California Road Charge

Public Private Roads Project

APPENDIX

H

Plug-In Device and Geolocation Report

Appendix was prepared to follow industry guidelines and standards for accessibility and usability to the extent feasible. If you have difficulty accessing content in this document, please contact Caltrans for assistance.

Public/Private Roads Project

Plug-In Device and Geolocation Analysis

Version: 2.0
Date: February 9, 2024
Task: 6.b.1, Agreement 82A0002

Prepared By:

Document Control

File Name:	CA-PPRP_6b1_Device_Location_v6.docx	
Version Number:	2.0	
Created By:	WSP	January 31, 2023
Reviewed By:	Caltrans	February 9, 2024
Modified By:	WSP	February 9, 2024
Approved By:		

CONTENTS

Figures

Tables

Acronyms & Terms

Acronym / Term	Definition
GB	Gigabyte
OBD-II	On-board Diagnostics, Version Two
PRIME	Platform for Road Charge Innovation and Mobility Evolution
TCA	Transportation Corridor Agencies
ODO	Odometer

1. Background

California has led several successful road charge projects to explore the general feasibility of the concept of a “tax per mile” or “road charge” as a potential replacement for the gas tax to fund transportation in the future. In 2017, the state conducted the largest pilot in the nation to date with 5,000 drivers reporting miles for nine months using low-tech (odometer reading) and high-tech options (plug-in devices and smartphone apps). In 2021, California conducted the Four-Phase Demonstration to further test the potential of road charge for use in the state, specifically studying how four business models – pay-at-the pump/charging station, usage-based insurance, ridesharing, and autonomous vehicle shuttles – might create an easy taxpayer experience to report miles and pay road charges.

1.1 PUBLIC/PRIVATE ROADS PROJECT OVERVIEW

The California Road Charge Public/Private Roads Project (the Project) is the third step in California’s study of how a road-charge program could function within the state. Although the two previous road charge projects explored differentiating between in-state and out-of-state miles, they did not investigate the delineation between public roads, private roads, and tribal lands within California in order to determine a “taxable mile”.

The Project’s key goals and objectives included the following:

- Demonstrating the technical, budgetary, and political viability of a geolocation-based system that differentiated between travel on public versus private roadways, as well as travel that traverses tribal lands, while addressing privacy concerns and equity considerations.
- Demonstrating the process and viability of a road charge program administered by a California tolling entity, serving as a commercial account manager.
- Enhancing the functionality of the Platform for Road charge Innovation and Mobility Evolution (PRIME), to provide support for interregional interoperation.

A key component of the Project was the staging of a six-month live demonstration pilot (the Pilot), which helped to evaluate the feasibility of a statewide road charge program across rural and tribal communities. The Pilot, which ran from April 2023 to September 2023, showcased the ability of using global positioning system (GPS) technology to report the miles traveled by vehicles; to delineate whether those miles were traveled inside or outside the state of California; and for mileage found to be in-state, to also determine if the mileage was driven on public roads, private roads, or roads that traversed federally recognized tribal land.

This differentiation of mileage by road type is critical to ensuring that a road charge is assessed solely for mileage driven on roads maintained by the state of California (i.e., roads determined to be both in-state and publicly maintained). By accurately differentiating between public and private roads, the technology ensures that California motorists who predominantly reside on private lands in rural and tribal communities would be exempt from paying for their mileage driven on privately maintained roads in those regions. This approach could further the fairness principle of road charge for motorists by creating a mechanism to automatically differentiate these miles in a safe and secure way. Furthermore, by focusing on these rural and tribal communities, key issues and considerations can be captured to ensure that all communities have a say in California's transportation funding future.

A commonly expressed public concern with road charge programs is related to the costs and complexities of having to build, deploy, and operate a potentially new and independent infrastructure to support a mileage taxation program. To help mitigate such concerns, by exploring the potential for leveraging existing infrastructure and resources in the administration of a road charge program, Caltrans partnered with California's Transportation Corridor Agencies (TCA) to also incorporate within this Pilot a small, 35-participant sub-pilot demonstration. This sub-pilot assessed the viability of a tolling agency serving as a third-party commercial account manager in a road charge system, by exploring ways in which TCA's existing toll transaction processing, account management, and financial reporting systems could be leveraged to support a more economical and familiar way to report, assess, and collect a road charge.

1.2 REPORT GOALS

This report aims to present the primary findings, lessons learned, and associated recommendations gleaned from the six-month live Pilot, specifically with respect to the use of an after-market OBD-II plug-in device to collect miles, and administer road charges. Discussion topics will include how the device works, the accuracy and effectiveness of the collected GPS data in differentiating miles, inherent device limitations, and associated recommendations. While this report is primarily written to address the immediate needs of California stakeholders, the findings included herein most likely apply to any road charge program that incorporates the option of utilizing OBD-II plug-in devices as a mileage reporting option for participants.

It should be noted that the objectives of the Public/Private Roads Project are in alignment with the research goals that Caltrans outlined in its application for the federal grant the state received from the Surface Transportation System Funding Alternatives (STSFA) program. More specifically, Caltrans' grant application requested funds to test the ability of current GPS technology to successfully support the accurate differentiation of public versus private roads, with the concurrent goal of filling the previous void in both the state and national body of knowledge in this space. An independent, objective evaluation to determine the extent to which the Pilot successfully fulfilled the criteria laid out within the grant application, is provided in the Task 7.a.3 Independent Evaluation deliverable (Final Report Appendix A).

1.3 REFERENCED DOCUMENTS

Following is a list of documents that are referenced within this report, and that can be referred to for further details and a more comprehensive understanding of the findings presented:

- California Road Charge Public/Private Roads Project
    - Task 3.a.2 Concept of Operations
    - Task 8.b.3 Final Report and Appendices
    - Task 7.a.3 Independent Evaluation (Final Report Appendix A)

  - Task 6.c.1 Pilot Operations Plan and Closeout (Final Report Appendix D)
  - Task 3.b.1/3.b.2 Pilot System Report (“PRIME Batch Updates”, Final Report Appendix E)
- California Road Charge Pilot
  - Final Report, 2017
- California Four-Phase Demonstration
  - Comprehensive Report, 2022

2. Introduction to Plug-in Devices

2.1 OBD-II PORT

Before discussing aftermarket On-board Diagnostics, Version Two (ODB-II) devices, this report will first provide a bit of background on the OBD-II port itself. Every vehicle sold in the United States since 1996 is equipped with an OBD-II port. The port's primary purpose is to ensure that there is a standard means by which an owner, mechanic, or emissions testing service can connect to a vehicle using a standardized port to obtain diagnostics data from a vehicle. For example, this diagnostics data can be used to determine the general emissions health of a vehicle. Data obtained from a vehicle via this port, such as vehicle speed, fuel use, and odometer readings, is directly relevant to a road charge program.

One caveat is that a few electric vehicles, such as the Tesla Model 3 and Model Y, have been provided an exception to the OBD-II support mandate that has been in place since the mid-nineties; such vehicles do not in fact include an OBD-II port from the factory. It is unknown if more electric vehicle manufacturers will also seek exceptions to the OBD-II port.

2.2 OBD-II DEVICES

The device used for the Pilot is similar in nature to many aftermarket OBD-II devices. The device is designed to be plugged into the OBD-II port of a vehicle and is equipped with an onboard central processing unit (CPU), memory, operating system, cellular transmission card, GPS receiver, and various other sensors. These attributes are used in concert to collect and process data from the vehicle into which the device is plugged, and then transmit the information wirelessly to the device-listener gateway of the device manufacturer.

The specific device used for the Pilot has the following specifications.

Physical Characteristics	Value
Dimensions	L=43mm, W=46mm, H=23mm
Weight	32.1g
Environment	IP64
Temperature Range	-40 C to +85 C
Humidity	0% - 95% (SAE J1455)
Shock, Vibration and Heat	SAE J1455, SAE J1211

Certification	Value
Carrier Certifications	FCC, PTCRB, AT&T Certified
Environment Certifications	RoHS Compliant

Electrical Characteristics	Value
Supply Voltage	24V (min. 8V to max 32V)
Current Consumption	<4 mA Average (sleep mode), <100 mA @ 12VDC (data upload)
Voltage Protection	Over and Reverse Voltage, Load Dump (J1113/11), Short Circuit, Transients (ISO 167502), ESD (J1113/13)
Current Protection	Internally Protected

Vehicle Communications	Value
Protocol Support	ISO 15765, GMLAN, FNOS, ISO 9141-2, J1850 PMW, J1850 VPW, KWP-2000, ISO 14230-4
Multi-CAN Communication	Simultaneous HSCAN = MSCAN (Ford and others) or Simultaneous HSCAN + SWCAN (GM)
Protocol Detection	Automatic vehicle protocol recognition
Ignition ON/OFF Detect	Automatic wake-up from sleep mode on IGN ON / Automatic sleep mode on IGN OFF

Wireless	Value
Cellular	4G LTE CAT1 (LTE Band 2, LTE Band 4, LTE Band 5, LTE Band 12) 3G UMTS/HSPA 850, 1700, 1900 MHz
Carrier Support	AT&T Certification, T-Mobile, Rogers, Telus, Bell, Telcel, Movistar
Band Support	LTE Band 2 Uplink 1850-1910 MHz/Downlink 1930-1990 MHz LTE Band 4 Uplink 1710-1755 MHz/Downlink 2110-2155 MHz LTE Band 5 Uplink 824-849 MHz/Downlink 869-894 MHz LTE Band 12 Uplink 699-716 MHz/Downlink 729-746 MHz UMTS Band 2 Uplink 1850-1910 MHz/Downlink 1930-1990 MHz UMTS Band 4 Uplink 1710-1755 MHz/Downlink 2100-2155 MHz UMTS Band 5 Uplink 825-849 MHz/Downlink 869-894 MHz
COMM	TCP/IP, UDP, FTP, SFTP, HTTP, HTTPS
SMS	Point-to-point MO and MT SMS cell broadcast
Bluetooth	Bluetooth 4.2, BLE, Dual-mode support, multi-phone, Secure Simple Pairing (SSP), Serial Port Profile (SPP)
WiFi	802.11 b/g/n 2.4GHz (Transmit data from device to back-end via WiFi connection)
Antenna	Internal built-in Cellular, WiFi, and Bluetooth
FOTA	Firmware-Over-The-Air update for configuration and device firmware

GNSS	Value
Satellite Channels	Acquisition 118 / Simultaneous Tracking 40
Constellation Support	GPS, GLONASS, Galileo, BeiDou
Antenna	Internal built-in
Cold Start/Hot Start	Cold Start <32 seconds TTFF Sensitivity – 145 dBm Hot Start < 1 second Sensitivity – 160 dBm
Data Acquisition Rate	Typical 1 Hz
Accuracy	Position <2.0 m CEP (open sky 1 Hz tracking)
A GPS	Full A GPS Support

Accelerometer / Gyro	Value
3-axis Accelerometer	X, Y, Z output
3-axis Gyrometer/Gyroscope	X, Y, Z output
Output resolution	+/- 2,4,8,16 g (12 bit resolution)
Auto=Normalization Algorithm	Self-Calibrating, Auto-Normalization of the data to the vehicle's direction-of-motion

Data Processing	Value
Microprocessors/OS	Cortex A7/Linux OS (platform for customer-specific device level development)
Memory (RAM/Flash)	~250Mbytes / ~500Mbytes
Security (Hardware)	Hardware Security Module
Security (Software)	TLS 1.2 or AES-128 (provides industry standard security for device and transmitted data)
Installation	Self-installed (10 sec or less)
Data Collection Interval	Configurable (1 Hz Max)
Back-up Power	SuperCap (10 F) (supports real-time disconnect events at extreme automotive temperatures and 500K cycles)

2.3 OBD-II RAW DATA COLLECTION

Figure 1 shows the high-level processing steps by which the raw travel data required for the accurate calculation of road charges was collected from the vehicle and flowed to the pilot system platform.

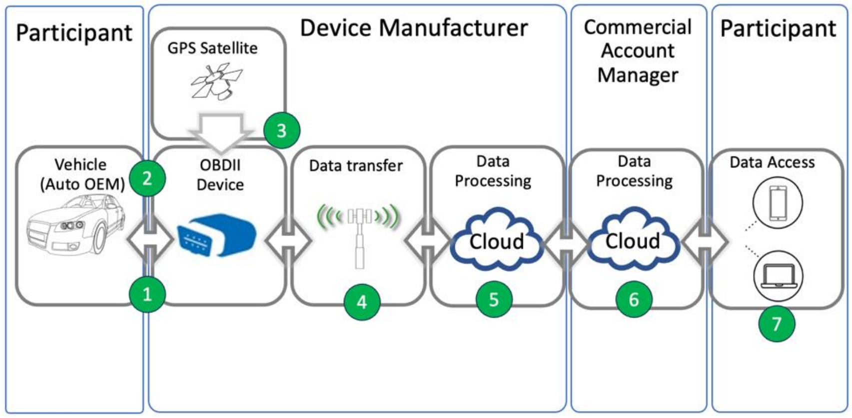

Figure 1: OBD-II Data Collection

The steps involved in the flow of data from the OBD-II devices are as follows:

1. Device sends a data request to the vehicle via the OBD-II bus. Types of data typically requested include the odometer reading, vehicle speed, fuel used, and battery voltage level.

2. Vehicle returns on the bus the data that was requested by the device.

3. Device determines GPS location information every X seconds, where the technology vendor can configure that interval of X.

4. Device sends collected vehicle data to Device manufacturer's gateway via cellular technology every X seconds. X is configurable. Data can be sent via appropriate industry standard protocols (TCP, UDP, FTP, HTTP or HTTPS).

5. Device manufacturer's gateway processes and stores the raw data into usable information that can be consumed by external entities, such as a road charge program's commercial account manager.

6. Commercial account manager pulls the data from Device manufacturer's gateway on a daily basis (or more often as needed), and then processes, stores and formats the data for consumption by road charge program participants.

7. Participants then use their computers and/or mobile devices to log into their personal accounts on commercial account manager's system to view their curated travel data.

The OBD-II device features an extensive message set for transmitting a wide variety of vehicle data to the vendor's back-office system. The specific message types used in support of this pilot are described in Table 1.

Table 1: Device Message Types Used for the Pilot

Message Type	Description
TRIP-START	Sent at the beginning of a trip, as indicated by an "ignition on" event being detected on the OBD-II bus. Payload includes device-assigned trip number, as well as the time, latitude/longitude position, and odometer reading at the outset of the trip.
TRIP-END	Sent upon device's algorithmic determination that a trip has ended. Payload includes device-assigned trip number; the time, latitude/longitude position, and odometer value at the end of the trip; and the device's determination of the distance covered by the trip.
GPS	Sent at regular intervals throughout a trip to communicate vehicle's GPS location at trip waypoints. Payload includes device-assigned trip number and latitude/longitude position.
CONNECT	Sent when device is physically plugged into vehicle. Payload includes time, latitude/longitude position, and vehicle-provided odometer reading at the time of the plug-in.
DISCONNECT	Sent when device is physically disconnected from vehicle, either intentionally or accidentally. Payload includes time, latitude/longitude position, and approximated odometer value at the time of the disconnect.
HEARTBEAT	Sent at regular intervals, to indicate device is still online as intended.

It should be noted that with the exception of the GPS message type, the above device messages are transmitted via User Data Protocol (UDP), which is a relatively streamlined and performant (speed-wise) protocol for wireless communication of data. The performance advantages of UDP are largely due to the fact that it foregoes "hand-shaking" (i.e., UDP does not ensure upfront that a firm connection with the targeted endpoint has been established prior to transmitting the data). The tradeoff for eliminating this time-consuming hand-shaking step, however, is that data packets transmitted via UDP can sometimes get lost. Additionally, data packets can also be received out of order, which is particularly troublesome for a road charge application that depends upon TRIP-START, GPS, and TRIP-END messages to be received in an orderly sequence. This relative "unreliability" of the UDP communication protocol is noteworthy, in light of several of the observations and finding recounted in Section 4 below.

Similarly, the contrasting circumstances under which the device transmits a TRIP-START message versus a TRIP-END message should also be noted. A TRIP-START message is actively triggered by the device's detection of an "ignition-on" event on the vehicle's OBD-II bus. However, the device is not able to similarly leverage the OBD-II bus to trigger the sending of a TRIP-END message, because once the vehicle ignition is turned off, the bus becomes disabled and is no longer available to facilitate the device's detection of a similar "ignition-off" event. As a consequence, the device employs an algorithm whereby if no successful response is received from the OBD-II bus over a 3-minute interval, the device assumes the ignition has been turned off and transmits the TRIP-END message accordingly. Again, this passive/algorithmic triggering of a

TRIP-END message is noteworthy, as it informs several of the observations and findings cited in the following sections.

2.4 RAW TRAVEL DATA TO ROAD CHARGE

The pilot team collaborated with the device manufacturer to configure an automated interface with their company's hosted gateway. Messages received from the devices and captured by the gateway were subsequently automatically forwarded to an Amazon AWS repository that served as a device message cache for the Pilot. This message cache was monitored by the pilot system for the receipt of TRIP-END messages, at which point an aggregation process was triggered to reference all device messages associated with the trip to create a trip record curating all parameters required to accurately determine the road charge for the trip, to include:

- The times, locations, and odometer readings at the beginning and end of that trip, respectively.
- A set of waypoint locations traversed by the trip.
- The device's determination of the overall distance travelled during that trip.

After the creation of a trip record, the distance covered by the trip was then differentiated by road type and a road charge for the trip was calculated, as depicted in Figure 2. This process is further detailed in the Task 6.c.1 Pilot Operations Plan and Closeout deliverable (Final Report Appendix D, hereafter "Closeout Report"). The process to calculate road charge from the plug-in device can be summarized as follows:

- Each GPS waypoint traversed by the trip was computationally assessed against a Caltrans-approved map-set, to determine whether the waypoint was located inside or outside the state of California … and for those waypoints found to be in-state, to subsequently also determine if it was located on a public road, a private road, or a road that was located on tribal land.
- The distance between consecutive trip waypoints was then calculated, with the resulting inter-waypoint distances being summed-up and allocated accordingly to the public road, private road, tribal road, and out-of-state road distance sub-totals for that trip.
- Mileage determined to have taken place on public roads was assessed a per-mile road charge. A road charge was not applied to any other road type, as such roads are not maintained by the state of California.
- For all mileage captured on behalf of a given vehicle (regardless of road type), the fuel used by the vehicle to achieve that distance was approximated by leveraging the vehicle's EPA MPG rating. The state tax that would be required to purchase the fuel to drive that distance, was then credited back against the road charges assessed for that same mileage.

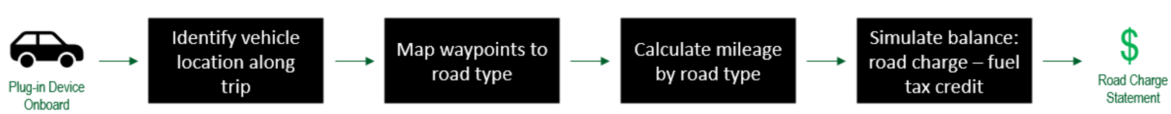

Figure 2: Calculation of Road Charge from Plug-in Devices

3. Highlights of Pilot Results

The Closeout Report (Final Report Appendix D) provides a detailed breakdown of the results for the live demonstration portion of the Project. These results include:

- The individuals who participated, segmented by demographics as well as by their inclusion within the following participant cohorts: representatives of the rural community (Rural Cohort); members of federally recognized tribes (Tribal Cohort); and TCA accountholders (TCA Cohort).
- The vehicles from which data was collected for the demonstration.
- The timing of, as well as the regions traversed by, the Pilot's vehicle trips.
- The distances driven by pilot participants, broken down by road type (in-state public, in-state private, in-state tribal land, or out-of-state).
- The resulting simulated fees that were assessed in association with the collected mileage information.
- The following sections provide a high-level summary of the above results, covered in much greater detail within the Closeout Report.

3.1 PARTICIPANTS

As shown in Table 2, there were 289 individuals enrolled into the pilot, and of those, 283 subsequently completed the installation of their OBD-II plug-in device such that travel information could be collected from their vehicle, thereby becoming "active participants" in the Pilot.

Table 2: Pilot Participants

Cohort	Enrolled	Installed Device and Actively Participated
Rural	238	234
Tribal	16	15
TCA	35	34
TOTAL	**289**	**283**

3.2 VEHICLES

Table 3 provides a breakdown by fuel type and by EPA MPG rating of the 283 vehicles from which data was collected during the Pilot's live demonstration. There were 14 fully electric and 2 alternative fuel (E85 flex fuel and Compressed Natural Gas) vehicles that took part, with the balance being made up of mostly gasoline vehicles and some hybrid, and diesel vehicles.

Table 3: Pilot Vehicles

Fuel Type	# Vehicles				MPG /MPGe	# Vehicles			
	Rural	*Tribal*	*TCA*	**TOTAL**		*Rural*	*Tribal*	*TCA*	**TOTAL**
					0-15	16		1	17
					16-25	122	7	16	145
Gas	212	14	32	258	26-35	58	5	10	73
					36-45	5	1		6
					> 45	11	1	5	17
					0-15	3	1		4
Diesel	8	1		9	16-25	3			3
					26-35	2			2
Electric	13		1	14	36-45	1			1
					> 45	12		1	13
Other	1		1	2	16-25	1		1	2
TOTAL	**234**	**15**	**34**	**283**		**234**	**15**	**34**	**283**

3.3 DEVICE MESSAGES

Over the course of the live demonstration's six-month period, the Pilot platform collected over 92.4 million discrete messages from the devices installed in the 283 participants' vehicles, as summarized within Table 4 below.

Table 4: Device Message Types and Counts

Message Type	6-Month Message Count
CONNECT	1,432
DISCONNECT	1,172
TRIP-START	167,019
TRIP-END	166,724
GPS	91,543,249
HEARTBEAT	532,089
TOTAL	**92,409,081**

3.4 PILOT VEHICLE TRIPS AND MILEAGE

The trips taken and the mileage driven by Pilot vehicles are summarized in Table 5. The average daily trips taken by each participant, as well as the average miles driven during each trip, were generally in line with the "rule of thumb" values observed within various other road charge pilots, both in the state of California and in other states (i.e., 3 trips per day, and 10 miles per trip).

Table 5: Trips and Miles for Pilot Vehicles

Trip Metric	Cohort			All Participants
	Rural	*Tribal*	*TCA*	
Vehicles	234	15	34	283
Total Trips *	121,456	10,889	20,138	152,483
Avg Total Trips / Participant	519	726	592	539
Avg Daily Trips / Participant	2.8	4.0	3.2	2.9
Total Miles Driven	1,238,453	107,740	176,803	1,522,996
Avg Miles / Trip	10.2	9.9	8.8	10.0

* Excludes zero-distance trips (see Section 4.1)

3.5 DIFFERENTIATION BY ROAD TYPE

The results of the road type differentiation process for each of the participant cohorts are shown in Figure 3. For the vast majority of the Pilot mileage information that was captured, this differentiation process was found to effectively identify the road types that were traversed during any given trip, and to accurately suballocate trip mileage accordingly. However, as outlined in Section 4.5, several scenarios were encountered during the Pilot that interfered with the system's ability to accurately differentiate the miles collected for a given trip, resulting in the miles in these scenarios having to be allocated to an "undifferentiated" category (versus a category such as public road, private road, etc.). It should be noted, however, that less than 4 out every 1,000 miles collected on behalf of the pilot, actually fell into this undifferentiated category.

For a more detailed breakdown and analysis of these road type differentiation results, see the Closeout Report.

Figure 3: Differentiation by Road Type

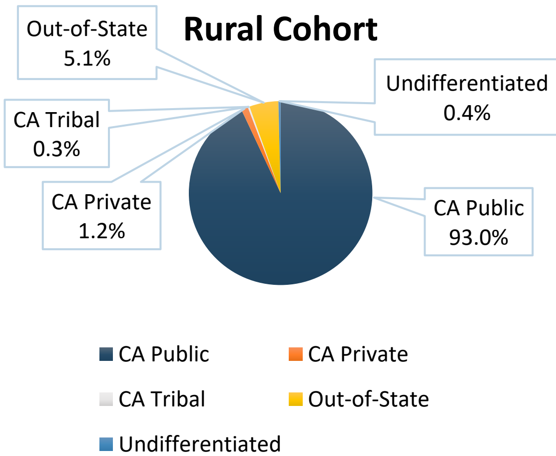

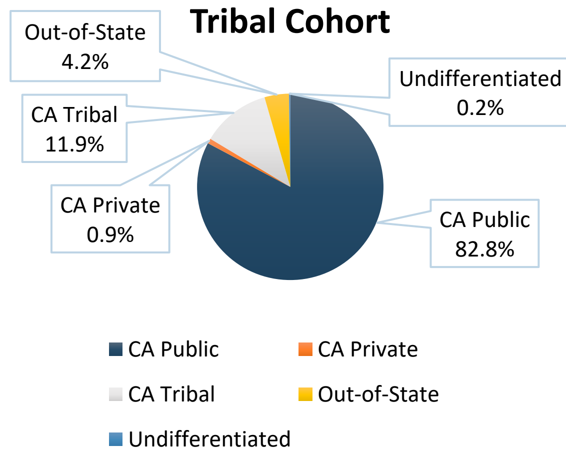

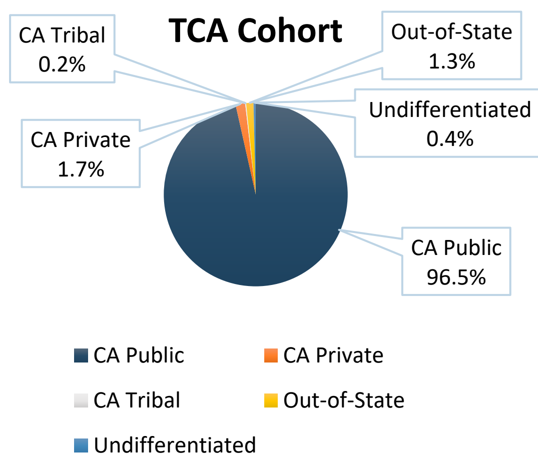

4. Pilot Findings

While analyzing plug-in device data, the project team identified a series of insights ranging from distance calculation inaccuracies to issues with specific trip messages. This section individually describes each finding and proposes a corresponding recommendation for future road charge projects using plug-in devices.

4.1 SUSPECT ZERO-DISTANCE TRIPS

As introduced in Section 2.3, the TRIP-END message sent by the device includes the device's determination of how far the vehicle actually travelled during the trip. The device derives this distance calculation from the vehicle's VSS (Vehicle Speed Sensor), an approach that has been previously shown to lead to highly accurate determinations of distance travelled. However, a seemingly very high number of TRIP-END messages for the pilot were found to contain a distance value of zero. More specifically, approximately 13,400 (8%) of the trips collected for the six-month Pilot were reported to have zero distance. Although it is possible for a given trip to legitimately involve a reported zero km travelled, as when the ignition gets turned on and subsequently turned off without travelling very far, this count of zero-distance trips appeared to be exceedingly high.

4.1.1 Analysis

In addition to the distance travelled by the vehicle, the TRIP-END message also provides the GPS positions (i.e., latitude and longitude) at both the beginning and the end of the trip. For each of the 13,400 zero-distance trips, the distance between the starting and ending latitudes/longitudes was calculated, yielding the following observations:

- The generally accepted error tolerance for GPS is approximately 16 feet. If we allow for such an error on both the starting and ending assessments of location by the device, a calculated distance between the reported starting and ending locations of less than 32 feet could potentially be attributed to GPS error.
- Of the 13,400 total zero-distance trips collected, approximately 6,200 of them had a distance between their reported starting and ending locations that exceeded 32 feet; i.e., for 6,200 of the trips for which zero distance was reported by the device, we can safely assume that the car did in fact move during the trip.
- The distance travelled is reported by the device in kilometers, and the Pilot system captured these kilometer values with 3 digits of precision (i.e., 0.001). Due to rounding, then, a trip during which the vehicle actually travelled less than 0.0005 km might still be reported by the device as having travelled zero distance (i.e., rounded down to 0.000 km). However, the aforementioned 32 ft threshold equates to 0.0097 km, a value greater than 0.0005 and which would not be rounded down to 0. Therefore, it's not feasible to attribute this observed device behavior to rounding issues for any of the 6,200 zero-distance trips with a calculated start/end locational gap exceeding 32 feet.

- In summary, then, it would appear that for 6,200 (4%) of the trips that were collected on behalf of the six-month pilot, there was an irreconcilable inconsistency in the trip information being reported by the device for the trip (i.e., for such trips, the GPS information reported by the device indicated that the vehicle travelled a distance that should've been reported as a value greater than zero, and yet the device failed to do so).
- In addition, 279 out of the 283 devices taking part in the pilot reported at least one such trip, and therefore this suspect behavior cannot be attributable to a handful of malfunctioning devices.

4.1.2 Suspect Zero-Distance Trips: Recommendations

The project team collaborated extensively with the device vendor to investigate this issue. To-date, the vendor's engineering team has not been able to identify any potential cause for this device behavior. In fact, the vendor indicated that this is the first time that this phenomenon has been brought to its attention on behalf of any of their other previous device-related engagements.

Moving forward, for future road charge initiatives that offer the OBD-II plug-in device as a mileage reporting option, it is recommended that logic be implemented to more closely examine trips for which the device reports zero distance travelled, rather than simply assume that such trips are outside the scope of a road charge. At a minimum, zero-distance trips for which the distance between the reported starting and end points exceeds a specified threshold should be flagged for further investigation.

It should also be acknowledged that the current generation of OBD-II plug-in devices was specifically designed for the insurance industry and was subsequently updated to address the needs of emerging fleet-management systems (see Section 5.1). Neither of these domains (i.e., insurance nor fleet management) require the precise assessment of miles driven, which is obviously an absolute necessity for road-charge applications. It is therefore highly recommended that a prerequisite for moving forward with the use of OBD-II plug-ins on behalf of road charge applications would be for road charge practitioners to collaborate with OBD-II device vendors to tailor the device functionality and message set to the unique distance-critical needs of the road-charge application.

Eventually, all road charge programs that use an OBD-II device must incorporate some type of 'true-up' process. The true-up process would be fairly simple and easy to implement. During enrollment, an initial true odometer reading is captured. This true reading can be obtained by a third party visually reading the odometer or the participant self-reporting. At any time a 'true-up' is needed, another true odometer reading is obtained. The difference between the initial reading and the subsequent reading would be the actual mileage. The actual mileage would be compared to the current total of all captured mileage and if the difference is over the allowed threshold, an adjustment to the account for the vehicle can be made.

4.2 ODOMETER DATA ISSUES

4.2.1 Background

As outlined in Section 2.3, four of the device messages utilized by this pilot contained an odometer value for the vehicle: CONNECT, TRIP-START, TRIP-END and DISCONNECT.

The CONNECT message is transmitted by the device when the plug-in device is physically installed into a vehicle's OBD-II port. To obtain the value of the odometer to insert into the CONNECT message, the device first communicates with an external third-party service to determine if the vehicle make is capable of fielding requests for its odometer value across the internal vehicle network and, if so, to determine the model-specific parameter that should be included by the device in issuing such a request. The device then puts a request for the odometer value onto the vehicle's OBD-II bus, awaits a response from the vehicle on that same bus, and then includes this vehicle-provided odometer reading within the transmitted CONNECT message.

Similarly, on behalf of assembling the content for the TRIP-START message sent at the outset of every trip taken by the vehicle, the device once again obtains the odometer value from the vehicle by placing a request on the OBD-II bus, and then using the value returned on the bus by the vehicle as the odometer value transmitted within the TRIP-START.

In contrast to the CONNECT and TRIP-START messages, at the time that the TRIP-END and DISCONNECT messages are assembled, the car has been turned off and the OBD-II bus is no longer available; i.e., the device cannot solicit the odometer reading from the vehicle on behalf of a TRIP-END or DISCONNECT message, because at that time the OBD-II bus is not available to facilitate the request. Additionally, because a non-trivial amount of network resources are required for the vehicle to respond to a device's request for odometer value at any point in time, the device vendor has, to-date, made the design decision to avoid placing undue burden on the vehicle's internal network by refraining from periodically requesting odometer value during a trip. As a result, rather than obtaining the odometer value from the vehicle, the device is forced to only approximate its value whenever it's transmitted via a TRIP-END or DISCONNECT message, using the following process:

- Throughout the course of a trip, the device updates an internal "current odometer" value.
- At the trip's outset, the current odometer value is set to the beginning odometer value for the trip, as obtained from the vehicle.
- As described in Section 4.1, the device continuously calculates the distance driven thus far for the current trip.
- Throughout the trip, the device periodically updates the internal current odometer value by adding the device-calculated distance for the trip thus far to the trip's starting odometer value.
- Whenever the device sends a TRIP-END or a DISCONNECT message, it sets the ending odometer in that message to the value of the internal current odometer value.

4.2.2 Unreliable Odometer in CONNECTs and TRIP-STARTs

If the device places a request on the OBD-II bus to solicit the odometer value from the vehicle on behalf of a CONNECT or a TRIP-START message but the vehicle's response time is excessive, the device will still transmit the message. However, the odometer value in that message will be set to 0.

This behavior was observed for both CONNECT and TRIP-START messages. More specifically, 17% of all CONNECT messages transmitted to the pilot system were found to contain an odometer value of 0. And of the 170,000+ trip records that were created on behalf of the pilot, 293 had a TRIP-START message with an odometer value set to 0. In addition, this latter phenomenon was much more likely to occur on the very first trip taken after a device had been installed, as was the case for 230 vehicles, or 81%, of the devices involved in this Pilot.

The reason for such a high number of vehicles providing an odometer value of 0 for the first trip after plug in is due to the fact that there is a specific process for activating a device the first time it is plugged into a vehicle and it was not followed by every participant during the pilot. The device must go through a series of diagnostics checks as well as determine the vehicles protocol. For this reason, it is always recommended that the first trip not be a trip at all but a 2 minute 0 mile trip where the vehicle just idles while it performs it's diagnostics. When this recommended 'first trip' process is followed, the percentage of vehicles that provide odometer correctly on the first trip is almost 100%.

4.2.3 Odometer Gaps between Consecutive Trips

The ability to identify "missing miles", or mileage that for any reason is unaccounted for by the road charge system (e.g., participants intentionally unplugging their device to temporarily evade road charge), is a compliance-related area of high interest to any operational road charge program employing an OBD-II plug-in device. (The topic is addressed in significant detail within Section 4.4.) In an effort to detect missing miles, the project team compared the starting odometer of every trip collected to the ending odometer of the previous trip; if a non-trivial "gap" existed between these two values, that would theoretically indicate that the vehicle somehow moved between those trips by a distance that was for some reason not accounted for by the road charge system.

In undertaking this analysis, the project team discovered an alarming number of trips for which the trip's starting odometer differed from the previous trip's ending odometer. More specifically, this condition was true for 20,024 trips, or almost 12% of the trip records collected for the Pilot. Moreover, as shown in Figure 4 below, for 38% of these trips in which such an "odometer gap" existed, the trip's starting odometer was actually <u>less than</u> the previous trip's ending odometer. Although some of these discrepancies could potentially be rightfully attributable to participants either intentionally or unintentionally disconnecting their plug-in devices (and therefore represent instances of successful identification of previously unaccounted-for mileage that should now be subjected to road charge processing), the sheer number of these occurrences indicated that something else was going on here. Said differently, it's not realistic to assume that participants uninstalled and reinstalled their devices over 20,000 times in six months.

As explained in Section 4.1, the ending odometer value contained in a TRIP-END message is only a calculated approximation of what the device believes the ending odometer should be, as opposed to an actual reading of the odometer as provided by the vehicle at the end of a trip. Additionally, all odometer values reported by the device have zero digits of precision; i.e., they are integer values. The device vendor's engineering team has confirmed that these two factors do lead to situations in which a difference between the odometer values across consecutive trips can in fact occur. The engineering team also reported that when this does occur, the difference in the two odometer values should never be greater than 2km (plus or minus).

For the purposes of this report, the "Trip Odo Gap" for a given trip is defined as follows:

"Trip Odo Gap" = (trip starting odometer – previous trip's ending odometer)

Figure 4 then depicts the frequency of "Trip Odo Gap" values that were observed on pilot trips for which the value was non-zero. For example, the figure shows that 1123 trips were observed for which the difference between the trip's starting odometer and the previous trip's ending odometer was greater than 2 km (i.e., indicating that somehow the trip started off more than 2km away from where the previous trip ended).

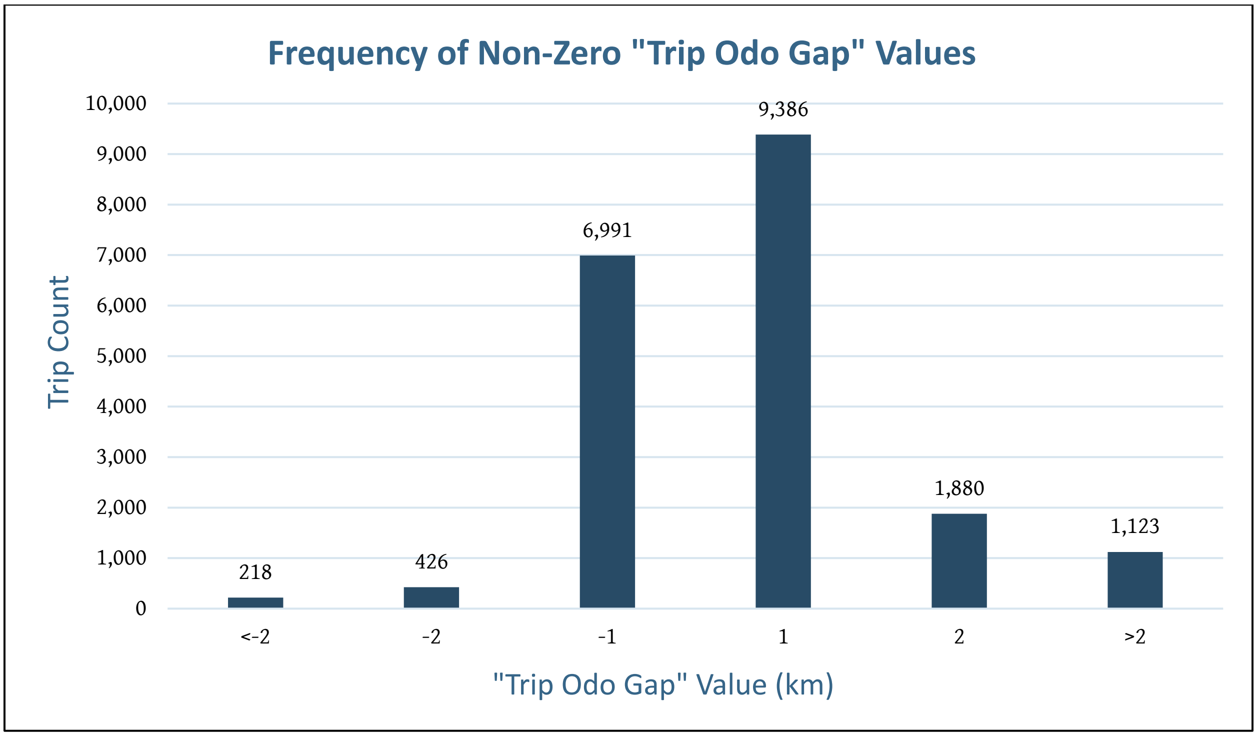

Figure 4: Frequency of Non-Zero ODO_GAP Values

Figure 4 reveals that 18,683 (93%) of the 20,024 trips exhibiting a non-zero "Trip Odo Gap" had an odometer difference that was found to be in the range from -2km to +2km, and therefore could be "explainable" due to the above factors cited by the device vendor.

At the same time, Figure 4 also reveals that 1,341 (7%) of the trips exhibiting a non-zero "Trip Odo Gap" had an odometer difference that did exceed the vendor-explainable threshold of 2km (i.e., the 218 trips shown with a "Trip Odo Gap" value less than -2 km, plus the 1,123 trips with a gap value exceeding 2 km). And again, it's unlikely that the gaps observed for these 1,341 trips were all attributable to the participant actually uninstalling and re-installing their device. It would therefore appear that there was an as-yet unexplained inaccuracy in the device's capture of trip information for at least a sizable share of these 1,341 pilot trips. Furthermore, the device manufacturer has to-date not been able to provide a plausible explanation for these unanticipated findings.

4.2.4 Odometer Issues: Recommendations

The odometer-related issues brought to light by the project team indicate that the odometer values reported by the device are at least somewhat unreliable. Because the device-reported odometer values were not directly involved in the calculation of road charges for this Pilot, the observed unreliability in odometer values did not impact the Pilot's accuracy in being able to differentiate mileage by road type and assess appropriate fees.

However, the unreliability would appear to be very relevant to any attempt on the part of future road charge initiatives to leverage the device-reported odometer values in pursuing the identification of missing miles. At minimum, any such effort should ideally involve further collaboration with the device vendor to determine whether non-nominal (i.e., exceeding 2km) "Trip Odo Gap" values do in fact point to authentic "missing miles", or whether they might sometimes alternatively be attributable to some as-yet unidentified processing issue going on in the device.

An effort should be undertaken to test full out the Odometer value obtained from an OBD-II device vs. the odometer value observed from the vehicle's dash. Because the odometer that can be observed visually is the actual measure of record, testing that is performed should be used to determine the true accuracy of the odometer obtained from an OBD-II device and then work with the device manufacturer to improve accuracy.

4.3 UDP SHORTCOMINGS

As described in Sections 2.3 and 2.4, three device messages play a role in the processing of road charges for a given trip: the receipt of a TRIP-START marks the beginning of the trip; the receipt of a series of GPS messages provide the waypoints traversed by the trip; and the receipt of the TRIP-END marks the completion of the trip.

In an ideal scenario, these trip-defining messages for any trip would be received by the road charge platform in that same order: TRIP-START; followed by multiple GPS messages; followed by a TRIP-END. Furthermore, these trip-defining messages would be received in their entirety, prior to the system's receipt of any device messages related to a subsequent trip; i.e., trip messages on behalf of two separate trips would not be interspersed. Lastly, trip-related messages would never get lost; i.e., all trip-related messages transmitted by the device would be received by the road charge platform.

However, utilization of the UDP communication protocol for the transmission of TRIP-START and TRIP-END messages can disrupt this ideal scenario in a variety of ways, in that neither the receipt of any given device-transmitted message, nor the order in which multiple messages are received, is ensured by the protocol.

4.3.1 Overlapping Trip Message-Sets

Envision a scenario in which a given vehicle makes back-to-back trips; we'll call them Trip #1 and Trip #2. As the TRIP-START and TRIP-END messages are transmitted via UDP, it's entirely possible for the road charge platform to receive the TRIP-START message for Trip #2, prior to receiving the TRIP-END message for Trip #1.

Because of the potential for a road charge platform to encounter such "overlapping" trip message-sets from the device, it's critical that the platform be designed in such a way that the processing to be performed on behalf of multiple trips can be executed in parallel. With respect to the above scenario, the fact that Trip #2 messages begin to be received by the road charge platform before all of the Trip #1 messages have been received should in no way impact the proper differentiation of mileage and assessment of road charges for Trip #1.

By way of example, the pilot system accomplished this by caching all incoming TRIP-START. GPS, and TRIP-END messages immediately upon receipt, and configuring a separate "event-listening" process to specifically look for the receipt of TRIP-END messages. Upon receipt of a TRIP-END, a multi-threaded "trip record creation" process was then triggered on behalf of the device-assigned trip number contained in the TRIP-END to gather all cached messages tagged with the matching trip number, extract all road-charge-related information for the trip, and persist that information into a unique trip record.

4.3.2 Tardy GPS Messages

Due to the idiosyncrasies of UDP, it's also possible that the TRIP-END message for a given trip can be received by the road charge system prior to the system's receipt of all GPS messages associated with that same trip. In other words, the system should not assume that it's safe to proceed with the road charge processing of a given trip immediately upon the receipt of the TRIP-END for that trip, as to do so would run the risk of some of the trip's waypoints being overlooked in that process.

This phenomenon was observed for 1,479 (1%) of the trips collected for the six-month pilot. Figure 5 presents a breakdown of these trips by the amount of time that elapsed between the time of receipt of the TRIP-END and the receipt of the last GPS message for the trip.

Figure 5: Delay of Receipt of Last Trip GPS Message, Past TRIP-END Message

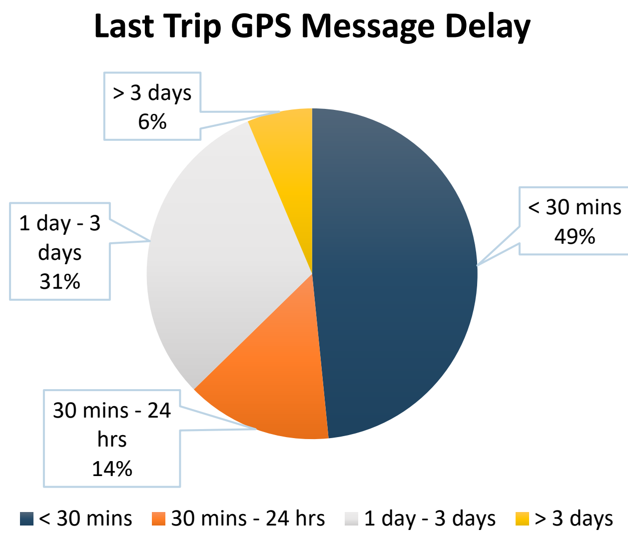

To mitigate the risk of "losing" these tardy trip waypoints, it's recommended that a road charge system incorporate a delay into the timing of trip processing, past the receipt of a TRIP-END. In other words, the system could essentially start a timer upon the receipt of a TRIP-END message and wait until that timer reaches a configurable threshold before undertaking the process of road charges for that trip. It should be noted that length of delay used, however, needs to be traded off against the extent to which that delay impacts the timely exposure of trip charges to the participants via the system's UI, as well as the timely delivery of monthly summary reports.

By way of example, this pilot utilized a delay of 30 minutes post-receipt of the TRIP-END, before trip-processing was initiated. In hindsight, Figure 5 reveals that this duration enabled the pilot system to completely capture tardy GPS messages in 49% of the instances in which they occurred.

4.3.3 Missing TRIP-ENDs

Although it's a relatively rare occurrence, it's also possible that a trip's TRIP-END message is never received by the road charge system due to UDP's lack of guaranteed message delivery. For this pilot, TRIP-END messages were never received for 0.02% of the trips collected over a six-month period.

As described in the sections above, the receipt of the TRIP-END message sets in motion the sequence of actions undertaken by a road charge system to process the trip. Due to the possibility that the TRIP-END will never be received, consideration should be given to a road charge system incorporating functionality to actively monitor the gaps in time between the receipt of GPS messages for the trip. Whenever that gap is observed to exceed a configurable "max TRIP-END wait time" threshold, the system should assume the trip has been completed and initiate the road charge processing for the trip, exactly as if the TRIP-END had been received.

Upon examining the 99% of pilot trips for which all GPS messages were successfully received by the system prior to the receipt of the TRIP-END (i.e., the complement to the trips addressed by Section 4.3.2), the project team found that the longest gap in time between receipt of the last GPS message and the receipt of the TRIP-END was a little over 3 minutes. It would therefore make sense for a road charge system to initially configure the above-referenced "max TRIP-END wait time" threshold to a value of approximately 5 minutes, and then monitor the number of times such a setting leads to falsely assumed completed trips and increment this configured threshold accordingly.

4.3.4 UDP Shortcomings: Recommendations

The Pilot findings and related recommendations associated with the mitigation of UDP-transmitted device messages that are detailed above may be summarized as follows:

- A road charge system should only use TCP as the transfer protocol to ensure messages always get to their intended location
- A road charge system should be designed in such a way that the processing to be performed on behalf of multiple device-reported trips can be performed in parallel.
- A road charge system should incorporate a configurable delay into the timing of trip processing, past the receipt of a TRIP-END, before such processing is initiated.
- A road charge system should also incorporate functionality to actively monitor the gaps in time between the receipt of GPS messages for the trip, and should that gap exceed a configurable threshold, proceed with the road charge processing for the trip as if the TRIP-END had in fact been received.

4.4 IDENTIFICATION OF MISSING MILES

In the interest of ensuring compliance, any operational road charge program will necessarily need to include some combination of process and/or functionality to pursue the identification of "missing miles"; i.e., mileage that for any reason is unaccounted for by the road charge system. For programs that offer an OBD-II plug-in device as a reporting option, missing miles would result from a participant either intentionally (to temporarily evade road charge) or unintentionally (e.g., hitting the device with their knee) unplugging their device, taking several trips, and then plugging their device back in. In such a scenario, the distance driven during the trips while the device was unplugged would constitute missing, and therefore untaxed, mileage.

4.4.1 Leveraging the Device to Find Missing Miles

Because device messages provide the pilot system with odometer values, it's reasonable to assume that those values might be leveraged by the system to look for missing miles.

For example, both the CONNECT and the DISCONNECT messages contain odometer values at the time that those events occur. Since the disconnection of the device from the OBD port triggers the transmission of a DISCONNECT and a subsequent re-connection of the device would result in the transmission of a CONNECT, it would make sense for the system to compare the odometer values in those two messages to determine if a gap exists, which could very well be associated with missing miles for that vehicle. However, as discussed in Section 4.2.2, the odometer value in CONNECT messages is unfortunately unreliable.

Similarly, since the TRIP-START and TRIP-END message both contain odometer values, it would seemingly make sense for a road charge system to compare the odometer value in each TRIP-START message to the value in the previous TRIP-END message to determine if a non-trivial gap might exist, in which case there could potentially be missing miles between the two trips. Again, as discussed in Section 4.2.2, the odometer value in TRIP-STARTs is also unreliable. We also learned in Section 4.2.3 that this pilot identified an alarmingly high number of such inter-trip odometer gaps; a number sufficiently high to cast significant doubt on the supposition that such gaps might actually represent missing mile scenarios.

4.4.2 Missing Miles: Recommendations

As long as a road charge platform has rigorous reasonableness checks in place for the device-reported odometer values it's receiving, either of the two above approaches might be used by the system to assist in the initial identification of "candidate" instances of missing miles.

In light of the above concerns, however, it's most likely inevitable that true compliance can only be achieved by the program mandating periodic "non-tech" (i.e., non-device-related) verifications of the odometer, for every vehicle utilizing a plug-in device to report mileage to the system. Such periodic non-tech verification might be accomplished via a participant upload of an odometer photo, or perhaps by visual verification of the odometer by a government entity (e.g., DMV) or certified third-party (e.g., gas stations).

Additionally, the system should maintain an independent (i.e., independent of the device-reported odometer values) history of these verified odometer readings for each vehicle, such that when each successive reading is obtained, it can be compared to the sum of the previous reading and the device-reported distance the vehicle has travelled since the previous reading was obtained.

4.5 UNDIFFERENTIABLE MILEAGE

As described in Section 2.4, the process by which the pilot system differentiated mileage by road type entailed the computational assessment of every trip waypoint against a Caltrans-approved map-set to determine whether the waypoint was located inside or outside the state of California … and for those waypoints found to be in-state, to subsequently also determine if it was located on a public road, a private road, or a road that was located on tribal land. The pilot did encounter several scenarios in which the mileage reported by the device for a given trip could not be successfully differentiated.

4.5.1 Data Pipeline Issues

As reported in Section 4.3, the utilization of UDP for the transmission of device messages led to several "data pipeline" issues that sometimes interfered with the pilot system's ability to properly capture all GPS waypoints associated with a trip. As described, the pilot system employed several strategies to mitigate the impact of these issues. Nonetheless, on rare occasions these issues with the flow of device data prevented the system from being able to differentiate any of the mileage for a trip. This phenomenon was observed for 877 out of a total of 171,000+ Pilot trips, or 0.05% of all collected trips, resulting in just 0.2% of all collected mileage being undifferentiable due to data pipeline issues.

4.5.2 Issues Related to Map-Sets & GPS

Beyond trips impacted by issues related to the device data pipeline, there were other pilot trips encountered for which the locational coordinates reported for a waypoint (or set of waypoints) could not be reconciled within the context of the specific version of the map-set that was being used at the time. In such situations, the waypoint was tagged as being undifferentiable, and the calculated distance between such waypoints was categorized as undifferentiable.

Table 6 shows a breakdown by pilot month of the trips having waypoints that failed the differentiation process, as well as the associated mileage that was therefore determined to be undifferentiable. Again, this was a relatively rare occurrence, with only 0.7% of the Pilot's collected trips involving at least one waypoint that failed differentiation, resulting in just 0.2% of all collected mileage being undifferentiable due to this issue.

Table 6: Trips and Mileage Due to Differentiation Failure, by Pilot Month

Pilot Month	Trips Involving Waypoints Failing Differentiation	Resulting Undifferentiable Mileage
April	0	0.0
May	1,233	3,151.8
June	22	96.5
July	15	109.8
August	7	0.9
September	0	0.0
TOTALS	**1,277**	**3,359.0**

The inability to reconcile a given set of GPS coordinates against a given map-set could be tied to an issue with the values of the coordinates as reported by the device or, alternatively, to an issue with the particular map-set being referenced. With regard to the former potential cause, at least a portion of the observed differentiation failures may very well be attributable to GPS fix error. This is an unavoidable byproduct of the fact that GPS receivers utilize radio waves to determine positional coordinates, and radio waves are subject to interference from a variety of sources (e.g., atmospheric phenomena, geographic and man-made obstructions, radio transmissions in nearby bands, etc.)

However, examination of the temporal pattern that is apparent in Table 6 would indicate that another factor is in play. This issue did not occur at all in the first month; it spiked in May; significantly dropped off over the next 3 months; and then once again did not occur at all in the last month of the pilot. The system employed several iterations of map-set versions, against which the GPS points were differentiated, over the course of the six-month pilot. The sudden onset of this issue in May indicate that a particular map-set version that was deployed in May was somewhat defective, in that trips began to be encountered for which waypoint coordinates could not be reconciled. And this is actually the case. In May, an update

As the maps-sets continued to evolve, subsequent iterations apparently "cleaned-up" whatever defects were introduced in May, to the extent that all trip waypoints in the last month of the pilot (September) were successfully differentiated.

The map-sets that were used (from March 14, 2023) for this pilot were as follows:

- State boundaries: https://www.census.gov/geographies/mapping-files/time-series/geo/cartographic-boundary.html.
- CA road network (all): https://caltrans-gis.dot.ca.gov/portal/home/item.html?id=836c3a3e2e7f40b8b36b24addf01cf14.

- CA road network (public): https://dot.ca.gov/programs/research-innovation-system-information/office-of-highway-system-information-performance/functional-classification.
- CA land ownership: https://gis.data.ca.gov/datasets/CALFIRE-Forestry::california-land-ownership/about.
- Tribal land boundaries: https://www.census.gov/cgi-bin/geo/shapefiles/index.php?year=2022&layergroup=American+Indian+Area+Geography.

From March 23, 2023 there was an update to the tribal boundary map. The update was triggered by the inaccuracy of the tribal land boundary map, and some test trips reportedly happening in tribal land were not correctly labelled. The updated map-set included the following update.

- The tribal land boundary map was updated to https://www.blm.gov/services/geospatial/GISData.

On May 8th, 2023, an update was triggered by an issue that the state boundary map did not cover water areas, causing trips made on bridges or along the coastline to be labeled incorrectly. The updated map-set included the following update.

- The CA state boundary map was updated to https://data.ca.gov/dataset/ca-geographic-boundaries. The boundary map for other states remains unchanged.

On May 10th, 2023, an update was triggered by the lack of map coverage outside of U.S. and incorrect labelling of trips made in Canada or Mexico. The new map-set included the following update.

- The boundary map for other U.S. states (other than CA), Canada provinces and Mexico states is added to the map-set from https://www.sciencebase.gov/catalog/item/4fb555ebe4b04cb937751db9.

4.5.3 Undifferentiable Mileage: Recommendations

Adoption by a road charge system of the recommendations captured at the end of Section 4.3 should help to minimize mileage that becomes undifferentiable by virtue of UDP-related issues with the flow of device data into the system (i.e., caused by device data pipeline issues).

Additionally, the system should also proactively/diagnostically monitor the frequency with which trip waypoints are failing the system's differentiation process, as a potential indicator of the need to further refine map-set files that are being used by the system.

4.6 ALTERING FREQUENCY OF LOCATIONAL REPORTING

As reported in Section 2.3, the OBD-II plug-in device incorporates a GPS receiver that determines the vehicle's locational coordinates at waypoints traversed during a trip. The frequency with which this locational assessment occurs during a trip is a configurable parameter, referenced in the following content as "GPS frequency".

4.6.1 Background

At the Pilot's outset, the Pilot devices were to be pre-configured with a GPS frequency of 60 "waypoint captures" per minute (i.e., an inter-waypoint interval of one second). Sixty days into the Pilot, however, the project team discovered that the device vendor had inadvertently set the GPS frequency for all pilot devices to only 12 waypoint captures per minute (5 seconds between waypoints). A process was therefore undertaken to alter every deployed device's GSP frequency to the proper value of one waypoint every second, but the time required to fully implement and test this reconfiguration unfortunately amounted to an additional 60 days.

As a result, GPS information for trip waypoints was collected at five-second intervals over the first four months of the Pilot, and was collected at one-second intervals for the last two months.

4.6.2 Storage Implications of Increased GPS Frequency

The most obvious impact of this mid-pilot change in GPS frequency was the increase in storage required to persist the GPS locational information for trip waypoints over the last two months of the Pilot. For example, consider any two months of the pilot for which waypoints were collected at five-second intervals during the former month and at one-second intervals during the latter month. If we were to assume that the total mileage collected for all participants, as well as the average speed driven by all participants, was exactly the same across both months, then we would expect to see a five-fold increase in the number of waypoints that would need to be persisted for the latter month.

To examine this further, Table 7 shows a breakdown by month of all the waypoints that were stored and the mileage captured during the Pilot. The table provides further insight by also deriving the following information for each month:

- **Waypoints per Mile**: Average number of waypoints that were stored for each captured mile during the month.
  (Calculation: Waypoints / Miles)
- **Seconds per Mile**: Given the amount of time between each waypoint capture (i.e., the inter-waypoint time setting for the month), this calculation is the average amount of time taken to travel from the first to the last waypoint in a one-mile span during the month. For example, if the waypoints were captured every five seconds during a month, and on average 20 waypoints were stored for each mile during that month, then on average it took 100 seconds to drive one mile during that month.
  *(Calculation: Waypoints/Mile * GPS Frequency)*

- **Average MPH**: Knowing the average seconds per mile (the previous column), we can take its reciprocal to determine the average miles travelled per second, and then multiply by 3,600 to determine the average miles travelled per hour for the month. *(Calculation: { 1 / Seconds per Mile } * 60 secs per min * 60 mins per hr)*

Table 7: Waypoints Stored and Mileage Collected by Pilot Month

Pilot Month	GPS Frequency (seconds)	Waypoints Stored	Miles Recorded	Waypoints per Mile	Seconds per Mile	Average MPH
APR	5	2,278,627	97,143	23.5	117.3	30.7
MAY	5	6,928,234	299,175	23.2	115.8	31.1
JUN	5	6,776,618	283,785	23.9	119.4	30.2
JUL	5	6,775,084	280,655	24.1	120.7	29.8
AUG	1	32,779,386	289,621	113.2	113.2	31.8
SEP	1	29,824,168	272,617	109.4	109.4	32.9

In Table 7, the waypoints-per-mile values for the first four months (all with waypoints captured at five-second intervals) averaged out to a value of 23.7, and the corresponding values for the last two months (during which the waypoint-interval was changed to one second) averaged out to 111.3, representing an increase by a factor of 4.7. We can also see that on average, while the average speed in the last column was relatively consistent across all six months (as would be expected), the participants did drive slightly faster on average during the last two months of the pilot, which helps to explain why the observed increase in waypoint storage did not more closely approach a factor of 5.0. This is attributable to the fact that for two devices configured with identical GPS frequencies, travelling along an identical one-mile segment of road, the device travelling at a relatively higher speed will collect fewer waypoints than the one driving at a relatively slower speed.

4.6.3 GPS Frequency: Recommendations

In light of the positive correlation between GPS frequency and both storage capacity and processing time (i.e., increasing the number of waypoints captured every minute will result in the collection of additional waypoints, which raises both storage requirements and processing time needed for differentiation), further research is warranted to investigate an "optimal" GPS frequency for the device. Decreasing GPS frequency would involve longer intervals between determination of waypoint locations, resulting in a beneficial reduction in both storage and processing time for a road charge system. However, the accompanying reduction in waypoint "granularity" (i.e., fewer waypoints collected for a given trip, with those waypoints being spread out more) will likely have at least some level of deleterious impact upon the system's accuracy in identifying both road type and state line boundaries, potentially resulting in less accurate differentiation of the mileage collected by the system. The goal of the recommended research, then, would be to identify the longest interval between the device's capture of waypoint location (i.e., the lowest GPS frequency) that would be accompanied by an "acceptable" loss of accuracy in the system's mileage differentiation process.

For example, three vehicles might be provisioned with OBD-II plug-in devices, with the GPS frequency of those devices configured such that waypoints are captured at one-second, five-second, and 30-second intervals, respectively. Then those three vehicles could be driven multiple times over the same route, with well-defined segments of the route comprising public, private, and tribal roadways, and with the length of each of those road segments being independently and precisely measured. An analysis of the resulting differentiated trip records collected by the road charge system might then be undertaken to more thoroughly investigate the relationship between decreased device GPS frequency (i.e., longer inter-waypoint intervals) and the accompanying loss of accuracy in the differentiation process.

4.7 PROVIDING TRIP MAP VIEWS

All participants in this particular Pilot were required to install a GPS-enabled OBD-II plug-in device, in accordance with the Pilot's requirement that all mileage collected for the Pilot be differentiated by in-state versus out-of-state, as well as by road type. Within an operational road charge program, however, privacy concerns on the part of participants will most likely necessitate the offering of alternative mileage reporting options that do not convey any locational information for the trips captured. For those participants in such future programs that still opt into a GPS-reporting option, their primary motivation in doing so will obviously be to avoid paying a road charge when they're not driving on roadways maintained by the state, and it can therefore be assumed that they will likely have high interest in reviewing the operational program's ability to accurately differentiate their specific trips.

It would seem that the best approach to facilitate a participant's "audit" of a road charge program's differentiation would be for the system's customer portal to make available a map view of that participant's trips, that visually distinguishes the segments of each trip that were and were not assessed a road charge. However, a system's provisioning of such a map view capability is accompanied by a surprising number of non-trivial costs and considerations:

- Although GPS drift (i.e., the difference between the vehicle's actual location and the location as indicated by GPS) introduces only a relatively minor inaccuracy in the device's capture of each waypoint within a trip, a map view approach that visually renders the trip by simply drawing straight lines between the (slightly inaccurate) waypoints will appear to be highly "jaggy". During the first four months of the pilot, the GPS capture frequency was set at 5 seconds. When the GPS points that were captured were drawn directly onto the map, it did cause some trips to show up as "jaggy". When the frequency was increased to every 1 second, there was an expectation that the "jaggy" lines would be reduced. The reality was that they were, but not by much.

- Therefore, a smoothing algorithm was employed to filter out the "noise" introduced by the inaccuracies due to GPS drift, before the trip is plotted on the participant's map view. Such smoothing algorithms have been found to be computationally expensive.

- There is a subscription cost associated with the third-party mapping services that are required to visually render a map view of trips. Some of these services will suspend the subscription cost, if the service provider is granted permission to use the plotted data for their own purposes, but such a cost-saving option is not viable for government agencies administering a road charge program.

- The provisioning of a map view for participants to visually audit their trips will shine a spotlight on the unavoidable anomalies associated with GPS positional tracking, and will undoubtedly be accompanied by an increase in the number of participant inquiries that must be fielded by a program's customer support team, which will translate to higher customer support costs for the program. The increase in inquiries is dependent on the number of participants that actually take the time to view their trips, which typically is fairly low.

4.7.1 Trip Map Views: Recommendations

If future road charge programs offer the ability for participants to review a map view of their trips within the system's customer portal, it is recommended that strong consideration be given to positioning this capability as a value-add, additional-fee premium feature, to offset the non-trivial costs incurred by the program in order to provide the feature.

5. Inherent Limitations of Devices

Even after a road charge system accounts for the findings described above by implementing each related recommendation, geolocation devices have inherent limitations due to their history and design.

5.1 OBD-II DEVICE DESIGN

As described in Section 2 the OBD-II port has been an industry standard of all vehicles since 1996. Since the standardization of the OBD-II port itself, the ‘language’ that the port uses – called a protocol – has gone through numerous changes. In the beginning, each manufacturer decided what protocol that it wanted to use. This resulted in many different protocols existing at the same time. In 2012, a new protocol was introduced that has since been used by all manufacturers in the United States. This helped companies using an OBD-II device for obtaining vehicle data to reduce costs by no longer having to create multiple versions of an OBD-II device to read and interpret the many types of protocols in existence.

The OBD-II port was originally designed to be used by automotive experts to determine the effectiveness of the vehicle’s emissions, but that didn’t stop other industries from leveraging the OBD-II port and the data it makes available for other purposes. The first of these was the insurance industry. In the mid-2000s, some insurance companies had the idea of plugging an OBD-II device into the port to use the data collected to track driving behavior. The driving behavior information was used to create a pricing model that would allow insurance companies to more closely align actual driving behavior to a risk profile. It was determined that driving behavior is over 2x more effective in determining risk than the next closest risk factor.

In addition to the insurance industry, the fleet management industry also found that an OBD-II device was a highly effective tool in providing a valuable service to a market that had stagnated. Prior to OBD-II devices, fleets managers who wanted to track their vehicles were paying $40 to $50 (or more) per month per vehicle to keep tabs on where their vehicles were and what they were doing. A few enterprising companies decided that the market was ripe for disruption and leveraged after-market OBD-II devices to bring fleet management services down to $15 to $25 per vehicle per month. This opened up tracking services to a whole new set of customers who had previously been priced out of the market.

The third industry to use the OBD-II port for a purpose beyond emissions testing was road usage charging (RUC). While the insurance and fleet management industries spent time with the OBD-II device manufacturers to develop hardware and firmware that was specific to their needs, the RUC industry just leveraged devices that were designed for the insurance and fleet management market. Although there are overlaps in the data collected between the three distinct industries, the precision requirements of the data collected between the first two industries and RUC is vastly different. Insurance and fleet management business needs can be fully met with data that is imprecise and prone to gaps. This is not the case with road usage charging. Mileage needs to be at least 95% accurate; if location information is collected, that information needs to be just as precise.

5.2 DEVICE COSTS

One factor that must be accounted for in administering a tax is the cost of administering the collection, processing and management of the tax. A RUC pilot or program is no different. For a road charge program that leverages an aftermarket OBD-II device, the associated costs of the device must be accounted for. Below is a list of all the costs associated with a device.

5.2.1 Device

The device that was used for this pilot is described in detail in Section 2.2 of this document. The cost of this type of device is dependent on the features and capacities of the device, as well as the volume of devices purchased. The cost of the devices used for this pilot was $95, which is near the upper range of what these devices typically cost (from $37 to $125).

5.2.2 Data Transmission

To use the data obtained from a device that is plugged into the vehicle, the device must be able to transmit that data to a back office for processing. There are two methods of transmission: cellular (device -> back office) and Bluetooth (device -> participant smartphone -> back office). You would only use one method at a time. Each method has benefits and drawbacks, and each has different cost considerations.

Cellular Costs

There are various costs associated with using cellular technology as the means for transmitting data from the device to the back office.

- **Cellular chip** – The chip is required to be included in the build of the device. Generally speaking, cellular chips at low volumes (<1 million per year) will range from $30 to $40 per chip.
- **SIM** – Not all cellular chips require a SIM any more as the technology is now handled via software or firmware. If a SIM is required, it can cost between $0.50 to $1.00 per device.
- **Data plan activation** – All wireless carriers require a SIM to be activated on their network. The range of activation fees is $0 to $30 per device.
- **Wireless data plan** – All wireless carriers require every single node (SIM) on the network to have an individual wireless plan even if all the SIMs are under the same account. Wireless plans have a fairly large range in terms of costs, but generally they range from $.045 to $1.00 per month for up to 2MB of data.
- **Wireless plan management** – Wireless plans are grouped into tiers. Data tiers range from Kbytes to Gbytes. An example of monthly ranges are: 0 – 1Mb, 1 – 3Mb, 3 – 5Mb, 5+Mb. For RUC programs, the 0 – 1Mb tier is usually large enough to handle typical driving patterns. As long as the device sends data volumes in the agreed upon tier or less, the monthly amount remains static to the agreement in place. However, if the device sends more than the agreed-

upon amount, an overage amount per extra kilobyte is added to the base amount. There is a small range between the base amount plus overages and the cost for the next tier where it makes sense to accept the overages. However, there is an exact threshold in which moving to the next tier is considerably cheaper than accepting the base plus overages. Managing each wireless plan to optimize which tier it should use for pricing each month is critical in ensuring that no extra expenses are incurred that are unnecessary. Wireless plan management usually requires a fraction of a resource to build the optimization models and then a fraction of a resource to review and oversee the running of the models to ensure that optimization continues to work as designed. For the pilot, no wireless plan management was performed, but in a real program this cost would need to be accounted for. Wireless plan management can cost between 20% – 30% of a technology vendor FTE (Full Time Equivalent) per year. In dollars, that would be between $10,000 and $60,000 per year. (For this report FTEs costs (salary plus benefits) are expected to be between $50,000 to $200,000 per year.)

Bluetooth Costs

Even though Bluetooth was not used by this Pilot as the data transfer mechanism, it is good to mention it here for comparison to get a complete picture of the associated costs for plug-in devices. There are various costs associated with using Bluetooth technology as the means for transmitting data from the device to the back office.

- **Bluetooth chip** – The chip is required to be included in the build of the device. Bluetooth chips at low volumes (<1MM per year) can range from $0.50 to $2.80 per chip.
- **Bluetooth data management** – Unlike a cellular enabled device that can send the data collected from a vehicle directly to the back office, a Bluetooth enabled device must use a customer’s smartphone. This means that a mobile app is required that can act as the data collector and processor of the raw data from the device and then use the customer’s wireless plan to push the data to the back office. A small team of developers that have expertise in writing native apps for Android phones and for Apple iOS phones is required. Between 1 to 3 FTE’s within the technology vendor’s staff would be needed for Bluetooth development and support.

5.2.3 Logistics

Because the devices are a tangible object, they must be physically moved from the location where they are manufactured to the home of a person who will install the device into their vehicle. The movement of the device can be handled in several ways, but the typical method is to use a company that specializes in logistics. These companies have the expertise in accepting large quantities of items and distributing them using the most efficient means directly to the end location where they are needed. Below are capabilities needed by a logistics company and their associated costs:

- **Warehousing** – *accepting and managing individual devices from the manufacturer as well as returns from individuals*. For this Pilot, warehousing was handled internally and thus did not have a separate charge. Typically, warehousing costs are included in the handling per shipment charges for small items or items that will not be in inventory for a long time.

- **Device Management** – *tracking exactly where a device is at all times and manually associating a device to the account it will be shipped to using the device management tool (typically a web-based interface).* For this Pilot, device management was handled internally and thus did not have a separate charge. Typically, device management costs are included in the handling per shipment charges.
- **Receiving** – *receiving returned devices, unpackaging them and getting them ready for refurbishment*. For this Pilot, all receiving was done internally so there was not a separate charge for this. Typically, receiving costs are included in the handling per shipment charges.
- **Refurbishing** – *flashing the device to purge all internal data and testing it to ensure it is ready for re-use*. This is done using special equipment from the manufacturer or it is handled by the manufacturer itself, in which case they also handle the receiving. For this Pilot, there was no need for refurbishing as no devices were needed to be reused. Typically, a manufacturer will charge between $2.50 and $4.50 for refurbishment and then ship the device back to the account manager handling the devices.
- **Shipping** – *the physical movement of the device from the warehouse to the participant*. This requires access working with one of the major shippers for sending out devices to participants and then receiving them when the participant leaves the program. Major carriers that could be used include USPS, UPS and FedEx. In general, FedEx is the cheapest when sending multiple or bulk items to a commercial address, and USPS is the cheapest when sending individual items to a residential address. Shipping costs range from $3.80 to $7.90 per item for normal delivery. Overnight shipping can cost over $35 per item. For this Pilot, shipping costs averaged close to $3.80 per device for outbound and return shipping.
- **Packaging** – *materials (boxes or bags, and labels) used to ensure a safe and damage free shipment*. Packaging materials can run from an inexpensive non-branded padded envelop ($0.24) to a fairly expensive custom designed box and shipping envelop ($6.50). For this Pilot, non-branded envelopes ($0.24) and 4x6 labels ($0.032) were used to ship devices. Envelopes with pre-paid labels on them were also sent to participants to use to ship devices back, which incurred an extra $0.48 and $0.064 per device.
- **Collateral** – *information on how to install the device and where to find the OBD-II port*. Collateral can run from an inexpensive document printed on an office printer on 8.5 x 11 printer paper ($0.002) to an expensive custom designed and printed and cut set of instructions ($3.25). For this Pilot, a simple one-sided document was printed at a local print shop ($0.25).
- **Collateral Management** – *tracking of inventory of packaging materials and collateral and ordering more if and when the levels dip below the order threshold*. Collateral management is usually included in the handling per shipment charges.
- **Handling** – *packaging up the device into a box, putting the box into a shipping package, creating and printing labels manually or automatically, and getting the package to the shipper*. Cost of handling usually starts with a base annual fee regardless of number of items. This base cost can range from $5,000 to $15,000 per year. In addition to the annual base cost, there is a per-item cost that covers warehousing, device management, collateral management and handling. This can range from $1.80 to over $5.50 per device depending on volumes. Once volumes become large and consistent enough, base costs are waived. For this Pilot, handling was all done internally, so there were no direct handling charges.

Table 8: Device Logistics Costs Overview

Category	Frequency	Low	High	Cost per Device in this Pilot
Warehousing	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Device Management	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Receiving	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Refurbishment	Per device per RMA	$2.50	$4.50	n/a (did not do it for this pilot)
Shipping	Per device per shipment	$3.80	$7.90	$3.80
Packaging	Per device per shipment	$0.24	$6.50	$0.24
Labels	Per device per shipment	$0.032	$0.040	$0.032
Collateral	Per device per shipment	$0.002	$3.25	$0.25
Collateral Management	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Handling	Annual base + Per device	$5,000 + $1.80 p/d	$15,000 + $5.50 p/d	n/a (handled internally – not treated as a separate cost)

Table 9: Device Logistics Costs for this Pilot

Category	Cost per device in this Pilot	Frequency of cost	Logistics costs for pilot
Shipping	$3.80	3	$11.40
Packaging	$0.24	3	$0.71
Labels	$0.032	3	$0.096
Collateral	$0.25	1	$0.25
Total per-device logistics costs for Pilot ->			$12.37

5.2.4 Gateway

For data from devices to be received in a central data repository, the data needs to be captured and processed. The combination of hardware and software functionality that provides this service is called a Gateway. Gateways accept and process millions of messages at an extremely high frequency. They are load balanced and can automatically adjust and scale their processing capabilities higher or lower dynamically based on message load and frequency. Because of the complexity of a Gateway and the high availability requirement, Gateways tend to be quite expensive and costly to maintain. Building a Gateway can take between $250,000 to $1 million and at least two full-time resources to maintain.

For this Pilot, the Gateway was provided by the OBD-II device manufacturer at a cost of $4.00 per month per device, and a one-time fee of $17,876.50 for technical support for 12 months.

5.2.5 Hosting

After data has been processed by the Gateway, it must be stored. Storage can take multiple forms based on the design of the system, but in general for road usage charging it will be split into two primary storage areas: raw data and processed data.

Raw Data

The raw data coming from the devices is identified and stored in its raw format into logical groupings. For this Pilot, there was only one logical grouping as all the data pertained to the same pilot. But in a real-world scenario where the Gateway is handling multiple businesses, pilots or programs, each pilot will have its own group of raw data. Data is stored in its raw format so that if re-processing needs to take place, it can be done from the point of time at which it was captured in its original format. For this Pilot, the project team used Amazon S3, which is a cloud-based storage service that is efficient and fast. 581.7GB of storage was used for six months at a total cost of $15.08 dollars, which amounts to about $2.51 per month and about $0.008 per vehicle per month. Total storage included all travel records during the pilot. Records were never purged, but in a normal program there would an algorithm to purge old data that was past the data retention date.

Processed Data

This data is in the format where it can actually be used for operations and reporting. The processed data serves as the source for:

- Customer portal (web and mobile)
- Monthly statements
- Customer support portal
- Accounting and finance
- Agency reports and extracts
- Analytics

The storage of the processed data can take many formats, but generally falls into two categories: structured data and unstructured data. Both have their benefits in terms of function, performance and ease of use. For this Pilot, the project team used Snowflake, which is a cloud-based storage system that includes database and processing capabilities that allowed us to store the data relationally. 674.9 GB of data was stored for 6 months at a total storage cost of $18.52, which amounts to about $3.08 per month and $0.01 per vehicle per month. Total Snowflake processing costs for the 6 months was $14,996.93, which amounts to about $2,499.48 per month and $8.83 per vehicle per month.

Snowflake calculates processing costs by "credits". Low-impact processing that use few compute cycles and run quickly incur fewer credits, and high-impact processing that take large amounts of compute cycles incur many credits. This Pilot used 4,100 Snowflake credits for the six-month pilot, which is an average of 683.33 credits per month. This computes to approximately $3.66 per credit (i.e., $2,499.48 monthly processing cost, divided by 683.33 in monthly processing credits).

A sidenote about the use of Snowflake: It could be considered a bit of an overkill for a road charge type application. Amazon S3 would work just fine and cost considerably less. In order to quantify exactly how much less would require some side-by-side comparison.

5.2.6 Compliance

One of the major components of a live RUC program will be compliance. Although not incorporated into this Pilot, it is important that compliance is mentioned in this report as a cost of a program that includes plug-in devices, as the ongoing support costs for handling enforcement and fraud detection are non-trivial.

As pilots transition from purely voluntary initiatives into mandatory programs, there will be a large increase in the number of participants that will do whatever it takes to evade road charges. Resources would need to be dedicated to pursuing fraud as well as the systems needed to host and process fraud detection algorithms that need to be continually monitored and updated.

5.2.7 Cost Charts for Devices

Table 10 shows a high-level summary of all the associated costs for plug-in devices in a pilot or program with low volumes (less than 50,000 devices).

Table 10: Components of Device Costs

Category	Frequency	Low	High	Cost per Device in this Pilot
Device	1x	$37.00	$125.00	$95.00
SIM	1x	$0.50	$1.00	n/a (Included by Device manufacturer)
Data Plan Activation	1x	$0	$30.00	n/a (included by Device manufacturer)
Wireless data plan	Per device monthly	$0.045	$1.00	$.90
Wireless plan management	Monthly	$1666 (20% FTE / 12 mos)	$2,500 (30% FTE / 12 mos)	n/a (did not do it for this Pilot)
Warehousing	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Device Management	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Receiving	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Refurbishment	Per device per RMA	$2.50	$4.50	n/a (did not do it for this pilot)
Shipping	Per device per shipment	$3.80	$7.90	$3.80
Packaging	Per device per shipment	$0.24	$6.50	$0.24
Labels	Per device per shipment	$0.032	$0.040	$0.032
Collateral	Per device per shipment	$0.002	$3.25	$0.25
Collateral Management	Monthly	Covered in handling		n/a (handled internally – not treated as a separate cost)
Handling	Annual base + Per device	$5,000 + $1.80 p/d	$15,000 + $5.50 p/d	n/a (handled internally – not treated as a separate cost)
Gateway (if outsourced)	Annual Base + Per device	$17.6K + $2.50 p/d	$30K + $4.50 p/d	$9.26
Gateway (if handled in-house)	1x + annual	$250,000 + 2 FTEs	$1,000,000 + 3 FTEs	n/a (handled by Device manufacturer)
Hosting – raw data	Monthly	$0.021 per GB	$0.025 per GB	$0.026 per GB ($0.008 / device)

Hosting – processed data	Monthly	$0.025 per GB	$0.027 per GB	$0.027 per GB ($0.01 / device)
Processing data	Monthly per credit	$2.50	$4.10	$3.67 per credit ($8.83 / device)
Compliance	Annual	.25 of an FTE	5 FTEs	n/a (did not do it for this Pilot)

Table 11 breaks down the monthly per-device cost for this Pilot.

Table 11: Calculation of Per-Device Cost for this Pilot

Category	Cost per device in this Pilot	Frequency of cost	Six-month cost
Device	$95.00	1	$95.00
Wireless data plan	$0.90	6	$5.40
Shipping	$3.80	3	$11.40
Packaging	$0.24	3	$0.71
Labels	$0.032	3	$0.096
Collateral	$0.25	1	$0.25
Gateway (if outsourced – per device)	$9.26	6	$55.56
Hosting – raw data	$0.008	6	$0.048
Hosting – processed data	$0.01	6	$0.06
Processing data	$8.83	6	$52.98
Total per-device cost for 6 months ->			$221.50
Total monthly cost per device ->			$36.92

Notes: Total and monthly costs do not take into account resource costs for operational device management (i.e., once the device is installed) or any other overarching management expenses.

5.2.8 Cost Recommendations

A cost of $36.92 per device is typical of a small-scale pilot. It is not typical of a large-scale program that runs at high volumes. Device, gateway and processing costs all are relatively high at low volumes. At scale and with a product and system designed specifically for road charging, many of these costs would be reduced dramatically.

Considerations for reducing or spreading out costs:

- Allow vendors to charge for additional services that can be provided by the device beyond base road charge functions.
- Charge participants for the use of the device.
- Allow additional state and local based funding mechanisms to leverage the same data that is collected for a road charge program.

5.3 DEVICE MANAGEMENT

Plug-in devices are a tangible, manufactured electronic apparatus that require constant care and attention. The following is a list of activities that a business will need to undertake in order to effectively leverage an OBD-II plug-in device within their business model. Each of these activities requires a resource or resources dedicated to that activity.

Purchasing

OBD-II devices are a custom made-to-order electronic apparatus. This means that orders need to be submitted to the device manufacturer at least 4 months in advance of needing the device. In certain circumstances such as during COVID, orders needed to be placed 12 months in advance due to supply chain constraints. Purchasing needs to account for the services needed by the device to ensure the right type of device is ordered. During the ordering process a decision needs to be made on the wireless carrier that is to be used so that the devices can be manufactured with the correct wireless carrier technology installed.

As soon as an order is placed it needs to be tracked from order date to delivery date. After delivery is made, each device needs to be provisioned. This is the process of activating the device onto the wireless network that was chosen during the purchasing process. Usually, the provisioning process is handled by the manufacturer, but in certain circumstances they may allow it to be done by the company purchasing the devices.

Issues Management

Devices are highly complex and are full of interesting and exciting features. But they come with a plethora of management and oversight needs. Below is a list of sample needs that are typical of a plug-in device-based program. This list is in no way exhaustive and is only shown to illustrate the types of issues that will occur during a pilot or program.

Compatibility

While all vehicles manufactured since 1996 are required to have an OBD-II port (with the exception of a couple Tesla electric vehicle models), the OBD-II device is not always compatible with every vehicle. Incompatibilities can be due to idiosyncrasies with the protocol that the manufacturer used or other issues with how data on the OBD-II bus is requested and obtained. The result is that for such vehicle models, data cannot be obtained or relied upon, making certain individual vehicles or whole model years incompatible. These incompatibilities must be documented and tracked to ensure that additional vehicles that are similar are not allowed to participate in the program.

Interference

There are specific guidelines addressing where the OBD-II port is required to be located within the vehicle. However, this does not mean that all locations are ideal for the insertion of a permanently installed OBD-II device. There are instances of ports being behind access doors or located in such a way that the device will be in danger of being kicked. In these situations, the type of vehicle must be documented and tracked to inform the participant of the potential interference so that they can make an informed decision before participating.

Firmware Updates

Occasionally the device firmware needs to be updated to account for vehicle-specific issues or new issues that have been identified. The firmware update is accomplished by sending a Firmware Over-The-Air (FOTA) command to the device to inform the device to check in and download and install the new firmware. On these occasions there is a specific testing process that takes place between the device manufacturer and the vendor using the devices in the field. The vendor needs to dedicate a resource to ensure testing is complete prior to the FOTA being sent. Once the FOTA is sent, the vendor resource needs to monitor the FOTA to ensure that it was recognized, and that the update process started and completed without issues. If issues arise at any point during the FOTA process, the vendor resource needs to manage them to a satisfactory conclusion.

Data Collection Anomalies

As data is flowing from the vehicle to the device to the Gateway to the processed data storage, anomalies can occur. They can occur at any of the steps. There will be standard reasonableness checks on the data every step of the way. Any time an anomaly is detected, a resource in charge of anomalies is informed and begins the process to identify the source, reason, impact, and solution to rectify the anomaly.

Lost Devices

Devices are expensive and need to be tracked at all times. In most cases, the device is where it needs to be – in inventory, en route to a customer, or in the participant's vehicle. Sometimes though, the devices get lost. This can be due to a participant selling their vehicle and forgetting to remove the device. Or it can be due to an accident and the device gets removed. Or maybe it gets unplugged by someone curious as to what the device is and they just throw it away or don't plug it back in. Regardless of the reason for a lost device, it can happen and there needs to be a resource in charge of tracking down lost devices and taking appropriate action.

Upgrades to Technology

There are many highly complex chipsets within a plug-in device, to include Bluetooth, GPS, Wi-fi, Memory, CPU, Accelerometer, Gyroscope, Cellular modem chipsets. Each of these chipsets is designed to last a fairly lengthy time; however none of them lasts indefinitely and will need to be upgraded if they no longer work with technology they need to integrate with. The primary chipset that can fall out of synch with external integration points is the cellular chip. In the past 10 years, prevailing cellular technology in the U.S. has evolved from 2G to 3G to 4G. Each time the cellular technology has changed, the plug-in devices that use the old technology must be swapped out with devices that have the new technology as the older technology is not compatible. When a device swap is required, a resource needs to be in charge of managing the entire swap process which can be very customer intensive.

5.4 VEHICLE CONNECTIVITY

There are a number of known issues with how OBD-II devices interact with certain vehicle types. Below is a list of these known issues that should be taken into account when setting up and operating a pilot or program.

- **Battery Drain** – The device is plugged into the OBD-II port and stays there indefinitely. It does draw a small amount of power from the port even when the vehicle is turned off. This small amount is not enough to drain a healthy battery to the point of the vehicle not being able to be turned on, but statistics show that up to 30% of the vehicles on the road have batteries that are 3 years old or older which means that they are at or past their life expectancy. This means that even a small amount of extra draw on a battery that is not fully healthy could push it past its ability to start the vehicle. Another potential battery drain issue is with hybrid vehicles. The battery that they typically use is much smaller than a standard car battery and thus can get drained much quicker, especially if the battery is past its life expectancy.
- **Electric Vehicles With OBD-II Ports** – Most EVs have OBD-II ports, but because EV's do not have gasoline powered engines and thus no emissions, there is no emissions-related data available on the port. This left the option open for OEMs to put whatever data they wanted to be available on the port, or no data at all in some cases. For such vehicles, it is safest for a pilot/program to not expect any data from the port on an EV and instead configure the device to operate in a mode known as GPS-Only, which uses only GPS distance calculations to determine distance travelled by the EV.
- **Electric Vehicles Without OBD-II Ports** – There are some EVs that do not have an OBD-II port. These vehicles either need to install an aftermarket adapter that allows for an OBD-II device to be plugged in or they must use an alternative mileage reporting option. These vehicles need to be tracked and handled appropriately during the enrollment phase of a pilot.
- **GM Vehicles with OnStar** – GM's OnStar service is the oldest and most widely known and used telematics service. It has been in place since 1996 (the same year that the OBD-II standard was implemented in all vehicles). There is a known issue with the monthly vehicle diagnostics report that active OnStar customers can run. The issue is that if a device is plugged into the OBD-II port, the vehicle diagnostics report will not provide the data that is needed to complete the report. The vehicle "thinks" that there is a mechanic working on the vehicle (why else would there be a device in the OBD-II port?) and stops performing all of its diagnostic routines until the "mechanic" is done with their work. There is a simple work-around for this by having the participant remove the device prior to running the diagnostics report, but then they must be diligent (and potentially reminded) to plug the road charge device back in subsequent to the generation of the report.
- **Aftermarket Products** – Many people take great pride in customizing their vehicle to match how they drive as well as their personalities and make it completely their own. In many cases, this customization can include aftermarket devices and apparatus that need to tap into the vehicle's electronic harness in some form or fashion. In most cases that will have zero impact on the normal functionality of the OBD-II device. However, there will be some cases that the aftermarket device does impact the OBD-II device, and this type of negative interaction needs to be documented and managed appropriately.

- **BMW Security** – BMWs are an extremely popular vehicle, and this popularity makes them a favorite among car thieves. Thieves at some point in the past 10 years figured out how to bypass BMW's security system by plugging in a nefarious OBD-II device. To combat this BMW has written software into their vehicle's platform to activate the security alarm if an aftermarket OBD-II device is plugged into the port after the vehicle has been turned off. There is a firmware fix from the manufacturer that solves most of the instances of the alarm being set off, but not all. This type of negative vehicle-device interaction needs to be documented and managed appropriately.

- **Vehicle being serviced** – The OBD-II port was designed for mechanics to connect to in order to determine how well the emissions systems within a vehicle are functioning. For a mechanic to service a vehicle and obtain the emissions data, they need to connect a diagnostics tool to the port. If there is an after-market OBD-II device in the port, they will need to remove it. When they completed servicing the vehicle, the mechanic would need to plug the after-market device back into the vehicle. However, with all things that are handled manually, mistakes can be made and in this case, the device might remain somewhere on the mechanics work bench and subsequently be lost. There are algorithms in place that would quickly notice that the device was missing, but there would be a fairly lengthy process to determine when the device was lost, where it was likely lost and what actions need to be taken to remedy the situation. During all of this time of course, there would be no mileage being collected by the device. This type of vehicle service disruption needs to be documented and managed appropriately.

5.5 DATA COLLECTION

The primary purpose of the OBD-II device within a road charge pilot or program is the collection of differentiable mileage data. This requires that the device be plugged in at all times during normal operation indefinitely. It also requires that the GPS data that is collected and used must meet a certain level of accuracy for it to serve as a reliable source for differentiation. Several known factors can keep a device from providing the fidelity of data that is desired as well as impact the costs and customer experience. Below is a list of these considerations.

- **GPS Accuracy** – The accuracy of a GPS chip depends on various factors such as environmental conditions and the quality of the GPS receiver. In optimal conditions with a clear view of multiple satellites, modern GPS chips can achieve accuracies within a few meters. Consumer-grade GPS devices, such as the one in an OBD-II device, typically provide accuracies ranging from 3 to 10 meters under favorable circumstances. However, obstacles like tall buildings, dense foliage, vehicle type, or adverse weather conditions can impact accuracy, leading to potential errors in location determination. This translates into the need to set expectations on how precise the device-provided GPS readings are and what type of information can be determined from using those GPS readings. As an example, such inaccuracies might interfere with the ability to differentiate between a highway and a frontage road.
- **GPS Fix Issues** – In addition to general GPS accuracy concerns, there are times when the GPS sensor is not able to obtain a fix quickly enough in order to provide a realistic reading. This can cause algorithms that are expecting a specific level of precision to provide wildly

inaccurate positioning on a visual map. This translates to a customer potentially seeing one trip as it should and another trip showing their location in the middle of the ocean, which is not a favorable customer experience. The good news is that GPS Fix issues happen very rarely. During the Pilot, it happened in only 1,277 of the trips out of 121,456 which is only 1% of all trips, but 1,233 of those were due to a bad map set, so realistically only .036% of the trips had GPS Fix issues. But this case does illustrate the point that it can happen, and it needs to be monitored and addressed when it does.

- **GPS Collection Frequency** – The frequency at which GPS data is collected has an impact on cost and performance. More data does not always equal better. The higher the collection frequency the higher the need for storage and processing. Storage increased nearly five-fold when the Pilot switched from five-second to one-second GPS collection. Processing time also increased. Additional testing would be needed to determine the optimum collection frequency. One observable fact as the result of the increased frequency for this Pilot, was the time it took to 'draw' a trip on the visual map viewed by the participant. Similar-distance trips that were previously fully rendered in seconds, took at least a minute or two to render after the change in collection frequency. Long trips could take even longer. This extra time proved to be a frustration to some Pilot participants.
- **GPS Blind Spots** – Generally, people know that GPS involved a wireless capturing of data of a satellite location in orbit around the earth and a GPS receiver within a device (i.e., a phone or an OBD-II device). Information from multiple satellites is collected to determine exact location. Three or more satellites are needed for accurate location identification. Intuitively, people do recognize that the wireless links between the GPS receiver and the satellites can be affected and, in some cases, severed completely. The cases where it is severed completely would be considered a blind spot. The most common blind spot would be a parking garage. If someone regularly parks in a garage that is mostly concrete and steel, this would directly affect the fidelity of their trips collected by the road charge system, as all of the start and end points and any points driven within the garage would be compromised. Urban canyons also would be considered a potential blind spot. Cities such as New York and Chicago with a high density of very tall buildings will have GPS blind spots that impact trip differentiation accuracy.
- **Inadvertent Disconnects** – Participants in a pilot know that they must keep the device plugged in at all times, but as discussed previously there are some vehicles in which the OBD-II port is located such that the device is susceptible to being "knocked out of" the port – either by a foot or knee. Some Ford models have the port located near the parking brake, and it has been shown that drivers of these vehicles that regularly used the foot parking brake also regularly knocked out the device. Every single one of these disconnects is potentially a trigger for a true-up being required in a road charge pilot or program.
- **Intentional Disconnects** – As is the case with every pilot involving an OBD-II port, there will be times that a participant intentionally unplugs the device. In most cases, this is benign. The disconnect could be for routine maintenance on the vehicle. It could be for a vehicle sale. It could be for work being done on the vehicle. Knowing the difference between benign disconnects and ones that are not is important.
- **Differentiation** – The accuracy of differentiation algorithms is only as good as the source of the maps. As was shown in the differentiation issue that occurred in May during the pilot (see Section 4.5.2), bad map data can lead to incorrect information. If maps are incorrect or even

just slightly inaccurate, the ability to differentiate between public and private roads can be substantially impaired.

5.6 CUSTOMER INTERACTION

One of the points that needs to be brought forward as a discussion point with any road charge pilot is that for a participant to be involved in the pilot, they must take action and do something. Participation in road charge pilots, and ultimately programs, is not a passive undertaking. For many revenue collection methods (e.g., paying for gas at the pump, or buying a shirt at a store), a participant need take no special action on their part to pay the taxes on the transaction. This is not the case with road charging. Participants must actively enroll (no one can do it on their behalf). Participants must wait for, receive, unpackage and install an OBD-II device (after finding out exactly where it plugs into their vehicle). And of course, they participant must pay their mileage fees. None of these actions are onerous or overly burdensome, but they are actions nonetheless and thus, how the product is designed and implemented must take participant experience into account. A flawless experience is the goal.

5.7 DEVICE LONGEVITY

After-market OBD-II devices have been around almost as long as the OBD-II port itself. It is reasonable to infer that they will remain around as long as the OBD-II port remains a required standard every vehicle must adhere to. With a need to change the funding model from fuel-based taxing to mileage-based taxing, the OBD-II device is one source of mileage data that can be leveraged. However, with the changing landscape of vehicles from internal combustion engines to electric vehicles, there is a very good chance that the OBD-II standard will change as well. This means that there is very likely a limited life left of the OBD-II port. How little life left? Well, there is no hard and fast end date – in fact there isn't even discussion of phasing it out. There is some discussion of possibly changing the data available on the port to support electric vehicle information and the physical port size and location have room for improvement, but these are just high-level discussions, and nothing is even remotely set yet. Also, because it takes so long for standards to propagate to all new vehicle makes and models, the soonest the OBD-II port would be phased out would be ten to twelve years – and that is if they mandated the phase out tomorrow.

Another potential time-based issue is the fact that since the inception of the OBD-II port and the subsequent creation of OBD-II devices, there has been a constant turn-over of device manufacturers. Since 2014 there have been at least four major players in the space that no longer make devices for the OBD-II market. There remain a few primary players, but they will only make devices as long as it is profitable to do so.

Despite the length of time it may take for the OBD-II port and their associated after-market devices to remain a viable source, this time based fact should play a role in the design of a road charge program.

6. The Bottom Line

The previous sections may make it seem that the notion of using an OBD-II device for a road charge pilot is dead on arrival. This actually is not the case; far from it. The plug-in device does have some drawbacks, and it was the goal of this report to call out the deficiencies in a way that have not been articulated previously. It was not to disparage the device nor tank its use for good, but to bring to light the deficiencies so that plans can be made to address the shortcomings.

The device does have some solid capabilities that cannot be provided any other way:

- As shown in section 5.5.2 and Table 6 above, differentiation using a location-enabled device is doable and it is highly accurate.
- The distance captured by the device using its own internal algorithm is very precise and can be relied upon as a source of distance travelled.
- Data is standardized and normalized across the broadest spectrum of vehicle years, makes and models. No other data collection method is even close to the number of vehicles that are supported.
- The amount of data that is provided from the device has many other potential uses beyond just road charging.
- Devices can provide location data in support of differentiation at a frequency higher than any other data collection method. Nothing else even comes close.
- The cost of the data is high, but it's actually not that much higher than other data collection methods. With a bit of engineering design, a less expensive alternative could likely be created.
- The data is clear from any ownership disputes. The data belongs to the driver, and they have a say in how it gets used.
- Devices allow for a "set it and forget it" style of revenue collection. Other data collection methods can get close to this but cannot match the simplicity of device-based approach, once the device is plugged in and successfully sending data.

So, what should some of the next steps be, with regards to the use of OBD-II devices within future road charge initiatives?

- First and foremost, it should be definitively stated that cost and support issues aside, a device with GPS technology could most definitely be used "to successfully support the accurate differentiation of public versus private roads", fulfilling a specific research goal that was identified by Caltrans within its application for STSFA funding (see Section 1.2).

- A precursor for moving forward with the use of OBD-II devices in support of operational RUC programs, would be for RUC practitioners to collaborate with OBD-II device vendors in order to tailor the device functionality and message set to the unique distance accuracy requirements of the RUC space. For example:

- Mutually explore moving from UDP to TCP in all device communications, to avoid pipeline issues.
- Investigate assembly of the entire trip in the device, with subsequent transmission of the entire via the more robust TCP protocol (rather than sending disparate trip components piecemeal via UDP, as was done for this Pilot).
- Have the device more frequently solicit ACTUAL odometer readings from the vehicle throughout the trip, to manifest better support for compliance.

- Address all aspects of the per-device costs, as the costs associated with this Pilot would not scale upwards and most likely prove to be prohibitive for an operational program. For example:
  - Snowflake is a relatively expensive option for meeting the transactional needs of a RUC program, and a more cost-effective alternative could most likely be employed.
  - Consider use of Bluetooth versus cellular communication for the device, which could help reduce device costs by as much as 30%.

- Further analysis of the tradeoffs between GPS frequency and accuracy/fidelity of differentiation would need to undertaken.

- As noted with the Conclusion section of the Closeout Report (Final Report Appendix D, Section 9.2), further analysis should be performed to investigate the ROI of undertaking differentiation to identify private road usage. It is quite possible that the cost of differentiation is higher than the benefit yielded, to either the participant or to the agency.

- Investigate leveraging using data from OBD-II devices that are already in vehicles for other products such as usage-based insurance and fleet management. Using a device from other established products could help alleviate the need for a road charge program to bear the full cost of the device, data plans and support. While this could prove to be a potential avenue of data, it might prove to be more effort than it is worth as only about 30% of insurance customer choose a telematics-based insurance product and of these, only a small percentage require a device to be plugged in full time. As for fleet management, this would only be a good source of data for commercial vehicles as there are no fleet management products for individual consumers.