FUTURE THREATS TO ITS NETWORKS AND CAV INFRASTRUCTURE

F-Secure

AUTHORS

Role	Name
Lead Author	Vic Harkness
Author	Joel Clark
Author	Rich Perry
Author	Terry Ip
Author	Imran Mughal
QA	Thomas Kane
QA	Keith McDonnell
QA	Corey Forbes
Project Manager	Grant Day
Project Manager	Sarah Phillips

ABSTRACT

The aim of this report is to examine the security of Intelligent Transport System (ITS) networks. It has been written as part of a work package aiming to secure future ITS systems, run by the Centre for Connected and Autonomous Vehicles in cooperation with Zenzic and Innovate UK.

Within this report is an overview of the various technical aspects which make up ITS networks. Different components (such as vehicles, sensors, and data stores) are described, and their security examined. Much of the technology proposed for use within ITS networks remains theoretical; as such, this report deals largely with literature on the subject.

It was found that although (or perhaps due to) a lot of technology remaining theoretical, significant consideration to security had not been given. Many documents simply state security is out of scope, or will occur at another level within the networking stack. Some of this has been done by design. The use of 802.11p OCB mode to facilitate the highly dynamic ad-hoc networks necessitated for ITS networks negates many of the standard security seen in 802.11x protocols. This report provides commentary on where security vulnerabilities may lie within a large range of technologies, as well as suggestions as to how they may be fixed (where appropriate).

It has been proposed that various Connected/Autonomous Vehicle (CAV) testbeds should be placed around the UK for vendors to test new technologies on. These are likely to draw a high level of interest from a wide variety of attackers. Recommendations have been made on how these testbeds might be secured, and in general how ITS network security may be tested at scale.

The particular types of threat actors who may be interested in ITS networks are described within this report. Their resources and goals are examined to assess which techniques they may attempt to utilise. Based upon this, example attack paths within ITS networks based on a variety of attackers with varying end-goals are described.

Recommendations as how to how threats may be identified and mitigated against at scale are made. In support of this, methods of tamper detection are discussed. As data is already signed within the network, a method of encapsulating signatures is recommended. This would allow for the point at which data was tampered with to be more easily identified, so that remedial action may take place.

Finally, considerations for testing security within ITS networks at scale are discussed.

CONTENTS

ABBREVIATIONS

Abbreviation	Meaning
AA	Authorisation Authority
APN	Access Point Name
AT	Authorisation Ticket
BSM	Basic Safety Message
BSSID	Basic Service Set Identifier
CA	Certificate Authority
CA Basic Service	Cooperative Awareness Basic Service
CAM	Cooperative Awareness Message
CAV	Connected/Autonomous Vehicle
CNI	Critical National Infrastructure
CPA	C-ITS Certificate Policy Authority
CRL	Certificate Revocation List
DEN Basic Service	Decentralised Environmental Notification Basic Service
DENM	Decentralised Environmental Notification Message
DoS	Denial of Service
DSRC	Dedicated Short-Range Communication
EA	Enrolment Authority
ECTL	European Certificate Trust List
ETSI	European Telecommunications Standards Institute
IBSS	Independent Basic Service Set
IP	Intellectual Property
ITS	Intelligent Transport System
LDM	Local Dynamic Map
LNM	Local Node Map
LTE	Long-Term Evolution wireless broadband communication
ML	Machine Learning

Abbreviation	Meaning
OCB	Outside the Context of a BSS
OSI	Open Systems Interconnection
PKI	Public Key Infrastructure
RHW	Road Hazard Warning
RSU	Roadside Unit
SLAAC	Stateless Address Autoconfiguration
TC	Trusted Component
TLM	Trust List Manager
VMS	Variable Message Sign
VPN	Virtual Private Network
V2X	Vehicle-to-any
WLAN	Wireless Local Area Network

1 BACKGROUND

1.1 PURPOSE OF THIS DOCUMENT

This document consists of F-Secure's contribution to the Connected and Autonomous Vehicle Cyber-Security Feasibility Studies program. This program comes from a partnership between The Centre for Connected and Autonomous Vehicles (CCAV), Zenzic (formally Meridian Mobility) and Innovate UK, part of UK Research and Innovation.

The aim of the program is to gain technical insight in to how the future of intelligent road infrastructure for connected and autonomous vehicles may be created and supported securely. In support of this, F-Secure set out to:

- Find ways to measure and maintain cyber-physical resilience and identify vulnerabilities and threats to connected vehicle networks
- Determine methods to identify and mitigate complex threats
- Explore security considerations for roadside infrastructure from a representative CAV testbed, such as identifying previously unknown attack paths, allowing detection efforts to be focused where they would be most effective.

1.2 APPROACH TAKEN

F-Secure had initially hoped to have access to a representative example testbed for use during this project. However, this was not possible. As such, this report takes a theoretical approach; standard documents related to the ITS infrastructure have been examined, as well as papers proposing contributions to the ecosystem.

This report takes a largely European-centric approach. Proposed measures for EU systems are examined and commented upon. As appropriate, US-centric technologies have been mentioned.

This report makes use of the terms CAV (Connected/Autonomous Vehicle) and ITS (Intelligent Transport System). CAV refers specifically to a vehicle, whereas ITS refers to the general concept of intelligent vehicles and all the technology that goes along with them. ITS-S (ITS station) refers to a single ITS node, which may or may not be a vehicle. As such, all CAVs are ITS-S, but not all ITS-S are CAVs. The term "CAV network" and "ITS network" however can be considered equivalent, as they both refer to the overarching network within which vehicles and other nodes sit.

2 INTRODUCTION

As technology advances, our worlds are becoming more connected. Our homes contain numerous devices which allow us to socialise, shop, work, all without leaving our sofas. Increased connectivity has the potential to ease the burdens of daily life, and has been appearing in a large range of areas. As part of the increase in connectivity, larger and larger amounts of data are being generated by a forever growing army of “things”. Manually processing this deluge of data became impossible, and so humanity turned to autonomy to do it.

An area ripe for the developments described above is transportation. People have long dreamed of cars which take the hassle out of driving; you and your family get into the car, punch in a destination on a screen, and enjoy a leisurely board game whilst your car delivers you to your desired location. In some dreams, the car would fly. Whilst we have not yet cracked this aspect of the dream, the self-driving car is fast becoming a reality. Driver assist functionality has flourished, blurring the boundaries between human operation and autonomy. Trials of fully autonomous vehicles have already begun.

In support of more pervasive autonomy in transportation, ITS (Intelligent Transport System) roadways have been proposed. Upon these ITS roadways, ITS vehicles will drive. ITS sensors will sit on gantries, monitoring road conditions. ITS road signs will automatically update themselves. ITS roadside units will receive data from every aspect of this ITS ecosystem, and transport it onwards to ITS data hubs for storage and analysis.

Data will flow constantly around the ITS network. It will be used not only to control traffic, but for decision making in individual autonomous vehicles. Even non-autonomous vehicles may partake of the rich data collected within ITS networks. Vehicles partaking of this data are referred to as Connected/Autonomous Vehicles (CAVs). The numerous sensors included in roadways will be able to identify and track non-CAVs, so that their presence may still be incorporated into traffic models.

This report aims to examine how security will occur in ITS networks. The vast ecosystem of vehicles, sensors, and other devices generate large amounts of data. As many aspects of future ITS networks have yet to be deployed, this report takes a literature-based approach to examining the security of ITS networks.

The report contains an overview of the proposed technical components of ITS networks, and descriptions of their security postures. Suggestions are made as to how their security may be improved. The proposed networking protocols are described, and their security examined. The proposed message-signing certificate system is also described and examined.

After providing a description of the proposed ITS networks, the types of threats it is likely to face are examined. Different types of threat actors are described, alongside their capabilities and goals. This information is then used to demonstrate example attack paths which may be used to compromise various targets within the ITS network.

Recommendations for threat mitigations are made in this report, including how the security of ITS networks should be handled at scale. In particular, the security of ITS testbeds is examined. These testbeds will consist of sections of ITS roadway on which vendors may test new technologies. As testbeds are likely to draw the attention of attackers, methods for maintaining the security of trial vehicles and testbeds are described.

Finally, considerations for testing ITS infrastructure at scale are provided, as well as suggestions for how this may be done.

3 CAV ECOSYSTEM OVERVIEW

3.1 IN-VEHICLE SYSTEMS

3.1.1 OVERVIEW

Vehicles are no longer expected to be used solely as a means of transportation. They guide us to our destinations using maps. They advise us on the best route to take using traffic information. They entertain us using audio and video. In the future, they will ease the burden of driving altogether by being able to autonomously navigate to our destinations. To facilitate this, vehicles contain complex internal networks. These take data from various sources, including:

- Transmission from external sources
- Integrated sensors
- Input from users

To manage all this data and the large-scale internal networking, a large amount of computational power is required. This will only become more complex as cars become more intelligent, and eventually autonomous. In parallel to this, cars will be able to take more and more actions without direct input from human operators. This makes the internal networks of vehicles an increasingly attractive target to malicious actors.

The reasons for wishing to compromise an internal car network are numerous. An attacker may wish to cause the car to take malicious actions. They may wish to steal data from the car's internal network. They may even wish to use the car as a point of entry into the CAV network, pivoting to more critical nodes.

Due to their complexity, the internal computing systems of a vehicle resemble a standard network. Individual computers may drive different aspects of the internal system, such as handling the infotainment system, or the car's lights. These systems require network connections to one another, as you would see in a standard connected computer network. Because of this, the process of compromising a computing system within a vehicle, and then pivoting to another, resembles the process seen during a standard network compromise.

To reduce the risk posed by a compromised vehicle, it is important that safety-critical functionality of vehicles within the CAV network do not implicitly trust all data being received from other vehicles. For example the network should not implicitly trust the data coming from individual vehicles, the vehicle should not completely trust data coming from its sensors, and vehicle internal systems should not completely trust infotainment data being received remotely from third parties.

Some key pieces of the internal vehicle systems are highlighted below.

3.1.2 VEHICLE SENSORS

Modern cars contain a large range of sensors. This number will only increase as vehicles move towards acting more autonomously, as more data is required to make safer decisions. If an attacker was able to feed malicious data to one of these sensors, they could cause the vehicle to take unexpected actions.

Beyond this, attackers can leverage the methods the internal system uses to communicate with sensors. Tyre pressure sensors are a common target for such attacks, as the rotation of the tyres forces the use of wireless communication. Researchers have been able to deconstruct the protocol used to transmit tyre pressure data and craft their own messages to be sent in the correct format [1].

As we move towards vehicles with greater levels of autonomy, the reliance upon sensor data will become more concerning. A greater number of sensors means a greater range of opportunities for an attacker to subvert the operation of the target vehicle. The use of object recognition in driverless cars is particularly fraught with danger. Adversarial examples have been used to cause the Machine Learning (ML) models in cars to incorrectly identify road signs, interpreting a stop sign as a speed limit sign [2]. Images of people have been projected on to roads, where they are interpreted as real people by vehicles (leading to emergency stops) [3].

3.1.3 ENTERTAINMENT SYSTEMS

In-car entertainment systems have been advancing in commercial vehicles for some time now. The complexity of these systems has grown from a simple car radio, to TV screens in the back of headrests, to touch-screen interfaces for drivers. As vendors have raced to provide the next piece of functionality to users, the interconnectivity of the entertainment systems has also grown.

In modern systems, a driver may be able to control the climate of the vehicle, manipulate the lights, and get live traffic updates. Entertainment systems can connect to a phone using Bluetooth or similar technology allowing the driver to access their stored media. The system can be used to control car peripherals (such as lights or temperature). In the future, it is expected that drivers may be able to receive tourist information about points of interest as they drive past them [4].

If an attacker could gain entry to a vehicle’s entertainment system, they would be presented with a large amount of potential systems to pivot to because of this high degree of connectivity. Whilst some degree of trust is essential to enable the flow of data, lessons have been taken from standard computer network security. For example, firewalls limit which components may communicate with each other. Despite this, vehicle entertainment systems provide an excellent point from which an attacker may move to more interesting aspects of the vehicle.

3.1.4 APPLICATIONS

In the envisioned future of intelligent vehicles, it will be possible to install various applications on to vehicles. These applications can be for a variety of purposes, from safety to entertainment [4]. Two applications which have been set out so far are the CA Basic Service and DEN Basic Service. The exact functionality of these two applications is described later within this report (see sections 3.7.2 and 3.7.3 respectively). At the basic level they are used to translate data from the vehicle’s internal network into messages suitable for external transmission, and vice versa. It is unclear at this time what other applications may be developed in the future, or how they will be managed.

As these applications interface with the vehicle’s internal network to exchange relevant data, they will likely make prime targets for attackers. If a vehicle’s owner could be convinced to install a malicious application into their vehicle, an attacker may be able to take various malicious actions. Likewise an attacker who has gained access to another aspect of the system may be able to leverage a legitimate application to take further actions, or even compromise another aspect of the ITS infrastructure.

The limitations on how applications may be used remain unclear. It is possible that applications could provide an interface for remote administration of stations by vendors. Vehicle owners could use these applications to purchase addons or receive updates for their vehicles. The remote administration of vehicles has already been seen in Teslas, whereby the battery capacity of their electronic vehicles can be manipulated via remote updates[1]. If this is the case, applications may present a particularly high-value target for attackers.

3.1.5 VULNERABILITIES AND RECOMMENDATIONS

If a vehicle places complete trust in all aspects of its internal system, an attacker may more easily be able to compromise it. In general, attacks on vehicle sensors fall into the following main categories:

- Attacks aimed at sensors which will cause the car to take erroneous actions.
- Attacks which use sensors as a vector to gain access to internal car networks.
- Attacks which cause a vehicle to disseminate incorrect information to other network nodes.

Any attack on the internal sensor network has the potential to be very harmful to the safety and security of a vehicle. It is important that internal systems do not completely trust the data which is being provided by sensors. For example, if a camera is detecting a person in the road in front of the vehicle but the sonar sensors are not, some form of additional checking should take place before any responsive action is taken. Conversely, if all sensors are presenting similar reports, more trust could be placed in the reports.

However, an attacker may be able to subvert all sensors simultaneously. This would be a complex attack, requiring a much higher level of skills and resources than compromising a single sensor. Removing absolute trust in the data from each individual sensor would reduce the attack surface significantly, and dramatically increase the skill level required for an attacker to be able to carry out viable attacks, thus lowering the overall likelihood of compromise.

As a dedicated attacker may still be able to compromise a single aspect of the internal network, different components should not trust each other as default. The use of simple firewall rules (i.e. only allowing certain components to communicate directly with each other) acts as a first line of defence, but can be subverted. Because of this, data should be checked for validity prior to its processing. If a component expects data containing only numbers to come from a sensor and suddenly it is receiving special characters, this should be detected and acted upon. An alert could be displayed to the driver[2], and data from that sensor should internally be considered untrustworthy.

[1] https://www.washingtonpost.com/news/innovations/wp/2017/09/11/as-hurricane-irma-bore-down-tesla-gave-some-florida-drivers-more-battery-juice-heres-why-thats-a-big-deal/

[2] A careful balance needs to be struck when displaying alerts to drivers. If many alerts are displayed frequently, a driver will likely become ambivalent to their presence. Likewise, frequent requests to seek technical assistance would probably be disregarded. Removing access to functionality which is detected to be behaving maliciously may lessen the threat faced by a compromised sensor, but is likely to annoy the driver. The most appropriate techniques for alerting drivers to errors is out of scope of this report.

Due to the use of some sensor data in safety-critical functionality (such as autonomous driving), the potential harm from a sensor malfunctioning should be carefully considered. If a vehicle's cameras, sonars, and other spatial awareness sensors are all malfunctioning, autonomous driving should automatically be disabled. If, however, a tyre pressure sensor is malfunctioning, it may be appropriate to display a warning to the driver but allow the system to continue driving.

Similarly, a high degree of trust is required in the functionality of applications. Applications should perform validation checking on data coming from both internal and external sources to reduce the risk of compromise. The sources of applications should also be carefully examined. An "application store", similar to those used by mobile phone providers may be an appropriate choice. Applications would need to meet pre-defined security levels before they can be added to the store. The store could also perform checks on the integrity of the vehicle prior to allowing installation of certain applications, such as ensuring the car has not been jailbroken[3]. Similarly, application vendors who receive data from vehicles back to their own networks should thoroughly sanitise received data.

[3] Jailbreaking refers to the practice of subverting the security restrictions placed on a device by the manufacturer, often done to overwrite the software placed on the device by the manufacturer with one's own. Whilst this can be done to give the user more granular control over the system, it also subverts numerous in-built safety systems.

3.2 PERSONAL STATIONS

3.2.1 OVERVIEW

The ETSI standard for ITS networks presents the personal ITS station [5].These allow for ITS application and communication functionality to be used via a handheld device. There are several scenarios where a personal station may be used [6]:

- Information and navigation service for low speed users such as pedestrians, cyclists and wheelchair users.
- Information for users transferring between different modes of transport, such as within a car park, a bus station or other transportation hub.
- Multi-modal traffic information service, for use by transportation support staff to relay information to users.
- Multi-modal navigation service providing real time guidance for different modes of transport and re-routing capabilities, similar to current navigation applications available on mobile phones.
- Community activities where a group of vehicles are travelling to the same destination and the group can track the lead vehicle. This may be used in events such as cycling or running races where the station can be used to ensure competitors follow the correct route.

These devices may also provide a human machine interface (HMI) to other ITS stations, such as the touch screen interface currently available in newer vehicles. Limited information is available on these devices at the time of writing, and real-world implementations are hard to find. It is anticipated that they will take the form of a tablet or smartphone with the capability to communicate over 802.11p.

3.2.2 VULNERABILITIES AND RECOMMENDATIONS

Due to their smaller form factor, personal stations may pose an attractive target for thieves. Measures should be in place to deter theft, such as increasing the difficulty of using the device after it is reported missing or stolen. This would reduce the ability for a thief to resell or use the device.

Other security measures to lock down the device and reduce its attack surface should also be considered. If the ITS functionality is integrated into a multi-function device, there may be several vectors through which a threat actor might be able to remotely compromise it. It is recommended that threat modelling and subsequent testing of such a device is built into the product development from the initial design process. Likewise, standard security practices for networked devices should be used.

3.3 ROADSIDE INFRASTRUCTURE

3.3.1 OVERVIEW

To facilitate the future of ITS roadways, large amounts of roadside infrastructure will be required. Not all vehicles will be intelligent or autonomous, and yet will still need to be incorporated in to traffic models. Roadside detection capabilities could be used to identify and track these vehicles. This could take the form of CCTV cameras, RADAR detectors, or other detection systems. The data from these sensors, as well as from cars, will need to be passed on to other systems to be collated. To handle this, Roadside Units (RSU) will be used [5]. These are described in more detail later within this report (see section 3.4). These RSUs will also be able to provide data to other aspects of the roadside infrastructure, such as variable message signs (VMS)[4].

3.3.2 VULNERABILITIES AND RECOMMENDATIONS

CCTV systems have a long history of being insecure. They are often configured to be accessible by anyone to facilitate easier administration. Users of Shodan[5] have long found and shared the video feeds of insecure camera systems. A search of the tag "CCTV" on Shodan returns numerous results[6]. This means that an attacker may not even need to leave their home to gain a foothold within the ITS network. They could simply browse Shodan until they find a suitable target. Once the attacker has gained access to the operating system of a camera, they could then attempt to pivot to other aspects of the network (such as an RSU) or begin to manipulate the data being transmitted by the camera.

As data generated by cameras and other sensors will be used to inform traffic control decisions, it is important that the data is reliable. The methodology for generally securing hardware devices is out of scope of this report. However, it is important that external connectivity to these devices should be limited. For management of individual systems, administrators should have to use a Virtual Private Network (VPN) connection to connect to a specific machine (which has permission to communicate with the sensors) and administrate them from there.

As with data travelling within the internal network of vehicles, different nodes within the network should not implicitly trust each other. Any system which is receiving data from a roadside sensor should perform validity checking of data. Connections from unexpected locations should be disallowed, and strong authentication should be used throughout. This is especially true of aspects of the roadside infrastructure which expects to receive data, such as variable message signs.

[4] Road signs which can have the displayed symbol or text updated remotely.

[5] Shodan.io is a website which allows users to search for insecure devices which are connected to the internet.

[6] https://www.shodan.io/explore/tag/cctv

3.4 ROADSIDE UNITS

3.4.1 OVERVIEW

RSUs are designed to collect and collate data from other aspects of the ITS network, including vehicles and sensors. There have been numerous suggestions as to what is then done with this data. Some parties believe that RSUs should solely collate data and transmit it to the data hub. Others believe that the RSUs should be able to make complex decisions based on the information received, and act upon it. This would include allowing RSUs to directly command the movements of autonomous vehicles. In practice, it is likely that the functionality of RSUs will fall somewhere between these extremes.

RSUs may also be used to pass data from the backend infrastructure back to individual vehicles. The process of using RSUs to transmit pseudonyms certificates used to sign messages (see section 3.8) has been described [7]. They could also be used to pass traffic data to vehicles, providing more localised information than is available from commercial map applications. A typical RSU will be comprised of several components which allow it to communicate to both vehicles and upstream systems such as data hubs, as well as locally process the data received. The actual type of device used for each of these components may differ between RSUs depending on the implementation.

In general terms an RSU ITS-S (ITS station) gateway, according to ETSI specifications [5]), is the interface between roadside infrastructure and the internal ITS network within the RSU. Roadside infrastructure can include variable message signs located at the roadside, or mounted on gantries, as well as sensors such as inductive loops embedded in the road surface. These might exist as serial to device servers, which allow the connection of several devices over serial connections (e.g. RS232, RS422, RS485 etc) to an ethernet network.

Internal networking within an RSU is commonly connected using an industrial network switch [8]. An industrial switch differs from other enterprise grade devices in that they are designed for harsher physical environments. They will still support many of the features expected of enterprise equipment such as MAC address filtering, remote management and network access controls.

Within RSUs sits an ITS-S border router. The primary purpose of the ITS-S border router is to allow connectivity between the RSU and a data hub. Data hubs are repositories of ITS data, and are described in section 3.5. Data transfer may occur using an LTE modem, DSL router or other type of perimeter network device. As such, data may be transmitted through networks which are not part of the ITS network. This limits control over the transfer medium, increasing the risk of data being manipulated in some way. The transfer medium used will depend on several factors such as the availability of services within the location. This is discussed further in section 3.6.

The above described RSU components are commonly housed in a weatherproof cabinet. RSUs are located approximately 500m apart[7], resulting in hundreds of units for longer stretches of highway. Due to the potential cost and logistical implications of using higher security locks, the cabinets are typically secured using standard panel locks.

[7] The 500m value stems from in-person discussions with Surrey University

3.4.2 VULNERABILITIES AND RECOMMENDATIONS

If an RSU has been compromised, then it is possible that malformed or false data may be transmitted from the RSU or forwarded from a vehicle. This can lead to a flawed decision-making process, affecting routing and congestion control. As RSUs forward data between a variety of sources, a compromised RSU could lead to malicious data being shared pervasively within an ITS network. Different aspects of RSU security have been broken down below.

3.4.2.1 RSU PHYSICAL HARDENING

Roadside units are typically located in cabinets within close proximity of roads and motorways. In urban environments, these will be in busy environments with both vehicular and pedestrian traffic, as well as public and private surveillance systems. In more rural areas, such as near motorways, it is possible that the RSU cabinets are concealed by natural surroundings if installation closer to the highway is not feasible. This can allow a threat actor increased opportunity to obtain prolonged physical access to an RSU with limited chance of detection.

The use of secure locks can reduce the risk of an attacker from gaining covert physical access. In the Midlands Future Mobility testbed [8], the RSUs are protected by standard panel locks for which keys can easily be obtained. Locks should not be easily bypassed with common tools and should not display key bitting codes on the cylinder.

It should be assumed locks can be compromised in more isolated environments, and as such, physical monitoring methods such as CCTV or tamper detection sensors should be considered. Detected intrusions may be automatically cross-referenced with maintenance schedules or confirmed with field engineers out of band to reduce the occurrence of false-positive alerts.

3.4.2.2 NETWORK HARDENING

With physical access to the cabinet, it is possible to connect to the network switch located within. A threat actor can take advantage of any unused network ports to connect a laptop, or leave a small dropbox device connected for prolonged access. These dropboxes may be battery powered, or if there are spare outlets, they can be powered by the supplied power within the RSU instead. It is possible from this position to obtain a foothold on the network and access the components located in the RSU. Making it more difficult for an attacker to connect to a network can reduce the impact of gaining physical access to an RSU.

The use of MAC address whitelisting should be considered. It can be trivial to bypass, and MAC addresses are commonly found on labels stuck onto equipment [8]. However, it serves to prevent opportunist lower skilled attackers, such as hobbyists and hacktivists, from connecting to the network.

Further hardening can be performed by locking down unused network ports. As with MAC address whitelisting this can be trivial to bypass using additional network hardware, but it can increase the likelihood of a physical network intrusion being detected.

3.4.2.3 SIGNING DATA

If an attacker is able to gain a foothold on the network, they may be able to transmit data from the RSU. As the network is expected to provide an assumed level of trust, especially if connected via a private MPLS or VPN connection, such data may be assumed to be somewhat trustworthy when implementation decisions are being made.

The use of cryptographic signatures can be used to ensure data integrity is preserved. Such an approach should be used throughout the entire flow, from the vehicle to the data hub. This would maintain a chain of custody as detailed further in section 5.1.3 of this report.

3.4.2.4 DEVICE HARDENING

If configuration settings have not been changed from the default values, it would be trivial for an attacker to access the network and management interface of ITS components used in the RSU. This could be through common vectors such as weak or default passwords, leaving unnecessary services exposed, or the use of outdated firmware which is vulnerable to attack. There have been several incidents in the past where roadside infrastructure, such as temporary signage, has been compromised due to weak configurations. This has allowed attackers to access the management using the default or easily guessable credentials[8].

The hardening of device configurations can increase the effort required to compromise systems. These can include the following measures:

- Disabling or changing default accounts and passwords.
- Disabling unused services.
- Updating and patching systems to the latest stable release.
- Encrypting on board storage to prevent offline data retrieval.

For large scale deployments, the use of centralised management systems can decrease administrative overheads but can introduce new risks if not appropriately protected. Only the management system should be able to connect to the devices, with robust spoofing detection methods in place. A secure VPN connection should be required for administrators to connect to the management system.

Devices used to access RSUs for maintenance, such as field engineer laptops, should also be considered when assessing the risks to an RSU. Common configuration issues, such as allowing local admin access to engineers so they can install software when out in the field, can increase the impact of a compromise. It is also common to see these laptops used for personal use, which may lead to laptops being compromised with standard malware. This can increase the risk of a compromise depending on the type of use and sites frequented.

3.4.2.5 NETWORK SEGREGATION

Individual components within the network are unlikely to support features such as end to end tunnelling. The reliance will be on WAN connectivity to provide the appropriate network routing to upstream systems, such as centralised data storage solutions. Thus, a threat actor with a foothold within the RSU would be operating from a zone with an assumed level of trust. If unsigned clear text protocols are utilised, this could

[8] https://jalopnik.com/how-to-hack-an-electronic-road-sign-5141430

allow traffic to be intercepted or modified. An attacker could potentially move laterally to other RSUs or to the CAV testbed infrastructure in order to perform attacks, such as injecting false data or attacking the application and data storage systems in order to disrupt the road infrastructure.

Reducing the types of traffic flows permitted between devices and systems can reduce the impact of an attacker gaining a foothold on the network. Firewalls can be implemented at the perimeter of network segments or on endpoints to ensure that only predetermined use case flows are permitted. For cloud-based deployments, the use of appropriately restricted network security groups can reduce the risk of an attacker compromising the data storage systems from locations such as an RSU. Using automated deployment technologies such as Terraform, CloudFormation and Azure Resource Templates reduces the administrative overhead of maintaining the principle of least privilege in a scalable environment.

3.4.2.6 MONITORING

Monitoring of systems and data can help to detect anomalous activity. To do this effectively, it is important to look at each system to determine which threats it is likely to face. This can be done initially through activities such as threat modelling and mapping these risks to frameworks such as Mitre ATT&CK to identify the most relevant use cases. Example threat models for the ITS network can be seen in section 4.5.3 of this report.

It should be noted that threats continually evolve and the process of developing the monitoring use cases should not be a one-off exercise. Once the use cases have been implemented in the monitoring system, it is important to regularly review them to ensure that they are still appropriate and sufficient for the risk appetite of each organisation. Additionally, proactive monitoring and hunting for potential threats should be performed on a continuous basis. This includes looking for indicators of compromise based on published intelligence on specific threat actors.

3.5 DATA HUB

3.5.1 OVERVIEW

Data hubs will collate all information gathered within the ITS network. Data from ITS stations will be collected by RSUs and sent onwards to the data hub(s). The segregation of these data hubs is unclear at this point. It will perhaps be the case that different regions will host their own independent data hubs. Testbeds (see section 3.10) will most likely have dedicated data hubs.

The level of detail stored will likely vary depending on the size of the data hub, and its exact applications. A data hub for a region of the UK would be servicing a large area, and so be collecting massive amounts of data. It may not be appropriate to collect data on every vehicle and sensor every second. Conversely, in the smaller testbed data hub a high degree of granularity is both viable and desirable.

When a vehicle is driven through the testbed, as much data will be captured as possible. The vehicle itself will be able to log all data transmissions it had received during the trial, along with timestamps. The testbed will record timestamped data from all sensors on the testbed. This data will then be made available for the vendor to download once the trial has completed. Using these datasets, the vendor will be able to create a high-fidelity simulator of the testing environment. Simulations can be used by the vendor to digitally re-run trials with various tweaks in place, or even provide preliminary testing to other vehicles. Collected data may also be used in the training of Machine Learning (ML) models.

The way in which data stored on the data hub can be accessed will likely be vendor-specific. It has been suggested that a web portal be used. Companies running trials would be able to register with the application and use it to gain access to data relevant to their trials. The trials data, accessible through the web application, would likely be stored by a cloud hosting provider.

3.5.2 VULNERABILITIES AND RECOMMENDATIONS

The data hub of a testbed would likely be a high-value target for attackers. There are various reasons for an attacker wishing to target it, including:

- The deleting of competitor data to disrupt their testing and delay product releases
- The theft of competitor data to learn about their latest developments, or assess their product development lifecycle
- To use as a route to attack other targets within the testbed (e.g. vehicles)
- To cause reputational or financial harm to the operators of a testbed
- To steal vendors’ Intellectual Property (IP) Vulnerabilities and recommendations

Due to the volume of data which needs to be processed as part of an ITS, data hubs are likely to be cloud based. This infrastructure is accessed by RSUs and other internal systems in order to upload and potentially retrieve data. External commercial and academic entities also require access to the data for the purposes of analysis and research, meaning internet connectivity is required. Each accessing party represents a potential attack path via which a threat actor may gain access to the central system within a testbed.

Data access is expected to be enabled through a website application which can be accessed by both testbed owners and third-party organisations. Any vulnerabilities which exist in the website applications can lead to the compromise of the stored data. Because of this, standard website application best practices should be followed. Furthermore, some limitations may need to be placed upon the directional flow of data. Access to this application should not provide a direct route to push data back out to an RSU, as this may provide a vector to affect routing and congestion control.

Due to the data hub being a high value target, the operators must ensure the system's security is very robust. Recommendations on how to ensure the security of testbeds is discussed later within this report (see section 6.4 onwards). However, the security of the website and the cloud-hosted data should be carefully evaluated as any other online application handling sensitive data would be. This includes regular penetration testing through established vendors, as well as secure configuration reviews of cloud storage.

It is important that robust logging of any changes to the data within the data hub be put in place. In the event that data is in some way compromised, it would make it easier to understand how it happened and put remediations in place to prevent similar occurrences. These logs should be analysed regularly, perhaps using automated tooling, to help to detect any so-far unnoticed breaches.

3.6 TRANSMISSION MEDIA

ITS infrastructure can make use of several methods of data transportation to facilitate V2X communications. The following sections provide an overview of the transmission mediums that are likely to be utilised within an ITS network, including testbeds.

3.6.1 LTE

Typically for RSUs, Long-Term Evolution wireless broadband communication (LTE) based connectivity with a private Access Point Name (APN) is used, particularly in environments where provisioning a physical line is not feasible. The use of a private MPLS network can isolate the RSU from the internet, reducing its remote attack surface.

Vehicles also have LTE connectivity. The APN used by a vehicle for internet access is also likely to be private. As with the RSUs, this prevents a vehicle from being accessed directly from the internet. The primary use of this connectivity is typically to communicate with the car manufacturer to retrieve over the air (OTA) software updates. The vehicle is also likely to provide detailed telematics in order to allow analysis in the event of an investigation[9]. There are several scenarios that may trigger an investigation, such as a road traffic accident, theft, or even to decide if a fault occurs commonly enough to warrant a product recall.

In the event of an accident, on vehicles with systems such as eCall[10], a phone call will automatically be made to emergency services when certain sensors are triggered, or if a passenger manually activates it. Once activate, the in-car microphone is switched on and vehicle data, such as sensor information and location, is transmitted.

3.6.2 LEASED LINES AND DIGITAL SUBSCRIBER LINES

Depending upon available connectivity at the site of an RSU, a testbed operator may decide to share an existing connection in the vicinity with the local authority that manages the highway infrastructure. These may be sites where CCTV, ANPR and safety cameras are in operation. These connections, especially when using DSL based connections, can be more economical for transmitting large amounts of data. In these instances, the data will have to traverse a third-party network in order to access the data hub, which may carry its own set of risks to consider as discussed further within the Vulnerabilities and Recommendations section below.

3.6.3 WIRELESS AREA NETWORK

Communications between vehicle to vehicle (V2V) and vehicle to infrastructure (V2I), also known overall as V2X, will occur on wireless communications. This communication will be similar to wi-fi, except this will be peer to peer rather than based around centralised access points. This is enabled by using the 802.11p standard in OCB mode. This is discussed further in section 3.7.4 of this report.

[9] Car telematics data was recently used as evidence in a murder investigation: https://www.bbc.co.uk/news/uk-wales-51466273

[10] eCall is a system which can automatically call the European single emergency number (112) if an accident has been detected. . It is able to provide geographic location details, and uses a standard SIM card to transmit data.

Vehicles may also broadcast wi-fi access points to allow passengers to connect to the internet or control in-vehicle entertainment systems. Personal ITS stations may be connected to this. This access is likely to share the same LTE connection used by the vehicle itself.

3.6.4 VULNERABILITIES AND RECOMMENDATIONS

While the use of private APNs can reduce the risk of compromise from external sources located on the internet, their compromise can lead to an attacker potentially obtaining access to the network used by other RSUs and similar ITS-S systems. This attack might be leveraged by first gaining access to the LTE modem. If not appropriately secure, then it can be possible to obtain the APN along with the username and password. There is a high likelihood that these credentials can then be used with a SIM card obtained from an RSU to gain access to the CAV infrastructure.

Where connectivity is shared with a third party, the attack surface is increased. A compromise of the third party could result in the compromise of the data transiting the network. It is therefore essential that data transported across such a network is encrypted. This could be through the use of transport layer security or a VPN tunnel. Likewise, third parties should perform appropriate sanitsation of any data received from ITS stations.

3.7 CONNECTING PROTOCOLS

3.7.1 OVERVIEW

The ITS infrastructure makes use of multiple protocols to transmit data, as well as various transmission mediums. Details of the key protocols and transmission mediums are described below.

3.7.2 CAM

Cooperative Awareness Messages (CAM) are a type of message which may be sent by an ITS station (such as cars or RSUs) to other stations nearby. CAM are used to provide a “heartbeat” for the system, allowing nearby stations to maintain a knowledge of each other’s locations and statuses. This can be useful for the orchestrating of more efficient road usage, as well as for accident avoidance. To increase the speed at which these messages can be received and acted upon, they are not encrypted. They should however be signed by the transmitting station.

CAMs can contain various pieces of information, as set out within the relevant ETSI standards [9]. This can include status data such as the current time, the position of the station, the state of motion, and which peripherals are active. They may also contain attribution data such as the dimensions of the station, and in the case of road vehicles, the type of vehicle and its role in road traffic.

CAM generation is governed by the Cooperative Awareness basic service (CA basic service) on each station. This CA basic service controls the generation, transmission, and receival of CAM. It sits in the facilities layer of stations, and links into various other aspects of the system. For example, the CA basic service takes data from interfaces with the station’s internal network to construct CAM and may pass data from received CAM to the Local Dynamic Map (LDM). From here, other applications within the station’s facilities layer may retrieve data.

As CAM are designed to share information only to the stations which may immediately need it (i.e. those in the immediate vicinity), they only have a single-hop transmission range. Receiving stations will process the contents of received CAM to examine the relevance of the data. This can be based on the locality of the transmitting vehicle, or other pieces of information.

The ETSI standards for CAM [9] provide the facilities for a large range of information to be transmitted in CAM, as well as for future modifications to be made. The standards require that some confidence be placed in the data being transmitted. Data such as the current heading and the speed of the vehicle must be provided to the CA basic service with at least 95% accuracy. If this cannot be provided, it will not be included within the transmission. It is unclear how this confidence value is calculated.

Some checks take place when a CAM has been received by the CA basic service. The permissions indicated in the transmitting station’s certificate are compared to those being claimed within the CAM. This takes two forms: a check that the transmitting station has permission to be transmitting CAM at all, and a check that the station has permission to transmit as the role that it is claiming to be (e.g. as an emergency vehicle). This is likely done by comparing the self-declared station type within the message to the one described within the signing certificate.

Whilst the standard documents state that CAM should be signed using certificates and that the above permissions should be checked, it does not state the practical mechanics for doing so. Techniques for verifying the integrity of messages are also not described, as security is handled at the layer above CAM; security headers encompass the CAM data packet.

The general practice of signing messages consists first of the sender first calculating a cryptographic hash[11] of the data to be transmitted (in this case, the CAM). The resultant hash value is then encrypted by the sender using their private key. The encrypted hash value is transmitted alongside the message. The receiving party decrypts the hash using the sender's public key and recalculates a hash of message (the CAM). If the hash of the message matches the decrypted hash, then the integrity of the message has been maintained. If the hash values do not match, the integrity has not been maintained. This could be through a transmission error, or malicious tampering. Either way, the message should be disregarded.

3.7.3 DENM

Decentralised Environmental Notification Messages (DENM) are another type of message which may be sent by an ITS station (such as cars or RSUs) to other stations nearby. The exchange of DENM between stations is intended to inform the inbuilt Road Hazard Warning (RHW) application. This application aims to facilitate the more efficient flow of traffic in roads, and to enable better road safety. Although it was developed with safety purposes in mind, it can be extended to provide data to any application type which may call upon road condition information. Like CAM DENM are not encrypted prior to transmission, but they are signed.

DENM can contain various pieces of information as set out within the relevant ETSI standards [10]. This can include information about an event which has been detected. Events may be anything relating to road safety including ice on the road, the presence of a broken-down vehicle, or even a driver suffering from a heart attack. Alongside details of the event, DENM also contain a geographic area of relevance and other pieces of metadata.

The generation of DENM is governed by the Decentralised Environmental Notification basic service (DEN basic service) on each station. This DEN basic service controls the generation, transmission, and receival of DENM. It sits in the facilities layer of stations, and links into various other aspects of the system. The DEN basic service may interact with other applications on the facilities layer. Applications may make requests to the DEN basic service to transmit data, and to receive incoming data from it. The DEN basic service may also interact with other facilities layer entities such as the Local Dynamic Map (LDM).

Varying types of DENM may be generated, depending on their intended uses. These are as follows:

- New DENM: A report of a newly detected event
- Update DENM: An update to a previously detected event
- Cancellation DENM: Signals the end of an event previously detected by the same system
- Negation DENM: Signals the end of an event, as detected by a different station

As the events DENM report on are geographically bounded, so are DENMs. Although DENM contain data to identify the station they originated from, they are not tied to their originating systems. DENM are

[11] A one-way deterministic function which is used to calculate a mathematical "hash value". Hashing two identical pieces of data will result in identical hash values.

automatically forwarded between ITS stations within a geographic area, including to those newly entering the area of relevance. Upon receiving a DENM, the DEN basic service will automatically parse the contents. If the receiving station is within the area of geographic relevance of the message, it will automatically repackage and retransmit the message. Due to how the DEN basic service operates, the transmitting station identifier of the original transmitter is overwritten with that of the forwarding station. The message is also re-signed with the forwarders' certificate.

Each DENM must contain at least one "trace" [10]. Traces consist of a series of geographic waypoints leading up to the event being described by the DENM. In the case of an event being reachable from multiple directions, multiple traces may be included. When a vehicle receives a DENM, it may compare the route(s) described in the trace to its own intended path to assess the message's relevance. Exact methodologies for using traces are not prescribed by the ETSI standards document [10].

Similarly to the CA basic service, the DEN basic service requires that 95% confidence be placed in safety-relevant data (such as directional heading and speed) before it may be transmitted. The ETSI standard for DENM [10] states that a 95% confidence interval[12] applies to such data. Again, no clear mechanism for how this value may be calculated is provided.

As with CAM, DENM are signed. The integrity checking process is again not described, but assumedly would be done similarly to CAM. When a DENM is received, the permissions declared within the message are compared to those afforded by the signing certificate. If there is a mismatch, the received message is disregarded. The exact mechanism for performing this check is not described.

3.7.4 ITS-G5

ETSI standards recommend that for transmission in the ad-hoc networks formed by vehicle-to-vehicle communications, the 5.9GHz frequency band be used [11]. For this purpose, the ITS-G5 standard is proposed. Based on the OSI model[13], ITS-G5 sits across the physical layer and data link layer (which can further be divided into the medium access control and logical link control sublayers). The physical layer and medium access control sublayer make use of the existing IEEE 802.11p standards, whilst the logical link control is based on ANSI/IEEE Standard 802.2 [11].

Traditionally, there are two main methods of configuring 802.11x networks: infrastructure mode, and IBSS (Independent Basic Service Set) mode[14]. In infrastructure mode, all data travels through a single access point. Whilst this can be set in mesh mode to support multiple access points in a single WLAN (Wireless Area Network), all data would still need to be transmitted through access points. In ad hoc mode, systems in the WLAN may communicate directly with each-other. Although no central access point is involved, devices still need to be "joined" to the WLAN so that network parameters and transmission particulars can be agreed upon. Packets being transmitted in such a WLAN still contain what is known as a BSSID (Basic Service Set Identifier) to mark them as being part of the WLAN.

The ETSI standards recommend the use of 802.11p for peer-to-peer networking in a newly facilitated mode. This does not make use of either of the aforementioned architectures, and is instead based around the newly introduced "dot11OCBActivated" parameter. When this parameter is set to "true" in an 802.11p

[12] Confidence interval is a statistic which may be calculated based on previously observed data. The previously observed data is used to calculate a range within which plausible data may fall.

[13] A model which aims to provide a technology-agnostic description of how the various "layers" of the data transmission process interact.

[14] Also known as ad-hoc mode.

network, free communication between devices is enabled Outside the Context of a BSS (OCB) [12]. Devices no longer need to be joined to a WLAN, and the BSSID which would govern which systems messages would be transmitted to is set to a wildcard character. This means that messages are broadcast freely, with any device in range able to read them. Due to the lack of formal network joining, none of the standard authentication or association capabilities of 802.11p are utilised.

The decision to use this peer-to-peer networking method has been taken to facilitate the speed of operation required for ITS transmissions. Although RSUs could act as stationary access points alongside the road, it has been suggested that vehicles may not be in range of a given RSU for long enough for comprehensive handshaking to occur. It is also desirable that messages can be received and decrypted in very low time frames to enable real-time traffic control[15].

Within intelligent roadway areas, a massive number of messages will be transmitting at all times. Each ITS station is intended to send ten CAM per second, and various DENM as required. Due to the lack of base station governing the network, standard decongestion techniques such as time-slicing will not be used. ITS stations will monitor the congestion of the ITS-G5 channel by examining how many messages are being transmitted, receiving signal strength, and packet loss rates [13]. The results of this are used by the CA basic service to govern the frequency with which CAM are transmitted [9]. However, the ETSI DENM standard does not set out how the DEN basic service makes use of congestion control [10]. For this reason, it may be assumed that it does not.

Due to the 802.11p/OCB configuration used by ITS-G5 transmissions, the bandwidth available for CAM/DENM dissemination is severely limited. The European Union has provided permission for ITS-G5 transmissions to make use of the 5,470MHz to 5,925MHz ranges [11]. This is broken down into the following bands:

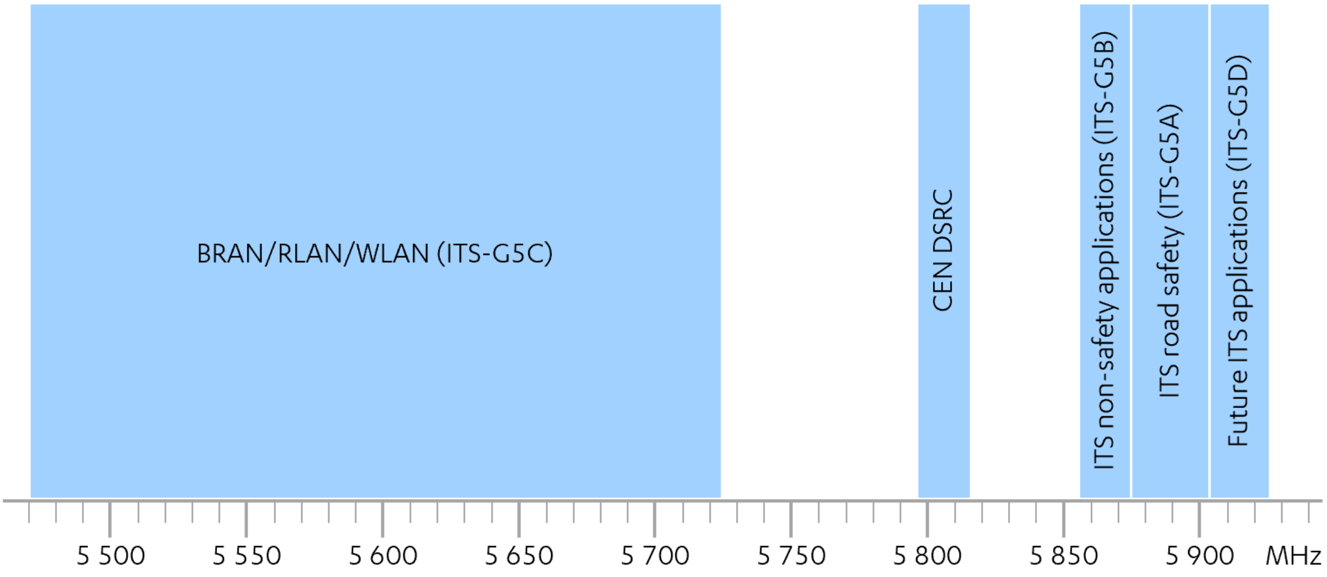

[15] During discussion of the technology with Surrey University, it has been proposed that RSUs could take a more involved role in traffic control. All vehicles would send messages to their closest RSU, which would then decide what actions each vehicle should take to enable optimal traffic flow. The RSU would then send commands back to vehicles. In the case of connected vehicles, alerts of these commands could be displayed to the human operator. In the case of autonomous vehicles, these commands would be taken as important inputs into the internal decision-making process. To facilitate this, Surrey University has suggested that the time from the moment an RSU transmits a message to when it has been received, parsed and actioned upon by a vehicle should be a period of no more than 20ms. This value appears to stem from the proposed latency required for safe vehicle teleoperation [28].

The purpose of these bands is described below:

Frequency range (MHz)	Band name	Purpose
5,470 to 5,725	ITS-G5C	For use in broadband radio access networks (BRAN), radio local area network (RLAN) and wireless local area network (WLAN)
5,855 to 5,875	ITS-G5B	For non-safety related road traffic applications
5,875 to 5,905	ITS-G5A	For safety-related road traffic applications
5,905 to 5,925	ITS-G5D	Reserved for future usage

By far the largest band available here is the ITS-G5C band. However, transmitting in this band requires dynamic frequency selection techniques to be used to avoid co-channel interference [11]. This is not possible in the 802.11p/OCB mode used by CAM/DENM, as transmission is not directed by a coordinating base unit. Therefore, transmission of CAM/DENM is limited to the narrow bands of ITS-G5A and ITS-G5B.

3.7.5 DSRC

Dedicated Short Range Communications (DSRC) is the US equivalent of the EU's ITS-G5. It makes use of similar operational procedures and operates in the same frequency bands as ITS-G5 [14], and contains similar functionality to ITS-G5. For example, DSRC uses Basic Safety Messages (BSMs) to share safety-critical information in vehicle-to-vehicle communications. Two types of BSM are shared [15]:

- BSM Part 1: These messages provide a "heartbeat" for vehicles, containing information such as a vehicle's size, location, and kinematic state. They are sent regularly at fixed intervals.
- BSM Part 2: These messages are sent on an ad-hoc basis, when an event has occurred which may be relevant. For example, vehicle brakes being activated.

BSM Part 1 and 2 can be seen as equivalent to CAM and DENM.

As this paper is aimed to provide advice on systems within the UK and the EU, DSRC and other US-only technologies are considered to be out of scope.

3.7.6 VULNERABILITIES AND RECOMMENDATIONS

In general, both the CAM/DENM protocols and the ITS-G5 transmission standard have not been designed with security in mind. Within the standards documents of these technologies, it is stated that security mechanisms are out of scope, or to be considered elsewhere within the ecosystem. This report has had to take a somewhat high-level approach to evaluating the security of these technologies due to the lack of a representative implementation. As such, this section will highlight a range of the more obvious vulnerabilities within them. Deep dives in CAM/DENM/ITS-G5 should take place in future work packages to ensure that all other issues may be discovered and remediated prior to any safety-critical application of them.

The lack of congestion control mechanisms within the DEN basic service has already been touched upon in section 3.7.3 of this report. Combined with the narrow bandwidth afforded by ITS-G5, Denial of Service (DoS) attacks present a serious threat to the ITS network. In the congestion control mechanism of CAM the current congestion of the channel is polled, and message frequency is lessened if congestion levels are high (a minimum transmission frequency is set which this mechanism cannot push beyond). DENM does not contain a similar feature. Whilst the CAM mechanism helps to ease congestion, it may well prove insufficient.

Accidental DoS appears to present a significant threat to ITS networks. If a large number of vehicles were present, CAM heartbeats would lead to some congestion. If an accident was to occur, this large number of vehicles may begin to transmit a large amount of DENM to report on it. All vehicles within the range of geographic relevance will automatically forward onwards these received DENM. Although ITS stations will not forward on a DENM relating to a given event if they have already recently done so [10], they will need to receive and parse each message to evaluate this. If the frequency of received DENM exceeds the processing capabilities of the DEN basic service, a vehicle may be rendered unable to receive and process more important messages. This may be especially risky if the DEN basic service processing capabilities of an RSU could be completely consumed, if the idea of having RSUs direct traffic is realised.

In general, the packet storm resulting from a large number of messages being transmitted simultaneously may prevent stations from receiving what may be relevant information. As previously mentioned, no interference remediations are in place as little coordination takes place in the transmission of CAM/DENM. It may be possible for a large numbers of vehicles in a small geographic area to act as jammers for each-other.

Likewise, intentional DoS is a considerable risk. The above conditions could easily be manually simulated; a large amount of meaningless messages within the ITS-G5 frequency ranges could be sent to jam legitimate communications. Due to the severe bandwidth limitations of ITS-G5 (as discussed in section 3.7.4), jamming of the ITS-G5 network presents a significant risk. As no peer relationships are established between nodes, nodes could not re-negotiate transmitting on a different frequency or using a different encoding mechanism as one would be able to do in traditional networks.

To increase network congestion, a malicious actor could even just replay legitimate CAM/DENM at a very high rate. This would have the added benefit of tying up the processing capabilities of the CA/DEN basic services on stations, as all received messages are processed to determine relevance.

Replay attacks within ITS networks could be used not only for the jamming of devices, but also for targeted disruption. In the envisioned future of ITS networks, it is possible that CAM/DENM could be used to send commands to autonomous vehicles. Consider the scenario where a human-operated vehicle is leading a fleet of autonomous vehicles. A person steps into the road in front of the human-operated vehicle, and so the driver performs an emergency stop. A message to perform an emergency stop is simultaneously transmitted by the human-operated vehicle, propagating along the autonomous fleet. An attacker is able to record this message. Once the pedestrian has left, the fleet begins moving again. The attacker then replays this message and causes the fleet to stop at their will. This could at best cause harm to the operating company due to delays, and at worst could cause a serious road accident if a vehicle following the fleet was not also able to emergency stop in time.

The simplest way to avoid replay attacks is through the inclusion of a nonce value in messages. A nonce value is a random number which gives a level uniqueness to a message. If a system receives a message containing a nonce value which it has previously seen, it may assume that the message is a duplicate and so disregard it. It is important that nonce values are both random and have a large range of potential values to avoid either accidental nonce collisions, or a malicious party from being able to deliberately “use up” large amounts of nonce values.

Although CAM/DENM do not contain nonce values, they do make use of timestamps. In ITS applications timestamps are considered as the amount of time since the start of 2004, in milliseconds [16]. In CAM, this value is modified prior to transmission. The value is divided by 65,536, and the remainder transmitted [9]. In DENM, the complete value is transmitted [10]. The time difference within which data is considered valid is likely to be application specific (if it is checked at all). In CAM, even if checking of this value did take place, a value could be replayed every ~65.5 seconds. Because of this, it is recommended that a dedicated nonce value be added to CAM. This could also be extended to DENM, although the need is less pressing as the timestamp value could be used for liveliness detection.

Within the ITS network, messages are not signed with a vehicle’s certificate per se. They are instead signed with a pseudonym certificate, which are cycled through by vehicles regularly. This may provide some protection against replay attacks, depending on how their invalidation is handled. An overview of how this system works as well as security considerations are discussed in more detail within section 3.8 of this report.

It may be possible for malicious actors within the ITS network to pseudo-masquerade as nodes. As discussed, DENM are designed to upon receival automatically be forwarded onwards to other stations within a geographic area of relevance. In doing so, the sender identifier value within the DENM is updated to that of the forwarding station, and the message is resigned using the forwarding station’s certificate. This feature could be useful to a malicious actor for obscuring their actions.

Although the malicious behaviour detection system within the ITS network is poorly defined, consider a scenario where a station is labelled as malicious if it has sent more than five malicious messages within a short time period. If a malicious actor had control of two ITS units, they could each send out three malicious DENM. Nearby stations would automatically retransmit these DENM, for a total of six malicious DENM messages. These stations could then automatically have their pseudonyms revoked, disallowing them from actively participating within the ITS network for a period of time. An attacker may also use the same principle to forward on their malicious messages without legitimate attribution being easy, essentially using other stations as proxies.

As stated in section 3.7.4 of this report, ITS-G5 makes use of the new 802.11p OCB mode. In this mode, nodes are not formally joined to a network; they instead send out directionless transmissions, with the onus being on receiving nodes to interpret the messages and evaluate their usefulness. Because nodes are not joined to a network, standard network security practices such as authentication are not used. To protect against the risk of unauthorised users acting within the network, messages are signed with certificates.

Depending on the point at which certificates are interrogated for validity, resources on the DEN and CA basic services could be used up by malicious parties. Even though messages bearing invalid certificates would be disregarded, an amount of processing will be required prior to this. In traditional networks, devices will disregard packets from unauthenticated nodes very early on. In this application, it may be that a station receives a message, passes it to the relevant basic service, which then performs some processing prior to disregarding the message. If the process for invalid certificate detection is not lightweight, this would provide yet another DoS vector to attackers.

Aside from the authentication issues, implementations whereby the CA/DEN basic services automatically parse all data received are worrisome. It may be possible for an attacker to fuzz[16] the basic services to find data which meets the formatting requirements of messages, and yet causes a crash in the CA/DEN basic service. Because of this, robust sanitisation of received data should be performed. There is no silver bullet for this; the more processing that is done on messages prior to their being processed and acted upon increases the time overheads of the system. However, in security-critical systems such as ITS networks, this should be an acceptable trade-off.

[16] Fuzzing refers to the practice of sending a system large amounts of varying (often pseudo-randomly generated) data with the aim of causing the target system to take some action outside of what is normally expected, such as crashing. This can be very useful for vulnerability discovery.

3.8 CERTIFICATE MANAGEMENT

3.8.1 OVERVIEW

It is important that the authenticity and integrity of messages being transmitted and received by ITS stations can be verified. This is especially the case when messages contain information which may be safety-critical, such as positional information. It has been proposed that each ITS station should have a certificate which it may include in messages to prove their authenticity.

The European Commission has proposed a security architecture, featuring a Public Key Infrastructure (PKI) alongside the use of frequently changing pseudonym certificates, to maintain the security of the system whilst balancing the privacy requirements of users [7]. This proposed architecture remains just that; the report states that it should not be considered as a statement, and that the accuracy of the data within cannot be guaranteed.

At the highest level of the PKI will be the European Certificate Trust List (ECTL), which is managed by a central Trust List Manager (TLM) entity. Membership to this TLM is appointed by the C-ITS Certificate Policy Authority (CPA). Below this will be a set of Root Certificate Authorities (Root CAs), with the report setting out the various security measures needed to facilitate this.

Below the Root CA will be two parallel systems, the Enrolment Authorities (EAs) and the Authorisation Authorities (AAs). These systems will be used for the authentication of ITS stations and will participate in the provisioning of the tokens which allow access to the ITS networks.

Manufacturers of ITS vehicles will be able to register themselves with EAs. Once approved, they may then register individual vehicles with the EA. During registration, the vehicle will be provided with various pieces of data. This includes:

- A globally unique canonical identifier.
- A canonical public/private key pair.
- Various certificates enabling the trust chain.
- Contact information for the EA/AA the vehicle is being registered to, including their network addresses.

At the same time, the EA will be provided with a matching set of information to enable authentication of later communications.

Once a vehicle has been registered with an EA/AA pair, it will be able to request its enrolment certificate from the EA. This enrolment certificate will then be used to acquire enrolment credentials as part of further authentication requests. This includes the renewal of the enrolment certificate, as well as in the requesting of Authorisation Tickets (ATs) required for communicating within the ITS network. The proposed methodology for managing ATs is described at a later point within this report.

Within the PKI, public keys will be used for encryption of initial enrolment, and of authorisation requests/responses. This is done to ensure the confidentiality of the data being used to authenticate. Messages should also be signed to assure the authentication and integrity of the data. The following schemes are to be used:

- ECIES_nistP256_with_AES128_CCM
- ECIES_brainpoolP256r1_with_AES128_CC

It may be necessary to update the certificates of stations for a variety of reasons. This should also be done using the encryption standards listed above. Certificates may be transmitted to stations in a variety of manners. This includes through ITS roadside infrastructure, through networks managed by others (such as cellular data or a WLAN), through wireless connections at electric vehicle charging stations, or through the onboard vehicle diagnostics port when the vehicle is at a garage.

The private key used in certificates at every level will expire after a varying period of time (e.g. for a vehicle, three years). Replacement certificates are to be created and distributed three months prior to expiry, to allow time for them to propagate around the system. Because of this, systems may have a maximum of two valid certificates at a time.

3.8.2 AUTHORISATION TICKETS

Vehicles which through the above described mechanisms hold enrolment credentials may request ATs. These ATs are used to sign communications to other ITS stations, proving that the transmitting station is trustworthy.

ATs were chosen to be implemented due to privacy concerns surrounding the use of a single certificate by each vehicle. If a single unique identifier was to be used at all times, a malicious actor may be able to track a single vehicle, or masquerade as a given vehicle for a long period of time. These ATs, also known as pseudonyms, were devised to provide a balance of privacy and security; a balance of unlinkability and of accountability for vehicles [17]. The use of these pseudonyms ensures that only a small number of bodies have access to the canonical information surrounding a station (i.e. the EA and the AA). Other stations will not be able to view such information.

In order to ensure the validity and origin of transmission over ITS, various security headers are used and signed with the inclusion of a message digest. This digest is encrypted using the private key component from a stations' corresponding pseudonym certificate. The public component may be used by neighbouring nodes to verify the sender.

Pseudonyms are designed to be changed frequently by a vehicle, even during a single trip. The European Commission has again provided guidelines on how this may be done [18]. It is recommended that each vehicle should have access to a fixed size pool of 100 pseudonyms. These would be drawn from randomly with equal probability, until the pool is emptied. A slight discrepancy appears here in management practices. The ETSI recommends that at this point a vehicle should be required to request a new set of pseudonyms [17]. The European Commission however states that the pool of pseudonyms may be reused, so long as they are still within their validity period [18].

Each pseudonym in the pool is valid for a period of 1 week. However, depending on implementation, each pseudonym may be cycled through multiple times during its validity window. Interestingly, the European Commission highlights the need for these pseudonym pools to be preloaded into a vehicle [18]. Vehicles may be pre-loaded with pseudonym pools for up to three months in advance. The terms of this do not appear to be specified. It is unclear if this is done automatically, or if it must be requested through a separate mechanism.

When a station cycles to its next pseudonym, it must also update its IPv6 address. Otherwise, it would be trivial for an attacker to track a vehicle despite the pseudonym changes [17]. To facilitate this, it has been recommended that Stateless Address Autoconfiguration (SLAAC) be used, as described in RFC 4862 [19].

RFC 4862 provides a framework for devices in a network to automatically configure themselves using a mixture of information being advertised by a router, and from locally available information. This allows for nodes on the same network to communicate directly, without having to pass data through said router. Although this is not quite the set-up created through the use of the 802.11p/OCB configuration, some principles from this RFC may be usable. For example, it describes how a node which has selected a new IPv6 address may advertise to its neighbours to avoid collisions.

The ETSI standards state that alongside this IPv6 change, "*the ITS-S Ipv6 layer shall modify the MAC address of each virtual interface which implies a change in the Ipv6 addresses associated to the virtual interfaces*" [17]. The wording of this point is unclear and could lead to variations in interpretation. It does however seem to be saying that the MAC address should be changed once the IPv6 address is. The exact mechanism for how a new MAC address may be selected is not specified. It may be performed in a similar manner to how the IPv6 address is changed.

3.8.3 CERTIFICATE AND PSEUDONYM REVOCATION

On occasion, it may be necessary to revoke the certificate associated with a given station. The revocation of a vehicle's EA registration would disallow it from receiving ATs. The lack of valid pseudonym means that other ITS stations would not trust and therefore not process data coming from this source. This would be an example of passive revocation, whereby an ITS station is removed from the network by no longer being able to receive valid credentials to participate.

The exact mechanism for revocation (as well as the threshold for revocation) varies between different types of certificate holder. The behaviour which will cause a certificate to be revoked varies within the PKI, as does the threshold at which a certificate is revoked. A concise set of terms for revocation does not yet appear to have been decided upon.

It has been proposed that to facilitate the revocation of certificates, a Trusted Component (TC) be used [20]. These would be a hardware component within the ITS station, which contains the certificate and associated pseudonyms. In the event that a certificate or set of pseudonyms need to be removed from the trust chain, a revocation command could be sent through the PKI down to the ITS. This would cause the TC to irrevocably wipe the data in question. As the ITS station would no longer have access to the certificate, it would no longer be able to engage with the rest of the ITS network. This technique relies upon the absolute security of the TC; in practice, attackers would likely be able to find a way to compromise the hardware[17].

In the ETSI/European Standards models, pseudonyms themselves cannot be revoked, although a single pseudonym may be blacklisted. The way in which pseudonyms may be revoked from a vehicle is through a trickle-down effect. If a vehicle has had its certificate revoked within the PKI, it will no longer be able to

[17] An example of people jailbreaking Teslas can be seen here: https://www.vice.com/en_us/article/y3mb3w/people-are-jailbreaking-used-teslas-to-get-the-features-they-expect

request new pseudonyms. However, due to the three month preload period, there may be a lengthy period of time before the effects of a revoked certificate are seen in practice.

3.8.4 VULNERABILITIES AND RECOMMENDATIONS

The key to enabling the above described certificate architecture is that the PKI remain secures. To facilitate this, the European Commission has set out robust specifications for how the PKI must be secured [7]. This report covers a wide variety of security methods, from the choices of encryption algorithms to how to physically secure the root CA's storage building. It is important that the security of the PKI be given very high priority. If an attacker was to compromise any of the high-level certificates within the trust chain, they may be able to wreak havoc in the ITS network. Alternatively, they may be able to collect sensitive data unnoticed for a long period of time. As such, care should be taken to follow the European Commission's recommendations, including for the auditing of the PKI to detect potential breaches or malicious actions.

Due to the distributed nature of the PKI, it should be ensured that it is secured at all levels. Any point of failure will have a knock-on effect to the security of any devices or organisations in the downstream. For example, if a vehicle vendor's certificate was somehow compromised, the attacker may have access to all cars developed by that vendor. To reduce the risk of this, regular auditing of all high-level PKI participants should take place. This could be facilitation by the consortium described in section 6.2 of this report.

The benefits of requiring stations to obtain a new certificate every three years is unclear. As the previous certificate is used to request the new one, if an attacker had compromised the old certificate they would be able to retain that foothold even when a new certificate is introduced. This however assumes that the old certificate has not been identified as compromised, and so revoked in some way; in this scenario, the new certificate may be loaded via some other mechanism. As such, the reasoning for changing certificates should be evaluated and the mechanism potentially re-examined.

The use of RFC 4862 in the cycling of pseudonyms may provide a vector for the Denial of Service (DoS) of a given ITS station. As there is a degree of pseudorandom selection of new IPv6 addresses (due to nodes not being assigned configurations by a single point), there is a risk of collision. The use of IPv6 increases the available address space and reduces the chance of a collision, but the RFC still addresses the risk. When a node selects a new configuration, it advertises to its neighbours its proposed new IPv6 address. If any neighbour is already using this address, it will indicate this in a response. If an attacker were able to remain physically close to a node (e.g. a box taped to an ITS station), it could automatically respond to each of these broadcasts claiming to already be using that address. This could hamper the ability of an ITS station to join and participate in the network.

RFC 4862 may also act to strip some of the anonymity of cycling pseudonyms. When a proposed new address is advertised, the station must surely make use of some other address so that it may receive responses. If one could match the address used in these advertisements to the one being proposed, they may be able to track a single vehicle through pseudonym changes. Additionally, if other pieces of identifying information (such as MAC address) are not cycled at the same time, an attacker could simply use this to track pseudonym changes.

Further vulnerabilities may lie in the way in which pseudonyms are generated. The mechanism for pseudonym generation is unclear. The generation of pseudonyms based upon a certificate will be actioned by the EA/AA pair. It is important that these pseudonyms cannot be used to infer the original certificate used in their generation. Likewise, it is important that related pseudonyms cannot be identified.

Each ITS station is provided with a pool of 100 pseudonyms per week. What happens once this 100 has been used remains open to debate. However, there seems to be an agreeance that a station may be preloaded with up to three months' worth of pseudonyms in advance. Whilst the use of these preloaded pseudonyms will be limited by their valid date range, it is still a large number of pseudonyms to pre-load a vehicle with.

In the proposed model, ITS stations may be detected as acting maliciously (the exact terms of how this would be done are not defined). If a station receives multiple reports, its certificate may automatically be revoked. However, as vehicles are not absolutely identifiable, it may only be possible to revoke the pseudonym that it is currently using. This issue is somewhat abated by the proposal to include Trusted Components (TCs) in stations. This would allow for an entire pool of pseudonyms to be revoked in one batch. However, as previously discussed, it is highly likely that attackers would be able to compromise TCs after an amount of time.

If a malicious actor was using their ITS station to send malicious data into the network, it would likely be detected and have its pseudonym revoked. So, the malicious actor would just switch to the next certificate in their pool. Depending on the terms of re-use, they may be able to request an additional set of 100 after these are all “burned”. Alternatively, they could purchase parts from write-off vehicles[18]. As these vehicles may be preloaded with up to three months‘ worth of pseudonyms, an attacker who is able to purchase the hardware responsible for dealing with these components could use this system for a period of up to three months to funnel malicious data in to the network.

The above technique of purchasing components from write-off vehicles can also be used as part of a Sybil attack[19] designed to deny others access to ITS networks. The viability of this will depend on if any stringent checks are set in place prior to a pseudonym being revoked. If for example it took three “malicious activity” reports for a pseudonym to automatically be revoked, an attacker could easily generate three messages using separate pseudonyms in a short space of time. This could be done via forced pseudonym cycling, or through the use of multiple devices. Various checks could be put in place to try and counter this, such as checking how many prior reports a pseudonym has placed, or the physical locations of the devices producing the reports. Due to the anonymising nature of the pseudonyms system, this type of attack would be hard to counter. However, some method of certificate revocation for vehicles which have been written-off should be considered.

Whilst the use of pseudonyms is good for user privacy, it is bad for system security. Without pseudonyms, an attacker who was able to compromise a data hub would be able to perform complex pattern of life analysis on a target based on a single vehicle. With pseudonyms, it would be very difficult for an attacker to consistently identify a single vehicle in the dataset. A dedicated attacker could follow a car to track pseudonym changes and match them up, but at this point they’re already physically following a car. A car could be tracked using compromised CCTV data based on its registration plate, but again, this increases the complexity of the attack.

However, with this level of privacy protection it is much harder to identify and take action against malicious behaviour within the ITS network. While the proposed use of TCs aims to reduce this risk whilst maintaining security, it is almost a certainty that these would be jailbroken after some time. Some technique for

[18] Anecdotally, intelligent vehicles contain more proprietary parts and so are harder to get repaired. It is often cheaper to declare them a write-off that to obtain the specialist parts required and pay for work to be done at a specialist garage.

[19] In a Sybil attack, an attacker creates a large number of fake identities to gain a disproportionate amount of influence within a network.

revoking entire pools of pseudonyms simultaneously is required, but the best way of implementing this is as of yet unclear.

The ETSI/European Standards models for pseudonym revocation mean that pseudonyms may not contribute as much towards the prevention of replay attacks as hoped. In a replay attack, a legitimate transmission is recorded and resent later. As a legitimate message is used, it may be accepted by the target as valid and processed. Changes in pseudonym have the potential to render replayed messages invalid. However, if a pool of pseudonyms is cycled through and reused, pseudonyms are likely not being invalidated as a matter of practice. If pseudonyms were invalidated upon cycling, this would reduce the window within which valid messages could be replayed. A single pseudonym may be blacklisted, but this seems to only be done if it has been associated with malicious behaviour.

3.9 DATA STORAGE

3.9.1 LOCAL DYNAMIC MAPS

3.9.1.1 OVERVIEW

Local Dynamic Maps (LDMs) provide a way for ITS vehicles to combine static map data with dynamic information on their surroundings. These maps can then be used by other aspects of the internal system, such as the route-calculating functionality of an autonomous vehicle. The data contained with the LDM is split up in to four layers:

- Permanent state: The static data which makes up the basic map. This includes roads, intersections, and points of interest.
- Transient static: Data which has a small area of relevance, connected to the permanent state layer data. This includes information on which lanes may turn in what direction, and for ITS stations would include the location of RSUs.
- Transient dynamic: Temporary data which might affect road conditions or routing plans. This includes weather information, and the locations of roadworks.
- Highly dynamic: Fast changing data which may only be relevant in the moment. This includes the positional data of other ITS stations, including vehicles.

The data contained within the LDM of a vehicle will be updated rapidly as information is gained from a variety of sources. A third-party vendor of map data may wish to provide updates to the data within the permanent state and transient static layers. Semi-frequent updates will need to be made to the transient dynamic layer, again perhaps by a third party. The highly dynamic layer will constantly be updating as the vehicle receives data from its local neighbours.

3.9.1.2 VULNERABILITIES AND RECOMMENDATIONS

The way in which third party data may be delivered to vehicles is unclear at this time. Regardless of the delivery method, the receiving vehicle should perform input validation on the data to ensure that malicious code has not been inserted. Vehicles should also have mechanisms to check the authenticity of the data provider. For example, public-key cryptography[20] could be used. Approved third parties should sign their data with their private key. Vehicles would be preloaded with copies of the public keys paired to these keys. Vehicles would therefore be able to verify that the data is coming from the expected provider. This would help to reduce the chance that an attacker is able to poison the LDM of a system with malicious data.

The data contained within the highly dynamic layer of an LDM will be coming from a source that is harder to authenticate; other vehicles and ITS stations. Due to the implementation of 802.11p/OCB in use for CAM/DENM transmissions, it may be difficult for an individual vehicle to determine if it can trust the data that it is being sent. Mechanisms should be put in place to prevent malicious actors within the ITS network

[20] Public-key cryptography refers to the use of a pair of keys to provide security in data transmission. An entity holds a private key, and releases a public key. People may sign messages to the entity using their public key, meaning only the owned of the public key may decrypt it. Alternatively, a person may sign their message using their private key, allowing people to use the public key to verify that the sender entity is who they claim to be.

from inserting malicious data into the LDMs of other stations, where it could lead to dangerous actions being taken.

3.9.2 LOCAL NODE MAP

3.9.2.1 OVERVIEW

The Local Node Map (LNM) is a system which sits alongside the LDM, enabling some of its operations. It maintains a list of information on neighbouring stations. This includes data on how neighbours may be identified such as their MAC address or IP address, as well as metadata on how fast the vehicle is travelling, and in what direction [5].

Little information is available on the nature of the LNM. However, it will likely contain the results of the pseudonym-rotation MAC/IP addresses changes described in section 3.8.4 of this report.

3.9.2.2 VULNERABILITIES AND RECOMMENDATIONS

Depending on the implementation used, the LNM has the potential to somewhat compromise the privacy afforded by temporary pseudonyms. By combining the LDM and the LNM, a station may track all neighbouring stations. This includes their MAC/IP address, their vehicle type, and their exact positions. If a neighbouring station was to cycle to a new pseudonym and therefore MAC/IP address, this would be reflected in the LNM. Through the data stored in the LDM and LDM, it would be trivial to match the old information to the new.

If the LNM is highly temporal, this is less of an issue. To keep track of the frequent changes, one would have to remain in transmission range of the target station. However, if a record of all the changes to a neighbouring station's identifiers (perhaps even those of the host station) was stored, an attacker may wish to compromise this data. If data from multiple vehicles could be obtained, an attacker may be able to match up the pseudonyms between logs and gain a comprehensive locational history for a target vehicle. Depending on the way in which pseudonyms are generated based on a certificate, it may also be possible for an attacker to utilise historic pseudonym data in an attempt to generate new valid pseudonyms.

Even if ITS stations do not habitually log LNM data, an attacker could potentially induce the system to do so. If a malicious application could be installed on to a target's system through some means, it could constantly poll and log the LNM data. It could even automatically transmit this data back to the attacker.

The ability of the LNM to track a single station between pseudonym changes has the potential to help identify malicious abuse of the pseudonym system. By comparing the geographic data of two received messages coming from different pseudonyms, it may be possible to assess the likelihood that two messages have come from the same vehicle. This could be used to detect pseudonym cycling as a method of gaining influence within a network or avoiding malicious behaviour bans. Likewise, it could be used to detect if multiple devices are transmitting from a single vehicle.

To counter this, attackers may spoof the geographic location of their malicious messages. A “pseudocar” entity could be created which follows the malicious party's real car. Spoofed messages could be crafted to appear to come from this pseudocar. To counter this, data within the LNM could be augmented with sensor data (e.g. sonar data or image processing data) to verify the physical existence of vehicles. As with many areas of cyber security, a technological arms race may arise surrounding this.

3.10 TESTBEDS

3.10.1 OVERVIEW

Once a vendor believes their technology may be ready to enter the market, it should be taken to a consortium-approved testbed for integrated trials. These testbeds will consist of a representative sample stretch of a future smart highway. This will consist of a stretch of road with all of the expected advanced infrastructure in place. Large-scale data logging will take place in this testbed. A vehicle vendor will be able to drive their vehicle along a stretch of road, and receive back all data relevant to the test. This includes timestamped data on what vehicles were present, what data was recorded by RSUs, the status of traffic lights, etc.

These testbeds will not be entirely isolated from the real world. Stretches of existing roadway will be upgraded to meet the testbed requirements. Because of this, regular road users will be able to use the roadways whilst testing is being conducted. This is representative of the expected CAV roadway scenario, whereby not all road users will have upgraded to networked vehicles. Due to the presence of civilian road users, the safety and security of systems must be demonstratable prior to the commencement of testbed trials.

Testbeds will provide a space for vendors to not only prove the efficiency of their technologies, but also to test their security in a more realistic environment. Pre-testbed trials will likely not provide as much coverage as testbed trials should do. For example, a system may behave differently when it is receiving erroneous data from one sensor at a time compared to every sensor simultaneously. During testbed trials, unintentional interactions may also be discovered. It is possible that data coming from a given gantry sensor may just happen to be in a format which causes an error state in the piece of technology being tested.

Because of the risk of “teething issues” for new vehicles, some sections of testbed should exist which do not have regular road users on them. This roadway could take the form of the testbed operator’s campus, a short piece of dedicated testing road, or even on sections of publicly accessible road which can be shut to regular road users for periods of time. These pieces of road would also serve as testing grounds for more full-on tests which are likely to expose unexpected behaviour, such as the aforementioned complete sensor failure test scenario.

It is important that testbeds do not just help vendors create security standards which are “good enough”. Although detecting and remediating against the simpler vulnerabilities will help to prevent unsophisticated attackers from targeting the ITS network, it will not stop more advanced threat actors. The cyber security testing performed within the testbed should aim to test vehicles with the resources which may be expected from a highly sophisticated threat actor.

The testbed should have a dedicated cyber security team which is on-hand during all trials, as well as numerous automated systems. These automated systems could make use of Machine Learning (ML) techniques to ensure a high degree of coverage is possible. For example, a dedicated vehicle fuzzing system could be developed by the in-house team. This team would be able to seed the model with data specific to the target vehicle, which would then be adapted by the ML of the fuzzing system to create a highly targeted fuzzing system. The use of an automated system in this context helps to ensure maximum coverage of trials, more so than would be possible by the human test team. Automated systems could be used to perform large variety of simulated attacks during trials.

The use of more heavy-duty testing techniques should not (at first) take place on main roads. Prior to any vehicle (or potentially other piece of technology) being deployed on roads with other users, failure states would need to be thoroughly tested. For example, during a DoS attack trial, the vehicle should still be able to operate safely. This is especially crucial in autonomous vehicles which are not intended to have human operators. Performing an emergency stop would not be sufficient; the vehicle should be able to safely operate without whichever sensor(s) or communication element(s) that it has been denied access to by the DoS attack, or safely remove itself from the road in some way.

To ensure the robustness of the security assurance provided during testbed trials, large amounts of detection and logging capability is required. Issues with autonomous cars may not be quite so obvious as a car driving off the road. The car might quietly be producing errors internally, which go undetected. Unnoticed errors may multiply and become more serious. It is important that the data collected within the testbed environment is examined carefully not only for errors, but for any anomalies. Dedicated tooling to automate some of this analysis could be created and maintained by the governing consortium.

3.10.2 VULNERABILITIES

Testbeds are likely to draw much attention from attackers. A full discussion of security considerations for both vendors and testbed operators before, during, and after trials is provided in section 6.4 onwards of this document.

4 THREATS TO ITS INFRASTRUCTURE

4.1 OVERVIEW

The future of ITS infrastructure will likely be considered as part of the UK’s Critical National Infrastructure (CNI). Because of this, it will fall target to a wide range of attackers, with hugely varying levels of skill and resources. Different attackers will have different end goals. This section of the report aims to:

- Highlight areas of the ITS networks which may be targeted by attacker.
- Identify key types of attackers which may target ITS infrastructure.
- Describe the goals which attackers may work towards.
- Model how attackers may gain an initial foothold within the network and move through other resources before reaching their goals.

4.2 ITS INFRASTRUCTURE ASSETS

4.2.1 INFRASTRUCTURE NETWORK

Due to the lack of representative testbed infrastructure for basing this work upon, a representative hypothetical ITS network has been used. The below diagram describes this hypothetical ITS network. It incorporates the ITS components proposed in the literature, and supporting equipment required to operate and monitor events.

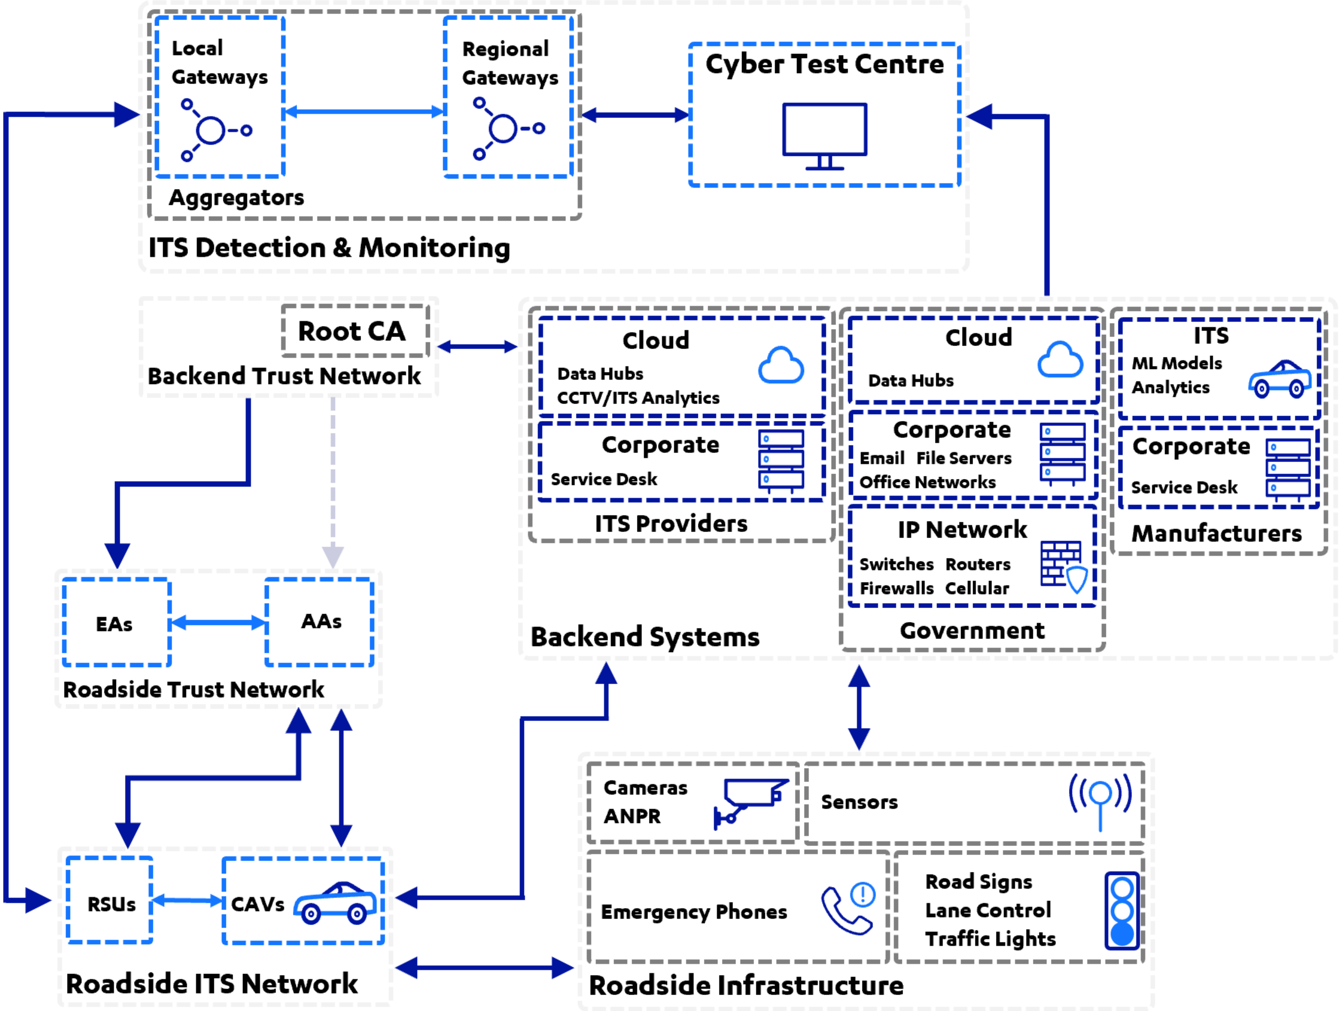

This diagram has been used for mapping threat paths within future ITS networks and testbeds.

4.3 THREAT ACTOR TYPES

4.3.1 OVERVIEW

As the ITS infrastructure and associated testbeds will form part of a highly network system, it is important to consider the different threats the system will be exposed to. Due to the complexity of the ITS network, it will draw the attention of a wide variety of attackers. These attackers will have differing levels of skills and resources, and will be interested in different end-goals. Threat actors can be categorised in a hierarchical structure, with high-risk better funded actors on top, and lower-risk less funded actors below:

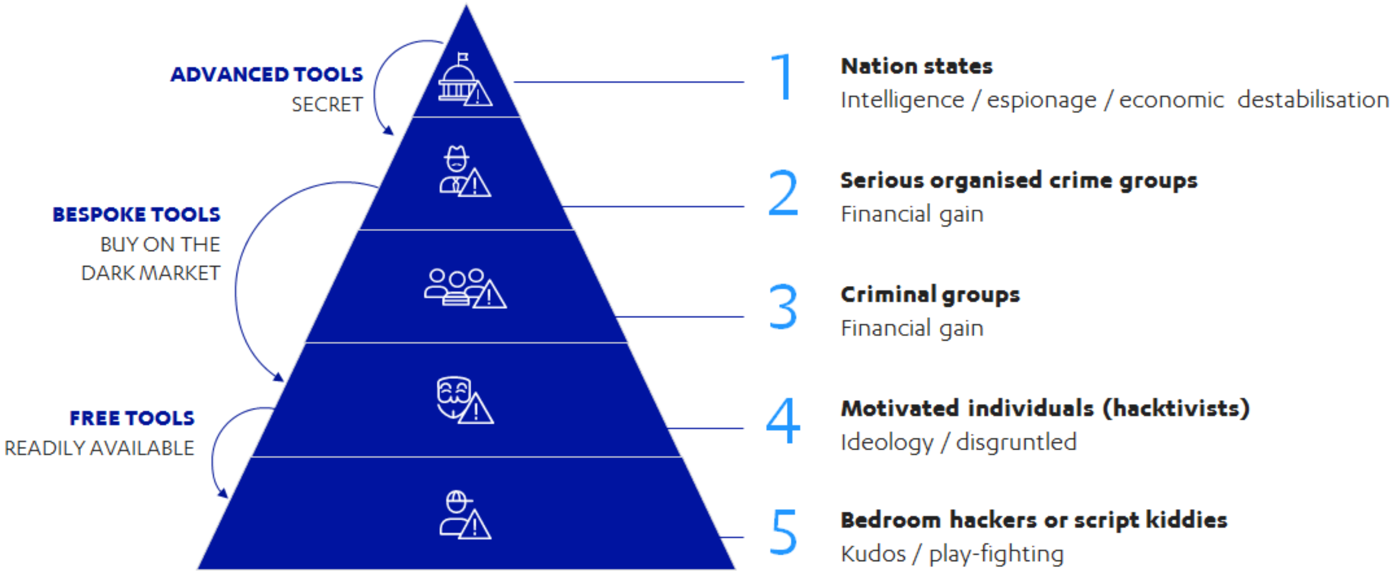

Although actors such as nation states and serious organised crime groups pose a higher risk than motivated individuals or bedroom hackers, these actors are much less frequently seen. Generally, more sophisticated attackers with large amounts of resources are few and far between. Script kiddies, however, are everywhere. Endless tutorials and tools are available online with low barriers of entry to their use. Whilst a nation state actor represents a significant threat due to their high levels of skill and extensive resources, the large number of unskilled "bedroom" hackers or low-skilled motivated individuals can pose a significant threat. After all, only a single attack needs to succeed to cause significant harm.

Although all of the threat actor types described above are likely to target ITS infrastructure, vendors should consider their individual risk appetites. It may be the case that a vendor decides that one particular threat is of such low risk that it would not make good financial sense to specifically protect against it. Vendors should seek external advice from risk assessment experts prior to making any such decisions.

4.3.2 ADVANCED THREATS

Advanced threats, often referred to as Advanced Persistent Threats (APTs), are located at tiers 1 and 2 of the threat pyramid. These represent highly skilled, well-funded, organised groups with intent to commit advanced intrusions for monetary gains or political leverage. F-secure has determined that advanced actors are identified by the following characteristics:

- High levels of funding, usually backed by adversarial governments.
- Well resourced, with advanced offensive security tooling, and collections of publicly available and bespoke exploits.
- Prioritising persistence, evading detection, and utilising specialised rootkits.
- Specialising in long-term operations focused around intelligence gathering, espionage, IP theft, large-scale disruption, and large monetary gain.
- Tendency to target critical national infrastructure, or other areas with potentials for large-scale harm.

Examples of advanced threat actors include:

- Organised crime and cybercrime groups, focused around leveraging misconfigurations in high value systems to install ransomware, commit fraud and other financially motivated activities.
- Nation state actors such as APT32, whose primary goal is to steal IP from automotive manufacturers.

APT32: OCEANLOTUS GROUP

One key threat actor group which may pose a threat to CAV networks is the cyber espionage group designated by FireEye as OceanLotus Group or APT32 [21]. The group is understood to operate for the benefit of Vietnamese state interests.

Active since at least 2014, the group's initial focus was political influence in Vietnam and Southeast Asia . Many of the tactics, techniques and procedures which have come to characterise APT32 were first observed in attacks targeting Vietnamese journalists and dissidents. Since 2016, APT32 has been observed targeting private sector organisations in manufacturing, technology, and banking. More recently, the group has begun targeting the automotive sector. It is believed that their economic espionage seeks proprietary business information and intellectual property in support of the country's emerging automotive sector.

APT32 is known to gain access to target networks using malicious attachments on spear phishing emails. The group is highly skilled at using social engineering methods to trick the victim into enabling the legitimate functionality which allows them to achieve their objectives. APT32 has proven itself quite adept at operating stealthily; in engagements where F-Secure's Incident Response (IR) team has encountered the group, they are often found to have been present and undetected in the target network for lengthy periods. The group is well resourced with a significant number of signature malware and backdoors [22]. These are often used in conjunction with commercially available offensive tools such as Mimikatz and Colbalt Strike.

Last year APT32 was believed to be responsible for hacking into the networks of several automotive manufactures including BMW, Hyundai and Toyota [23]. It is not known what propriety business information was accessed in these attacks. APT32 is also believed to behind the breach of Japanese Toyota dealerships resulting in the compromise of 3.1 million items of data pertaining to Toyota customers [24].

4.3.3 TECHNOLOGICALLY COMPETENT GROUPS

Technologically competent groups are found between the 3rd and 4th tiers of the threat pyramid. They encompass several types of individuals and smaller groups, motivated by varying reasons. Actors in this group will have the necessary skills required to identify and exploit a number of vulnerabilities present within ITS networks, without the large financial or technical backing as available to APTs. Actors that fall within this tier of the pyramid are usually identified by the following characteristics:

- Hobbyist level funding, though some groups may be better equipped with advanced tooling.
- Prioritising financial gain, disruption, or reputational impacts.
- Primarily financially motivated, though may be politically inclined.
- Focused on identifying misconfigurations and utilising public exploits for quick rewards.
- Typically, attacks committed originate from jurisdictions which don't enforce cybercrime legislation.

Examples of advanced threat actors include:

- Groups like Anonymous, whose members are motivated primarily by showing technical prowess or causing disruption in order to gain mainstream notoriety.
- Hacktivists, such as environmental groups that don't like motorways, or don't like the idea of a new motorway infrastructure with the potential to impact the environment.
- Researchers, who whilst not setting out to negatively impact the network or its users may cause reputational harm to companies upon vulnerability disclosure, or through the public release of tooling.
- Industry competitors who may attempt to disrupt testing or impact performance of products produced by competing manufacturers. They may try to steal IP on CAV technology during trials or in active operation.
- Insider threats, typically rogue employees, that either accidently or purposely sabotage the security of system components to negatively impact the security posture or operation of the system.

4.3.3.1 LOW-COMPETENCY, UNFUNDED INDIVIDUALS

These types of threat actors found in tier 5 of the threat pyramid. They are typically characterised as low skilled and opportunistic. They will be in one of three common categories of threat actor:

- Low level of technical knowledge. Actors in this group usually operate by using premade and automated tooling with the aim to exploit exposed services.
- Motivated out of curiosity or an attempt to get fame, though individuals may be looking to commit crime and defraud companies, such as gain access to paid upgrades on cars for free.
- Non-technical parties who may not understand the security impact of their actions.

Examples of low-competency actors include:

- Script kiddies, reliant on premade tooling. They will likely only exploit remotely accessible systems. Their motive is generally to gain respect from their peer groups and other hackers.
- A person who heard from a friend that they could get a free upgrade on their car if they just plug a USB stick into it and press a few buttons.
- Vandals, typically motivated to break open physical infrastructure enclosures to steal equipment or cause disruption for minor monetary gain.

4.3.4 INSIDER THREATS

Insider threats are a specific class of threat characterised as individuals operating in organisations with access to the ITS infrastructure. Typically, these will be employed by an organisation or government entity, either as a member of staff or as a contractor. They possess knowledge of and have access to internal systems and processes. Insider threats are more able to identify weaknesses which can be leveraged to bypass security controls. Typical motives for an insider could include revenge, malice, coercion, bribery, or ignorance.

A malicious insider could be a disgruntled employee attempting to cause disruption and reputational damage to the organisation. This may be motivated by disagreements between the insider and the organisation or a colleague. Common insider attacks stem from contract workers who did not have their remote access removed once their contract was terminated.

Bribery, radicalisation or blackmail may be used by an external entity such as cybercrime groups or APTs to coerce someone into performing malicious actions on their behalf. The motivation will vary depending on the entity. APTs may even have the resources to have a stooge gain employment legitimately within the target organisation as a method of gaining persistent access to resources.

It may also be possible for weaknesses in a system to be the result of an accident or misconfiguration. In this instance, no ill-intent has been directed at the organisation, but an exploitable vulnerability may be exposed. An example of this could be an employee copying code from Stack Overflow[21] and adding it to a product's code base without fully inspecting it.

[21] A website where people commonly share code excerpts and help each other with programming problems.

4.4 ITS ASSETS

4.4.1 OVERVIEW

The ITS infrastructure and its associated stakeholders are made up of physical and non-physical assets. These assets are targets which may be sought after or exploited by threat actors. In this paper, assets have been separated into categories which represent the different components that make up the ITS infrastructure.

4.4.2 PHYSICAL ENTITIES AND ASSETS

Physical assets are a clear target for threat actors. Although their compromise may not be the goal of an attack, they are likely to be targeted in order to achieve a foothold onto the network. The following non-exhaustive list enumerates the physical components that may be found in the ITS and supporting backend networks:

ITS Gateway – Regional nodes used to collate or aggregate data with roadside equipment and usage of the ITS network. Distributed detection mechanisms could also forward malicious event and logging data through these nodes.

ITS RSU – Roadside units which form the bulk of the fixed infrastructure communicating with CAVs. Designed to provide up-to-date event and congestion data to vehicles on the road, or potentially direct traffic flow.

CAV/ITS Vehicle – Connected and Autonomous Vehicles are vehicles designed to take advantage of the intelligent highway infrastructure to optimise available routes, passenger safety, road congestion, and the driving experience.

ITS Application – Applications which run on an ITS station, including the basic service components inside a CAV used to communicate with the ITS network. These applications are anticipated to be approved by manufacturers and extend the functionality available to the end user.

ITS Charging Stations – Some electronic charging points will be equipped with ITS stations capable of providing secure or stable connections for expansive data downloads or system updates.

ITS PUs – Personal Units consisting handheld equipment, including monitoring and maintenance devices, as well as mobile phones. PUs will be equipped with components which allow them to communicate with ITS interfaces on vehicles and other systems.

ITS Garage – Maintenance and repair garages equipped with facilities to interface, diagnose and repair ITS compatible vehicles. Likely to be brand specific due to the specialised parts used in CAVs.

Related Services – Backend Services support the operations, analytics, data collection, and monitoring of the network. These may not directly make up part of the ITS network. Backend networks will be owned and managed by different stakeholders, including:

- Government owned network lines which event and traffic data flows through.

- Manufacturer networks, which take metadata from vehicles (such as telematics), and provide data to vehicles (such as infotainment, or other subscribed addons).
- Industry networks, such as the Midlands Future Mobility (MFM) testbed, or test systems used by security researchers.

Backend Servers – Any large holding of ITS data. For example, this could be data collected at a testbed during a trial, traffic metadata held by road design organisations, or a CCTV feed repository.

Cloud-Hosted Services – Enabling ITS related services which are hosted in the cloud. This could include website portals which allow for access to data, distributed automated monitoring systems, or automatic testing systems.

Cloud-Hosted Servers – Similar to Backend Servers, these will be data repositories deployed on cloud infrastructure for greater availability and accessibility.

ITS Testbeds – A subsection of the ITS roadside network deployed as a test environment, where manufacturers can test their pre-production CAVs and other pieces of ITS technology.

ITS Test Centre – Centralised Facility for monitoring and managing the ITS network. Security alerts and operational data are all forwarded to these systems to provide comprehensive oversight into the network’s state.

ITS Trust Network – The ITS network makes use of device-level certificates within an overarching PKI to provide authentication and validation for node communications. Aspects of this will be contained within dedicated physical infrastructure to ensure secure storage. The PKI consists of various elements, including:

- Root CA – Authority designated as the ultimate decider of authenticity of entities in the system.
- Enrolment Authorities – Authority used to onboard and generate identities for CAVs and other ITS systems communicating in the network.
- Authorization Authorities – Once enrolled, entities request tickets from the AA. These are used to authorize access and provide pseudonym identities to vehicles, so they may utilise the network without sacrificing their privacy.

4.4.3 DIGITAL ASSETS

Although non-physical assets are not as visible a target as physical assets, they are a serious target for threat actors. Through the theft of data, an organisational attacker could gain a competitive advantage. Through the destruction of data, the efforts of other organisations could be hampered. Through the insertion of malicious data, ITS nodes could be caused to take malicious actions.

Data represents a large amount of the non-physical ITS assets. Due to their potentially high impact or large monetary value, data will be a high-value target for threat actors. Some of the data regularly broadcast within the ITS network is described below:

Sensor data – Data produced by various sensors on roadside equipment as well as from those onboard CAVs, e.g. sonar detection data of nearby vehicles.

Regional events data – Data associated with temporary conditions on the road, e.g. the presence of black ice on the road ahead.

Congestion data – Data describing the current road capacity, e.g. details of a road blockage.

Weather data – Meteorological data regarding weather conditions for a localised area.

Routing data – Data describing the route being taken by a vehicle.

Identifying information – Data maintained by ITS nodes on the identity of surrounding vehicles, e.g. the contents of the Local Node Map.

Telemetry data – Data used by manufacturers to track the actions of a vehicle, e.g. the locational data and the status of peripherals.

The following data will also be included within the ITS network, although more often at rest:

Testbed results – Data which has been collected both by sensors within testbeds, and by ITS nodes themselves during trials. Performance metrics and other sensitive data may be included. This data will be made available by the testbed operators to test participants via an online portal.

Financially-sensitive data – Data related to billing for different features and services contained within the ITS network. This may be relevant to applications, or for the purchasing of firmware upgrades from manufacturers.

Intellectual property – The firmware contained within nodes, or any other proprietary information within the ITS network. This could include information relating to scheduled future trials.

The above data may be used by vehicles, both autonomous and human-operated, to shape driving behaviours. For example, an alert that there is black ice on the road ahead would inform a driver's choices, causing them to slow down. Because of this, it is safety-critical that data integrity is maintained.

4.4.4 REPUTATIONAL RISK

Alongside the above physical and digital assets, vendors within the ITS ecosystem may also consider the security of more difficult to quantify assets: their reputation, and financial status. A loss in one will lead to a loss in the other. As part of protecting their assets, organisations will also want to protect these.

Maintaining the reputation of a company in part relates to the public's perception of them. Negative press surrounding vulnerabilities or data loss would cause an organisation to suffer reputational losses, as the general public would place less trust in them. Likewise, the government may stand to lose public trust if the ITS network is seen as insecure. Although reputation can be managed through good public relations, the ideal way to build and maintain a good reputation for ITS ecosystem participants is through the delivery of secure products.

If a company suffers reputational harm, financial harm will most likely follow. The stock prices of a company could drop, or sales. A company which has suffered significant reputational harm could even lose access to markets. If an ITS technology vendor was found to have massive amounts of vulnerabilities in existing technologies and to have been subverting standard ITS vulnerability detection processes, it may be appropriate to blacklist them from involvement within the network.

4.5 THREATS TO ITS

4.5.1 OVERVIEW

A large variety of nodes are connected within the ITS network, each with its own individual risk profile. The connectivity between nodes increases the attack surface exponentially, creating a tiered attack surface which may provide entry points to attackers of varying skill levels. A vehicle node would likely have a high barrier of entry to compromise due to extensive physical hardening. A CCTV camera may have remote administration enabled, using weak credentials. If an attacker could compromise one of these nodes, they could attempt to pivot to the other via the ITS network.

4.5.2 KEY TARGETS AND GOALS

This section discusses some of the key targets and goals which an attacker may have. The list is in no ways exhaustive, as new end goals are likely to be constantly emerging. It also remains to be seen what types of attacks are worthwhile. For instance, whilst it may be possible to induct vehicles into a spam botnet, this may not be seen in practice due to compromised vehicles having much more useful applications.

4.5.2.1 TRAFFIC THROUGHPUT RATES

An attacker may wish to manipulate the rate at which vehicles travel through an area of road. This could be done by:

- Simulating traffic jams in some roads to cause bottlenecks in others.
- Denying vehicles access to the ITS infrastructure so they do not feed into congestion statistics.
- Manipulating data which may be used by autonomous vehicles to inform driving style, such as weather data.
- Manipulating traffic lights to alter directional traffic flow.
- Masquerading as a dangerous goods vehicle which is unsafe to overtake, causing autonomous vehicles to become “stuck” behind it.
- Creating fake road accident reports.
- Manipulating map data to create fake road closures.
- Manipulating speed sign data to slow traffic.
- Physical or digital damage to RSUs to prevent their operation.

Because of the complexity of ITS networks, this list is by no means exhaustive. The potential attack vectors available to an attacker are endless.

4.5.2.2 MANIPULATION OF AUTONOMOUS VEHICLES

The threat to autonomous vehicles is much greater than that faced by connected vehicles; there is no human in the loop to correct malicious actions when they occur. There are various reasons why an attacker may wish to manipulate an autonomous vehicle; some examples are given below:

- Replaying of unlock signals or blocking of lock signals to facilitate theft.

- Physically harm the driver.
- Use the vehicle as a kinetic weapon against a person, object, or building.
- Harm the reputation of a vendor by causing accidents.
- Gain fame/infamy for "hacking" autonomous vehicles.
- Cause public mistrust in autonomous vehicles.
- Cause road accidents to deny access to areas of road.
- Highjack a vehicle temporarily for other purposes.
- Deny targets from transportation.
- Track a target.

4.5.2.3 TESTBEDS

Testbeds will provide a place for vendors to test new technology. Because of this, testbeds are likely to draw the attention of more advanced threat actors. They may also draw the attention of less skilled attackers, who are hoping to gain some "quick wins" from the less secure technology being tested. Reasons to target testbeds include:

- Gaining insight into competitor developments before they are formally announced.
- Theft of competitor IP.
- Disruption of competitor tests to delay releases.
- Compromise of a vendor's new technology with the hopes of gaining access to the rest of their network.
- Compromise of a vendor's new technology with the hopes of gaining access to the rest of the ITS network.
- Manipulate the machine learning models of vehicles which may still be in training.

4.5.2.4 TECHNOLOGY VENDORS

Vendors within the ITS supply ecosystem are likely to be targeted by attackers for a variety of reasons. These include:

- Journalists wishing to get previews of upcoming technology.
- Attackers wishing to steal IP directly from vendors.
- Attackers aiming to gain a route into all technology associated with a vendor (e.g. every vehicle by a single manufacturer).
- Identification of key employees to target in other attacks.
- Targeting of production systems.
- Theft or other financial gain.
- Ransoming of corporate systems.
- Hampering the performance of competitor systems.

4.5.3 EXAMPLE ATTACK PATHS

Endless attack paths are available to attackers within the ITS network. A wide variety of entry points are available, with an almost infinite number of paths available to reach a given goal. Access to a certain path may only be limited by the attacker's skills, resources, and time. An attacker could be physically co-located with their entry point, or gaining access remotely. They could find several different ways to navigate the network, or only one.

Examples of how an attacker may gain entry to and move around the ITS network can be seen in the sections below.

4.5.3.1 RSU COMPROMISE

Vehicles and RSUs are attractive targets for attack due to the fact they are easy to physically access, as well as communicate with wirelessly. Access to vehicle or RSU technology could be achieved maliciously through theft and intrusion, or through legitimate purchases. Once an attacker has obtained a foothold by compromising a vehicle or an RSU, it is then possible to perform attacks to cause disruption or facilitate crime and fraud.

To gain access to the RSU, an attacker may wish to move through a compromised vehicle, or to aim directly at the RSU. The exact choice of path will depend on the preferences of the attacker, their skill levels, and the security of the devices.

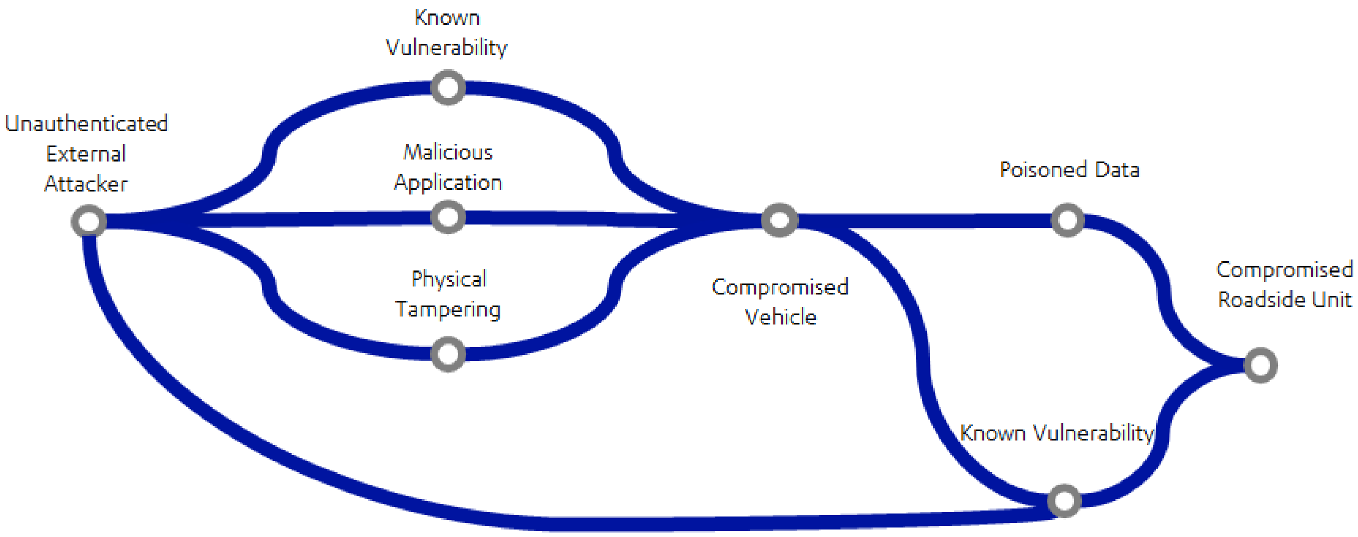

Example scenario

Bob lives in a small village, which has a section of smart roadway running through it. He is upset by the speed at which the cars travel through the village, and wants to slow them down. He has read online that the RSUs placed every 500m along the roadside control the speed at which cars drive, and so he decides to target them.

Bob spends some time learning about hardware security online. One night he takes his laptop and a crowbar to an RSU which is hidden behind some trees. The RSU is in a locked cabinet, but Bob manages to lever it open with the crowbar. The physical tamper detection system of the RSU detects the unexpected opening of the door and sounds an alarm. An alert is sent to the ITS network requesting investigation. The loud alarm startles Bob, and he returns home.

Based on his experience with the RSU, Bob decides to try and use other devices to get to the RSU. His neighbour has a new intelligent vehicle, which connects to the RSUs. Bob does some research on his neighbour’s new car, and discovers the car broadcasts a hotspot which drivers can connect their phones to. Bob reads a web forum and finds unofficial reports saying the technology used for this hotspot has an abusable authentication vulnerability.

As Bob is within the hotspot’s range, he can exploit the vulnerability to connect to the car from his bedroom. Using his knowledge of network security, he is able to pivot through the car’s internal networks until he finds the service which allows it to interface with the RSU. After some trial and error, he learns how to send status update messages to the nearby RSU. Bob uses these to send the RSU many reports that the road is covered in black ice. The RSU collates these reports, and does not correctly identify that they have all come from a single vehicle. The RSU updates its model with the presence of black ice and decides that the speed limit for that area should be lowered.

The dynamically updated road signs along the village now display half the previous speed limit. Autonomous vehicles automatically adhere to them, as do most manually driven cars. Bob is happy once again.

4.5.3.2 DISRUPTION

As with any transportation ecosystem, the ITS environment is incredibly complex. There are several vectors that could allow disruption. The diagram below shows just some of the possible disruptive end goals that an attacker may have, and how they may reach this point within the network. Overall, the motives can be summarised in the following categories:

- Obtain competitive market advantage.
- Cause reputational damage.
- Reduce national capability as part of a wider attack.

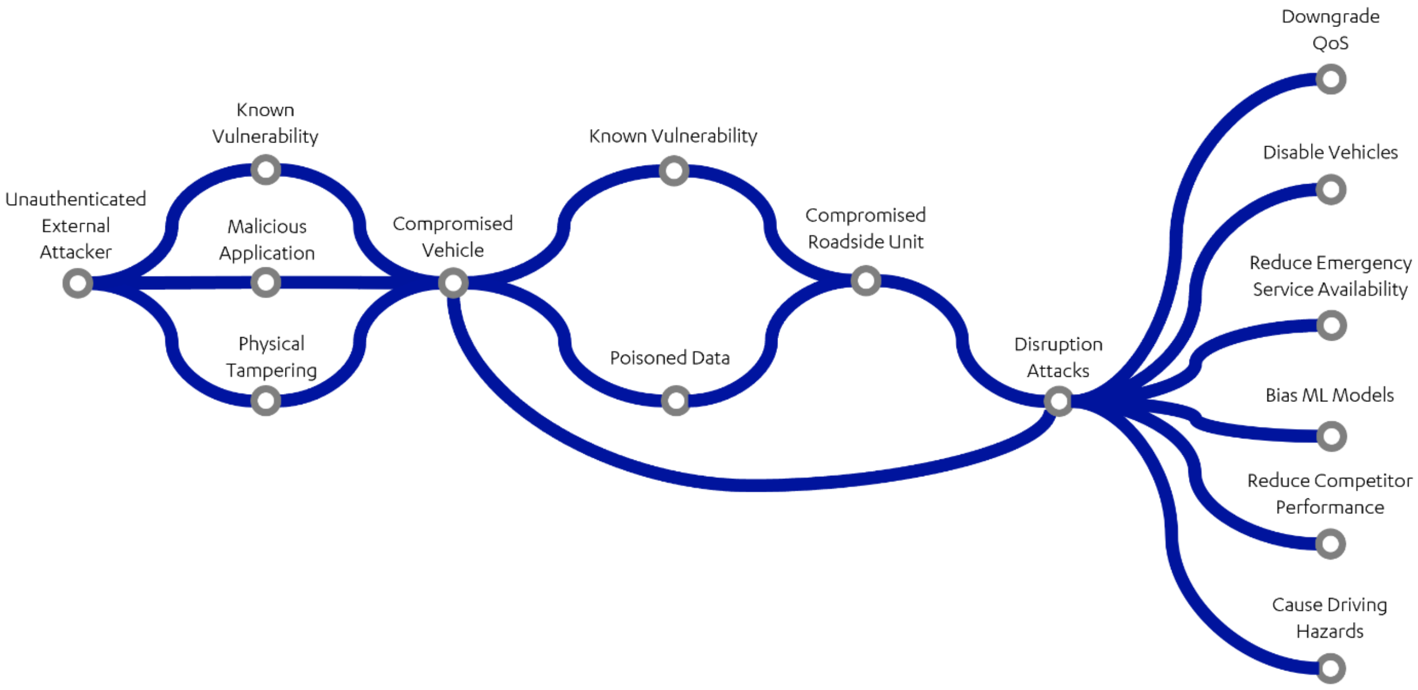

Example scenario

Company A has had extensive success elsewhere within the automotive business, and so decides to enter the autonomous vehicle market. Company A spends a significant amount of time developing their system, ensuring it meets all accreditation requirements, and launches the new vehicle. This new vehicle is more advanced than those of their competitors, and quickly gains market share. This upsets Company B, who lose money as customers switch to Company A's products when upgrading.

Company B is so upset about the success of Company A, they decide to punish them. If only they could cause Company A to suffer significant reputational harm, customers may return to Company B. Company B decides to hire a team of black hat hackers. This team is instructed to disrupt the performance of Company A's vehicles using whatever means necessary, with the one caveat that attacks cannot be linked back to Company B. They offer the hacker team large amounts of money to ensure the job is done.

The group of hackers use some of their extensive financial reserves to purchase a vehicle from Company A. The team spends several months extensively examining the vehicle, and finds a large range of vulnerabilities within it. Over the longitudinal research period some vulnerabilities are discovered independently, and fixed by Company A. The hacker team has a catalogue of attacks at their disposal, and so are unaffected by this.

The hacker team wants to gain access to as many vehicles as possible, so decides to use a malicious application to gain access to the vehicles. The hacker team subcontracts a team of black hat application developers. This team creates a new application which adds new control screen themes. Instead of having to listen to a robot voice tell them where to go, a character from the driver’s favourite cartoon provides guidance. An animated avatar is displayed which the user can interact with via the touchscreen.

The application proves very popular, and is installed in a large proportion of Company A vehicles over the following months. At this point, the black hat team begins to act. They instruct the subcontracted developer team to push malicious code to the application in an update. Application updates are not scrutinised to the same level as initial installations, and so the malicious code goes unnoticed.

The malicious code is designed to detect the brand of vehicle it is installed on. If it detects it is on a Company A vehicle, it begins to disrupt the operation of the vehicle in several ways. Firstly, it begins to send out a large amount of junk data within the vehicle’s internal network. This causes a small yet noticeable amount of lag in the vehicle. Many users notice the lag, but put it down to their vehicle aging; it does not seem worth investigating.

Next, the malicious code hooks into the starter mechanism. The vehicle no longer starts on the first press of the “on” button. The user must press it several times before the vehicle starts. Customers begin to complain about this on social media. Company A notices this and begins trying to find the source of the problem.

The hacker team decides that more kinetic disruption is required. Malicious code is pushed via the application to a small number of vehicles. This code is designed to periodically shut down the sensors of the vehicle. The vehicle had been designed with assumptions that only one or two sensors may be in a fault state at any given time; not all of them at once. This small subset of affected vehicles begin to behave erratically. One vehicle just happens to have had its sensors deactivated when a pedestrian stepped in front of it. The driver was not paying attention, as they had trusted the vehicle to take appropriate actions if needed.

The story of a Company A vehicle striking a pedestrian makes the news. Further reports emerge surrounding the recent performance issues experienced by Company A vehicle users. Despite press conferences vowing to compensate affected drivers, the reputational damage to Company A has been done. Customers begin to move away from Company A vehicles, back to Company B.

4.5.3.3 BACKEND INFRASTRUCTURE

As we move further into the infrastructure, other vectors are available to provide a potential attack path through which an attacker may compromise backend networks. Backend networks may provide access to cloud-based core systems such as the data hub. Access to this network may be more feasible from more trusted systems such as operator workstations or from a vulnerable perimeter end point hosted in cloud infrastructure.

An attacker may choose not to enter the ITS network from any vehicle or RSU node, but through a vendor. Vendors may choose to maintain a two-way connection with ITS vehicles to push updates, and receive telematics data from the vehicle. Because of this, the compromise on a vendor network could be used to gain malicious access to a vehicle, and vice versa.

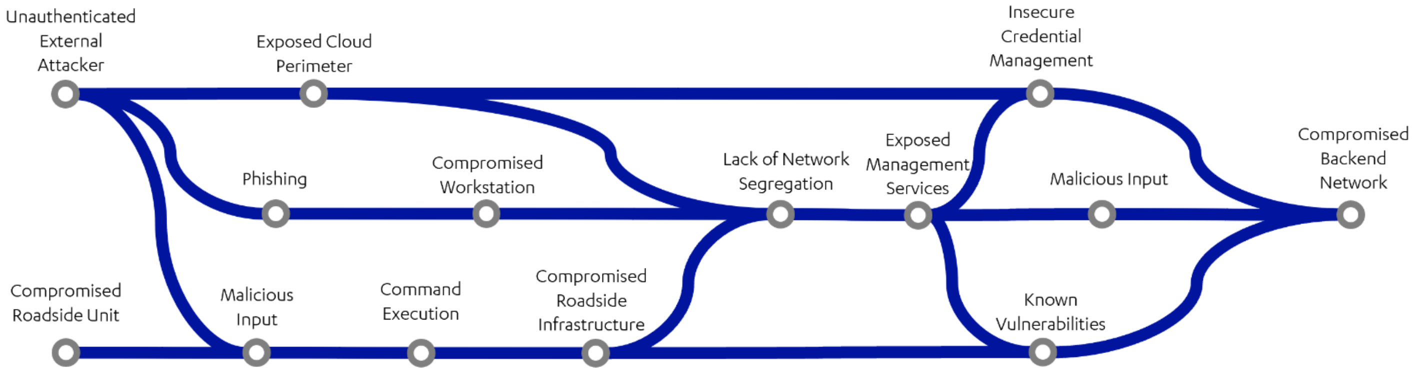
Unauthenticated External Attacker
Exposed Cloud Perimeter
Insecure Credential Management
Phishing
Compromised Workstation
Lack of Network Segregation
Exposed Management Services
Malicious Input
Compromised Backend Network
Compromised Roadside Unit
Malicious Input
Command Execution
Compromised Roadside Infrastructure
Known Vulnerabilities

Example scenario

The country of Zerhone spends a lot of time comparing itself to its neighbour, Nerund. Nerund has a highly developed ITS roadway system, with many internationally acclaimed CAV developers based in the country. Zerhone decides it would like to have more developed ITS roadways too, but does not have the resources to fund the extensive CAV development work Nerund has done. Zerhone decides corporate espionage would be a much more cost-effective measure for developing its technical capabilities.

The Zerhone government engages APT-77, a Zerhone-backed hacking group. This group is tasked with stealing as much information on CAV as possible from companies in Nerund. The APT-77 team decides that their first move should be to gain access to the internal networks of Nerund-Auto, the country's leading vehicle manufacturer. A spear phishing campaign is launched, and APT-77 soon finds itself with footholds within various support functions of Nerund-Auto.

APT-77 discovers Nerund-Auto does not make use of network segregation. The group can move from an HR employee's computer to one belonging to the tech support team. This tech support team has access to the research and development (R&D) team networks. Once APT-77 gains access to Nerund-Auto's R&D network, they are able to steal large amounts of research work, including plans for unreleased vehicles.

The Zerhone government is very happy with the information gathered by APT-77. They ask APT-77 to see if their foothold within the Nerund-Auto network could be leveraged to gain access to a given car. The group quickly discovers that through the remote upgrade functionality implemented in all Nerund-Auto products, they may push malicious code to arbitrary vehicles. The Zerhone government provides APT-77 with details of a Nerund-Auto brand vehicle owned by a key Nerund politician. The APT-77 group confirms they can push code to this vehicle. As the code comes from the trusted manufacturer, it is not evaluated; the APT group is able to cause the politician's car to take arbitrary actions.

4.5.3.4 TESTBED COMPROMISE

As vendors trial new technological developments, they are likely to draw the attention of sophisticated actors, wishing to gain a competitive advantage (be that at the company level, or at the nation state level). They may also draw the attention of smaller hacking groups who see them as easy targets, as new technology may not have robust security measures yet in place.

Attackers may make use of numerous techniques to cause disruption in testbeds to meet their goals. Some example attack routes may be seen below.

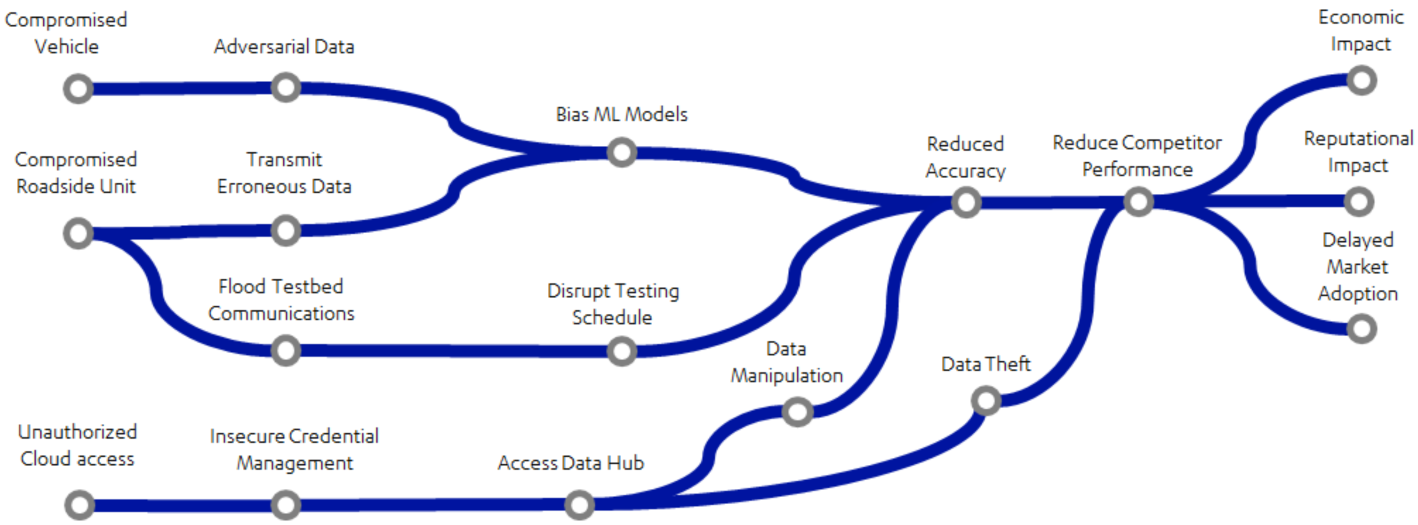

Example scenario

In this scenario, a team of grey-hat researchers decides they are feeling more black-hat than usual. They have heard about new autonomous vehicles test beds, and read that data collected during trials is used to populate simulators. The researchers are quite interested in machine learning, and are familiar with various techniques for subverting object recognition systems, such as those used by autonomous vehicles. They decide to see how far they can push the testbed.

Some of the researcher group head to the location of a testbed. They set up their laptops in a coffee shop near a road intersection, and use Wireshark[22] to capture packets. Because of the lack of encryption in 802.11p/OCB, the researchers quickly begin to identify different types of nodes within the network, and their functions. An ambulance goes through the intersection with its sirens blaring. The intersection‘s traffic lights set themselves to red allowing the ambulance to have priority. The hackers easily pick out the data packet responsible for this, and re-play it again. To their delight, the traffic lights turn red every time they do.

The next day, the researchers decide to scope out the physical assets of the testbed. They work out the areas of road included, and take note of what sensors are transmitting. They also take an interest in road signs in the testbed road signs noticing there are “slow down for children” signs here, next to a school. The

[22] Wireshark is a tool which lets you capture and analyse data packets that are being transmitted in your network

researchers have identified the autonomous vehicles currently being tested via their transmissions, and notice they slow down once the sign is within line of sight of its sensors.

The researchers decide they will investigate machine learning poisoning within testbeds. They craft some stickers and place them in identical positions on every national speed limit sign within the testbed area. They leave these stickers for a period of many months; no human really pays attention to them, they're written off as being the result of kids messing around.

After several months, the researchers return to the testbed. They then place the stickers once reserved for the national speed limit signs on to the "slow down for children" signs. Unbeknownst to the vehicle vendors, the vehicles had been updating their machine learning models to interpret those stickers as being a key identifier of a national speed limit sign. This has been perpetuated through the use of testbed data to train other vehicles. When their sensors detect the researcher's stickers on the slow down signs, they speed up to the national speed limit.

5 THREAT MITIGATIONS

5.1 ENSURING DATA INTEGRITY & VALIDITY

5.1.1 THREATS TO DATA INTEGRITY

The ITS environment is a complex set of interconnected systems which generate a large amount of data, much of which comes from external sources. Each node of the network is generating data and sending it to others. Since a lot of data is generated by systems which are outside of the receiver's control, there are risks associated with the use of that data.

The ITS-G5 specification mandates the signing of messages sent by nodes on the network. This ensures the data has not been tampered with during transit, and that the only nodes which are able to communicate with the infrastructure are those which are trusted. This mechanism may prevent low skilled outside threats from injecting malicious data, but it is still vulnerable to abuse. It may be targeted by insider threats who have the knowledge for generating messages which do not deviate from the spec, or from well-funded attackers who may have access to genuine keys used for generating valid signatures.

It may also be the case that during testing, specifications are followed less rigorously leading to some vehicle nodes generating data without signatures. It is also likely that testbed infrastructure itself generates supplemental data during testing which is to be used by manufacturers. This supplemental data may not be signed if it is deemed too complicated or unnecessary to implement the additional processing.

This means it is possible that there will be a large amount of data which is intended to be used later (by vehicle manufacturers or for providing information about the state of the road) for which the source of the data is no longer available.

5.1.1.1 ATTACKING MACHINE LEARNING MODELS

The data generated and stored by the ITS infrastructure, particularly testbed infrastructure, is likely to be used by manufacturers for training new CAV machine learning models. The data may also be used by the vehicles themselves for decision making.

An advanced attacker may wish to damage the reputation of a target manufacturer by causing their machine learning models to fail in a specified way. This can be achieved by biasing the target manufacturer's machine-learning models during the training process. In order to carry out this type of attack, the attacker would need to inject a large amount of specially crafted data into the system knowing it will be used later for training machine learning models.

It is possible that manufacturers could take advantage of malicious data of this nature by using it for adversarial training of their machine learning models. This is when the models are trained on data which is known to be malicious so that it can be recognised as such later.

5.1.2 IDENTIFYING MALICIOUS NODES

A key part of mitigating the risk from malicious data lies with identification. There are several research papers which discuss mechanisms for vehicles to determine the trustworthiness of other vehicles in the local area. For example, VEBAS, as described by Schmidt et al [25], uses a method of classifying vehicles as trustworthy, untrustworthy or neutral based on the output of behaviour analysis modules inside each vehicle. The MisDis protocol as described by Yang et al [26] recommends that vehicles maintain a secure log of all messages sent by other vehicles on the network, which can then be used to ensure the accountability of other nodes. However, this proposal seems difficult to scale and poses potential privacy issues.

Hasan et al [27] discuss several ways for vehicles to identify other malicious vehicles on the network. For our purposes however we are more concerned with ways which the infrastructure can detect malicious behaviour and identify untrustworthy data.

The main underlying mechanism for determining trustworthiness by vehicles is usually to compare the received data values to some expected value. Often this is done by performing statistical analysis on a series of data points to establish whether a new data point falls within an expected range. For determining location and speed data, the vehicles often correlate the received values to their own estimated positions and velocities in order to determine whether the values are plausible or not. The RSUs have the advantage of being fixed in place at a known position and can therefore use that information to verify that the received data is plausible. This information can be then be disseminated to the vehicles allowing them to update their internal models.

Associating trustworthiness metrics with individual vehicle nodes within the ITS network is problematic due to the use of pseudonyms (see section 3.8 for more information). These pseudonyms are designed to be difficult to associate with vehicles to protect user privacy. A balance will need to be struck in future networks between privacy and accountability for vehicles.

If any of the proposed V2V protocols for determining trustworthiness achieve widespread adoption, then they could also be used by the RSUs as part of a mechanism for determining trustworthiness of the received data. For example, if vehicles are maintaining a log of other vehicles in their vicinity which they have deemed untrustworthy then this information can be passed to an RSU to inform the RSU's data reliability assessment systems.

5.1.2.1 APPROACHES TO DETECTION

As discussed by Hasan et al [27] there are two main approaches to detecting anomalous and malicious data. These are entity-centric and data-centric.

An entity-centric approach focuses on determining whether the source of the data is trustworthy or not, i.e. determining whether the vehicle or sensor sending the data is misbehaving – intentionally or otherwise. This approach aims to establish a trust-level for the node sending the data. For a vehicle passing an RSU, this would need to happen very quickly and would likely depend on the validity of the message signature. This approach would require robust PKI with an effective revocation mechanism.

A data-centric approach focuses on verifying the validity of the data itself, usually using statistical methods to correlate the data with expected values. The expected range of values would be predicted using past data from other vehicles and sensors along with ground truths produced by the RSUs themselves.

5.1.3 MAINTAINING A CHAIN OF CUSTODY

A chain of custody for data may be necessary, where the data is received by an RSU and transported to another server for storage. This is likely to be particularly applicable to testbed infrastructure. For example, the Midlands Future Mobility (MFM) testbed will be receiving messages from vehicles under test along with data from sensors along the testbed. This data is then stored in the Chordant data hub [8]. The data must travel through untrusted, potentially malicious networks in order to be stored in the data hub so it will be necessary to ensure it arrives without being tampered. The most effective way to achieve this will be to use a VPN to tunnel from the RSU back to a trusted endpoint which is able to create an encrypted connection to the data hub. To further ensure the data has not been modified in transit, or by any intermediate node which was required to process the data, a chain of custody should be used.

A chain of custody requires that the receiving node, in this case the RSU, signs and timestamps the raw data with a trusted private key. This can be done in batches before being sent for processing or storage elsewhere. Once the data is sent on to the next node for processing, either by the data hub or an intermediate processing server, the receiving server can verify the integrity of the data by checking the signature. If the processing node makes changes to the data, these changes should be appended to the raw dataset and the whole dataset signed again. This process is repeated for any intermediate server which needs to process the data. The final receiving endpoint for the data should be able to recursively verify each signature along with the changes made by each signing node, thereby verifying the entire chain of custody back to the raw data as received by the RSU.

Maintaining a chain of custody ensures that any entity which uses the data for processing is using the original, untampered data.

5.1.3.1 HIERARCHICAL CONFIDENCE

Attackers who compromise an intermediate node may be able to generate valid signatures for the data. They would therefore be able to make malicious changes to the data and sign them with a valid key. However, it would not be possible for them to regenerate signatures from further down the chain, making it possible to detect where the malicious actor is in the chain. If the data is sent directly from the RSU to the data hub then the malicious/misbehaving RSU should be easy to identify.

It is obvious that a compromised RSU which can generate valid signatures would be difficult to detect based on signature verification alone. For this reason, further analysis by the entities which store/use the data should be done. This should consist of statistical analysis to ensure that the data falls within an expected range. The data from surrounding RSUs can be used to correlate the values. Robust tamper detection mechanisms (both physical and digital) should be placed on RSUs to provide an extra layer of data assurance.

If a certain intermediate node or RSU is found to be compromised or otherwise misbehaving, i.e. if it is found to be sending malicious/invalid data, then it should be considered untrustworthy and the data it has sent should be marked as such. This should also raise an alert in the test centre and the node should be marked for immediate investigation by an engineer.

5.1.4 APPLICATION SECURITY

Applications providing access to information stored within the data hub are accessed by multiple parties who are all utilising the testbed. Restrictions on IP whitelisting to reduce attack surface are not always going to be a realistic option for systems that are accessed by many commercial and academic institutions.

Applications and associated APIs should be tested prior to release and before each significant change is pushed to production. Ideally, the testing of applications should not be restricted to the end of the development process. A pragmatic approach to software development should be taken. This can involve the use of threat modelling to determine the potential vectors of attack for an application. The approach can be complimented with frameworks such as OWASP Application Security Verification Standard (ASVS). The outcomes from the threat modelling can be used to develop unit tests that can be implemented into a secure development lifecycle process.

5.1.5 PROTOCOL SECURITY

As mentioned in section 3.7.6, at several points within the descriptions of ITS protocols, security is mentioned as being out of scope. F-Secure recommends that this approach to security be reconsidered. Security should not be considered out of scope when the usage application has so much potential for harm. Whilst the protocols cannot provide absolute security, they should still contribute overall to the security of the system. Formal analysis of all protocols should take place, with discovered vulnerabilities being fixed prior to deployment.

6 TESTING CAV INFRASTRUCTURE AT SCALE

6.1 OVERVIEW

In the future, connected and autonomous vehicles will become commonplace. To enable this, intelligent roadway infrastructure will make up large swathes of the UK's road systems. Eventually, CAV/ITS infrastructure may be considered part of the UK's Critical National Infrastructure (CNI). Maintaining the physical and cyber security of ITS infrastructure will be of the utmost importance. In support of this, large scale testing will be required. This includes regular tests as part of standard operations for existing infrastructure, as well as thorough testing of newly developed pieces of the network.

To assist in this, the UK plans to make use of several testbeds. These testbeds will provide a real-world environment within which ITS infrastructure vendors (be that vehicle manufacturers, algorithm designers, sensor developers or others) will be able to test their contributions. Testbeds will contain a massive number of sensors, allowing vendors to gain highly granular trials data. This data will be recorded and held by the testbed operator, with access to relevant data provided to vendors.

Although testbed trials will provide an excellent way for vendors to gain some real-world trials data (and meet any test drive distance requirements which may come in to place in the future), they will also provide a way for the cyber security of technologies to be tested. Whilst extensive security testing and development should take place prior to testbed trials, the testbeds themselves should offer security testing during all trials. This is because technologies may behave differently when attacked in an isolated test environment compared to in a real-world environment.

While the testbed will help to improve the security of the ITS networks, they should not themselves become a weak link. Testbeds are likely to be targeted by a variety of actors (as described in section 4.3), with a wide variety of end goals. Testbeds will need to be able to protect not only themselves against this onslaught, but also vendors who are utilising their facilities to test CAVs and other pieces of ITS infrastructure. As trials will most commonly involve technology which has not been released yet, businesses will need to be able to trust that the testbed trials do not present an increase in risk levels.

The envisioned ITS networks are very complex, with many actors involved in the development, testing, and maintenance of them. To provide a degree of unity to the many involved actors, an overarching organisation entrusted with controlling the security of ITS networks should be created. This organisation would be involved in every aspect of the network and would contain stakeholders from a large range of involved parties.

In this section, recommendations are made in the following areas:

- Consortium membership: The proposal of an overarching organisation for managing ITS network security. This body would govern the operation of testbeds, the security standards required of the ITS network, and how they may be maintained. The way in which this consortium will act with different ITS ecosystem actors is discussed, and how the consortium may support in identifying and mitigating against emerging threats.
- Pre-testbed trials: What sort of testing vendors should aim to undertake prior to bringing their contributions to testbeds.

- Maintaining security during testbed trials: Guidance on how the security of testbeds may be maintained during trials, as well as that of the technology being tested.
- Security considerations post trial: Security considerations for vendors and testbed operators post trial.
- Considerations for testing ITS infrastructure at scale.

6.2 CONSORTIUM MEMBERSHIP

6.2.1 OVERVIEW

In support of nation-wide role out of secure test bed facilities, a governing consortium should be formed. This consortium would be both ratified by the government, and involved in the shaping of future UK law in the context of connected and autonomous vehicles. A wide variety of stakeholders should be included in this including legislative bodies, representatives of vehicle/test bed infrastructure vendors, cyber security experts, and advisors from academia. Whilst vendor representatives should be able to actively engage with the consortium, it is important that they are not placed in positions which could harm the overall impartiality of the organisation.

The consortium would exist to both drive the development of security measures, and their enforcement. They would be heavily involved in the shaping of the processes for the development of new technologies, and what security requirements they must meet before they can be brought to market. When a vendor wants to develop a new technology, they would be able to seek security advise from the consortium. The consortium would be able to help facilitate the testing of new technologies throughout their development cycles, including through testbed trials. The consortium would also periodically conduct testing on existing infrastructure, both in testbeds and in public roadways.

Ultimate decisions on the introduction of new technologies in to the ITS network should lie with the consortium. It would be up to the consortium to describe the security standards to be met, and to ensure that they are complied with. The consortium however should be nurturing in this role; instead of acting as a barrier for entry, they should be highly supportive of vendors and testbed operators. The consortium should not just exist to say no to vendors wishing to release new technologies; the consortium should exist to provide support to vendors in developing measures to meet the required security standards.

6.2.2 CYBER SECURITY TEST CENTRE

The consortium should run a cyber security test centre, which sits separate from standard testbeds. This would provide an environment for ITS security researchers employed directly by the consortium to operate in. It would also provide a base of operations for the consortium activities described within this report. As such, the cyber security test centre should ensure that it follows security best standards so it may lead by example.

6.2.3 AT THE LEGISLATIVE LEVEL

At the highest level, the consortium would create guidance on how to handle security throughout the entirety of the V2X ecosystem. Through collaboration with academics and cyber security experts, cyber security standards would be established alongside robust testing frameworks to verify products meet standards. Bespoke tooling could be created to facilitate testing. Cyber security companies would be able to seek accreditation to carry out approved testing within the V2X ecosystem.

The consortium would be heavily involved in helping the UK government to shape the future of ITS networks. It would engage with EU bodies to ensure the interoperability of all technologies. The consortium would also be able to help with the creation of public information campaigns, to help the general public adjust to the future world of prevalent CAV roadways.

It may be necessary for more specific quality assurance standards be created with the support of the consortium, to better ensure compliance of ITS technologies. Other standards such as ISO 27001 and ISO 27032 could alternatively be extended to provide explicit standards for ITS technology.

6.2.4 FOR ITS INFRASTRUCTURE

In the middle level, the consortium would test in-place infrastructure. This includes dedicated testbed infrastructure as well as ITS networks in the field (once these are set in place). The consortium would be required to accredit any new testbeds, and ensure that they are being ran in a secure manner. As such, they would be responsible for conducting regular audits. Likewise, regular testing of representative samples of CAV roadway infrastructure should take place.

As mentioned, newly created testbeds would need to be authorised by the consortium. If they could not meet their security standards, they would not be allowed to operate. The consortium would set out security standards for testbeds, and perform regular auditing to ensure compliance. If a testbed was no longer able to meet these standards, it should no longer be able to operate. ITS network component vendors would not be able to gain permission to deploy systems to the overarching network by running trials at such testbeds.

6.2.5 FOR ITS VENDORS

At the lower level, the consortium would deal directly with vendors. The consortium would be able to advise vendors on how to develop secure systems, and guidance on how security may be evaluated internally. From here they would be able to advise on how to obtain external verification testing. The consortium would also be available to provide Incident Response support to vendors should they suffer a breach. Descriptions of how the consortium may enforce security standards are described within later sections of this report.

Vendors would be encouraged to actively engage bi-directionally with the consortium. The purpose of the consortium should not only be to enforce behaviours. Vendors should be comfortable with seeking expertise from the consortium. If a difficult security issue arises within a vendor, they should be comfortable with seeking advice from the consortium on it without fear of the loss of sensitive information to competitors. Because of this, it is important the consortium be able to remain impartial when dealing with vendors.

6.2.6 FOR CYBER SECURITY VENDORS

As cyber security is involved heavily in the services provided by the consortium, it will be necessary for a large amount of cyber security consultancies across the country to be involved; it would be impractical for a single company to provide all services required. As such, the consortium would develop a set of standards for cyber security vendors to meet if they wished to be an accredited ITS cyber security vendor. This could be done through the creation of a dedicated accreditation scheme and associated exams, or through leveraging other testing standards available.

To facilitate the large-scale testing of ITS networks and vendors, some standardisation should be required. Companies seeking approval for performing accredited ITS network tests should be able to meet the testing standards set out by the consortium, with the consortium providing auditing of security consultancies regularly to ensure that quality standards do not drop. The consortium would work with cyber security and ITS network experts to assess what standards should be set in place prior to this.

Standard tooling should be developed, which is suitable for use by a variety of vendors. The route to the creation of such tooling is not currently obvious. In the future, the consortium may wish to commission the creation of specific tools. Alternatively, existing cyber security vendors may wish to develop future-looking tooling or modify existing tooling to be suitable for use with ITS infrastructure. It may also be the case that existing ITS networking vendors (such as vehicle manufacturers) wish to collaborate on the creation of testing tools. One of the consortium's duties would be to facilitate the creation of new tools by enabling knowledge sharing, and to provide impartial auditing of the results.

6.3 SUPPORT IDENTIFYING AND MITIGATING AGAINST COMPLEX THREATS

6.3.1 OVERVIEW

The cyber security landscape is constantly evolving. Over time, the threats facing product developers and system operators will change. It is important that all parties involved within the ITS network maintain a high level of awareness of emerging threats, and the forever adapting threat landscape. If one aspect of the ITS network is targeted, it is likely that others will soon face similar threats.

Actors within the ITS supply ecosystem may be wary about sharing details of threats which they face, for a variety of very understandable reasons. These include:

- Not wanting to be known for having suffered breaches, as this could lead to reputational harm.
- Concerns that competitors may be able to glean sensitive information about new developments from shared information.
- Wanting to maintain a competitive advantage by being aware of and remediating against threats which others have not yet discovered.

The supply ecosystem should aim to act, as is in the name, as an ecosystem. The term ecosystem implies that all aspects involved are supporting of each other, even if not through directly taking actions towards each other. The consortium should help to facilitate the indirect support of all parties involved within the ITS supply ecosystem. By helping each other to become more secure, the vendors reduce the risk that they in turn will face.

6.3.2 SHARING FROM THE VENDORS

To keep up with emerging threats, it is important that information is shared by vendors within the sector. Useful information could originate from a vendor in a variety of ways, including:

- Disclosure of unintended behaviours, such as bugs uncovered in third party libraries.
- Disclosure of breaches as they occur.
- Details of unsuccessful attack attempts.
- Incident response results following successful attacks.

The tendency for an organisation is to reveal as little as possible about incidents for fear of reputational damage. However, sharing information can benefit other organisations and as a result, the sector overall.

Disclosure of such information does not necessarily have to occur publicly and could instead be facilitated by the consortium. Information would be (where appropriate) thoroughly anonymised prior to sharing. A model such as that used by the government managed NCSC Cyber Security Information Sharing Partnership (CiSP)[23] could be used.

6.3.3 AWARENESS PROVIDED BY THE CONSORTIUM

The consortium should help members of the industry to maintain an awareness of emerging threats, and advise on how they may be mitigated against. This information would not only come from vendors, but also from work conducted by the consortium. Through expert analysis, the consortium should be able not only to alert members to emerging threats, but to provide explanations of how they might impact the industry and how they should best be addressed. News of these emerging threats should be used to adapt the advice and requirements set out by the consortium, if appropriate.

The consortium should be able to act as a medium for the sharing of potentially sensitive information between vendors. The consortium should be able to help vendors to share details of discovered threats with each other, whilst not compromising details of their intellectual property. If a member organisation discovers a breach, the consortium would be able to help them share information with relevant involved parties.

The consortium should be able to directly share knowledge with participants within the ITS supply chain ecosystems. As many different parties will be involved with diverse interests, there is no "one size fits all" approach to threat-intelligence sharing. It may be useful for the consortium to create a list of subject areas which may be of interest, which companies could then subscribe to. The presentation of requested information could take the form of a newsletter, or even an RSS feed.

It may also be necessary for the consortium to send out more targeted threat intelligence briefings. These could be built on information which should be acted upon as quickly as possible, such as news that a similar company has recently been targeted. These would ideally aim to help companies to address high risk vulnerabilities which have recently been discovered, whereas the newsletters or RSS feeds could provide a trickle-feed of useful information which could be actioned upon by companies as required.

Information to be shared with newsletters and targeted bulletins should come from a large variety of sources, including:

- Anonymised information from vendors.
- Research funded by the consortium.
- Research from partner organisations, such as security consultancy firms.
- Research from other parties which has been identified as relevant.
- The outputs of limited-attendance conferences ran by the consortium.
- The outputs of relevant public conferences.

[23] https://www.ncsc.gov.uk/information/cyber-security-information-sharing-partnership--cisp-

- Information gained from media such as news outlets, blogs posts, and journal articles.
- Details of publicly released vulnerabilities, including those informally described on blogs as well as those formally assigned CVE identifiers[24].
- Outputs of relevant threat hunting work.
- Monitoring of zero-day vulnerabilities[25] available on dark web markets.

It would be one of the duties of the consortium to summarise the available information, and to provide guidance on how to may be used.

6.3.4 BUG BOUNTIES

In support of generating information to be shared with vendors, the consortium could run a bug bounty program. Bug bounty programs work by announcing to the world that you will pay money to anyone who can uncover a vulnerability within your technology. Bug bounty programs will also often promote the work of the person discovering the vulnerability, once a period (often 60 to 120 days) has passed to allow time for the vulnerability to be fixed by the vendor. Bug bounty programs aim to disincentivise hackers from selling vulnerabilities on the dark web.

In the bug bounty program, security researchers would be encouraged to investigate the security of the ITS network devices themselves. Discovered vulnerabilities could be highlighted to the consortium (or to vendors themselves, should they run their own bug bounty program), in exchange for a financial reward. Through such a program, vulnerabilities may be discovered and fixed before they can be discovered by malicious actors.

6.3.5 SECURITY ADVICE

The consortium would employ experts on security within the ITS domain, or otherwise have access to them. Their expertise would be made available to vendors to help with the development of secure systems, as needed. Such people would likely work on a lot of sensitive material for a variety of vendors, and so would require extensive security vetting before being able to do so.

These security experts should be available remotely for advice (perhaps feeding into the consortium newsletters on occasion) but should also be deployable to vendor sites if needed. The experts should not only help to educate vendor staff, but also to assist with secure delivery more directly. For example, pair-programming exercises could be conducted with the expert and a vendor employee.

[24] The CVE (Common Vulnerabilities and Exposures) system provides a way of publicly listing vulnerabilities. Various databases of these are publicly available. These include vulnerabilities which have been disclosed to a vendor for fixing without exact details of the exploit being publicly available.

[25] A zero-day vulnerability refers to a vulnerability which as of yet does not have details publicly available. Due to the lack of awareness, vendors have yet to fix them. They are commonly sold online by back-hat hackers, and are often only publicly discovered through their use to compromise systems.

6.4 SECURITY CONSIDERATIONS PRIOR TO TESTBED TRIALS

6.4.1 OVERVIEW

CAV testbeds will provide a way for vendors to test their technology within a real-world environment, whilst ensuring that as much data as possible be available for post-trial analysis. Although it is likely that issues will be uncovered during the testing of new technology, it is important that contributions be tested thoroughly prior to being introduced into the test bed ecosystem. Whilst the most common piece of technology to be trialled in test beds is vehicles, other contributions may be tested. This could include new algorithms for V2X data transmission, new RSU hardware, or even a new gantry CCTV system.

6.4.2 PREPARATIONS BY THE TESTBED

Testbeds should not need to make many special preparations in advance of a new trial. As the testbed will consist largely of publicly-accessible road infrastructure, the secure operation of the infrastructure should be maintained at all times.

Prior to testing, the testbed should liaise with the vendor to understand their needs and to identify relevant consortium testing guidelines. As testbeds will likely be conducting multiple tests simultaneously, it is important that testing schedules are designed and adhered to so that aspects of testing are not overlooked. Likewise, it is important that the testbed operators ensure that their in-house testing team has the man hours available to conduct manual testing prior to agreeing to any schedule.

It has been suggested that trial participants be able to gain access to their trials data through a website application, post-trial. Prior to this website having access to any real client data, its security should be thoroughly tested. Internal testing should be conducted by the developers, as well as testing performed by an external penetration testing provider. It should be ensured that any vulnerabilities discovered to present a significant risk are fixed prior to formal launch of the application. The application should require additional security auditing and testing any time that significant changes to its code base are made, as well as annual penetration tests as a matter of practice. As the website will be publicly facing, it is likely to be frequently targeted by malicious actors.

When data is being collected and collated in the testbed data hub, it should already be being sanitised. It is likely that you would see attackers attempting Hail Mary[26] style SQL injection attacks[27], such as a drop tables command. Robust sanitisation at this level would prevent the integrity of data from being affected. However, a web portal provides an additional vector through which code injection attacks could be used. An attacker may anticipate that data collected during trials will eventually be displayed in a web portal, and so attempt to insert malicious JavaScript code into the data hub. When it is displayed to the end user, it

[26] An attack where you don’t have any feedback, or even knowledge on what is likely to work. You just throw the attack out there and see if it works.

[27] SQL injection attacks refer to the class of attacks whereby a malicious actor attempts to insert code using valid SQL syntax into the target database. If they have managed to format their command correctly, it may be executed by the database (often with harmful results)

would be executed. Because of this, a whitelist approach to input sanitisation should be used[28]. Conversely, it should not be possible for a website application user to insert special characters into any field within the website for similar reasons.

6.4.3 PREPARATIONS BY THE PARTICIPANTS

The exact format of pre-trials testing will be bespoke to the type of technology being tested. Whilst CAVs will make up the bulk of testing, other pieces of ITS hardware such as CCTV systems or handheld devices may be tested. A general set of security assurance techniques should be employed for all new contributions. This includes:

- Secure code review.
- Penetration tests of user interfaces.
- Product security reviews.
- Secure configuration review of any off the shelf software/hardware.
- Network penetration tests.

Whilst vendors should have their own internal security teams which are able to carry about the above tests, the results of such tests should be independently verifiable. To that end, external bodies (approved by the overarching consortium) should be contracted to carry out their own testing, using approved frameworks and tooling to do so. These test companies themselves should be audited annually to ensure that they are meeting testing standards.

The bespoke testing needs of different CAV network contributions should be examined by the consortium, with testing frameworks created for each. By using standardised tooling, vulnerabilities may be uncovered which had otherwise been unknown to the developers. For instance, the University of Surrey carried out formal analysis of previously described V2X revocation protocols using the Tamarin protocol verification tool[29] and were able to discover vulnerabilities which had not been found during prior manual inspection [20]. Similar testing devices for other aspects of ITS networks would be invaluable.

By performing robust testing on elements prior to their introduction into the test bed, it can be ensured that a layered approach to security is maintained. Without robust prior testing, an attacker may be provided with a route into the test bed infrastructure. An example scenario is described below.

A new car model is brought to the CAV test bed for trials, without prior testing. Unbeknownst to the test bed operators, this new car contains a simple vulnerability which allows a remote attacker to gain access to its ITS station router. An attacker exploits this, and from their privileged position are more easily able to pivot to other aspects of the infrastructure. They are then able to remotely switch off the various cameras recording cars in the test beds. The lack of important trials data means that the developer of the new car model must redo trials, as well as other entities which were testing technology at the same time. All

[28] Instead of explicitly disallowing characters which may be of use in different types of attack, a set of pre-defined characters should be allowed. This could include alphanumeric characters, and limited punctuation only. The use of whitelisting prevents attackers from using esoteric characters to bypass filters.

[29] Tamarin Prover: https://tamarin-prover.github.io/

companies involved suffer financial losses from having to redo tests, along with delays to their release schedules impacting their time to market.

Whilst it is important that strong cyber resilience is in place between different aspects of the system (making it harder for a single point of compromise to harm the rest of the infrastructure), this problem should have been identified prior to the insecure vehicle being introduced to the test bed. If the consortium had worked with the vendor to ensure that the new vehicle had been robustly tested, such an occurrence may have been less likely.

6.5 MAINTAINING SECURITY DURING TESTBED TRIALS

6.5.1 OVERVIEW OF TESTBED TRIALS AND GENERAL CONSIDERATIONS

Although a large amount of interoperation is required between the testbed and new pieces of ITS infrastructure, no element should place absolute trust in all others. Due to the sheer scale of the proposed future ITS networks, it is almost a certainty that at least one aspect will be compromised eventually. Whilst the interoperability of the devices would make it difficult to prevent knock-on effects, some isolation must be implemented to limit the potential for harm.

6.5.2 FOR THE TESTBED

It is important that the connection of potentially insecure new devices to the ITS testbed infrastructure are not able to induce malicious actions in other aspects of the testbed. Whilst it is likely that vendors will unintentionally provide insecure devices for testing, it may also be the case that a well-funded malicious actor submits a piece of technology for testing which is deliberately malicious in some way. This device could be designed to gather information about the testbed infrastructure, to target other devices being tested, or numerous other goals. It should not be possible for an attacker to leverage such devices to compromise the ITS network. Likewise, the testbed should be able to automatically detect potentially malicious activities such as this.

When a vendor brings new technology to the testbed for testing, it is likely to draw the attention of threat actors. These threat actors will have varying goals in mind, and varying levels of resources available to accomplish these goals. They may wish to disrupt testing of competitors to delay product releases. They may wish to steal sensitive data about the latest technological developments. They may simply want to see how far their skills will let them get into the network.

Whilst vendors should implement appropriate border security measures within their systems (i.e. not trusting all data that they receive), it is important that security measures go both ways. The testbed should be able to provide assurances to vendors that their testbeds will not substantially increase the risk faced by users of the facilities. The consortium should conduct yearly audits of the security of the testbed, in addition to regular penetration tests. The testbed should be held to the same security standards as the rest of the supply ecosystem. This includes the mandatory reporting of any vulnerabilities discovered within the testbed.

Because of the increased threats being faced by a testbed during a trial, the testbed operators should ensure that thorough logging of all actions takes place not only for participants, but also for the operators themselves. It may be the case that an attacker decided to target an RSU positioned within the testbed, rather than the vehicle being tested at the time. The testbed operators should be able to detect this immediately and react to it. This would take the form of both physical and digital tamper detection in key nodes such as RSUs.

A bespoke intrusion detection system would be of great use within the testbed. This would provide the standard network breach detection services seen in off the shelf tools, combined with physical tamper detection. For instance, an alert could be displayed if it is detected that the access door on of an RSU has been opened. From here human operators could take action as appropriate, from investigating the source of the alert through to removing the targeted node from the testbed network for maintenance if necessary.

As the testbed is more likely to be targeted during trials than at any other time, it may be necessary to have more security staff on hand than usual. The consortium should be able to assist with this, perhaps providing a framework for partnered security companies to send employees on short secondments to the testbed.

6.5.3 FOR THE PARTICIPANTS

When participants test new technology within the testbed, it should be assumed that the technology will be compromised in some way. Ideally, this compromise should be carried out securely by the testbed operators, with no real harm taking place. This is an idealised scenario. An attacker may wish to target vehicles currently undergoing testbed trials, as their security will assumedly still be lacking. Because of this, testing vehicles or other technologies should be completely stand-alone. They should not maintain a connection to vendor networks. This would limit the potential for damage to an organisation should their test device be compromised by a malicious actor during testbed trials.

The way in which cyber threats should be handled would require some agreement between the testbed operators and the testing vendor. A vendor may wish to continue trials, even if a breach or serious vulnerability is detected; they may wish to see what additional issues may arise in the continuation of testing. Conversely, the testbed operators may wish to halt all tests if a serious threat is discovered so that it does not have a knock-on effect on the rest of the testbed. It is likely that some middle ground will need to be struck, based upon the resilience of secure separation of the technology being tested, and the rest of the testbed.

6.6 SECURITY CONSIDERATIONS POST TRIALS

6.6.1 OVERVIEW

Once trials have been completed, both the vendor and testbed operators will be holding a lot of data which has come from potentially untrustworthy sources. The vendor will be able to gain even more data post trial completion by requesting sensor data from the testbed operators. Because of this, it is important that participants do not assume that security is “done” once the trial has been completed.

6.6.2 FOR THE TESTBED

Once testbed trials have been completed, the testbed will hold a large amount of trials data. This data will be made available for download by the vendor through a website application. The exact mechanism of this is unclear. However, it has been suggested that the vendor will be able to login to the website, and request all data relating to a given data/time range. It is unclear what permissions will be used here; if vendors will only be able to access data for certain date ranges and relating to certain nodes, or larger portions of data.

Depending on the sensitivity of data collected, a permission model may need to be inbuilt in to stored data. Whilst the onus would ultimately be on technology developers not to broadcast potentially sensitive data, even the metadata on vehicle DENM/CAM could be of interest. A potential solution may be for vendors to provide the testbed with a set of pseudonyms which they want data relating to. Or, that data may only be requested for areas during the times where their technology was present. A discussion would need to occur to ensure that the data needs of testing vendors are balanced with the security of other vendors who may simultaneously be running tests.

The data storage capabilities of testbed data hubs are undefined. It is unclear if all data ever collected during the lifespan of the testbed will always be available through the website application. Due to data storage limitations, it is unlikely this would be the case. To limit the potential for harm in the event of a website breach, the testbed operators may wish to only make relevant data available for as long as is required for vendors to download it. After this point, data should be moved to an offline storage location.

Appropriate data storage principles should be followed, including ensuring the physical security of storage medium as well as ensuring that data backups are stored. As such, the testbed operators should ensure that they are compliant with data handling regulations after the completion of trials. The consortium may be able to provide guidance on this.

Although any breaches or attacks should be detected during testing by automated systems, it is important that logs are inspected. It may be the case that during a busy period with multiple tests running simultaneously, an important alert was missed. By inspecting logs (manually or using dedicated tooling), these missed alerts may still be addressed. Depending on the severity of the post-trial discovery of an incident, the testbed may need to request incident response assistance from the consortium. If it is discovered that a breach did occur, the testbed operators should alert all potentially affected parties.

6.6.3 FOR PARTICIPANTS

After the completion of a trial, the vendor will wish to collect data from the local storage of the device being tested, as well as from the testbed infrastructure. This data should not be considered completely trustworthy. Returning technology and data received from the testbed should be considered “dirty” and handled as such. This includes the use of sheep dip[30] techniques to evaluate the safety of data prior to it being introduced to internal networks or used to populate ML models.

The exact format for trials data being downloaded from the website application has not been specified. It is likely to consist of a very large file, if not multiple files. Some of the onus should be on the vendors to ensure that they download this data securely; checking that they are visiting the correct website, and that the download is taking place securely using HTTPS. Checksum values could be provided by the website which could compared to checksums calculated for the downloaded data to ensure that the integrity of the data was not hampered during download, intentionally or otherwise.

Post trials, vendors should be invited to provide feedback to the testbed on any security issues discovered. This is important for multiple reasons: it allows the testbed operators to verify that any vulnerabilities in test pieces were not leveraged to attack other areas of the network, and it also allows vendors to gain assurance that vulnerabilities in their systems have not stemmed from vulnerabilities within the testbed. Likewise, information could be provided to the consortium to shape future pre-testbed testing best practices, or to help shape the development of future technologies across the industry.

[30] The use of a system which sits isolated from other systems to investigate the safety of data prior to it being included in a network. This could include checking for unexpected characters, performing virus scans, or the manual inspection of data.

6.7 INFRASTRUCTURE TESTING CONSIDERATIONS

Testing of the entire ITS networks at scale should draw on testbed infrastructure testing. The entire networks are, in essence, massively scaled up testbeds. The practice of using automated event detection should be scaled up, with alerts being inspected by dedicated human operators. If required, alerts could be escalated to someone suitably geographically located to take manual actions. Extensive logging should take place, with automated analysis of logs being standard practice. If required, more intensive manual examination could take place.

It is recommended that all new technology to be inserted in to the ITS networks be proved in testbeds prior to deployment; not just vehicles. Because of this, high standards of security within the networks should be maintained. That said, auditing of infrastructure should still take place. It would be impractical to test the entire network manually, and so smaller representative samples should be examined. The section being examined could be changed regularly to provide coverage, with any results which may require acting upon being disseminated by the consortium to all parties responsible for managing local regions of intelligent roadways.

To support the identification of any existing issues within the infrastructure, consortium-led wargames could be run. These could be entirely hypothetical, or could make use of representative sections of infrastructure (including testbeds). Security firms which are not normally involved with the ITS ecosystem should be brought into these exercises, to ensure that their skillsets are not being missed out on. This could potentially be linked into the bug bounty programs provided by the consortium to incentivise involvement. Likewise, ITS infrastructure could become involved in competitions such as Pwn2Own[31].

Testing of the overarching infrastructure may be expanded to include pieces of technology leveraged by the ITS infrastructure. For example, it has been proposed that council networks could be used to transport data from RSUs to data hubs. In support of this, pro bono security advice could be provided to councils. This would help to secure the overarching ecosystem, but also to make them a less attractive target to attackers in general.

In general, complete testing of the ITS networks would be impossible to conduct manually. Representative samples should be used for manual testing, alongside automated testing and monitoring systems for more wide-scale security verification.

[31] An annual contest in which hackers may try to compromise pieces of high-profile tech (using responsible disclosure practices) to win sizeable prizes.

7 CONCLUSION

F-Secure aimed to take a holistic view to the security of CAVs, and the overarching intelligent roadway networks. This approach was taken because existing literature was found to be quite isolated; it often deals with a single aspect of the solution, without considering the interoperability of different parts. These single aspect reports often consider security to be out of scope.

Numerous potential security vulnerabilities were identified by F-Secure during this project. F-Secure recommends that these should be tested for and addressed prior to any safety-critical deployments. The creation and examination of representative CAV testbeds will help to uncover these vulnerabilities.

F-Secure identified a need for overarching guidance on CAV and ITS security. A consortium should be formed with involvement from government, vendors, and the security industry. This group would aim to establish and maintain ITS security standards. Vendors would need to meet these to deploy devices within the intelligent roadway infrastructure. Likewise, security consultancies would need to be accredited by this group before they may provide approved security auditing services.

Finally, the need for thorough detection and logging capabilities within ITS networks is of very high importance. The identified threat actors likely to target intelligent roadways will have a broad spectrum of skills and resources available to them. Whilst the described security practices will help to prevent lower-grade attackers from compromising the networks, it is inevitable that a breach by a high-grade threat actor will occur at some point. Complex automated detection and logging will help to identify when such a breach has occurred, so that remedial actions may be taken as appropriate.

8 REFERENCES

[1] Ishtiaq Rouf, Rob Miller, Hossen Mustafa, Travis Taylor, Sangho Oh, Wenyuan Xu, Marco Gruteser, Wade Trapper, Ivan Seskar, “Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study,” 2014.

[2] Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li,Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song, “Robust Physical-World Attacks on Deep Learning Visual Classification,” *CVPR,* 2018.

[3] Ben Nassi, Dudi Nassi, Raz Ben-Netanel, Yisroel Mirsky, Oleg Drokin, Yuval Elovici, “Phantom of the ADAS: Phantom Attacks on Driver-Assistance Systems,” 2020.

[4] José Santaa, Fernando Pereñígueza, Antonio Moragóna, Antonio F. Skarmetaa, “Experimental Evaluation of CAM and DENM Messaging Services in Vehicular Communications,” *Transportation Research Part C Emerging Technologies,* 2014.

[5] ETSI, “ETSI EN 302 665: Intelligent Transport Systems (ITS); Communications Architecture,” 2010.

[6] ISO, “ISO 13111-1:2017: Intelligent transport systems (ITS) — The use of personal ITS station to support ITS service provision for travellers — Part 1: General information and use case definitions,” 2017.

[7] European Commission, “Certificate Policy for Deployment and Operation of European Cooperative Intelligent Transport Systems (C-ITS),” 2018.

[8] Midlands Future Mobility, “MFM Infrastructure: Cyber Security: Project Charter,” 2019.

[9] ETSI, “ETSI EN 302 637-2 V1.3.2: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service,” *European Standard,* 2014-11.

[10] ETSI, “ETSI EN 302 637-3 V1.2.2: Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service,” *European Standards,* 2014-11.

[11] ETSI, “Draft ETSI EN 302 663 V1.2.0: Intelligent Transport Systems (ITS); Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band,” 2012-11.

[12] J. Kenney, “Tutorial of 802.11p/OCB,” 20160.

[13] ETSI, “ETSI TS 102 724: Intelligent Transport Systems (ITS); Harmonized Channel Specifications for Intelligent Transport Systems operating in the 5 GHz frequency band,” 2012.

[14] Federal Communications Commission, “Dedicated Short Range Communications (DSRC) Service,” 22 April 2019. [Online]. Available: https://www.fcc.gov/wireless/bureau-divisions/mobility-division/dedicated-short-range-communications-dsrc-service. [Accessed 19 March 2020].

[15] United States Department of Transportation, “Vehicle Based Data and Availability,” 2012.

[16] ETSI, “ETSI TS 102 894-2: Intelligent Transport Systems (ITS); Users and applications requirements; Part 2: Applications and facilities layer common data dictionary,” 2014.

[17] ETSI, “ETSI TS 102 941 V1.2.1: Intelligent Transport Systems (ITS); Security; Trust and Privacy Management,” 2018-05.

[18] European Commission, “Security Policy & Governance Framework for Deployment and Operation of European Cooperative Intelligent Transport Systems (C-ITS),” 2017.

[19] RFC, “IPv6 Stateless Address Autoconfiguration,” 2007.

[20] Jorden Whitefield, Liqun Chen, Frank Kargl, Andrew Paverd, Steve Schneider, Helen Treharne, Stephan Wesemeyer, “Formal Analysis of V2X Revocation Protocols,” 2017.

[21] FireEye, “Cyber Espionage is Alive and Well: APT32 and the Threat to Global Corporations,” 14 May 2017. [Online]. Available: https://www.fireeye.com/blog/threat-research/2017/05/cyber-espionage-apt32.html. [Accessed 25 March 2020].

[22] MITRE, “APT32 Information,” 14 October 2019. [Online]. Available: https://attack.mitre.org/groups/G0050/. [Accessed 25 March 2020].

[23] B. Barth, “Reputed Vietnamese APT group hacks BMW, Hyundai: report,” 09 December 2019. [Online]. Available: https://www.scmagazine.com/home/security-news/apts-cyberespionage/reputed-vietnamese-apt-group-hacks-bmw-hyundai-report/. [Accessed 25 March 2020].

[24] Toyota, “Announcement regarding possible leakage of customer information at our Tokyo area sales outlet,” 29 March 2019. [Online]. Available: https://global.toyota/jp/newsroom/corporate/27465617.html. [Accessed 25 March 2020].

[25] Robert K. Schmidt, Tim Leinmuller, Elmar Schoch, Albert Held, Gunter Schafer, “Vehicle Behavior Analysis to Enhance Security in VANETs”.

[26] Tao Yang, Wei Xin, Liangwen Yu, Yong Yang, Jianbin Hu, Zhong Chen, “MisDis: An Efficent Misbehavior Discovering Method Based on Accountability and State Machine in VANET,” *Web Technologies and Applications: 15th Asia-Pacific Web Conference.*

[27] Monowar Hasan, Sibin Mohan, Takayuki Shimizu, Hongsheng Lu, “Securing Vehicle-to-Everything (V2X) Communication Platforms”.

[28] Huawei Technologies, “Use Cases, Requirements, and Design Considerations for 5G V2X,” 2017.