Ipswich Connected Vehicle Pilot

Lessons Learned Report

March 2022

Queensland Government

Translating and interpreting assistance

The Queensland Government is committed to providing accessible services to Queenslanders from all cultural and linguistic backgrounds. If you have difficulty understanding this publication and need a translator, please call the Translating and Interpreting Service (TIS National) on 13 14 50 and ask them to telephone the Queensland Department of Transport and Main Roads on 13 74 68.

Feedback

Please send your feedback regarding this document to: tmr.techdocs@tmr.qld.gov.au

Document control options

Contact for enquiries and proposed changes

If you have any questions regarding this document or if you have a suggestion for improvements, please contact contactCAVI@tmr.qld.gov.au

Version History

Version no.	Owner	Date	Nature of amendment
0.1 – 0.3	David Sulejic, Transport and Main Roads Jian Qin, Transport and Main Roads Justin White, Transport and Main Roads Merle Wood, Transport and Main Roads Zinah Tam, Transport and Main Roads Peter Chalmers, Transmax Shaleen Chand, Transmax Nigel Nielson, WSP David Alderson, WSP	06/07/2020	Draft creation
0.4	Miranda Blogg, Transport and Main Roads	20/10/2021	Draft final development
0.5	Geoff McDonald, Transport and Main Roads	21/10/2021	Draft final review internal
0.6	Aditya Bhalla, Kapsch Kash Munir, Cohda Ioni Lewis, Queensland University of Technology Chris Cheek, Roadtek	22/10/2021	Draft final review external
0.7	Miranda Blogg, Transport and Main Roads	15/12/2021	Final review

Departmental approvals

Date	Name	Position	Action required (Review / endorse / approve)	Due

Acknowledgements

The Ipswich Connected Vehicle Pilot was delivered by the Department of Transport and Main Roads, supported by Motor Accident Insurance Commission (MAIC), QUT's Centre for Accident Research and Road Safety – Queensland (CARRS-Q), iMOVE Australia, Telstra, Ipswich City Council and the Department of Infrastructure, Transport, Regional Development and Communications.

Contents

1 Introduction

The purpose of this document is to share learnings from the Cooperative Intelligent Transport System (C-ITS) pilot, delivered by the Cooperative and Automated Vehicle Initiative (CAVI) team. The team comprised Transport and Main Roads and the pilot vendors (Amazon Web Services (AWS), Cohda Wireless, INTEGRITY Security Services (ISS), Kapsch Trafficom, Queensland University of Technology (QUT) – Centre for Accident Research and Road Safety - Queensland (CARRS-Q), Transmax and WSP).

The report is broken into the following sections, with each section containing a sample of lessons learned – the situation, outcome and future recommendations:

- Vehicle positioning – lane or approach level positioning.
- Traffic Light Use-Cases – signal phasing and timing and map messages.
- Cellular Use-Cases – message broker, and back of queue, roadworks, road hazards and speed messages.
- Security Credential Management System – implementation by the roadside, vehicle and central stations.
- Equipment – roadside and field design, installation and commissioning.
- Central Facility – data repository, message broker, monitoring and maintenance.
- Security Credential Management System – implementation.
- Testing – all phases of testing.
- Data – data collection, analysis and surveys for the safety and system evaluation.
- Management – participant, change, asset, and project management.

Many of the issues require a detailed understanding of the system, standards and tools – a basic overview of the pilot has been outlined in subsequent section.

1.1 Pilot background

In 2016, Transport and Main Roads developed a C-ITS business case for South East Queensland examining the benefits and costs of 10 safety use cases that could be deployed by manufacturers in commercially available vehicles in the short and mid-term. Based on international pilots in Europe, USA, Japan, and South Korea, the cumulative crash reduction (fatal and serious injuries) of the 10 use cases was approximately 20%. Assuming a moderate penetration of C-ITS (in 70% of new vehicles by 2031) over a 30-year period from 2021, there is a potential to save \$3.40 for every \$1.00 spent.

Figure 1.1 – Crash reduction savings

To validate these benefits in Queensland, and Australia more generally, Transport and Main Roads embarked on a large-scale pilot of C-ITS – more commonly referred to as the Ipswich Connected Vehicle Pilot (ICVP). The pilot objectives were as follows:

- **validate** the safety impacts and user perceptions of the C-ITS use cases
- **demonstrate** technologies publicly and build public awareness and uptake
- **grow** the department's technical and organisational readiness for C-ITS technologies, and
- **encourage** partnerships and build capability in private and public sectors.

On the basis of safety merit and local issues such as roadworker deaths, five vehicle-to-infrastructure (V2I) use cases (plus the provision of speed limits) were selected for the pilot:

- **Advanced Red-light warning (ARLW)** - This warning alerts drivers there is a risk of driving through a red-light ahead or are in the conflict zone on a red-light.
- **Turning warning Vulnerable Road Users (TWVR)** - This warning alerts drivers to pedestrians or bicycle riders potentially crossing at the traffic lights.
- **Road Hazard Warning (RHW)** - This warning alerts drivers that there is a risk they are travelling at an unsafe speed for a hazard up ahead, such as water on the road, road closures or a crash.
- **Back-Of-Queue (BoQ)** - This warning alerts drivers there is a risk they are travelling at an unsafe speed for upcoming traffic queue.
- **Road Works Warning (RWW)** - This warning alerts drivers there is a risk they are travelling at an unsafe speed for upcoming road works, giving them time to slow down or change lanes. It also alerts drivers if they exceed the speed limit within the road works.
- **In-Vehicle Speed (IVS)** - This display provides drivers with information about the current speed limit (static, variable, roadworks and school zones).

Vehicle-to-vehicle use cases were not included in the pilot, however, two use cases – slow / slopped vehicle and emergency electronic brake light were analysed using a simulator and controlled tests with Lexus Australia and Transport and Main Roads.

A summary of the use cases is provided in Attachment A.

The business case assumptions are now five years old. More recent publications suggest slower C-ITS growth – such as Austroads' "Future Vehicles Forecasts Update 2031" (AP-R623-20), where standardised C-ITS in new vehicles is anticipated to be less than 50% by 2031. Positive cues are emerging from the Australasian New Car Assessment Program (ANCAP), who aligned with Europe and as such, include C-ITS in the 2025 five-star vehicle safety rating.

1.2 Pilot overview

The ICVP was deployed from September 2020 for 12 months, and included the following:

- 355 participants - their vehicles were retrofitted with a Vehicle Intelligent Transport System Station (V-ITS-S) and a display tablet referred to as a human machine interface (HMI).
- 29 Roadside Intelligent Transport System Station (R-ITS-S) installed at traffic lights.

- A cloud-based Central Intelligent Transport System Station (C-ITS-S) covering 300 square kilometres of Ipswich.
- A Security Credential Management System (SCMS).

Once installed, the participants used the equipment for nine months. Data collected throughout the pilot duration was used for a safety and a system evaluation. Self-report data to assess user experience and acceptance of the system and the use cases via questionnaires, individual interviews, and focus groups were also conducted. At the time of writing this report, all data collection, both objective driving data and self-report data, was complete however, the analysis and, thus, final evaluation was incomplete.

The pilot implemented a randomised controlled methodology with a participant control and treatment group. The treatment group is used to measure changes in the behaviour of a participant between a baseline period – where no C-ITS warnings are presented to the driver – and a treatment period – where C-ITS warnings are presented. The baseline and treatment periods were randomly allocated (50% start with baseline, 50% start with treatment). A control group who never received any C-ITS warnings was also included.

Data was used to explore the following research questions:

- Does the system improve or degrade safety for road users?
- Is the system acceptable to users and what is their willingness to use it?
- How accurate, timely, reliable and secure are the messages?

1.3 Pilot system

Per the Federal Chamber of Automotive Industries (FCAI) recommendation, European C-ITS standard were adopted. European Telecommunications Standards Institute (ETSI) message types and protocols used within the pilot include the following:

- **Continuous Awareness Message (CAM)** – location, speed and heading are broadcast by the vehicle every 100 ms.
- **Decentralised Environmental Notification Message (DENM)** – events such as hazard, queues or roadworks are broadcast by vehicle, roadside, or central station.
- **In Vehicle Information (IVI)** – a road network model that includes regulatory information such as posted speed limits are pushed by the central station.
- **Signal Phase and Timing (SPATEM)** – traffic light information polled by the field processor every 500ms and broadcast by the roadside station every 100 ms.
- **MAPEM** – the centreline of each lane on the approach to an intersection and traffic light signal group for that lane is transmitted by the central station to the roadside station, and broadcast by the roadside station every 500 ms.

Safety advice is provided to the driver as per the following value chain:

- **Observe the situation** – messages are generated, transmitted, and received by the V-ITS-S, R-ITS-S and C-ITS-S.
- **Assess the situation** – V-ITS-S use its CAM to assess the relevance of the received message.

- **Deliver the advice** – if relevant, the V-ITS-S uses preconfigured visual and/or audible warnings delivered to the driver via the display, also known as the HM.
- **Driver reads and reacts** – the driver receives the advice and takes evasive or alternative action. There is no automation, and the driver was always in control of the vehicle.

As illustrated in Figure 1.3(a), the ICVP uses a hybrid model with short-range ITS-G5 (5.9 GHz) and long-range (3G / 4G) communications. Short-range communications between the roadside and vehicle stations is used for low latency safety-critical use-cases such as ARLW and TWVR. Long range communication between the central station and field stations extends the reach of the system beyond the roadside stations, which is a viable for RHW, RWW, IVS, and BoQ warnings.

Figure 1.3(a) – Hybrid communications model

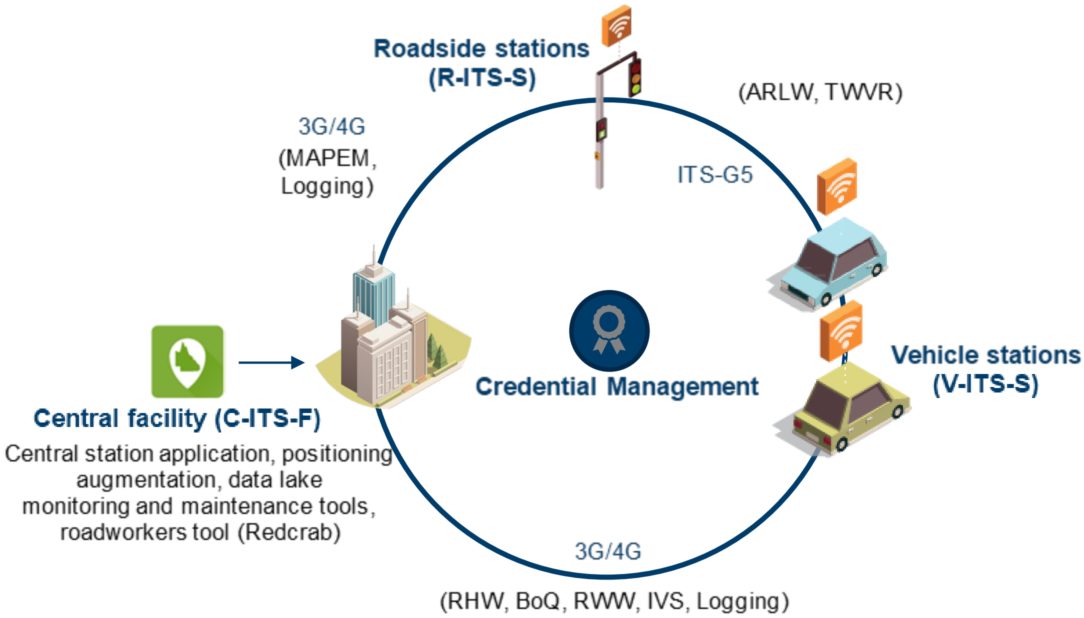

The central station message broker uses an Internet of Things (IoT) protocol – Message Queuing Telemetry Transport (MQTT) (ISO/IEC 20922:2016). Vehicles subscribe to topics – these topics represent tiles across the Ipswich pilot area and include the IVS, RWW, RHW and BoQ messages and a positioning augmentation feed from AUSCORS. Roadside stations also subscribe to topics that store MAPEM. The tile size has been optimised, as illustrated in Figure 1.3(b). The Australia tile and state tiles have been named according to the ISO 3166-2 standard (IE, AU-QLD for Queensland).

The Security Credential Management System (SCMS) issues certificate bundles to the different station types over 3G / 4G, which are used by the stations to sign transmitted and verify received messages. If the station has no valid certificates, it does not transmit messages. Due to the size of the statis speed data, the central station does not sign individual speed messages – rather it packages these up in each title.

Data generated or observed by the field stations is transmitted to the central facility message broker where it is saved in the Amazon Web Service (AWS) simple storage service (S3) data lake repository. The data is then used to generate dashboards and exception flags within the Tableau monitoring and reporting service (MRS). The central facility also supports remote configuration and access to the vehicle and roadside stations.

To enable RWW, an application was created based upon an existing product called Redcrab. The application allows the placement of different types of signs on a map so that it can be transmitted to the vehicles within the pilot. The agreed traffic guidance scheme or road works permit is entered into the Redcrab system prior to the works, and a mobile application (tablet interface) is used by the Contractor onsite / in the field to confirm or move the signage and activate and deactivate the roadworks.

Figure 1.3(b) – Ipswich tile layout for the message broker

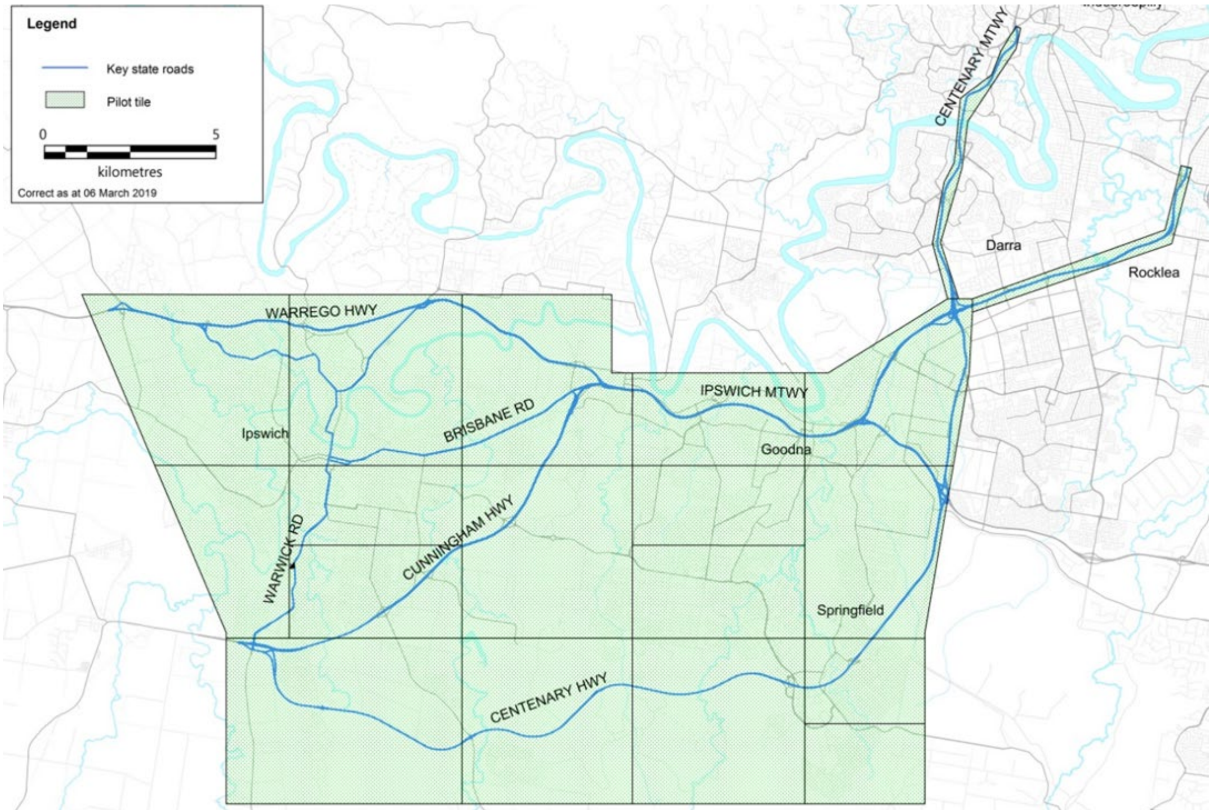

2 Vehicle positioning

Situation: No or false traffic light warnings caused by incorrect vehicle positioning

To support the traffic light use-cases, the vehicle station was required to meet the ETSI standards for lane level accuracy - 95% of the data should be within 1 metre of the centre of the lane. In an effort to meet this requirement, the vehicle station was enhanced with a ublox F9P dual frequency chipset plus cellular AUSCORS augmentation fed from a pilot deployed Continuously Operating Reference Stations (CORS).

Outcome:

A detailed analysis of single lane segments in the north-south and east-west direction indicates 99% and 98%, respectively, are within 1 metre of the centre of the lane, respectively.

Despite this, a driver can experience false or no alarms at a traffic light. For example, at a stop bar a positioning error of 1 metre could place the vehicle station in the conflict zone on a red-light, which would trigger a red-light warning.

Poor vehicle positioning can lead to the selection of an adjacent lane traffic light information – for example, a driver that travels through an intersection on a green light might receive a red-light warning because the vehicle station thinks the vehicle is in an adjacent lane with a red-light.

This condition may also occur when the positioning accuracy has been achieved, however, the MAPEM centrelines do not reflect vehicles' typical travel path, especially in wide and/or on curved lanes.

In the pilot there was approximately 1500 red-light warnings (from 350,000 trips through a traffic lights). Of these, 500 were stop-bar false alarms and 400 were adjacent-lane false alarms.

Recommendation:

Modify the vehicle station algorithms to allow for curves where swept path varies from centreline.

Modify the vehicle station algorithms to allow for error in positioning at the stop bars. Modify the implementation to include the vehicle's turning indicators.

Degrade the use-case such that it only provides warnings when the adjacent lanes have the same signal group.

Situation: No or false warnings caused by slow positioning location lock

Typically, the accuracy of the vehicle station's ublox chipset is based on GNSS (+/-10 metre accuracy) until the AUSCORS feed is available and the ublox can complete the correction calculation. In open sky this can take up to a few minutes to get a positioning fix. If a participant reaches an event before the correction is completed, they may receive no or false warnings.

Outcome:

The vehicle station was hard-wired with vehicle battery access to ensure the ublox chip was always powered and therefore retained satellite, timing and other GNSS information for start-up.

Recommendation:

Ensure that the vehicle solution retains critical positioning information for start-up.

Situation: No warnings due to poor implementation of the AUSCORS feed

Within the Ipswich area there are six CORS including the local Ipswich CORS. The CORS augmentation feed available through AUSCORS is updated every 2 seconds while the location is updated every 15 seconds. The pilot sought to implement two requirements:

- CORS topic subscription – when approaching a new tile, the vehicle station would subscribe to the relevant topic to receive the AUSCORS feed for the closest CORS.
- CORS outages – in the event of an outage, the vehicle station would reference up to two other failover topics with the nearest alternative CORS.

In implementation, the vehicle's position was jumping out of the pilot area for up to several minutes. As the vehicle moved between tiles or alternatively from the primary to failover topic, it was continuing to use the old CORS location.

Outcome:

To avoid switching, the central station topic was fixed to a static CORS and in the event of an outage, all primary topics were populated with the same backup CORS data.

Recommendation:

Enhance the central station so that the 15 second CORS location is repeated every 2 seconds. Adopt precise point positioning and retire CORS.

Situation: Density of CORS for the pilot positioning was thought to be adequate

Around Ipswich there are six CORS within 40 km of most of the pilot area that would be acceptable to use. At the time, there were concerns that 40 km would not provide a reasonable positioning solution.

Outcome:

Based on testing, the team determined that stations up to 40 km away were able to meet the 95% within 1 metre positioning requirement.

Recommendation:

CORS up to 40 km from the target area are acceptable for the needs of a C-ITS.

3 Traffic Light Use-Cases

Situation: No or late warnings caused by excessive end-to-end processing time

The pilot requirements for the traffic light use-cases stated a 95th-percentile end-to-end processing time of 300 milliseconds (ms) from the traffic light controller to the vehicle display.

Outcome:

As shown in the table below, the average and 95th-percentile message processing time for the individual components has been estimated for the traffic light use cases. Most of the processing time is related to the traffic light signal controller – these are older controllers that can only be polled every 500 ms and at intersections with higher number of detectors (more detectors, more competing data) the traffic light colour may not always be available. The processing time is also impacted by the vehicle station – the 95th-percentile of 249 ms is three-quarters of the target end-to-end processing time.

Result (ms)	Signal controller to roadside station	Roadside station (receive to transmit)	ITS G5	Vehicle station (transmit to present)
Average	300	25	5	162
95th Percentile	400	65	10	249

Recommendation:

Newer traffic light controllers are needed that can be poll the traffic light colour more regularly.

Vehicle stations should optimise their internal processing to reduce the internal message processing time.

A Real Time Operating System (RTOS) could be used to achieve deterministic message processing times.

3.1 Traffic light map

Situation: Human error developing MAPEM

The MAPEM creation included the use of the USA DOT ISD Map data tool. This process was human intensive, which can lead to errors and false red-light warnings.

Outcome:

Despite quality control processes, a large number of errors were found. These errors were identified over the test and pilot phases.

Recommendation:

Reduce manual steps, create validation tools, create an ETSI compliant tool rather than use the USA tool.

In the future, Transport and Main Roads NextGen signal controller has the potential to automatically generate part of MAPEM.

Situation: False warnings caused by V-ITS-S worse-case lane assumption at taper connections

According to the standards, MAPEM includes a taper from the adjacent through to turn-lane links. At the connection point where the vehicle may either continue to travel in the through lane or turn into the turn lane, the vehicle vendor assumed the worse-case signal group (that is, the lane with a red-light). This issue presented as a false red-light warning.

Outcome:

To ensure timely issue resolution, Transport and Main Roads removed the tapers, resulting in a workable but non-compliant MAPEM.

Recommendation:

MAPEM taper connections should be included PlugTests to ensure that the standards are unambiguous.

Situation: False warnings caused by V-ITS-S definition of the conflict zone

The vehicle station uses MAPEM to define the traffic light conflict zone. When the vehicle is in the conflict zone during a red-light, a red-light warning is triggered. Two implementation issues were observed:

- The vehicle vendor initially implemented a single stop bar location on the intersection approach, which could not support offset stop bars. This issue presented as early or late red-light warnings.
- At a clustered traffic lights (two closely spaced intersections), the vendor assumed a single conflict zone which included the queuing area between the conflict zones. This issue presented as a false red-light warning for vehicles queued between the two intersection.

Outcome:

To ensure timely issue resolution, Transport and Main Roads decommissioned the clustered traffic lights within the pilot area.

Recommendation:

Offset stop-bars and clustered traffic lights should be included Plug Tests to ensure that the standards are unambiguous.

Situation: False warnings caused by V-ITS-S use of adjacent lane signal data on unsignalised slip lanes

According to the standards, MAPEM does not require the inclusion of unsignalised slip lanes, however, the vehicle vendor assumed the signal information from the adjacent lanes. The issue presented as a false red-light warning on the unsignalised slip lane during the testing phase.

Outcome:

To ensure timely issue resolution, Transport and Main Roads added slip lanes to MAPEM and assigned a signal group of 0 (unknown), resulting in a workable but non-compliant MAPEM.

Recommendation:

The use of signal group 0 is in the standards but not facilitated in the new C-Roads guideline. This means that Transport and Main Roads will likely not use this solution in the future, and it will be up to the vehicle station vendor to manage unsignalised slip lanes.

Situation: False warnings caused by V-ITS-S movement assumption in a shared through-turn lane

On shared lanes with different traffic light control, MAPEM has multiple signal groups - for example, a red-light for the through movement and a green arrow for the turn movement. In the pilot, the vehicle vendor assumes the vehicle is travelling through the intersection – and hence, the driver will receive a red-light warning when turning on a green arrow if the shared through movement has a red light.

Outcome:

There were five shared lanes within the pilot. At one location, the turn was the predominate movement resulting in five times the number of red-light warnings compared to other pilot intersections.

To ensure timely issue resolution, Transport and Main Road's removed the MAPEM lane through-movement signal-group resulting in a workable but non-compliant MAPEM. Alternatively, the team could have descoped the lane by assigning a signal group of 0 (unknown), effectively removing the use-case service from the subject lane.

Recommendation:

A more mature fitment should include the vehicle's turning indicators so that the relevant signal group can be confirmed. Alternatively, warnings could be withheld when the signal groups are different in the same lane.

3.2 Signal phasing and timing

Situation: Late or false warnings cause by insufficient traffic light polling rate for SPATEM generation

According to the standards, SPATEM should be updated and broadcast every 100 milliseconds (ms). Due to the limited controller baud rate and the potential for other software issues, the traffic light colour is only polled by the field processor every 500 ms. At intersections with higher number of detectors (more detectors, more competing data) the traffic light colour may not be available, increasing the potential for late SPATEM updates and false warnings.

Outcome:

Testing confirmed the Eclipse controllers support a higher baud rate. At locations within excess of 14 detectors, Eclipse controllers were installed to improve the probability of accessing the signal colour. While this controller type also presents the opportunity to decrease the polling rate (from 500 to 100 ms), two software builds would have been needed – one for the Eclipse and another to support the other controller types, and as such, was not implemented.

Because the polling rate could not be fully resolved, from time to time a participant could see a late pedestrian warning immediately after the pedestrian crossing time had ended, and similarly, a red light warning immediately after the traffic light had changed from red to green.

Recommendation:

In the future, Transport and Main Roads NextGen signal controllers may accommodate the required polling rate for SPATEM updates.

Situation: Late or false warnings caused by updating the message time in an older SPATEM

The roadside station must receive a SPATEM that is less than 300 milliseconds old or it will discard it. As the Field Processor only polls the traffic light colour every 500 ms, it sends SPATEM to the roadside station every 100 milliseconds, timestamped with the time that is was sent rather than when the signal colour was polled.

Outcome:

The workaround means that SPATEM may have been old and the vehicle station would not have known to discard the message. From time to time a participant would see a late pedestrian warning immediately after the pedestrian crossing time ended, and similarly, a red light warning immediately after the traffic light changed from red to green.

Recommendation:

In the future, Transport and Main Road's NextGen signal controllers may accommodate the required refresh rate for SPATEM updates.

Situation: Interrupted red-light warning caused by incorrect use of minimum red time

According to the standards, the minimum red time is a mandatory field in SPATEM. Transport and Main Roads initially defined the value as the all-red time – typically 2 second – however, the vehicle vendor used this value to forecast the start time of the next phase. This interpretation presented as an interrupted red-light warning (starts the warning – then predicts green in a few seconds, so end the warning early, and restarts the warning when the prediction is incorrect).

Outcome:

Both Transport and Main Roads and the vehicle vendor have made changes. The minimum red time in SPATEM is now populated with over an hour (or 36001).

Recommendation:

The use of over an hour for the minimum red time in SPATEM is a workaround for the pilot. The time should be updated per the C-Roads profile.

Situation: False warnings caused by field processor failure to update SPATEM

During comms failures between the field processor and the signal controller, SPATEM generated by the field processor did not populate "unavailable", instead sending the same / old SPATEM for up to 10 seconds. In a number of instances, this presented as false traffic light warnings.

Outcome:

The field processor vendor included "unavailable" in SPATEM during comms failures.

Recommendation:

None - this issue has now been resolved.

Situation: At times, unnecessary pedestrian warnings were provided

According to the standards, SPATEM should only include pedestrian information when a pedestrian is present. In the pilot, pedestrian buttons were used as a surrogate for the pedestrian's presence. At times, these buttons can get stuck – calling the pedestrian phase in each cycle. Of greater impact was COVID, where the pedestrian buttons were forced on during day-light hours so that pedestrians did not need to touch the buttons.

In addition to the above, a driver could receive a warning if they were in shared lane and travelling below 40 km/h irrespective of their intent to turn left or right across a pedestrian.

Outcome:

The vulnerable road users use-case was the considered the least useful of the six-pilot use-cases.

Recommendation:

Confirmation of a pedestrian - advanced detection such as a "puffin" can be used to confirm the presence of a pedestrian. Currently, only a few of these types of treatments are deployed around Queensland.

Turn indicators – Turn indicators should be used to confirm the intended movement of the vehicle, noting that for the pilot this data was not available.

Situation: The V-ITS-S was not able to fully implement the no red-light warning when entering on yellow

The pilot requirements state that the vehicle station should not provide a red-light warning if the vehicle enters the traffic light conflict zone on yellow.

Outcome:

Around 10% of the conflict zone red-light warnings were related to vehicles who entered during the yellow, but it changed to a red-light while the vehicle was still in the conflict zone.

Recommendation:

The vehicle station should implement the requirement - as entering on yellow is legal.

4 Cellular Use-Cases

Situation: No or late warnings caused by excessive end-to-end message processing time

The pilot requirements stated 95th-percentile end-to-end message processing time of 4 seconds.

Outcome:

As shown in the table below, the average and 95th-percentile message processing time for the individual components has been estimated for variable speed limits (VSL), back of queue (BoQ), road hazard warnings (RHW), and road works warnings (RWW). For RHW and RWW, the central station pulls the raw data from the data source every minute – hence there is no data source processing time.

The BoQ and VSL source data – STREAMS Gateway – have an average and 95th-percentile that exceed the target of 4 seconds. The central station has an average and 95th-percentile of 2.2 and 3.1 seconds, respectively. The 95th-percentile vehicle station processing time – approximately 500 ms – is almost double the traffic light use cases.

Average:

Use Case	Data source (field to gateway)	Central station (receive to transmit)	3G/4G (transmit to receive)	Vehicle station (receive to present)
BOQ	3270	2205	170	95
RHW	unavailable	2205	170	114
RWW	unavailable	2205	170	125
VSL	3336	2205	258	10

95th percentile:

Use Case	Data source (field to gateway)	Central station (receive to transmit)	3G/4G (transmit to receive)	Vehicle station (receive to present)
BOQ	5510	3116	166	506
RHW	unavailable	3116	166	515
RWW	unavailable	3116	166	527
VSL	5826	3116	282	527

Recommendation:

Implement a direct interface to field devices to eliminate back-end processing delays.

Feeds should be designed as a push on change rather than pulled every minute.

A Real Time Operating System (RTOS) could be used to achieve deterministic message processing times.

Situation: Late, confusing or no cellular warnings due to map alignment errors

The road network model is used to generate the cellular use-case message trace. Other spatial data sets such as intersection, vehicle queue and variable speed limit (VSL) sign provide data to generate cellular warnings. The vehicle station determines the relevance of the hazard when travelling on the trace – if the trace or the location of the event is incorrect, the vehicle will either not issue or provide late or confusing warnings.

Often road geometry is captured when a road is built, but not updated when changes are made (a dual carriageway becomes a divided road, a new road segment is added, there is roadworks that results in a significant but temporary shift in the centreline, or an intersection becomes a roundabout). Furthermore, roundabouts are not accurately represented on the through or approach to the roundabout. Event locations, for example for a VSL sign, are sourced from external systems that often do not have the same data accuracy requirements as C-ITS.

Outcome:

Many examples of alignments issues were found during testing / pilot periods and manually updated.

Processes and procedures were developed to verify event spatial data before it was released to the production system.

Recommendation:

Improve the existing state-controlled road network model maintained by Transport and Main Roads.

Buy commercial probe data to validate the model.

Improve the state government owned state-wide roads and tracks data.

Improve the quality of road furniture and event locations by implementing data support systems with high-level accuracy.

Situation: No messages due to tile design limit within the V-ITS-S

The final configuration of tiles (size and layout) in the central facility was the result of a number of iterations to address the tile size capacity of vehicle station to download and process data and consideration of 'natural boundaries' between urban areas within the pilot area or the scoped roads such as the highway sections.

Outcome:

Some tiles were divided into multiple tiles to reduce the number of speed messages related manageable by the vehicle station. The vehicle station further filtered the content of the tile to a 500 metre radius around the vehicle's location.

Recommendation:

Tile creation in the central facility should become an automated process, taking into account the volume of data in each tile and the relevance of data to users within the tile.

Consider international trends for tile solutions, such as quadtree.

4.1 Back of Queue on Motorways

Situation: No or late BoQ warning caused by poor automated queue definition

The back of queue warning is based on STREAMS algorithms that estimate the location of the back of queue on motorways using detector speed, occupancy and volume. Once the speed/occupancy threshold for a queue is triggered at the "tail detector", STREAMS estimates the back of queue location between the tail and upstream detector.

The central facility algorithms assume the tail detector is the event location and builds a trace upstream from the event equivalent to 30 seconds of travel at the posted speed limit. To trigger a warning, the participant must be travelling on the trace in excess of 40 km/h.

A number of participants observed that the queue warnings were too late. Likely issues included the following:

- the actual back-of-queue is likely to be upstream of the event definition at the tail detector.
- the participant is already driving slowly and does not find the warning useful. This may be common for a participant who regularly enters the middle of a queue from an on-ramp.
- a queue is perceived by the participant but not by the system – road operators typically set the detector speed thresholds between 45 – 55 km/hr, but queues can be perceived at higher speeds.
- there are several existing STREAMS defect tickets related to the back of queue recovery algorithms where the variable speed limits remain for some time after the queue has dissipated.
- the configuration of queue detection zones does not appear to have been analysed for suitability or updated to align with real-world conditions or experiences.

Outcome:

The event definition was moved from the tail detector to the next upstream detector, which increases the chance that the vehicle receives the warning prior to the actual back of the queue which is prior to the STREAMS detected back of queue.

To ensure that warnings are not provided when the vehicle was already travelling slowly, the minimum speed was also increased from 40 to 60 km/h.

The existing STREAMS defect tickets were not resolved in the pilot period nor was configuration improvements. From ICVP data analysis it appears that compliance with the variable speed limit set at 60 km/h was poor – over 50% of the vehicle speeds where 10% over the speed limit. In part, this may be related to the absence of a queue due to defects or poor configuration. The speed differential between vehicles - some at 60 km/h and some well in excess of 60 km/h - may actually create a safety issue outside of just ICVP participants.

Recommendation:

To further ensure that warnings are not provided when the vehicle is already travelling slowly – say entering the queue from an on-ramp - the minimum speed could be increased from 60 to 70 km/h.

The existing STREAMS defects should be addressed and configuration analysed and updated as soon as possible.

4.2 Roadworks

Situation: Few RWW were generated and managed by roadwork contractors

RoadTek – one of Transport and Main Roads primary contractors for road maintenance - was trained to use the Redcrab tool to generate and manage real time, short and long-term roadworks messages. During the pilot period, they only had very few roadworks where Redcrab was used.

Outcome:

To maximise events, the team would identify roadworks and input these into Redcrab without the support of the on-road contractors. The roadworks would be field checked by the team every few days to confirm they were still active, and the signs had not moved – which meant at times, changes had occurred before the team could update the messages.

The team also engaged with Fulton Hogan, who undertook Redcrab training to generate and manage their roadworks messages within the pilot area.

Recommendation:

Where practicable, a number or roadworks contractors should be on-boarded to increase the generation of RWW.

In the future, an interface specification should be developed to allow roadworks industry to digitise their roadworks.

Redcrab should be enhanced and made available to industry for their use.

Situation: No RWW for major works involving realignment of the road

The roadworks warning is built from a base map using the existing road network model centreline. Major roadworks can result in significant changes to the centreline and hence there were times in the pilot where no roadworks speeds were issued to drivers.

Outcome:

Roadworks that temporarily resulted in a change in the centreline of the road were deemed out of scope for the pilot.

Recommendation:

A methodology to manage changes in the centreline of the roads for temporary works will need to be developed.

4.3 Road Hazards

Situation: QLDTraffic feed limitations addressed by the central station

QLDTraffic feed is used by the central station to generate a RHW. The feed has a number of limitations including the following:

- No way to subscribe to geofenced or specific topics – the entire feed had to be polled and geofenced for the City of Ipswich.
- No notification of new or cancelled events – the feed was polled every minute and any new or cancelled event inferred by comparing against the previous minutes data and change inferred as an update to a previous message.
- Location of events do not always align with a road and heading is a free text field - the location and heading had to be implied from the road network model within the central station.

Outcome:

The central station was designed to manage the limitation of the input, including improving the geospatial context.

Recommendation:

The above items should be considered in any enhancement of the feed by Transport and Main Roads Engineering and Technology who manage QLDTraffic.

The standards define information quality fields. Quality of the QLDTraffic event could be populated by traffic management staff. For example, quality could be rated low where an event is called in by the public but cannot be verified by traffic management staff. Similarly, the central facility could make some assumptions about the quality of the message based on the type, spatial and temporal information. For example, where the event location is not on a road in the central facility base-map, and must be moved to generate an event trace along a road.

Situation: DENM message cause-codes did not include all QLDTraffic generated events

The pilot project includes many QLDTraffic hazards that did not have a like-for-like cause code within the C-ITS standards. The team used cause code 1 – which, per the standards, is for traffic congestion – to include hazards such as yellow flashing traffic light and smoke.

Outcome:

The pilot used many of the QLDTraffic events, resulting in a usable but non-compliant DENM.

Recommendation:

The C-ITS cause codes should be considered in any enhancement of the QLDTraffic processes or systems.

4.4 Speed

Situation: Incorrect speed selection by the V-ITS-S at connecting roadways with small angles

The road network model is made up of centreline links that include the speed limit information. The vehicle station determines the vehicles position and displays the associated link speed. At link connections with small angles the vehicle vendor conservatively assumed that the vehicle would move to the link with the lowest speed. For example, when travelling along the motorway, a vehicle at an off-ramp would display the off-ramp speed.

Outcome:

Transport and Main Roads split the off-ramp into two segments and assigned the motorway speed to the first segment of the off-ramp (the segment connected to the mainline). A similar methodology was applied to the local street network with small angles at the intersections. While this method worked for the pilot, it is effectively an incorrect speed map.

Recommendation:

The vehicle station should have its own base map to track the vehicle path and improve its predictive path estimation.

Situation: Incorrect speeds selection by the V-ITS-S on parallel roadways

At parallel and close proximity links or elevated / stacked roads where the vehicle station is unsure which link it is travelling, the vehicle station chooses the lowest speed. This issue presents as incorrect posted speed limits. This was a common issue along the motorway, where the posted speed would change between the mainline and frontage road speeds.

Outcome:

Transport and Main Roads defined the lane-widths and lane-counts within the road network model so that vehicle station could determine what link they were on in relation to the centreline location. A similar methodology was applied to the local street network. White this method worked for the pilot, it is effectively an incorrect speed map.

Recommendation:

The provision of a base map with lane-width and lane-counts is workaround for the pilot. The vehicle station should have its own base map to track the vehicle route. The relevant route speed should continue to be displayed until the vehicle station can confirm that it is on a ramp or side street for example.

Situation: Incorrect speed definition in the road network model

Speeds in the road network model are not always correct/ are out of date; unsigned speed zones or roadways are common (for example, residential 40 km/h zone or default rural road speeds of 100 km/h); some roads having different speeds in each direction – which is difficult to represent on a single centreline.

Outcome:

Many examples of speed inaccuracies were found during testing / pilot periods, and manually updated.

Recommendation:

- Use computer vision to identify speeds from Transport and Main Roads video data.
- Buy commercial data where Transport and Main Roads is contractually allowed to publish the speed data.
- Source a lane-based road model to improve speed assignment.

Situation: Incorrect school speeds due to calendar maintenance issues

Creating and maintaining the school speed zone calendar is a manual task. Individual private schools have different school holidays, which change from year to year and need to be sourced from each school. Furthermore, different schools had different COVID closures, which were difficult to track, and often with little notice.

Outcome:

Errors in the school calendars did occur from time to time.

Recommendation:

To negate the need for calendars, automate the collection of the status of electronic school speed signs. Noting – not all school speed signs are electronic.

Situation: No speed limit data caused by a poor V-ITS-S message broker connection

On occasion, a vehicle station error would terminate the download process from the message broker, with the issue presenting as a blank speed on the HMI often lasting many tiles of travel.

Outcome:

The vendor issued new software with an improved broker connection, albeit mid-way through the pilot. Post implementation the latency improved from the 50th percentile to in excess of the 90th percentiles with less than 400 ms.

Recommendation:

The issue has been rectified.

5 Equipment

5.1 Roadside

Situation: Obstructed R-ITS-S antenna range

Antenna range should equal to or exceed the approach length defined in MAPEM. Range can be impacted by obstructions such as roadside infrastructure and trees, which can impact the timeliness of the use-case warnings.

Example 1: obstructions from existing road infrastructure

Example 2: Significant signal disruption due to obstruction from large trees

Outcome / Action:

Of the 101 approaches, 55 had a MAPEM extent of 300 metres - the remaining were shorter due to the presence of upstream intersections. Using two months of vehicle station received SPATEM (May – June 2021), the maximum antenna range for each intersection approach was estimated. All but one approach had an antennae range that exceeded the MAPEM definition.

Recommendation:

A review of the pole location and / or desktop models should be used check for obstructions prior to installation.

Situation: Substandard R-ITS-S mounting height

As per the vendor's specifications, the antenna should be mounted between 6 and 8 metres in the centre of the intersection. Using existing poles, 14 sites were mounted between 5 and 6 metres, and 15 sites could not be centrally located.

Outcome / Action:

Based on analysis, there was no correlation between the height and range of the antenna. Increasing the mounting height per the specifications was not likely to address obstructions from nearby trees, buildings or mounting infrastructure.

In general, the location of the device was expensive to maintain – requiring the intersection to be shut down to service the roadside station attached to the pole over moving traffic.

Recommendation:

Nearby obstructions are more critical than the mounting height.

Mounting location should consider ease of maintenance.

Situation: R-ITS-S were not connected

During the cable testing procedure, the ethernet cable needs to be disconnected in order to perform the test. However, some cables were left unconnected on one end after the tests, leaving the roadside station unpowered / disconnected.

Outcome / Action:

Resources needed to be re-deployed to connect the cables back so the power to the devices could be restored.

Recommendation:

LEDs should be visible from the ground and checked prior to completing installation.

Photographs should be taken of the devices/ connections within the cabinet for proof of correct installation.

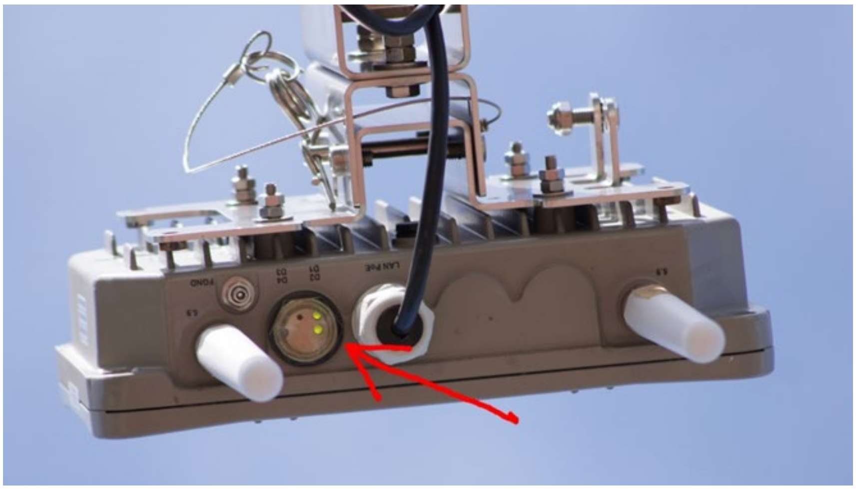

Situation: Grounding points were not utilised

During commissioning the grounding ports on the devices were not connected per the vendor's specifications.

Outcome / Action:

The installer returned to ground the roadside stations as per the installation guide.

Recommendation / Lesson learnt:

Use installation checklists.

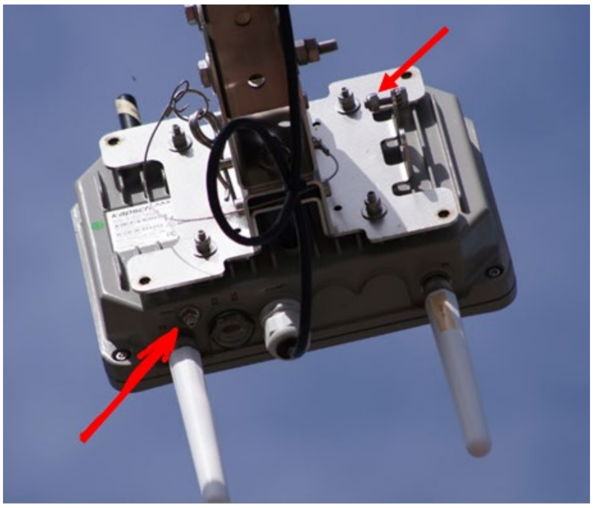

Situation: Different mounting standards between Europe and Australia

Australia has own standards for band, mounting arms and brackets. Direct comparison between EU Standard and Australian Standard is possible but time consuming.

Outcome / Action:

Long release cycles because of mapping between EU standards and Australian standards

Recommendation / Lesson learnt:

Production and Purchase in Australia for off the shelf items to prevent long release cycles.

Calculation of an extra time buffer for clarification of standard conformity. Selection of standard supplier which provides standard conformity.

Situation: V-ITS-S ignored R-ITS-S data with poor GNSS fix

Per the standards, the roadside station messages include a geo-location in the security header, which is tied to the certificates. If the message is generated at a location that is too far away from the reported geo-location (that is, the station has moved), the vehicle should ignore the message - the distance is defined by the vehicle.

In the pilot, two roadside stations reported persistent GNSS synchronisation issue and the traffic light messages were ignored by the vehicle station.

Outcome:

The vendor speculated that the harmonic of LTE 700MHz may have caused GNSS interference. The roadside station antennas were replaced with cellular compatible antennas, which rectified the issue.

Recommendation:

Cellular compatible antennas should be provided by the vendor for all roadside stations.

Situation: Dual-use export permit are required for 5.9GHz antennas

5.9GHz antennas are classified as dual use items and a general permit is needed for export. Antennas were the first dual-use items delivered to Australia.

Outcome / Action:

There was a delivery delay of 1.5 weeks.

Export permit requested => Future deliveries to Australia aren’t a problem anymore due general Export permit.

Recommendation / Lesson learnt:

Check if a general export permit is available at project start, especially when contractor/ client is outside of common trade agreements.

Situation: Mounting straps installation did not meet specifications

After the installation of the roadside stations, it was identified that some of the mounting brackets only had two rather than four mounting straps per the specifications signed off by a licenced engineer.

Outcome / Action:

RoadTek rectified the mounting brackets installation.

Recommendation:

Use installation checklists.

Situation: Poor traffic pole condition

Engineering and Technology Structures did not support mounting of any weight of device on traffic poles without structural analysis. On an existing pole, the analysis must consider the integrity of the pole such as the degree of internal corrosion.

Despite months of effort, neither CAVI nor Structures was able to confirm the integrity of the pole. CAVI attempted to procure a "pole shaker" – but they were unwilling to take the risk and sign the T&Cs. CAVI attempted to procure structural engineers – but they were unwilling to take the risk and sign-off such an analysis. Structures attempted to perform a polenoscopy – but none of the poles have access doors to permit an internal inspection – opting as a last resort to an external inspection.

Outcome:

Upon visual assessment the traffic poles were found to be in fair to good condition. Assuming a degree of pole degradation/aging (a design reduction factor of 0.95), the preliminary structural analysis for several standard pole types indicate that increases in the load are within the design capacity and risk of structural failure due to the installation of the device is low (see attached report). The additional loading imposed by the installation of the device has been estimated at 2% for Types A to C poles, and 25% for Type D poles (see CAVI typical designs). Type D (pedestrian poles) show the largest increase due to the installation of the extension pole to mount the device at the recommended height.

Recommendation

- That the regions use the result of the analysis for mounting of low weight ITS devices (< 10 kg) – noting the standard pole types and mounting arrangements captured in the analysis.
- That RoadTek replace corroded poles and / or missing grout pads when installing the ITS device.

5.2 Vehicle

Situation: HMI design framework could not be implemented consistently
Per the specifications, a framework for the triggers and escalation of the warnings (from low / white to a high / red) was developed to ensure a consistent approach across each use-cases. The central station configuration tool was then designed to support the remote management of the framework parameters for each use-case within the vehicle station. The framework was an excellent mechanism to approach the use-case design but was implemented inconsistently to manage concerns regarding the accuracy of the use case events - such as the presence of a pedestrian or validity of the roadworks.
Outcome:
The following general principles applied to all use cases: • Yellow is intended for "heads-up' warnings • Red is intended for "take-action" warnings
To manage concerns regarding the accuracy of events, the following elements of the framework were amended: • Event Zone - The red-light warning included a "beep-beep-beep" in the event zone, whereas, the pedestrian and roadworks warnings included a "boop" in the event zone. • Approach to event - The red-light warning was the only use case to implement a red warning on the approach. Even though other use cases triggered warnings on the approach, yellow warnings were used. The red-light, pedestrian and roadworks warnings included a "boop" on the approach, but other use cases had no sound on the approach.
Recommendation:
Frameworks for the HMI should consider the expected accuracy of the event.

Situation: V-ITS-S cable type and the length issues
Cable used for vehicle fitment was intended to support a range of vehicle types, but the selected cable resulted in a weaker 5.9GHz range.
Majority of the vehicles support the placement of the vehicle station underneath the driver seat; however, some EU cars require a boot installation.
Outcome:
3.8 m cable was required for install and not 10 m as originally stated in the tender documentation.
Recommendation:
Further work should have been done to optimise the design in advance of the installation.

Situation: Display Wi-Fi failure

In some circumstances the vehicle station lost wi-fi connectivity with the HMI. This issue presents as a black screen with no use-case information or warnings.

Outcome:

The vendor provided a wi-fi module firmware update, albeit in the last month of the pilot.

Recommendation:

A hardwired connection would resolve Wi-Fi or similar remote connection issues.

Situation: HMI heat deformation

During summer, the in-vehicle temperatures exceeded 90 degrees Celsius in a parked vehicle, causing some HMI screens and / or mounting brackets to deform. The HMI was intended to support in-vehicle temperatures – and the team took a lot of care selecting and further testing the product.

Outcome:

A total of 15 impacted participants' HMI / brackets were replaced. All participants were issued with sunshades, and few issues were reported post summer 2020.

Recommendation:

The HMI is not recommended for in-vehicle temperatures in Queensland.

Situation: HMI bracket glue failure

The HMI vendor changed the bracket design part way through the installation phase, and the glue failed on some devices.

Outcome:

The vendor provided an alternative adhesive and 89 participants were asked to return for reinstallation of the bracket.

Recommendation:

Vendors should be reminded that notification of product change is required in advance of product delivery.

6 Central Facility

Situation: Cloud operational cost grew as the system scaled

The introduction of cloud services to meet system requirement can cause significant increases in costs as systems scale.

Outcome:

A summary of the proportion of the cost across the different services is as follows:

- Data storage – 45%
- Extract transform and load – 40%
- Compute – 10%
- Redcrab – 2%

Some takeaways:

- The team met regularly to review the data management plan – storage and retention options – deleting redundant data sets such as raw and staging data from the cloud and moving others to lower costs storage options. Major changes in storage and retention resulted in as much as a 50% reduction in the monthly costs.
- Approximately 30% of the above costs were related to the cloud governance tools – CloudTrail and CloudWatch. As the Transport and Main Roads corporate team already pays for this service, collecting much of the same data, CloudTrail was turned off toward the end of the pilot.
- The team's implementation of extract transform and load was costly. Other tools may now be feasible, and the data lake folder structure could be restructured to improve the efficiency of the service.

Recommendation:

Engage AWS support early to optimise the design and minimise cost.

Assign a responsible person to monitor and report on cloud service costs and facilitate changes at the onset of the project.

Situation: Lack of an automated deployment process increases risk of redeployment issues

Continuous Integration – Continuous Deployment (CICD) should be implemented for all aspects of software development. CICD process automates manual steps and reduces associated human error and is commonly referred to as a pipeline. CICD was recommended and adopted for the central station but not the data lake and monitoring accounts.

Outcome:

Due to an issue outside of the team's control, the monitoring service was removed from production. It took the team 5 days to manually rebuild the monitoring service – effort which could have been avoided if a pipeline had been implemented.

Recommendation:

CICD should be implemented for all aspects of software development.

Situation: Inefficient data lake structure adversely impacts the analytics performance

The data lake storage structure can reduce the rate in which the extract transform and load process is performed, causing delays and increased processing costs.

Outcome:

Because the scale of the pilot was not well understood during the development phase, the ETL commonly failed throughout the pilot – requiring manual intervention.

Recommendation:

The storage structure can be much more efficiently organised – addressing the limit on the number of files that can be read in a second per folder by increasing the number of folders or flattening the file structure in the data lake.

Situation: The central station does not always choose the optimal trace for the cellular messages

The central station uses the road network model to generate traces or breadcrumbs into the events. Per the standards, up to seven (7) traces can exist into a single event. An automated process was developed for the trace generate; however, this process does not always choose the optimal path – for example, it chooses the seven (7) shortest rather than most likely trace paths.

Outcome:

Occasionally useful traces where not issued to the message broker, likely resulting in missed driver warnings.

Recommendation:

Implement a trace algorithm that considers road speeds, road classification and traffic volumes.

Situation: Use of Software as a Service enabled the team to focus on the transport applications

The central facility is comprised of over 40 AWS building block services.

Outcome:

The team was able to stand up a proof of concept using the AWS building blocks within 8 weeks. The use of these tools allowed the team to focus on the development of the central station applications, and similarly the use of cloud meant that CAVI did not have to focus on the storage or compute infrastructure and could scale if needed.

Furthermore, a number of AWS services including the MQTT broker, security service did not meet the requirements of the pilot. The team made a number of service feature requests, all of which were supported by AWS in time for deployment.

Recommendation:

Use SaaS so as to focus on the application rather than the building blocks.

Situation: Software as a Service limits regularly exceeded as the system scaled

Issues were identified as the system scaled to full production usage where Software as a Service limits caused system performance issues. From time-to-time the team failed to project the system growth against the limits causing issues that were largely related to the monitoring service extract of the data from the data lake.

Outcome:

The software was redesigned to work within limits.

Recommendation:

Implement the recommended monitoring to identify limit breaches early and allow fixes to be developed.

Situation: Documentation was delivered after some content providers had left

Some documents, for example, design and procedures were mostly completed after the project had been operational for some time. Staff who had developed aspects of the project and would have been ideal content providers had left the project without completing the relevant documents.

Outcome:

Staff documenting the system had to research the current system to learn enough to document the system. It is likely that some design details and procedures are still not documented.

Recommendation:

Include documentation as a requirement to pass testing phases.

7 Security Credential Management System

Situation: Late or no warnings caused by the V-ITS-S and R-ITS-S "software" certificate management latency

According to standards, hardware-based trust management is required, however, vehicle and roadside station vendors were reliant on third-party software that was not available during the testing phase. As a result, a software-based solution was implemented. This solution resulted in high latency and in some cases late or no warnings due to verification processing of messages from surrounding vehicles and/or roadside stations.

Outcome:

Due to COVID delays, Transport and Main Roads was able to test and implement hardware-based trust management, and latency issues were resolved.

Recommendation:

Use hardware-based security to ensure availability, latency and integrity of the message.

Situation: New implementation of "hardware" certificate management in the central station

According to standards, hardware-based trust management is required for a central station, however, commercially available cloud hardware security modules do not support the standards (IEEE1609.2) and are not able to conduct all security operations in the hardware.

Outcome:

AWS was engaged to provide additional protections for the software applications, with a hardware security module deployed in the cloud in a standalone account for the pilot.

Recommendation:

Continue to use hardware-based security in the central station until EU standard exists, and update as needed to ensure the standards alignment.

Situation: V-ITS-S did not verify the central station C-ITS certificates

According to standards, all C-ITS messages must be digitally signed, including in-vehicle information messages and decentralised environmental notification messages from the central station.

Outcome:

The vehicle station did not verify the messages as the product issues where too significant to resolve in the timeframes. Message security was still supported through the broker encryption and authentication process.

Recommendation:

Vehicle stations should be required to prove they verify signed C-ITS messages from any station type (including government or industry cloud) before they are enrolled with the production SCMS.

Situation: Corruption of V-ITS-S certificate management during abrupt power failures

In rare circumstances, the vehicle station was unable to complete security certificate download for signing of its messages. Signed messages are used between interacting cooperative vehicles to generate vehicle-to-vehicle use case warnings such as a hard-braking event. As vehicle-to-vehicle use cases were descoped, digitally signing cooperative awareness messages was not a critical safety issue for the pilot.

Outcome:

In the last few months of the pilot, a fix was issued by the vehicle stations' third-party security provider.

Recommendation:

The issue has been rectified.

Situation: Late or no warnings due to V-ITS-S certificate top up process

During the pilot, there was higher than acceptable vehicle latency across 1.2% of the data – this peaked on a Wednesday, which was the same day that most vehicle stations were bootstrapped and would request a weekly top-up of the certificate bundles. In rare circumstances, processing of messages was competing with the top-up request.

Outcome:

A fix was issued by the vendor, which allowed the request for certificates and the use-case algorithms to operate in parallel rather than consecutively.

Recommendation:

The issue has been rectified.

Situation: Expiration of the trust and revocation list, causing systemwide outages

The system experienced two outages – one when the certificate trust list expired before the pilot, and one where the revocation list expired after the pilot.

Outcome:

The vehicle and roadside stations did not recover gracefully and required the team to manually log into the devices to update the lists. New third party vehicle and roadside software since issued by the vendor does not require manual intervention. The security vendor has also implemented a new alert service notifying customer of upcoming critical dates.

Recommendation:

Automate alerts of critical dates to ensure good governance (now implemented by vendor).

Situation: Interpreting the standards for the pilot

There were a number of challenges implementing the standards for the pilot. The issues and outcomes are as follows:

- The standard makes no stipulation on what makes a root Certificate Authority trustable. The pilot used the well-established eCommerce PKI assurance standard "WebTrust".
- Multi-country deployment: Establishing and maintaining a system for multi-country deployment was excessive for a pilot, however, its exclusion would cause incompatibilities with the standard. The pilot established a "once-off" European Certificate Trust List to solve this.
- Station permissions: The standards do not fully define which stations should send out which message types or permissions. The pilot assigned permissions based on the use cases. This was then extrapolated and expanded to a perceived future production deployment.
- Vehicle roles: The standards provide no clarity on how to change roles (for example, when auctioning an ambulance) or how to switch roles quickly (for example, going from an unmarked/undercover police vehicle to "lights on"). While this was out of the pilot scope, the vendor recommend that a station have multiple Enrolments with the SCMS in order to achieve this.
- Revoking stations: The standards rely on "passive" revocation, preventing stations from retrieving further authorisation credentials, though not removing any existing credentials. All existing credentials must naturally expire before a station is removed from the C-ITS environment. While this was out of scope of the pilot, limiting the number of allowable off-line cached credentials would reduce the security risk to C-ITS – and further research is required.
- Stored certificates: stations may store a "pool" of credentials to use to interact with other C-ITS stations – though the standards don't stipulate these. Hence, the pilot pool was based on assumptions from international pilots.
- Communications with the certificate service: The standard assumes immediate response to all certificate requests. This forces requests to occur within peak driving hours and does not provide an even computing load across any given day. A smoothing algorithm was investigated – which is not standards' compliant and was not implemented.
- Unique identifier: The standards don't specify the nomenclature. The pilot used asset's serial number prefixed with a unique vendor code.
- Station registration: The standard does not establish a means for secure station registration. Due to the small number of stations the certificate system supplier provided a means for registering stations using email.
- Station compliance: The standards do not define the security posture or controls that the station should comply with.

A number of other interpretation issues were clarified by the vendors through their various Plug Test efforts in Europe.

Recommendation:

- Conduct further research into the privacy vs the number of vehicle certificates stored.
- Investigate the load spreading performance benefits proposed by the use of the alternate request mechanism described in v2.1.1 of the ETSI TS 102 941 standard.

8 Testing

Situation: Test practices, knowledge, and culture took time to embed, delaying the project

Test practice issues included:

- Vendors were provided but did not utilise each other's products in development – as a result, the FAT approved products could not be integrated without months of additional effort.
- Tests were developed in parallel with the development of the products, and hence vendors were not able to readily pass the test phases.
- Some tests were time consuming to validate or visualise – such as the position of the vehicle in relation to the event.
- There was a lack of automated tests for regression and scaling – while automated testing was adopted for some components, others were manual and time consuming.
- Some tests appeared to fail but were actually configuration errors, incorrectly specified test results, or defects in the automated tool that were being developed in parallel with the testing.
- It was difficult to overlap testing with the timing of live hazard events, leading to a reliance on the simulation.
- Client issues could not be readily repeated on the vendors bench, leading to longer ticket resolution.

Human resource issues included:

- A lot of specialised resources were required – traffic, data, information technology security for example.
- A lot of resources were still learning about C-ITS.
- The team had a heavy reliance on junior resources.
- Due to general program delays, tests phases were overlapped, which was complex for the team to manage.
- There was frequent context switching between tasks.
- For the documentation, evidentiary data was not collected consistently.

Outcome:

The tests practices were an iterative process and required sufficient time for development. Many bespoke tools were developed for testing – largely within the central facility – through the visualiser and monitoring tool.

A number of areas worked well:

- Backend protocols (Abstract Syntax Notation) and C-ITS sample data sets were provided early for vendor development and testing.
- Over 500 requirements were tested, ensuring a robust pilot level product – as evidenced by the participant feedback and the quality of the safety and system evaluation data.
- A multi-unit vehicle and roadside station bench - running 24 hours a day with data and dashboards - was used to determine the reliability and performance of the stations, and readily accessible to all vendors.
- Given the various locations of the vendors (Brisbane, Melbourne, Adelaide), the central facility solution was able to quickly deploy field test environments where needed to accommodate convenient testing.
- The release process has become routine and repeatable using automated testing and visualisation tools.

Recommendation:

- Specify tests in advance of development
- Automate as many tests as early as practicable
- Ensure configuration is cross checked before running the tests
- If working in different offices, ensure the configuration is identical
- Capture events that are difficult to find in the field for replay testing
- Develop a knowledge base for the system and test resources (such as Confluence).

Situation: Test bench data and compute costs were larger than the pilot proper cost

During the pilot proper, the test bench continued to be utilised for software releases. By the end of the 12-month pilot period, the 6-station test bench had exceeded the 355 vehicle / 29 roadside station pilot costs (cloud data storage and extract, transform and load).

Outcome:

As each station is designed to collect transmitted and received data, the scale of interacting stations over 24 hours on the bench was in excess of the scale of interactions in the field.

Recommendation:

- Turn the bench off when it isn't needed – as it was only needed for short periods of regression testing.
- Minimise the infrastructure on the bench – the bench.

Situation: Field test facility in a rural location had cellular and GNSS limitation

Mt Cotton test facility is located in a relatively rural area and was not ideal for use-case testing.

Outcome:

While the location was ideal for rural testing - noting parts of Ipswich are rural, the Ipswich testing was pulled forward where the cellular and GNSS coverage is better. The team did examine alternative sites, but these were generally located in unsuitable areas.

Recommendation:

Prior technical evaluation of the enabling infrastructure at a test site is recommended.

Situation: Human Machine Interface Useability and Ergonomics Testing (HUET) recommendations were limited

HUET was completed by third party experts in audio, vision, and human machine interface. Evaluations were conducted (three in total) at different times in the product development. Two of these evaluations were conducted in the pilot area, with the last one conducted two months prior to the dress rehearsal. Unfortunately, because of the ad-hoc nature of the RWW, RHW and BoQ warnings, it was difficult for the HUET driver(s) to find these events to test in the field. These warnings were part of the broader discussion; however, experts were not able to experience those three use-cases under real, expected conditions.

Outcome:

The experts recommended a simplified approach to the design that provided only essential messages that would require the driver's attention and exclude extraneous system messaging. The removal of the status and error messages meant the drivers were not aware of unit faults - only that the display was on or off – however, all health data was monitored through the back-office systems. More detailed use-case behaviours were not possible as the HUET tester only performed limited test drives. Furthermore, the desire to complete a final HUET before go-live was not possible due to the development turn-around times of four weeks.

The experts recommended several specific changes to generic elements of the product, such as the following:

- Removal of the status bar
- Removal of error messages
- Removal of "unknown" speed
- Removal of speed being compliant in roadworks zones (non-compliance only)
- Reduction in number of escalations of warnings
- Removal of brightness selection

Recommendation:

HUET is a useful exercise when considering the HMI's design and presentation of static and frequent messages, however, for infrequent messages, a desktop exercise should also be included to maximise the third-party experts' review.

Situation: Dress rehearsal limitations reduced the learnings for the pilot proper

According to FESTA guidelines, a dress rehearsal is a recommended component of a pilot project where a small sample of participants are used to represent the likely end-to-end experience in the pilot, prior to the pilot proper commencing. The dress rehearsal took place one month prior to the pilot commencement, with a target of 15 participants for a duration of 2 weeks. Deinstallations of dress rehearsal participants was occurring when recruitment for the pilot proper had already commenced.

Outcome:

So as to maximise potential public participation for the pilot, recruitment for the dress rehearsal was limited to Transport and Main Roads staff who were not directly involved in the day-to-day ICVP project tasks / activities. There was limited time available for recruitment and participants needed to meet ICVP eligibility requirements, which resulted in only a small participant pool of 6. However, the eligibility criteria was reduced from 3 hours of driving per week (ICVP criteria) to 2 hours per week – to increase the pool of eligible participants. For these reasons, the extent to which pilot processes could be tested were limited, and further testing outside of the dress rehearsal was required.

Recommendation:

A longer and larger dress rehearsal, which occurs well in advance of the pilot proper. The timing between rehearsal and pilot needs to be sufficient for the identification and implementation of learnings from the dress rehearsal into the pilot proper. For the ICVP, the dress rehearsal should have concluded at least one month prior to pilot recruitment.

Situation: A commission test track was established to ensure equipment was working after install

Because of the different installation considerations for vehicle make and model – for example the location of the station, and antennae, and power integration, all equipment was checked by the vendor on a short test track internal to the installation site.

Outcome:

The commissioning test track process ensured participant confidence in the system was high. A variety of issues were intercepted – including faulty equipment, expired certificates and cellular connection issues.

As the test track was located on the production environment, all Briggs related tests had to be geofenced from the data for analysis and dashboards.

Recommendation:

A commissioning test track should be utilised where participant buy-in is important to the outcome of the pilot.

Situation: Test data was present on the production system

Occasionally, problem vehicle stations were removed from vehicles and put on the commissioning test bench to investigate. On one occasion, the test team ran the station on an Ipswich simulation loop in the production environment. The simulation triggered a high number of red-light running events, which was presented in the daily production dashboards.

Outcome:

To remove the erroneous data from the dashboard, the monitoring team filtered out the stations under investigation. Ideally, the test team should not have used the production environment or limited the simulation to the Briggs geofence or a location outside of the pilot.

Recommendation:

Create a better process to register vehicle stations as test or live and create data fields to tag data as "test."

Situation: Frequent software releases were time consuming to test

Vendors typically took two weeks to release a new build in response to tickets, followed by two weeks with the team for testing on the bench and in the field. Regression was not uncommon – adding this time again to the schedule.

Outcome:

There were 24 vehicle releases, 15 central station releases, and 6 roadside releases. In an effort to stay within the original development schedule, the team descoped two vehicle-to-vehicle use cases (captured instead in simulator testing) and C-ITS credential management (later added back into scope for 5.9GHz messages). By the completion of the development phase, the schedule was in excess of 6 months late.

The team's goal was to lock down all releases for the pilot, however, given a number of defects observed during installation, there were three (3) additional vehicle station releases and one (1) central station release. The bulk update tool worked well – and really the only feasible way to manage the change on a large number of participant devices.

Recommendation:

Vendors to demonstrate issues is fixed and there is no regression prior to testing by the customer.

Situation: When new RHW messages where introduced, output to V-ITS-S was not tested

Despite the generation of 172 cause-code one (1) events in the central station (for example, flashing yellow traffic lights, and smoke), the use of this cause-code was not a requirement of the vehicle station, and hence never implemented. This issue wasn't identified until the last few months of the pilot. The development and testing of the new central station build was completed, however, the vehicle station was not tested as part of the change.

Outcome:

Due to the late identification, the vehicle station was not enhanced to include the cause code.

Recommendation:

Requirements and test results should be carefully linked, linking enhancements back to the requirements.

9 Data

9.1 Collection

Situation: Additional effort due to multiple data schema changes

The pilot requirements include the abstract syntax notation (ASN), which describe the data collected from each station type. There were three major ASN releases over the testing phase – with all station types updated accordingly, triggering the need for additional development and integration testing resources.

Outcome:

While the team sought to limit ASN changes, the initial requirements were a bold attempt at defining the data needs at the onset of the development phase. Throughout the testing phase, some parts of the ASN were found to be deficit. The last ASN update was delayed for as long as possible to ensure that there were no more changes before the pilot go-live.

Recommendation:

The level of effort required to adequately describe all data needs should not be underestimated.

Schema flexibility should be considered in the design to more easily allow for change. For example, the team limited the trace length to 500 metres, which did not need to be limited.

Including all parties (particularly the safety evaluation team, central facility team and monitoring and reporting team) in the ASN development early helped establish a consistent dataset and minimise the number of changes.

Situation: The R-ITS-S data comprised 90% of data in the pilot

Roadside stations operating 24 hours a day logs all the transmitted data, as well as the data received from the field processor (SPATEM), central station, and passing vehicle stations. As a result, the roadside station data comprised approximately 90% of all data collected in the pilot. While there were 355 vehicle stations, collecting the transmitted and received data also, the typical uptime of these devices was approximately 20 minutes a day.

Outcome:

The detailed roadside station data was used to generate a number of metrics including – the accuracy of the SPATEM traffic light colour against STREAMS; timeslip of the message in relation to the interacting stations; and latency of the messages. This would not be possible without detailed logging.

Recommendation:

In deployment, roadside station data could be limited to logging only on change – that is, the traffic light changes colour, rather than 10 times per second. Detailed logging history should be retained on the roadside station, and extracted when there is an issue requiring investigation.

Situation: Excessive internal V-ITS-S latency between the V-ITS-S baseline and control

For the system and safety evaluation, the latency between the vehicle station and display is measured. During the initial weeks of the pilot, excessive latency of the baseline / control group (participants who don't receive warnings) was observed skewing the vehicle station health report. This issue was not observed prior to go live, because the team had not tested the baseline data – all efforts focusing instead on the treatment condition.

Outcome:

As the baseline group does not display warnings, the vehicle station latency was using an erroneous timestamp – this defect was later corrected, with the prior month of data filtered out of the reports.

Recommendation:

Ensure that baseline condition is well tested prior to the pilot deployment.

Situation: In a naturalistic driving study, up to half the data in an event could not be used

From the original event sample, data was filtered for a range of reasons, including the following:

- TWVR – As part of COVID measures in the central business district, pedestrian buttons were permanently activated during the daytime hours over a five-month period. This measure was implemented so that pedestrians would not need to touch the buttons.
- RWW – In an effort to increase the sample size, CAVI staff would input C-ITS roadworks messages without the support of the roadwork's contractors. The roadworks would be checked by the team every few days to confirm that they were still active, and the signs had not moved – which meant at times, changes had occurred before the team could update the messages.
- RHW – At times the hazard events had cleared before the warnings were remove. Similarly, some event did not impact all lanes in the same manner, especially on freeways.
- BoQ – There were a number of use-case configuration and algorithm updates over the course of the pilot period, where the driver's experience may have been inconsistent.
- V-ITS-S health flags – several station health flags were found to be statistically significant in terms of the driver's behaviour within an event – such as the availability of the network and latency.
- Contextual data – some third-party contextual data such as traffic volumes and rainfall were not available for a given location or time within an event.

Outcome:

During the pilot, the team instigated a log to document the dates and times of events that could have an impact on the safety evaluation – such as the late termination of roadworks and back-of-queue configuration changes.

After filtering for these events, statistically significant flags, and missing contextual data, just over half of the C-ITS event data remained for the evaluation. At the time of writing this document, the team was further analysing the filtering methodology to minimise the loss of data.

Recommendation:

The degree of filtering is not uncommon for a naturalistic driving study. An alternative is a controlled study where there is more control of the environmental factors that would result in a reduction in the sample size.

Situation: The backend data sources need to be scrutinised

During the pilot, it was important to measure message latency to determine the timeliness of messages and use case warnings. The system design expected data source providers to identify each event with a unique identifier that is used to link logged message processing events as it progressed from the data source to the central station and to field stations such as the vehicle ITS station. For example, a variable speed limit sign change, the corresponding IVIM message published by the central station and its use by a vehicle station. The data sources were not adequately scrutinised to determine if source data events were uniquely identified.

Outcome:

Certain data sources did not use a unique identifier which made it difficult to match corresponding published C-ITS messages. For example, unique identifiers for traffic queue events were repeated daily. Hence, the latency between the sourced / backend data and the published messages was difficult to measure.

Recommendation:

Identify limitations in the backend data sources. Ensure that metrics can be developed with the data available. If not, invest in the development of backend data sources to ensure that these metrics can be measured against the pilot's requirements.

9.2 *Evaluation*

Situation: The control sample size was reduced to 10% given concerns with recruiting participants during a global pandemic and to maximise the public experience of the technology

The study design included both within-groups and between-groups comparison, such that there were three participant samples:

a) Treatment:

   i. TreatmentFirst: participants who experienced warnings for the first 6 months of the FOT (treatment period), then 3 months without any warnings (baseline period);

   ii. BaselineFirst: participants who did not experience any warnings for the first 3 months of the FOT (baseline period), then 6 months with warnings (treatment period), and

b) Control: participants who did not experience warnings for the entire 9 months.

The FESTA handbook mentions within-groups and between-groups separately, and did not stipulate that both were required. However, both were adopted to increase confidence in the findings. Later, due to various reasons including concerns associated with recruiting a sufficient number of participants during a global pandemic and to maximise public experience of the technology of those who were recruited, the relative size of the control group to the total sample was reduced from 20% to 10%. The duration of the baseline and treatment periods were unchanged.

Outcome:

Safety evaluation was not based on these groups. Instead, scenarios with participants who had warnings on were compared with participants who had warnings off. That is, data of the control group were grouped with the data from participants in baseline periods (even though the participant experiences were not entirely equivalent). Primarily, this was because there were periods that there were few participants who were not receiving warnings, as highlighted in the box in the simplified image below.

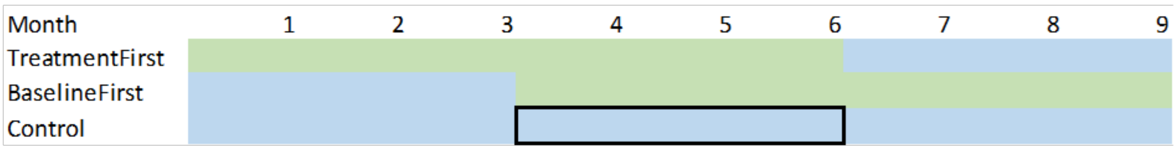

Meaningful insights have been produced from this safety evaluation methodology, without using the control group data uniquely. Despite this, the University partner recommends a control group that is adequately sized for statistical power, for some between groups comparisons including the control group.

Recommendation:

One potential option may to be consider staggering the treatment and baseline periods so that there is more consistent data over the pilot period. That is, for ICVP, periods 3 to 6 (in image above) could have had more baseline participants. However, this option may have implications for the duration of the FOT to ensure sufficient amount and quality of data for analyses.

Situation: Machine learning was not successfully calibrated due to data limitation

Machine learning technique was employed to predict the frequency of near crashes. A robust machine learning model requires training with high quality data. Typically, data collected from the same study will be divided, with some portion being used for modelling training and validation and the remaining for analysis. The driving data collected from control group would be the ideal candidate for modelling training, while the treatment group data would be applied to the model for comparison.

To train the machine learning model, the control data needed to be clearly labelled for any near-crash or crash incidents. Without any onboard video camera installed in the control vehicles, it was not feasible to determine near-crash situations by only examining deceleration data which was found to be quite noisy.

Outcome:

100-car data (from the naturalistic driving study based in Virginia) was sourced to train the ICVP machine learning model, as their data was annotated. The trained model may not be fully transferable, as 100-car collected kinematic data from triaxial accelerometer and ICVP collected GPS location-based data to derive kinematic measures.

Data smoothing methods applied may also be different, meaning that the thresholds suggested in US-based studies may not be appropriate for ICVP.

Exploratory analysis showed that the distribution of output was shifted towards a higher crash rate as compared with output using 100-car data.

Recommendation:

The machine learning model should be trained with data collected by the same technology to avoid systematic bias.

To allow the driving data being annotated for validation, video camera needs to be installed in a selection of control participants / test vehicles.

Situation: Data was difficult to join and quickly evaluate

Tagging datasets makes it easier to rapidly understand the driving behaviour. For any given event, there was a unique scenario id that allowed data to be rapidly linked. Spatial references were also extremely helpful to position data relative to other data (often inferred and joined using station name and time). Other tags that could have improved the ease of the data analysis include the following:

- A unique identifier for the participant ID was based on the participant code and the vehicle station name. Over time, some vehicle stations were replaced, and participant codes modified (spelling corrections for example).
- HMI "on" was not attached to the data and hence, there was a reliance on the external participant management database to confirm if the participant was in baseline or treatment groups.
- The traffic light signal group was only attached to the data when a warning was displayed. It was not possible to analyse events where a warning was not issues – a useful dataset to examine gamification for example.
- Data included both UTC and AEST timestamps, and these data sets were not clearly labelled.

Outcome:

At times, analysis and understanding of events took longer to ensure the right information was joined and irrelevant information filtered.

Recommendation:

- Ensure relevant data has a consistent header or tag with spatial, temporal and participant information.
- One unique identifier for the participant should be used through all datasets (where privacy / confidentiality are maintained).
- It is essential to test and validate that time is applied consistently across all parts of the system.
- When essential data tagging is missing, it is better to fix early and generate a new dataset which is more complete. Many of these would preferably been attached to the dataset when created however modifying the analytics datasets to join this data would have minimised the impact.

9.3 Surveys

Situation: Questionnaires were not collected post-deinstallation of equipment

The last of four questionnaires was conducted one month prior to de-installation, with only focus groups offering the means to explore post-deinstallation sentiments. A subsample of approximately 50 ICVP participants volunteered to take part in these focus groups.

Outcome:

While qualitative research methods offer in-depth insights, a final questionnaire conducted on all participants and based on the same measures as previously assessed in preceding questionnaires may have been useful to explore whether participants' user perceptions changed after equipment removal. A HUET expert suggested that an indicator of success was whether the technology is missed by the user after they no longer have access to it.

Recommendation:

Conduct self-report questionnaires at the following time points, at a minimum: (a) before installation; (b) soon after installation; (c) after condition switch; (d) before de-installation;
(e) after de-installation.

Situation: Broader public sentiment surveys in the treatment period not completed

External to the pilot participant surveys and focus groups, a public survey was conducted in 2018 to measure the public's perception and knowledge of C-ITS before planned public outreach activities. Around 2000 Queensland residents responded, and of these 25% indicated that they had knowledge of C-ITS (compared with 95% for automated vehicles).

The team did not obtain the necessary permission to deliver a general awareness campaign – approval is required by the whole of state government's Government Advertising and Communications Council (GACC). As a result, no additional public surveys – such as a post-campaign survey – were completed.

Outcome:

As one of the key objectives in the benefits realisation plan is to improve public awareness of cooperative vehicles, the baseline survey was still useful in understanding the sentiments of the public.

Recommendation:

Limit surveys as part of broader benefits realisation metrics, and/or limit to treatment period only surveys that seek information about the impact of the pilot on their understanding.

10 Management

10.1 Participant Management

Situation: Complexity of the project resulted in a number of ethics variations being required

As part of the collaboration with a University for participant-based trials (with the University overseeing all aspects of participant management), an approved ethics submission from a designated Human Research Ethics Committee (HREC) was required. However, due to the complex and agile nature of the ICVP, approved ethics submissions were superseded a number of times, approximately 30 variations had to be submitted for approval of changes to project protocols.

Examples of these changes included:

- COVID practices for the participant onboarding and vehicle installation
- Addition of Youi and Budget Direct (insurance companies) acknowledgement of potential customers' involvement
- Use of fleet vehicles that were not owned by the participant
- Reduction in driving time from 3 to 2 hours to increase participant recruitment during dress rehearsal
- Ability to issue letterbox flyers for recruitment
- Ability to issue vehicle sunshades to protect equipment during a heat wave
- Ability to issue roof racks to participants at the end of the pilot
- Any changes to consent documentation or training
- All correspondence to participants – such as troubleshooting, power down the device if the car was not going to be used for extended periods of time and the ability to ask participants if they were interested in continuing to use the equipment after the pilot

Outcome:

ICVP was run under an Agile project management methodology with 2-week sprints. This meant that there were continuous changes. Changes such as those above could not be enacted until approved, impacting the schedule.

The number of ethics variations was not anticipated in the design of the project.

Recommendation:

To reduce the number of ethics variations, ideally, the technology should already be built and tested prior to human research.

Situation: Recruitment efforts required a mid-plan correction

The original publicised target for the pilot was "up to 500 participants". Recruitment communication was largely issued in the first month and included the following:

- 1st weekend - Ipswich mall pilot launch, channel 7, 9 and 10 stories, two-page newspaper article
- targeted twitter/Facebook posts
- radio interviews
- insurance companies' newsletters
- multiple online newspaper articles.

These efforts resulted in approximately 400 registrations (both eligible and ineligible). Based on the expected conversation rate, only 25% of these were likely to be eligible after the initial and final screening tasks.

Outcome:

To supplement the communications effort, a paid letter drop across 66,000 homes was added to the recruitment plan. This and another round of newspaper and radio interviews saw these number of registrations rise to 1041 before registrations were closed. The greatest number of participants were recruited via "other" means (understood to include letterbox drop and traditional media), and social media. "Other" methods also had the greatest conversation of expressions of interest (EOI) to completing the pilot (38%).

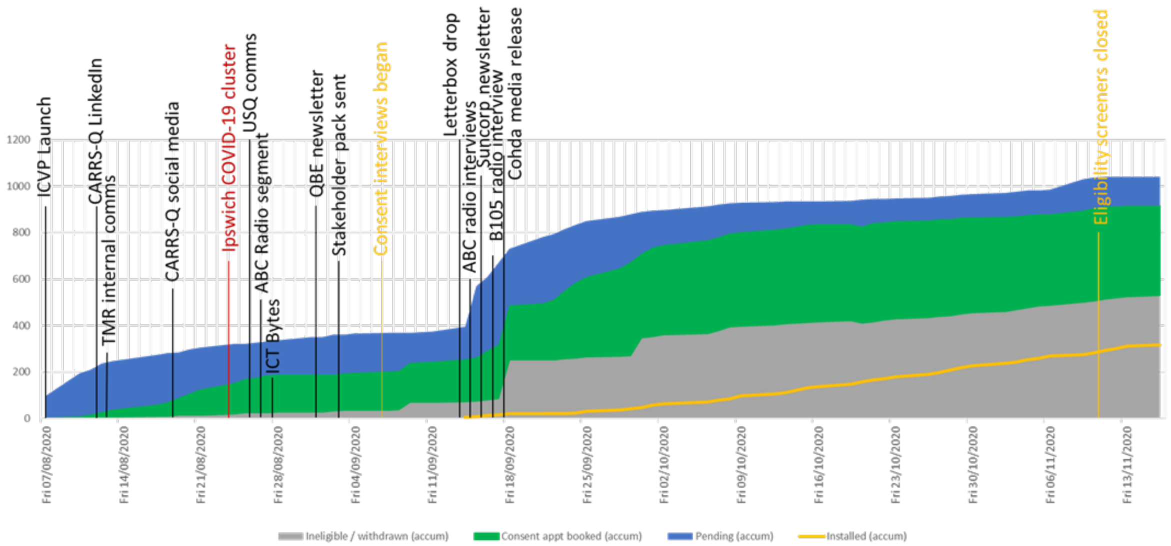

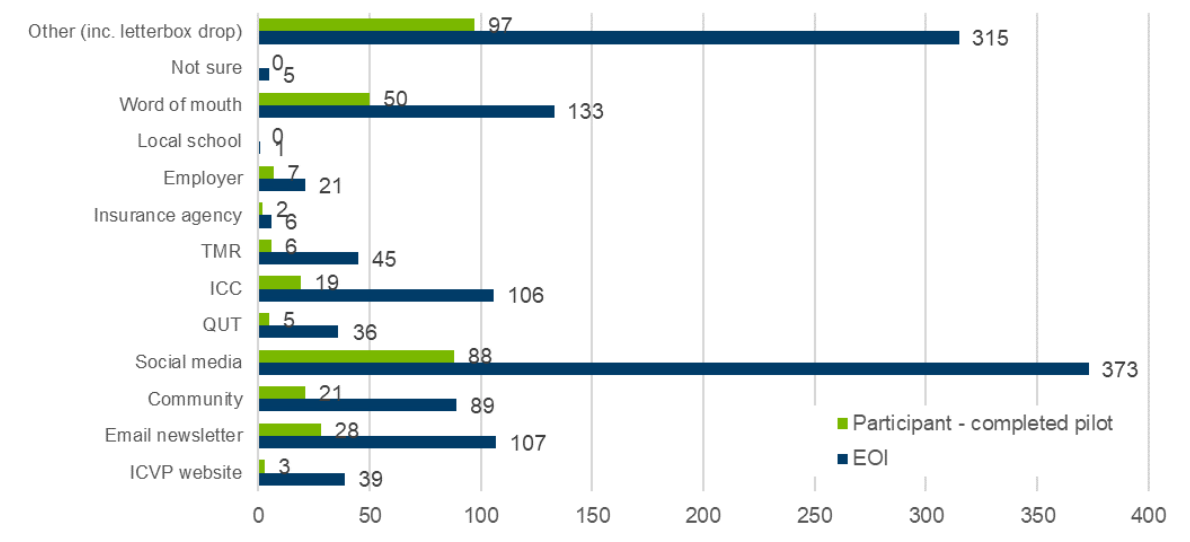

Recommendation:

A letterbox drop is effective for mass recruitment and may prevent mid-plan corrections.

Situation: Conversion of recruitment for the pilot was in line with findings from the dress rehearsal

The conversion rate from registration to participant in the dress rehearsal was around 30%. This was largely due to interested volunteers not fulfilling eligibility requirements.

In total, 1041 people registered, 85% where eligible in pre-screening, and 50% were eligible after more detailed screening. Most could either not be contacted or were not actually eligible in terms of meeting the participation criteria (e.g., they did not have comprehensive insurance, they had comprehensive insurance but it was not with one of the identified insurers whom had approved their clients partaking in the pilot, or they did not drive for up to three hours a week in Ipswich).

Further to this, there was tendency for people not to continue through with their consent interview after being deemed eligible after registration or they did not follow through with their installation appointment after attending the consent interview. In some instances, individuals simply did not turn up to an appointment, they cancelled before their appointment, or they attended their appointment but, for instance, their vehicle was not able to be fitted with roof racks.

Outcome:

Consent booking levelled off at around 400 participants per week throughout October 2020, with new consent bookings simply replacing cancelled bookings. To keep the project on budget and on schedule, it was agreed to complete the installs by the end of November 2020.

No new consent bookings were made from mid-November 2020. As a result of participant cancellations after this date, the number of participants reached 351 by the end of November 2020.

The conversion rate from registration to participant in the pilot was around 30%, consistent with the dress rehearsal.

Recommendation:

Recruitment efforts needs to consider how the eligibility criteria and onboarding process can reduce the potential pool of interested participants. Each screening and/or onboarding step requires relationship management to keep potential participants interested through the multi-step onboarding.

Situation: More staff resourcing was required to manage participant management activities due to overlapping or concurrent in-person, manual activities

With concerns due to COVID and its potential adverse impacts on recruitment, the attempt to retain consented participants, and to meet budgets, it was decided that the installation period would start only one week after the consent interview period commenced. This meant that there was an intensive period of in-person, manual, overlapping activities. Additional resources were required to complete these activities within three months.

Outcome:

Casual labour (up to 7 people) were hired to support the in-person, manual overlapping and concurrent participant management efforts.

Recommendation:

As early as practicable, explore digital solutions for participant and research management to reduce manual processes, including the capability for:

- Recording of relevant personal details from EOIs into records of those who subsequently become participants
- Online bookings for appointments, including amendment and automated confirmation of these bookings
- Automated administration of questionnaires and record of responses
- Automated administration and recording of incentives, and
- Ease of reporting of various statistics, for example, number of completed questionnaires, conversion rates of EOI to consented participants.

Situation: The estimated rate of 6 installation per day was achieved

Participants were required to attend a consent interview and an installation booking one week apart to ensure the availability and fit of correct roof racks to the participants' vehicles. This meant that the consent interview and installation appointments were managed concurrently. The participant consent team could achieve up to 8 consent interviews per day (but 12 if the activities were separate), and the installer team 12 installations per day (limited by the number of available rental vehicles).

Outcome:

The team averaged 32 installations per week (355 vehicles over 11 weeks) – which is equivalent to 6.4 per day. This was less than the capacity of the consent and installation teams. An average of four participants per week were no-shows (either for consent appointments or installation appointments) - with no notice given to the team. While the research team continued efforts to contact and/or replace any 'no-shows', their ability to do this was subject to participants' availability. (Note that 93% of participants who completed consent appointments moved on to complete installation appointments and, thus, became an ICVP participant).

Recommendation:

If C-ITS equipment does not require parts that differ between vehicle models (for example, different roof racks for different vehicles), consider combining the consent and installation appointments into a single appointment to maximise staffing resources.

Situation: Participant support services were fit for purpose

Over the 12-month pilot period, the following support was available to participants:

- Project / participant phone resourced 7am to 5:30pm, Monday-Friday (excluding public holidays) from commencement of recruitment throughout the FOT and continuing to end of focus group data collection post-FOT.
- After-hours automated phone service.
- Pilot project generic email.
- One day per fortnight of onsite support (e.g., maintenance appointments).

Outcome:

- Phone calls / emails: 1019; common requests: equipment issues, use-case questions, withdrawals, concerns regarding obligations.
- Onsite support: 200+; main categories: 89 display bracket recalls, 75 drop-ins for the collection of a vehicle sunshade, 47 display or VITSS related equipment issues, and 31 unrelated withdrawals.

Recommendation:

Despite excessive equipment issues requiring participants to return to the installation site, the process worked well and there were no participant withdraws related to these inconveniences.

Situation: The estimated participant attrition rate was better than expected

The FESTA handbook (Connected Automated Driving | The FESTA Handbook) notes an expected attrition rate of 10% for large pilots of similar scale to the ICVP. This attrition rate was adopted in the design of the pilot in terms of the sample size considerations and statistical power for analyses.

Outcome:

The pilot had an attrition rate of 8.7%, thus below expected rate as per FESTA guidelines. In the ICVP, most participants withdrew from the pilot due to a vehicle no longer working/no longer driveable, a crash (unrelated to the equipment), selling their car, or moving. These were evenly spread over the duration of the pilot, and there was no effort to recruit after the installation period was completed. Transport and Main Roads also wrote-off four (4) participants C-ITS stations – three of whom could not be contacted for uninstall. It also appeared that the in-person consent appointment and participant support services contributed to high retention.

Recommendation:

Expect a 10% attribution rate and 2% equipment write-offs.

Situation: Vehicle installation damage assumptions were within expectation

The team was concerned about possible damage due to the vehicle installation / deinstallation. A number of treatments were put in place, including, the third-party participant management – responsible for managing issues between the participant, installers and Transport and Main Roads, RACQ service for vehicle battery issues, and contingency for any minor damage at the discretional use of the installer.

Outcome:

Few vehicle damage issues emerged (2% of total participants) with rectification fully covered by Transport and Main Roads at a total cost of less than $5000:

- 2x vehicles with vehicle frame damage from the roof racks
- 1x vehicle with roof rack damage (participant's own roof rack)
- 1x vehicle with ute tray damage
- 1x vehicle with paint damage on roof from tape removal
- 1x vehicle with water leak due to the installation
- 1x vehicle general complaint about the potential for damage – asked to be removed from pilot post installation

All issues were managed without escalation and minimal financial impact.

Recommendation:

Ensure contingency is available to resolve vehicle issues.

10.2 Change Management

Situation: (System) JIRA was effective for defect management

Given that the pilot was an integrated system with multiple vendors, and the need to share learnings was one of the primary objectives of the pilot, a single instance of JIRA (a commercially available software tool) was adopted for defect management.

Outcome:

JIRA was a useful tool for the management of tickets, noting the following limitations:

- Transport and Main Roads JIRA - Vendors were required to access Transport and Main Roads JIRA which was difficult because of firewalls. Furthermore, the version of JIRA was lacking some tools that could have streamlined the workflows.
- Defect reporting - The team would manually re-label the ticket type (triage, enhancement or defect) when the defect was confirmed. There were several issues with this approach - labels were not always updated; when the defect was confirmed, it was difficult to determine the vendor's start date for service agreement tracking; and at times the team would add new devices to an open defect ticket stretching resolution times.
- Ticket details - At times the team would rectify an issue without logging a ticket or would not clearly document the outcome – making reporting and corporate memory more difficult.

Recommendation:

Ideally Transport and Main Roads should be using the enhanced JIRA tools for response time tracking. In the absence of these tools, ticket linking may have been a more effective way of tracking against contractual agreements – adding new ticket numbers for defects rather than re-labelling the ticket.

Situation: (System) Half of the defects were resolved within the service agreement timeframes

From approximately 400 JIRA tickets there were 140 defects within the 12-month pilot period. Of these, 70 defects were reported in the first four months to the end of the installation phase, and around 10 defects per month post the peak. Over half the defects were not rectified within contractual timeframes – with competing workloads / resource availability, limited diagnostic data from the device, or the complexity of the ticket being a common reason for delays.

Where practical, the vendor would restart or replace the device at their own cost but would not necessarily resolve the root cause. This meant that software-based issues could occurred again. By the end of the pilot there were three (3) unresolved issues that required more detailed data for analysis – noting tickets with pending data needs were not tracked against the agreements.

To manage resolution times, the team conducted regular meetings to discuss unresolved tickets and reflect on the outcomes of the previous months' resolution times.

Outcome:

Given the lack of contractual ramifications, service level agreements were useful for tracking performance, but generally did not impact the pace of ticket resolution.

Recommendation:

Contracts around the roadside stations need to be carefully considered as the installation labour cost is significant and was not the responsibility of the vendor.

Situation: (System) Any change should consider all impacted parties

There were two examples of poorly managed change:

- The team agreed to move to the latest version of the standards during the test phase but neglected to inform Lexus who were also using the test bed at the time.
- Road hazard types for flashing yellow traffic lights and smoke, were included in the central station scope. Despite the generation of 172 events (which would have translated to 220 vehicle warnings), the use of this cause-code was not a requirement of the vehicle station, and hence never implemented. This issue wasn't identified until the end of the pilot.

Outcome:

A change management process was implemented during the pilot – however, change requests should have been more closely tied to the requirements and transparent to impacted vendors or partners.

Recommendation:

Consider change management as a critical process for a complex integrated system even in a pilot.

Situation: (People) Limited knowledge sharing across other Transport and Main Roads areas
A staff survey was completed in late 2019 suggesting that from 250 responses around 42% viewed themselves as moderately to extremely knowledgeable in C-ITS, 33% had some knowledge and 24% had no knowledge.
Outcome:
Since the survey, the team has provided a range of activities to actively involve interested Transport and Main Roads staff including the recruitment of 100 guest vehicle equipment testers, external and internal technical presentations, pilot go-live information via various communication channels, and technical training for regional staff. The impact of these efforts will be measured in the treatment survey.
Recommendation:
Multi-jurisdictional pilot to involve operational staff from different regions should follow major pilot efforts in order to share learnings as soon as practicable.

10.3 Asset Management

Situation: Transport and Main Roads PDO preferred traffic light field processor and controller could not be integrated
Transport and Main Roads Program Delivery and Operations (PDO) noted a preference for a specific field processor and controller, however, the integration of C-ITS was not possible with some of these devices (see Signal Use-Cases SPATEM for details). PDO's concern led to some creative agreements, including a promise by the field processor vendor to replace failing devices at their own cost, and the need to return signal controllers that were removed for the pilot at CAVI's cost. These agreements were captured in the ICVP Asset Agreement.
Outcome:
Over 24 months, one field processor failed, however, it was presumably replaced at PDOs cost.
Recommendation:
Allow for budget to accommodate business as usual needs within a pilot.

Situation: No or false warnings caused by various non-CAVI led field changes

As the pilot was integrated with existing traffic light equipment, there were a number of non-CAVI led changes that resulted in no or false traffic light warnings at the pilot traffic lights. These included the following:

- Transport and Main Roads PDO changes to signal groups - Per the ICVP Asset Agreement, Transport and Main Roads PDO and CAVI teams agreed to manage their respective equipment at the 29 traffic lights and inform the other party of any changes. PDO updated the signal groups at one intersection but did not inform CAVI. As the pilot map was not similarly updated with the signal groups, participants experienced false red-light and pedestrian warnings.
- Transport and Main Roads PDO cancellation of signal equipment subscriber identity module (SIM) - Per the agreed CAVI Asset Management Plan, PDO was responsible for telecommunications at the test site, however, PDO cancelled the SIM because the SIM was unknown to them.
- RoadTek general maintenance works – Despite efforts to ensure that the pilot equipment was clearly labelled in the field, unrelated field processor works by RoadTek resulted in removal of the pilot software, and in a separate occurrence, an incompatible baud rate at one of the traffic controllers.

Furthermore, as some elements of the central facility service are managed by corporate Transport and Main Roads and there were a number of corporate led changes that resulted in the pilot administrative service outages such as an upgrade to AWS – resulting in the removal of the production account containing the monitoring service; and upgrades to Tableau, resulting in the pilot monitoring tool outage for up to 24 hours.

Outcome:

Most of these issues were one off and intercepted by the team as part of daily checks, however, it highlights the need for transparency of the asset agreements and the associated change management captured in these agreements.

Recommendation:

As C-ITS is integrated with the existing traffic light equipment, Transport and Main Roads PDO will ultimately be responsible for the operation and maintenance of the equipment, removing much of the multi-party coordination.

Ensure asset agreements are well circulated.

Situation: Loss of cellular use-cases caused by various third-party system outages

Given C-ITS is a highly integrated system, the adhoc nature of both planned and unplanned outages across these systems can impact of the availability of the service. Over the pilot period, outages include the following:

- CISCO Jasper SIMs management tool – an unplanned issue impacting cellular access. The team identified a workaround within the first few days of the issue being identified, however, some participants had to return to the installation site to complete the vehicle station commissioning. It took the vendor over 30 days to fully rectify the issue.
- Telstra Firewall – planned outage of 4 hours, during which time the message broker and hence all use-cases and supporting tools were unavailable.
- STREAMS Gateway (data feed for BoQ and IVS) – regular planned and unplanned outages between 1 and 14 hours, during which there is no BoQ or variable speed limit information.
- QLDTraffic (data feed for RHW) – unplanned outages between 2 to 48 hours, during which there is no RHW information.
- AUSCORS (positioning augmentation feed) – a planned and unplanned outage of several days impacting vehicle positioning.

Outcome:

Initially both unplanned and planned outages were identified by the team as part of daily checks, and the team was later added to the number of vendor's email lists for awareness of service interruptions.

Recommendation:

Ideally, in a highly integrated system, planned service updates should not impact the availability of the service, be serviced at night, or be serviced on an agreed schedule.

Situation: Pre-commissioned V-ITS-S certificates expired, causing commissioning delays

In some cases, the pre-commissioned vehicle station certificates had expired before installation.

Outcome:

Because the root cause of the issue was not immediately understood, the team did not give the vehicle station time to re-establish a connection and download the new certificate bundles. In concert, the stations were also attempting a firmware update, which further exacerbated the update. It took the team a couple of days to identify this issue, slowing initial installation efforts. A bench at the installation site was established to ensure correct operations prior to installation.

Recommendation:

The asset log should clearly list the product bootstrapping date or certificate dates to ensure the device is ready to operate.

Where these are likely to have expired, ensure a "re-commissioning" step to allow the device time to reconnect.

Situation: Poor asset management led to weeks of effort locating equipment

Four batches of vehicle stations and associated equipment were received over the course of the test and development phase – two large and two small batches. Equipment was located over four test batches, in six test vehicles, at the Briggs Street installation site and the Mary Street storeroom.

Outcome:

An asset tracking register was created, and an asset lead appointed. It was not uncommon for the team members to move test assets between any and all locations without tracking the asset. At the start of installation process, the team took several weeks to reconcile the assets against the asset register, and a similar level of effort was required again at the end of the deinstallation process.

Recommendation:

Perform regular asset audits and closely manage the movement of all assets.

Situation: Changes to hardware mid-installation causes issues

There were two changes to the product during the installation phase:

- The display unit vendor changed the bracket glue without notice.
- The third-party vendor changed their hardware on the second batch of units with limited notice.

Outcome:

The new bracket failed, resulting in the recall of 89 participants units.

The first batch of vehicle stations were sufficient for uninterrupted installation while the second batch was tested.

Recommendation:

Changes in the product mid-installation are undesirable.

10.4 Project management

Situation: (Scope) Lack of shared understanding of the scope and requirements

Despite working with all vendors for 12 months in the planning phase and co-developing the pilot requirements / specifications, there were regular misunderstandings during the pilot development, resulting in delays to the program.

Outcome:

Many measures were put in place including shared stand-ups, multi-vendor workshops, and use of JIRA (a commercially available software tool) for development and defect ticket management.

The most effective tool was co-location (a requirement of the contract). Parties who did not co-locate were typically not forthcoming with issues that could have been avoided with learnings or support from the boarder team. Some packages also had technical managers that were not from their organisation – this worked well in breaking down vendor silos.

Recommendation:

The development phase should include co-location of development staff.

Situation: (Schedule) Poor understanding of the development level of effort

The planning phase was procured using an early contractor involvement (ECI) model – with four successful station vendors. The intent of this phase was to co-develop the requirements and high-level design for the pilot tender, including the associated schedule.

Outcome:

Only one vendor stated that the pilot schedule was insufficient – and hence, did not tender for the pilot. Others did not understand that the ECI offered them an opportunity to guide the project in a way that would improve their success in the next phase. Most felt that the seed funding for the ECI was insufficient and limited their engagement to meet the budget.

The pilot development took an excess of 6 months longer than planned – effectively doubling the development schedule, impacting all parties' resources and costs.

Recommendation:

Transport and Main Roads' should work more closely with vendors to improve their understanding on an ECI prior to the adoption of an ECI model.

Situation: (Resources) Underestimated staffing resource needs

Most of the package staffing needs were underestimated during peak periods.

Outcome:

- Testing (12 months) – full time team of 5, supplemented regularly by the central facility team
- Participant management (36 months) – full time team of 3.8, and 6 casual staff who assisted with roster shifts during the consent and installation phase
- Safety evaluation (36 months) – full time team of 3, regularly supplemented by the central facility team
- Central facility team (36 months) – full time team of 3, and 12 at the height of development.

Recommendation:

If possible, consider a pool of resources that can be accessed during peak times so that the staff have the appropriate skill set for the work, and less busy work packages are not impacted.

Attachment A: Use Case Summary – Quick Reference

QUICK REFERENCE

Use-case	Example display	What is included in the pilot area?	What warning is displayed on approach to the event?	How is the warning triggered (generally)?	What warning is displayed in the event?	How is the warning triggered (generally)?	Priority	Common issues	Alternative display
In Vehicle Speed	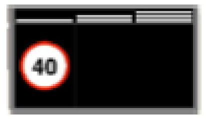	The lowest speed posted: - Static Posted Speeds - Variable Speeds - School Zone Speeds - Some Roadworks Speeds	None	There are no warnings.	None	There are no warnings.	The speed is always on in the pilot area.	***Note - the speed display typically changes when the vehicle has past the sign.*** - The speed map is wrong, missing a new road, or realigned due to road works, - The school calendar or times are incorrect, - The roadwork isn't included, - The variable speed limit sign is faulty, - The vehicle can't locate itself on the road, - The vehicle doesn't have cellular access and couldn't access the data, - The vehicle is out of the pilot area.	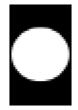 Speed unknown display
Road works		Only some roadworks - typically longer term roadworks.	 Boop	1) The distance to the roadworks is less than a generous stopping distance - approximately 20 seconds travel time; 2) The vehicle is travelling > 10km/h.	 Boop	The vehicle speed is 2km/h or more over the speed limit.	Roadworks have a higher priority than back of queue or road hazard.	- The roadwork isn't included, - The roadworkers have not cancelled the works when finished, - The roadworkers changed the speed limit or location and forgot to update the information, - The vehicle can't locate itself, - The vehicle doesn't have cellular access and couldn't get the data, - The vehicle is out of the pilot area.	
Road hazard		Debris on the road, a crash, or a flooded road. Some events are called in by the public and identified in the QLD Traffic Application.	 (the image doesn't change based on type)	1) The distance to the hazard is less than a generous stopping distance - approximately 30 seconds of travel time; 2) The vehicle is travelling > 40km/h.	None	None - the event is a point.	Lowest priority	***Note: Participants are encouraged to call QLD Traffic for road hazards or signal faults - call 13 19 40*** - QLD traffic hasn't yet identified or confirmed the event, - QLD Traffic has indicated the wrong location, - QLD Traffic has not cancelled the event, - QLD Traffic listed hazard isn't one of the pilot events or is out of the pilot area, - The vehicle can't locate itself on the road, - The vehicle doesn't have cellular access and couldn't access the data, - The vehicle is out of the pilot area.	
Motorway Back of Queue		State motorway queues less than 60km/hr. (Logan Motorway is not a state motorway)		1) The distance to the back of queue is less than an generous stopping distance - approximately 30 seconds travel time; 2) The vehicle is travelling > 60km/h 3) The alert is cleared when the vehicle is travelling < 40km/h	None	None - the event is a point.	Back of queue on motorways has a higher priority than a road hazard event.	- The queue is est. using detectors 500 metres apart -between these, the location of the queue is estimated, - The speed on the detector must be less than 60km/h - a queue may still be evident at higher speeds - The algorithm and operator takes time to confirm a queue, - The queue is between two detectors and not yet detected, - The detector is faulty, - The section of motorway is not configured to detect queues, - The vehicle can't locate itself on the road, - The vehicle doesn't have cellular access and couldn't access the data, - The vehicle is out of the pilot area.	
Advanced Red Light		State road signals - 29 traffic lights on Warwick, Brisbane, Limestone, East, Downs	  shrill -3xBeep	1) The distance to the stop bar is less than an est. safe stopping distance; 2) The vehicle is travelling > 30km/h.	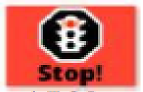 shrill -3xBeep	1) The vehicle crossed the stop bar during a red light	Red light warnings have a higher priority than roadworks or road hazard warnings (there are no back of queue warnings on arterials).	- A lane can have different signals for different movements e.g. left & through shared lanes. In these types of lanes, the through warnings are shown. There are half a dozen locations with this configuration in the pilot, - The traffic equipment has stopped communicating and the light colour is incorrect or unknown, - The intersection is not in the pilot or the lane arrangement of the intersection has been changed, - The vehicle can't locate itself in the correct lane ((particularly for vehicles with offset antennas).	
Turning warning for vulnerable users (pedestrians)		State road signals - 29 traffic lights on Warwick, Brisbane, Limestone, East, Downs	 (including left or right arrow) Boop	1) The vehicle is in a lane that can turn left or right; 2) The distance to the stop bar is less than an est. safe stopping distance; 3) The vehicle is travelling < 40km/h (turning).	 (including left or right arrow) Boop	1) The vehicle is in a lane that can turn left or right; 2) The vehicle crosses the stop bar when the pedestrian light is showing a green man or flashing red man.	Pedestrian warnings have a higher priority than red-light warning, roadworks, road hazard warnings (there are no back of queue on arterials).	- The vehicle is at a COVID traffic light where pedestrian buttons are activitated without pedestrian activity, - The vehicle turned left or right at speeds greater than 40km/h (no warning), - The vehicle travelled through but was in a lane that can also turn left &/or right (warning provided), - The traffic equipment has stopped communicating and the pedestrian light is incorrect or unknown, - The intersection is not in the pilot or the lane arrangement of the intersection has been changed, - The vehicle can't locate itself in the correct lane (particularly for vehicles with offset antennas).	

Pilot area https://www.google.com/maps/d/viewer?mid=1Xv3nXaYNHdRmpgGYvrg1UPdOMploVOkf&usp=sharing

file: G:\TSB\ITS Trial Project\8 - C-ITS Pilot\A - C-ITS Planning\4 - Deliverables\5 - C-ITS High-Level Design

Date edited: 16/12/2021