining Access to a NMEA 2000 Network

To be able to attack one of these CAN-based networks, the attacker just needs to be able to access the network. Examples of these include:

- physically add small device whose aim is to disrupt network.
- Putting a USB key into a PC on the vessel. If the PC itself, is connected to the NMEA 2000 network, then this is a way in.
- reflash or reconfigure an ECU.
- via a gateway, e.g. CANopen, IoT etc.
- Software backdoor added by a disgruntled employee of a device manufacturer.
- Software backdoor in open-source software and third party components that is then used in a NMEA 2000 device.

A study [11] discusses some issues with gateways from unsecured to secured domains based on CAN. Figure 3 shows some possible gateway applications that are common for NMEA 2000 and marine applications.

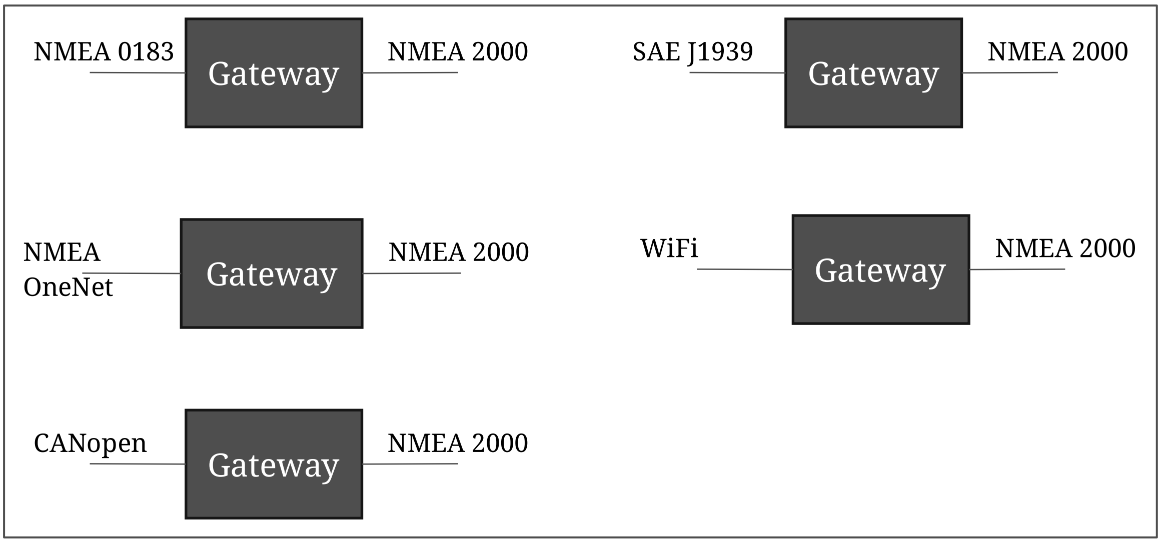

Figure 3 – Possible NMEA 2000 Gateways

5 New NMEA 2000 Device Types for Increased Cybersecurity

The NMEA has set up a Cyber Protocols Technical Committee to look at the susceptibility of the NMEA protocols. NMEA 2000 is the de facto marine protocol and therefore is the initial focus. As well as looking into mechanisms for protecting against NMEA 2000 vulnerabilities such as those mentioned in this paper, there are also two new types of devices that are being discussed:

- Intrusion Detection System (IDS) – aim to be able to protect a legacy network e.g. containing devices that do not have any protections against cyberattacks.
- Secured Gateway – since a gateway is one of the main ways for a malicious actor to gain access to a NMEA 2000 network, a secured gateway is needed.

5.1 Intrusion Detection System (IDS)

Efforts will be put into how to make normal NMEA 2000 devices cybersecure as well as defining the functionality of a new type of device; namely the Intrusion Detection System (IDS). An IDS is a device that can be installed in existing NMEA 2000 networks to monitor the network for cyberattacks. This has the benefit of being able to protect legacy devices released to market before cybersecurity issues were a consideration. The IDS can maintain a list of devices (even a certified device list) on the network and look out for unexpected activity. If an issue is found, the IDS needs to be able to alert to the existence of the suspected cyberattack. Alerts could be made across the CAN bus, by SMS or via an audible alarm. The NMEA 2000 specification provides the ability to send alerts to a NMEA 2000 Multi-Function Display (MFD). An IDS has been highlighted as a reasonably easy solution, which can be added to existing networks with no impact [11].

An IDS would have to be easy to operate by marine technicians and therefore an installation procedure would be necessary that can be run after the network has been successfully installed. This would also have to be protected by sufficient security so that only authorized personnel can run the installation procedure.

5.2 Secured Gateway

Mitigations in the installation of NMEA 2000 devices is seen as important so as to remove the possibility of a physical connection to be able to make an attack. Therefore, hiding away NMEA 2000 cable access points is a deterrent. As well as the protocol vulnerabilities highlighted in this paper, gateway devices are a major weakness (e.g. to protocols 0183, OneNet, CANopen, J1939, Wi Fi etc.) as they provide a way into the NMEA 2000 network. The top-level aim will be to ensure that a NMEA 2000 network becomes a secured domain and the management of information coming from other network domains based on some of the technologies shown in Figure 3 is important. It is imperative that malicious actors cannot gain access to the NMEA 2000 network. A discussion describes requirements for secured gateways [11] that takes information from an Untrustworthy Network Domain (UND) to a Trustworthy Network Domain (TND). Therefore, making NMEA 2000 a TND is an objective and the definition of requirements for a secured NMEA 2000 gateway is also underway.

6 Mitigation Against NMEA 2000 Vulnerabilities

Table 1 summarized cybersecurity vulnerabilities at CAN-, SAE J1939- and NMEA 2000 levels. This has been extended to include advice on the kind of mitigating strategies that can reduce or eliminate the vulnerability to a cyberattack. These are shown in Table 3.

Protocol Group	Cyberattack / Vulnerability	Impact	Mitigation
CAN	Janus	Message	Certification
	High Priority CAN ID DoS	Network	MURA, BLM, SCT, ADL
	Frame Spoofing	Message	SCT, AoM, WM, FP
	Relay / Man-In-The-Middle	Message	SCT, AoM, WM, FP
	Double Receive	Message	SC
	Bus Off	Device	HM
	Freeze Doom Loop	Network	MURA, OD
SAE J1939/ ISO11783	Address Claim Hunter	Device	PC, FP, Snap
	Transport Protocol	Message	Flow
	Commanded Address	Device	ComAddMon
NMEA 2000	NAME Instance	Device	CompComMon
	Fast Packet Sequence	Message	Flow
	Data Instance Hopping	Message	CompComMon
	PGN Timing Attack	Network	CompComMon, MURA, BLM
	Steganography in Fields	Message	Statistical Analysis
	Packet Sniffing	Message	Encryption

Table 3 : Categories of cyberattacks for NMEA 2000 along with mitigation strategy.

The mitigations are described further here:

- **Certification** – The NMEA 2000 product certification process helps ensure that the device under test has the correct CAN sample point.
- **MURA** – Message Update Rate Analysis can be used to detect whether a message transmission has not met its expected transmission period or time.
- **BLM** – Bus Load Monitoring can check whether the loading on the CAN bus exceeds its expected level.

- **SCT** – Secured CAN Transceiver such as the NXP TJA115x range can provide Bus Load Monitoring and Allow/Deny Lists.
- **ADL** – Allow/Deny List can provide protection a list of CAN messages that are allowed on a network or transmitted by a device or a list of CAN messages that are denied access on a network or transmitted by a device.
- **AoM** – Authentication of a message.
- **WM** – Watermarking, e.g. by Steganography, as a way of authenticating the transmitter of the message.
- **FP** – Fingerprinting could involve learning some of the physical characteristics of devices such as those based on CAN signal voltages or bit timing and then using these to detect an anomaly on the network [12, 13, 14].
- **SC** – Sequence Counter can be used to check that the data in a TP message is changing and also is in the correct order.
- **HM** - Heartbeat Monitoring can be away to monitor if a device is still alive. There is a field available for CAN Controller State which can be used to report when a device has gone bus off.
- **OD** – Overload Detection is a particular feature that some but not all CAN controllers will have.
- **PC** – Plausibility Checks can be used to ensure that the received fields are as expected as per the specification. An example of this is for the Address Claim Hunter protection in which checks can be made that parts of the NAME field are a valid combination.
- **Snap** – Network snapshot to store expected messages and timing.
- **Flow** – Transport protocol flow monitoring to ensure that the data in a TP message is also in the correct order. This can have similarities with the use of Sequence Counters.
- **ComAddMon** – Commanded Address Monitoring involves counting how often the Commanded Address service is sent to a particular device. A threshold can be set for which an excursion above results in the issuing of an Alert.
- **CompComMon** – Complex Command Monitoring involves counting how often the service is sent to a particular device. A threshold can be set for which an excursion above results in the issuing of an Alert.
- **Encryption** - is a method by which information in the fields of a NMEA 2000 message are scrambled so that the meaning is hidden and only devices that have the key can unscramble.
- **Statistical analysis** - to help discover unusual patterns in signal behaviour.

7 Conclusion or Recommendations

NMEA 2000 is a CAN-based higher layer protocol for marine electronic device communications. This paper has highlighted vulnerabilities of the NMEA 2000 protocol. These have been broken down into three levels; CAN, SAE J1939 family-related and NMEA 2000 specific.

For some vulnerabilities, the solution is straightforward and has been discussed in this paper. Many of the attacks can be stopped, whilst others can at least be detected and therefore provide the opportunity for an alarm to be raised (e.g. audible, SMS or a NMEA 2000 Alert PGN). Many of the vulnerabilities in NMEA 2000 devices are the ones that make it easy to set up or configure. The ease of marine device installation is an important feature of the NMEA 2000 protocol. However, the features that make configuration easier are also those that are extremely easy to exploit with a cyberattack.

This paper has suggested some solutions to detect and help prevent NMEA 2000 cyberattacks. This will be a continued area of discussion and research to ensure that solutions provided both meet cybersecurity and industry requirements. The ultimate aim is to highlight the existence of these issues, rather than to fully standardize approaches to deal with them. Operational ambiguity can be a strength when it comes to cybersecurity.

It is unlikely that this review will be the last version and additional vulnerabilities will almost certainly be identified. Based on the findings of this paper, recommendations for NMEA 2000 improvement are:

- The implementation of the mitigations proposed in this paper should be investigated. These could be features that manufacturers should consider implementing in their devices. Perhaps some could lead to improvements in the NMEA 2000 protocol specification.
- Some of the mitigations proposed in this paper may not be suitable for being included as improvements to the NMEA 2000 protocol. However, application notes providing detailed implementation and design advice may be appropriate to help make devices cybersecure.
- Manufacturer specific mitigations could be implemented. Rather than standardising every single detection and prevention method, NMEA 2000 device manufacturers could implement these as they deem appropriate. Ambiguity can be a good protection against a cyberattack. Additionally, if device manufacturers implement their own methods, this can help avoid a single mode failure or vulnerability across all devices within a network.
- New devices are needed such as secured gateways and Intrusion Detection Systems (IDS). These new devices could benefit from new features in modern semiconductors that are provided to improve the cybersecurity of devices. The NXP TJA115x is a new secured CAN transceiver with such features as Allow/Deny lists and Bus Load Monitoring. Many new microcontrollers now have features to help improve cybersecurity such as flash write protections and cryptographic co-processors to help with the calculation of encryption routines.
- The NMEA leads its industry for NMEA 2000 installation training. New methods of installation can help mitigate intrusion convenience of attackers looking to gain access to a system by way of physical connection. Physical deterrents such as limiting node connections, is a low impact form to curve the potential of security intrusion. Other forms that are in practice today, is installing cabling through conduit, implementing node access lock box for additional node installations.
- Installing NMEA 2000 Class 2 devices for mission critical applications helps avoid the loss of function when a network is under attack. Critical functions can remain in operation with the loss of one network.
- For Class type ships it is important to design these systems for a zero-trust environment. This includes careful evaluation of chosen network components. Select devices that are primarily registered as NMEA 2000 certified. Select cabling components that have met the NMEA cabling approval process. Junction boxes, tee device and active connection media should be under scrutiny to avoid adding a risk level into the network infrastructure.
- Manufacturers of products should regularly report back to NMEA of any software or hardware versioning changes to allow network security devices to perform regular audits of trusted equipment.

8 References

1. McLaughlin T and Quigley C (2018); “CAN Bus at Sea”, Professional BoatBuilder Issue No. 177, February/March 2019.
2. Tran, K.; Keene, S.; Fretheim, E.; Tsikerdekis, M. (2021) “Marine Network Protocols and Security Risks”, Journal of Cybersecurity and Privacy. 2021, 1, 239–251.
3. Kessler G (2021); “The CAN Bus in the Maritime Environment – Technical Overview and Cybersecurity Vulnerabilities”, the International Journal on Marine Navigation and Safety of Sea Transportation Volume 15 Number 3 September 2021 DOI: 10.12716/ 1001.15.03.05.
4. Tindell K (2021); “The Janus attack”, CAN Newsletter 4/2021.
5. Tindell K (2020); “CAN Bus Security - Attacks on CAN bus and their mitigations”, Canis Automotive Labs, Document number 1901, Version 07, Issue date 2020-02-14.
6. Murvay P.S. and Groza B. (2018); "Security Shortcomings and Countermeasures for the SAE J1939 Commercial Vehicle Bus Protocol," in IEEE Transactions on Vehicular Technology, vol. 67, no. 5, pp. 4325-4339, May 2018.
7. Daily J. (2018); “Introduction to SAE J1939” Cybertruck Presentation, page 124
8. Chatterjee R, Mukherjee S, Daily J. (2018); “Exploiting Transport Protocol Vulnerabilities in SAE J1939 Networks”, Colorado State University https://www.ndss-symposium.org/wp-content/uploads/2023/02/vehiclesec2023-23053-paper.pdf
9. Quigley C (2023); “*J1939-based networks vulnerability to Address Claim Hunter cyberattack”, CAN Newsletter, June 2023.*
10. Quigley C. (2023); “Beware CAN with No NAME” – Professional Boat Builder, Professional BoatBuilder Issue No. 204, August/September 2023.
11. Tindell K, Gardiner B, Maag J (2023); “Implementation Requirements for Secured Gateways”, CAN Newsletter 1/2023.
12. Cho K. and Shin K. (2016); “Fingerprinting electronic control units for vehicle intrusion detection”, In Proceedings of the 25th USENIX Conference on Security Symposium, SEC’16, pages 911–927, Berkeley, CA, USA, 2016. USENIX Association.
13. Avatefipour, O., Hafeez, A., Tayyab, M., & Malik, H. (2017) “Linking Received Packet to the Transmitter Through Physical-Fingerprinting of Controller Area Network”. Information Forensics and Security (WIFS Conference, Rennes, France.
14. Shin K. and Cho K. (2017); “Identifying Compromised Electronic Control Units via Voltage Fingerprinting”, US patent US 2019 / 0245872 A1.
15. NMEA 2000 Main Document, v3.000, March 2022 (www.nmea.org).
16. Campo M.T., Mukherjee S. & Daily J.; “Real-Time Network Defense of SAE J1939 Address Claim Attacks”, SAE International Journal Commercial Vehicles, vol. 14, no. 3, Aug. 2021, doi: 10.4271/02-14-03-0026.

Acknowledgements

The authors would like to thank the National Marine Electronics Association (NMEA) for its support during this research.

Contact Information:

Warwick Control Technologies
Unit 3, Block A7,
Coombswood Business Park East,
Coombswood Way,
Halesowen, West Midlands, B62 8BH
United Kingdom
chris@warwickcontrol.com

Digital Yacht Ltd
6 Farleigh Court,
Old Weston Rd,
Flax Bourton
Bristol, BS48 1UR,
United Kingdom
paul.sumpner@digitalyacht.co.uk

National Marine Electronics Association (NMEA)
International Marine Electronics Alliance (IMEA)
846 Ritchie Highway, Suite L4
Severna Park, MD 21146
Standards@nmea.org