Connected Vehicle Deployment Decision-Support Analysis and Stakeholder Impact Analysis:

Summary of Findings

1. Introduction

Connected Vehicle is a Federal initiative that envisions the use of wireless vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication to reduce crashes, congestion, and emissions. In the Connected Vehicle concept, vehicles would broadcast standardized safety messages using Dedicated Short Range Communications (DSRC) at 5.9 GHz, a frequency set aside for this purpose. These messages would share information on vehicles' speed, heading, location, and other attributes, and this information would be used for safety applications that provide drivers with warnings about imminent collisions. In addition, the data could be used by state and local transportation agencies to improve signal timing, road maintenance, and overall traffic flows.

Connected Vehicle has the potential to produce significant reductions in the societal costs of automobile crashes, which claim over 30,000 lives per year and cost billions of dollars in property damage and other impacts. The National Highway Traffic Safety Administration (NHTSA) has announced that it is pursuing a rulemaking process that may result in Connected Vehicle equipment becoming standard on all light-duty vehicles produced for the US market. The American Association of State Highway and Transportation Officials (AASHTO) is also sponsoring research on potential deployment scenarios for Connected Vehicle and how this deployment would affect their member organizations.

Although Connected Vehicle has great promise for safety and mobility, it would also entail considerable expense for adding the requisite equipment to vehicles and roadsides, and for ongoing telecommunications and maintenance costs. As such, the ITS JPO has sponsored a number of analytical efforts to estimate the potential benefits and costs of Connected Vehicle and to understand the implications of different deployment scenarios, including the Volpe Center analysis presented in the document.

This document is a final summary memo that recaps the work that the Volpe Center team has conducted over the past two years on an inter-related set of questions regarding Connected Vehicle safety benefits and location choices for roadside equipment. The specific area of focus was on V2I applications and whether these applications offer enough incremental benefits over a V2V-only approach to warrant the additional expense for roadside infrastructure. Additionally, a geographic analysis was conducted to understand how the quantity and location of roadside units would affect V2I safety benefits. The document presents information on the overall methodology and key assumptions that were used to generate estimates; discusses the key findings, including sensitivity analyses that were conducted; and highlights suggested areas for further analysis and refinement. This information is intended to provide ITS JPO with decision-support information as it considers different deployment options. The analysis and findings here are preliminary and are *not* structured as a formal regulatory impact analysis, but it is possible that this information, or forthcoming related analyses, may ultimately inform future rulemakings or guidance.

2. Methodology Summary

This section summarizes the overall approach for the analysis, the input data and assumptions that were used, and important limitations and areas for further refinement.

2.1 Overview of Approach

This analysis initially looked at the total potential for collision avoidance benefits of V2V and V2I applications over the time period from initial deployment to 2040. Estimated benefits in terms of crashes avoided were converted into monetary values using USDOT guidance on the societal value of injury prevention, with a $9.1 million value for preventing a fatality and lesser values for non-fatal injuries by severity level. Monetary values in future years were converted to present-value terms using OMB's recommended discount rate of 7%.

The analysis included safety benefits only, and did not include benefits from V2V/V2I applications that might enable congestion relief, reduced emissions, or other impacts. The analysis was also limited only to light-duty vehicles (passenger cars and light trucks).

Estimates of the types, frequency, and severity of crashes that could potentially be avoided through V2V and/or V2V applications were taken from NHTSA-sponsored research on crash scenarios and their applicability to Connected Vehicle. For each year in the analysis period, these estimates were then adjusted to account for rising vehicle-miles traveled (VMT) over time; a generally declining crash rate per VMT over time unrelated to Connected Vehicle; the growing deployment of Connected Vehicle equipment on vehicles in the fleet; and the level of V2I roadside equipment (RSE).

The analysis was then refined to generate more specific estimates of V2V/V2I benefits at different levels of RSE deployment and at different levels of safety effectiveness. This included gathering data on the spatial concentration of crashes to understand the impact of RSE placement on total safety benefits. Cost estimates for each scenario were also developed, using current working assumptions from the Connected Vehicle team. This allowed benefit-cost estimates to be compared across scenarios. In the subsections below, each of these modeling components is discussed in more detail.

2.1.1 Connected Vehicle Deployment Scenario

Drawing on informal NHTSA scenarios, onboard equipment (OBE) for Connected Vehicle was assumed to be installed on new vehicles starting in 2020, with a 3-year phase-in. That is, 35% of new vehicles produced in 2020 would have OBE, then 70% in 2021 and 100% in 2022. For simplicity, no distinction was made between calendar years and model years. The scenario further assumed that a small number of existing vehicles would be retrofit with OBE, starting in 2022. The retrofit would cover 5% of (otherwise unequipped) vehicles in model years 2022 and 2023, and 10% of vehicles in model years 2024 to 2026. For modeling purposes, factory-installed and retrofit OBE were assumed to be otherwise identical.

RSE was assumed to be installed over a 5-year period, with 20% of the total number of units installed each year. The Volpe team tested five RSE scenarios: no units; 50,000 units nationwide; 250,000 units; 500,000 units; and 2.3 million units. For illustrative purposes, some calculations were also done with an assumption of ubiquitous RSE availability from Day 1, to show the maximum potential benefit.

2.1.2 Fleet Turnover Model

The OBE deployment scenario listed above applies primarily to new vehicles. Older, non-OBE equipped vehicles would remain in the fleet for many years, other than the small number that would be retrofitted. Because (generally speaking) only equipped vehicles can avail of the safety applications, the benefit potential for Connected Vehicle would tend to rise over time as a growing share of the on-road fleet becomes equipped.

To estimate OBE penetration for each year of the analysis period, the Volpe team developed a light vehicle fleet turnover model based on previous work for the Corporate Average Fuel Economy rulemaking.[1] This model starts with a breakdown of the current fleet by vehicle age, and includes forecasts of new sales and scrappage (final removal from the fleet). Using these inputs, the model estimates the size and age composition of the fleet for each year of the analysis period, and tracks the number of OBE-equipped vehicles in each year based on the details of the deployment scenario. (As a simplifying assumption, the presence or absence of OBE was assumed not to affect sales or scrappage.)

The fleet model also estimates total VMT and VMT by vehicle age. The VMT adjustment is important

[1] See the CAFE regulatory impact analysis: http://www.nhtsa.gov/staticfiles/rulemaking/pdf/cafe/FRIA_2017-2025.pdf

because newer vehicles are driven slightly more per year than older ones, so the actual composition of miles driven – which is a proxy for crash exposure – is somewhat more weighted toward newer vehicles than the overall fleet composition alone might suggest. For estimating crash avoidance potential in a given future year, the relevant variable is the VMT (or share of total VMT) that is generated by equipped vehicles.

For simplicity, crashes addressable with V2V were assumed to be two-vehicle (rather than multi-vehicle) crashes. In that case, the likelihood that both vehicles are equipped is the square of the VMT-adjusted OBE penetration rate. For example, if OBE-equipped vehicles account for 20% of total VMT, the probability that both vehicles involved in a pre-crash scenario are equipped is 4% (0.20 * 0.20). Thus, a maximum of 4% of the total count of V2V-relevant crashes could be prevented in that year.

For crashes addressable via V2I, it is assumed that only one vehicle must be equipped, so the relevant figure is simply the VMT-adjusted OBE penetration rate, along with an adjustment for the presence of RSE (discussed below).

2.1.3 Subject Crashes

Subject crashes are defined as the subset of crash scenarios for which V2V/V2I technologies are applicable – i.e., crash types for which V2V/V2I collision avoidance applications are designed to operate and can at least potentially prevent the crash. Information on subject crashes was drawn from NHTSA-sponsored research, a summary of which appears below in Table 1.[2] Note that crashes involving driver impairment are excluded to be conservative.

Some types of crashes are addressable via V2V, others via V2I, and many are addressable by either or both. In order to avoid double-counting in cases where both V2V and V2I could address a particular crash type, the Volpe team followed the convention of the source material and designated V2V or V2I as either "primary" or "secondary" for tabulation purposes. This does not imply that one system or the other is more important; it simply designates one system as primary in the sense that crashes avoided through that system are counted first. Then, only the incremental crashes avoided through the secondary system are tallied. (Using an example from the table below, if V2V were considered the primary system, it can potentially address 39,000 out of 42,000 crashes in the "running stop sign" scenario. If it were to be 100% effective in doing so, only 3,000 crashes would thus remain for a V2I system to address. Conversely, if V2I were the primary system, it could potentially address all 42,000 crashes in this scenario.) In most of the scenarios analyzed, V2V applications were considered the primary system, because almost all Connected Vehicle concepts include OBE, whereas some do not include RSE.

[2] Najm, W.G. et al., "Pre-Crash Scenario Typology for Crash Avoidance Research," NHTSA report DOT-HS-810767, April 2007, and Najm, W.G. et al., Frequency of Target Crashes for IntelliDrive Safety Systems," NHTSA report DOT-HS-8113181, October 2010.

Based on further discussions with the sponsor and subject-matter experts, one change was made to the classification from Najm et al. (2007), namely that “running red light” crashes were changed from being potentially V2V-addressable to being V2I-only. The reasoning here was that, although two vehicles are typically involved and V2V is potentially viable, it is also necessary to know the red/green state of the traffic signal for optimal performance and to avoid excessive false alerts.

Table 1. Summary of V2V- and V2I-Addressable Crash Types

Crash Description	Total Annual Crashes (Baseline Year)	Addressable with V2V as primary	Addressable with V2I	Weighted Average Crash Cost ($2012)
Running red light	226,000	0	226,000	$32,557
Road edge departure/no maneuver	240,000	0	48,000	$87,098
Road edge departure/maneuver	54,000	0	9,000	$51,113
Pedestrian/maneuver	19,000	0	8,000	$85,649
Pedestrian/no maneuver	39,000	0	5,000	$210,786
Object contacted/no maneuver	61,000	0	5,000	$37,015
Running stop sign	42,000	39,000	3,000	$34,249
Object contacted/maneuver	32,000	0	3,000	$14,357
Rollover	3,000	0	1,000	$42,258
Control loss/vehicle action	89,000	89,000	0	$38,215
Control loss/no vehicle action	414,000	414,000	0	$85,358
Backing into vehicle	127,000	127,000	0	$6,984
Turning/same direction	195,000	195,000	0	$13,399
Parking/same direction	38,000	38,000	0	$16,699
Changing lanes/same direction	329,000	329,000	0	$15,045
Drifting/same lane	102,000	102,000	0	$18,908
Opposite direction/maneuver	9,000	9,000	0	$116,032
Opposite direction/no maneuver	102,000	102,000	0	$93,796
Rear-end/striking maneuver	80,000	80,000	0	$14,259
Rear-end/LVA	22,000	22,000	0	$15,227
Rear-end/LVM	190,000	190,000	0	$22,125
Rear-end/LVD	384,000	384,000	0	$15,385
Rear-end/LVS	906,000	906,000	0	$15,183
LTAP/OD @ signal	195,000	195,000	0	$29,870
Turn right @ signal	29,000	29,000	0	$9,650
LTAP/OD @ non signal	178,000	178,000	0	$30,540
SCP @ non signal	634,000	634,000	0	$34,646
Turn @ non signal	43,000	43,000	0	$23,372
Other - Rear-end	1,000	1,000	0	$42,258
Other - Sideswipe	2,000	2,000	0	$42,258
Other - Turn Across Path	1,000	1,000	0	$42,258
Other - Turn Into Path	1,000	1,000	0	$42,258

Note: Average crash costs are based on then-current USDOT injury values at the time the source document was prepared. USDOT has since issued updated guidance, which is used in the calculations below.

2.1.4 VMT and Crash Rates

VMT serves as a proxy for crash exposure: each additional mile driven raises the chance of a crash, however slightly. VMT has generally been increasing from year to year, although growth has flattened out since about 2008. Meanwhile, the prevalence of severe crashes, as expressed as a rate per VMT, has generally been declining over the past decades, due to improvements in vehicle safety technologies and societal factors. The safety analysis therefore needs to take these trends into account when estimating potential benefits in future years. Otherwise, the future safety benefit potential would be based on an under- or over-estimate the actual number of preventable crashes in those years.

For the purposes of this analysis, future VMT and crash rates per VMT were forecast using an existing Volpe Center model developed for FHWA. The model forecasts a relatively slow VMT growth rate of around 1% per year, while the crash rate per VMT declines at a similar rate. The net effect is that total crashes are forecast to level out or very slowly decline over time. (As will be seen in more detail below, the maximum potential safety benefit of V2V/V2I technologies levels out along with it; in other words, in the latter part of the analysis period, there are simply fewer crashes each year to prevent.)

These adjustment factors were applied to the initial estimates of subject crashes to produce a year-specific estimate of total V2V- and V2I-relevant crashes for each year in the analysis period. For simplicity, the share of crashes by particular scenario (e.g. running red light vs. running stop sign) were assumed to be fixed. It is possible that some types of crashes will become more or less common in the future, but there is little information on which to base a specific assumption in this regard.

2.1.5 Safety Effectiveness Rates

Safety effectiveness refers to the likelihood that a crash will actually be prevented, given that the vehicle(s) are in the relevant pre-crash scenario and other preconditions are met. Effectiveness rates are almost always lower than 100% due to technical limitations of the system as well as human factors limitations. Technical limitations include issues with sensors, telecommunications, and warning algorithms; human factors include inattentiveness, reaction time, comprehension of warnings, and ability to execute the necessary maneuver. In addition, some types of crashes may be essentially unavoidable despite advanced safety systems.

The Volpe Center team's analysis began with an assumed 100% effectiveness rate, which demonstrates the maximum safety benefit potential and is useful for illustrative purposes. Subsequent analyses included rates of 10%, 25%, and 40% in addition to the 100% level. Connected Vehicle safety technologies are still the testing phase and no definitive information is available on their actual effectiveness rates; however, tests with earlier generations of technology suggest that effectiveness rates on the order of 20% to 30% may be most realistic.

2.1.6 RSE Geographic Coverage

The DSRC telecommunications envisioned for Connected Vehicle have a range of approximately 300 meters. This is generally adequate for V2V interactions, since vehicles involved in a pre-crash scenario are generally within that range of each other, if not much closer. For V2I applications, one key limitation on safety benefits is that RSE cannot be everywhere, and thus only pre-crash scenarios that take place within 300 meters of a functioning unit can be potentially addressed. As such, the overall safety benefits of V2I applications depend on the number of units and the location of these units relative to where crashes (or pre-crash scenarios) actually occur. This, in turn, depends on how spatially concentrated crashes are – in other words, are they concentrated at a relatively small number of trouble spots that can be equipped with RSE, or are they spread thinly across the road network?

There is little information in the transportation literature on the spatial concentration of crashes in the United States. In order to estimate this factor, the Volpe Center team used geocoded crash data from the Missouri State Police, covering crashes of all severity levels statewide from 2002 to 2011. (No other states appeared to have comprehensive, high-quality crashes with latitude and longitude recorded, but the Volpe team continues to search for other sources of data.)

Using the historical Missouri data, an iterative geographic analysis was conducted in STATA to identify the optimal RSE locations relative to the crash locations – i.e. the locations that, for any given number of RSE units, yielded the maximum number of historical crashes within their 300m range. This exercise was repeated with an adjustment for crash severity. To account for temporal variation and reversion to the mean, the location algorithm was also tested using a holdout sample, i.e. by estimating the optimal RSE locations using only a portion of the data (2002-2008) and then testing these locations against the remaining data (2009-2011). One limitation of the geographic analysis is that RSE were not required to be located within the highway right-of-way, so some of the calculated optimal positions may not be technically feasible.

As noted above, various RSE deployment sizes were tested, from 50,000 units up to 2.3 million units nationwide. Missouri's share of total national VMT was used as a scaling factor to convert between statewide estimates and nationwide estimates. Missouri's VMT is just slightly higher than average, yielding a roughly 44:1 scaling factor for national totals (rather than 50:1 if Missouri were perfectly average).

The result of these calculations is a function relating RSE buildout size to the percentage of V2I-eligible crashes covered by those RSE. This is used to scale V2I benefits appropriately and test the benefit-cost impacts of smaller and larger RSE buildouts.

2.2 Benefit-Cost Calculations

Putting together all of the pieces from the discussion above, estimated V2V and V2I safety benefits for each year in the analysis period and system costs were calculated as follows:

- The number of crashes potentially addressable by V2V and V2I was taken from Najm et al. (2007) and the monetary equivalents were adjusted using updated USDOT injury values. These totals were scaled for VMT growth and changes in future crash rates as discussed above. These steps yield the theoretical maximum safety benefit potential for each year in the forecast period.
- The share of that overall potential that is actually achievable is scaled down based on the level of OBE and RSE deployment in each year of the deployment scenario.
    - For V2V crashes, this is drawn from the fleet turnover and VMT model. Total achievable benefits equal the theoretical maximum for V2V, multiplied by the square of the (VMT-adjusted) fleet penetration rate for OBE.
    - For V2I crashes, total achievable benefits equal the theoretical maximum for V2I, less the estimated crashes already prevented by V2V, multiplied by the (VMT-adjusted) fleet penetration rate for OBE. This total is then multiplied by the share of crashes that are estimated to occur within the 300m range of an installed RSE unit, based on findings from the Missouri geographic analysis.
- The above step yields total achievable benefits for a given year. This total is then adjusted for the safety effectiveness rate, using 10%, 25%, 40%, and 100% as illustrative values, to give a more realistic estimate of the level of safety benefit is likely to be achieved.
- Finally, the monetary values from each year are converted to present value using the 7% discount rate and are summed across the time period.
- Estimated deployment costs for OBE, RSE, and system operations and maintenance are likewise discounted and summed across the time period. The benefit-cost ratio is the total present value of benefits divided by the total present value of costs.

3. Modeling Results

The focus of this analysis was to understand (1) whether the potential safety benefits of V2I applications are large enough to warrant investment in RSE, and (2) if so, what level of RSE buildout would provide the most cost-effective approach. V2I can provide benefits over and above what V2V provides, both in terms of crash scenarios for which only V2I is applicable (e.g. single-vehicle crashes) and in providing additional coverage during the years when OBE penetration is low and thus V2V interactions are scarce. These questions are discussed below.

3.1 Incremental Safety Benefit Opportunity of V2I

The initial analysis calculated the safety benefit opportunities for V2V and V2I applications. The goal was to highlight the incremental benefits of V2I applications relative to a V2V-only baseline (or "V2V primary" in the terminology as described above). These illustrative calculations ignored the impacts of RSE placement and range – i.e. they assumed that RSE was ubiquitous – and further assumed that all applications had 100% effectiveness rates with respect to subject crashes. As such, these calculations highlighted the theoretical maximum safety benefits that could be achieved with V2V and V2I applications.

As shown in the chart below, the benefit potential of a primary V2V system rises steeply as a larger and larger share of the vehicle fleet becomes equipped. The benefit potential then levels out in the second half of the analysis period; this reflects the impact of the team's forecast of a falling rate of crashes per VMT and slow VMT growth. However, the overall benefit potential is very large: over 3 million crashes per year with an undiscounted monetary value of $100 billion per year or more at peak deployment.

By contrast, an incremental V2I system would have the greatest safety benefit potential in the earliest years of deployment, when OBE penetration rates are low, reaching a peak about 6 years into the deployment scenario. The benefit potential then levels off as V2V interactions become more common and the role of V2I becomes increasingly limited to V2I-only applications such as single-vehicle crashes. As with the V2V system, the benefit potential also gradually levels out over time reflecting the influence of the underlying crash rate. Again, however, the overall benefit potential is quite large, peaking at over 400,000 crashes per year with an undiscounted monetary value of over $30 billion. In Figure 1, the *incremental* benefits of V2I can be viewed as the vertical distance between the two curves.

Figure 1. Maximum Safety Benefit Potential (Undiscounted) for Primary V2V System and for Combined V2V and V2I Systems, By Year, 2020-2040

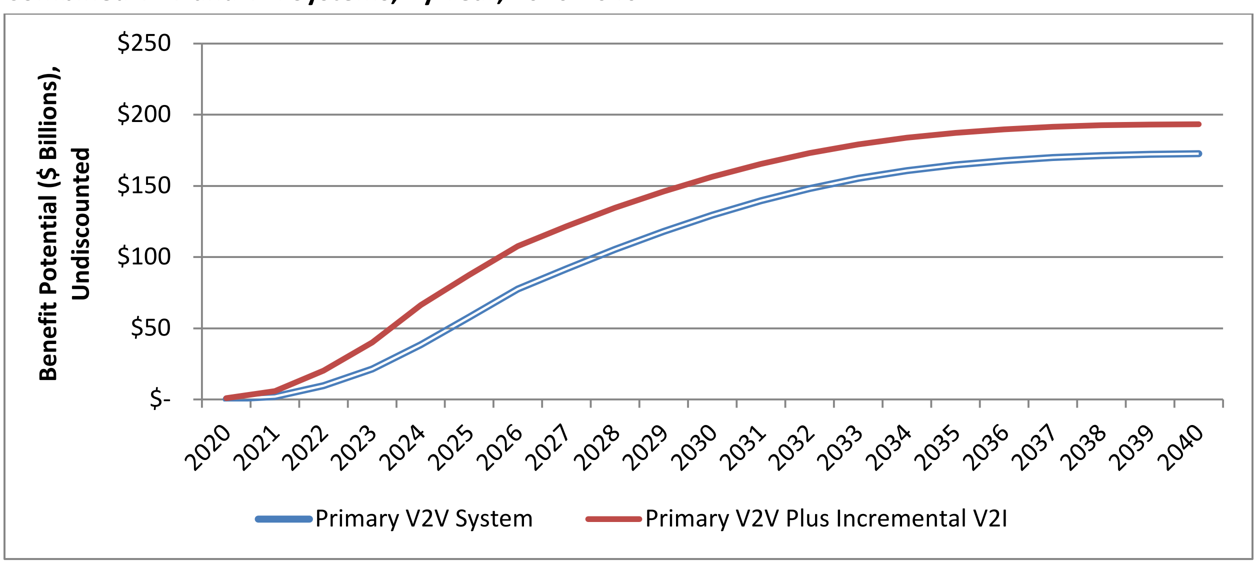

3.2 Crash Locations and RSE Placement

The safety benefit potential of V2I as calculated above is very large compared to the benefits that have been realized from other safety interventions, and would justify significant investment. The catch is that these benefits can only be realized to the extent that RSE is deployed and that imminent crashes actually take place within communication range of an RSE unit. Each new RSE unit increases the geographic coverage of the system and the likelihood that a particular pre-crash event will be addressable; however it also increases the costs of the overall system. This phase of the analysis was designed to study that tradeoff.

From the multi-year geocoded crash data from Missouri , the study team estimated the relationship between the number of installed RSE and the share of crashes that would be within the 300-meter communication range of those RSE. This is a function of the actual spatial concentration of crashes as they occur in the real world. A holdout sample was used to control for reversion to the mean.

The results of this analysis for Missouri, as shown in the chart below, suggested that crashes exhibit a moderate degree of spatial concentration and that this concentration shows some stability over time. Estimating the ideal geographic locations for RSE using the 2002-2008 data and then applying these same locations to the 2009-2011 data, a somewhat non-linear relationship was found between the level of RSE installed and the share of crashes (or crash costs, when adjusting for severity) covered. For example, the first 1,141 RSE units statewide – roughly equivalent to a national buildout of 50,000 units – would cover nearly 14% of all crashes. Yet because of the relatively long "tail" of the crash location distribution, it would require nearly 42,000 RSE units statewide (equivalent to 1.8 million nationally) to get to the 90% coverage level.

Figure 2. Share of V2I-Relevant Crash Costs Potentially Covered by RSE Units (300m Range), by Size of Statewide RSE Deployment for Missouri

Based on Geocoded Crash Data 2002-2011

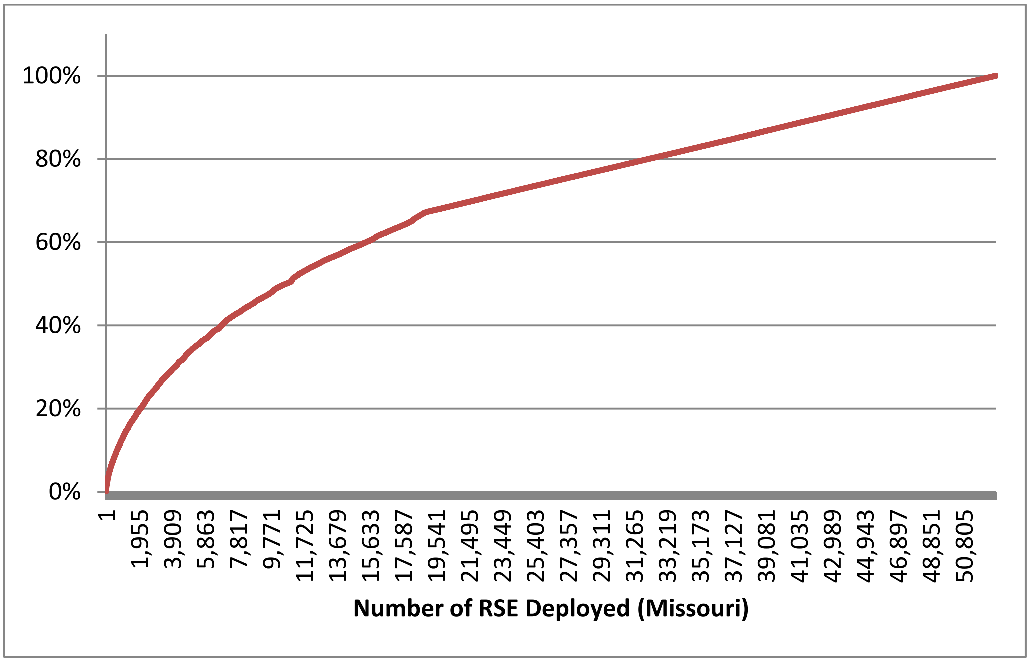

In order to provide a more intuitive feel for what a particular RSE buildout might look like, the project team also created a map of Missouri that shows the calculated RSE locations alongside historical crash locations. The map below is based on a scenario of 50,000 RSE nationwide, or just over 1,000 in Missouri. For this scenario, within Missouri, the geographic analysis placed 13 RSE in the central area of Columbia, a university town of approximately 110,000 people. These RSE were generally located at intersections and highway interchanges, covering locations with multiple crashes in the historical crash data. Many other crashes remain outside the collective coverage of the RSE.

Figure 3. Illustrative RSE Deployment Map for Columbia, Missouri: Calculated Optimal Locations in Scenario with 50,000 RSE Nationwide

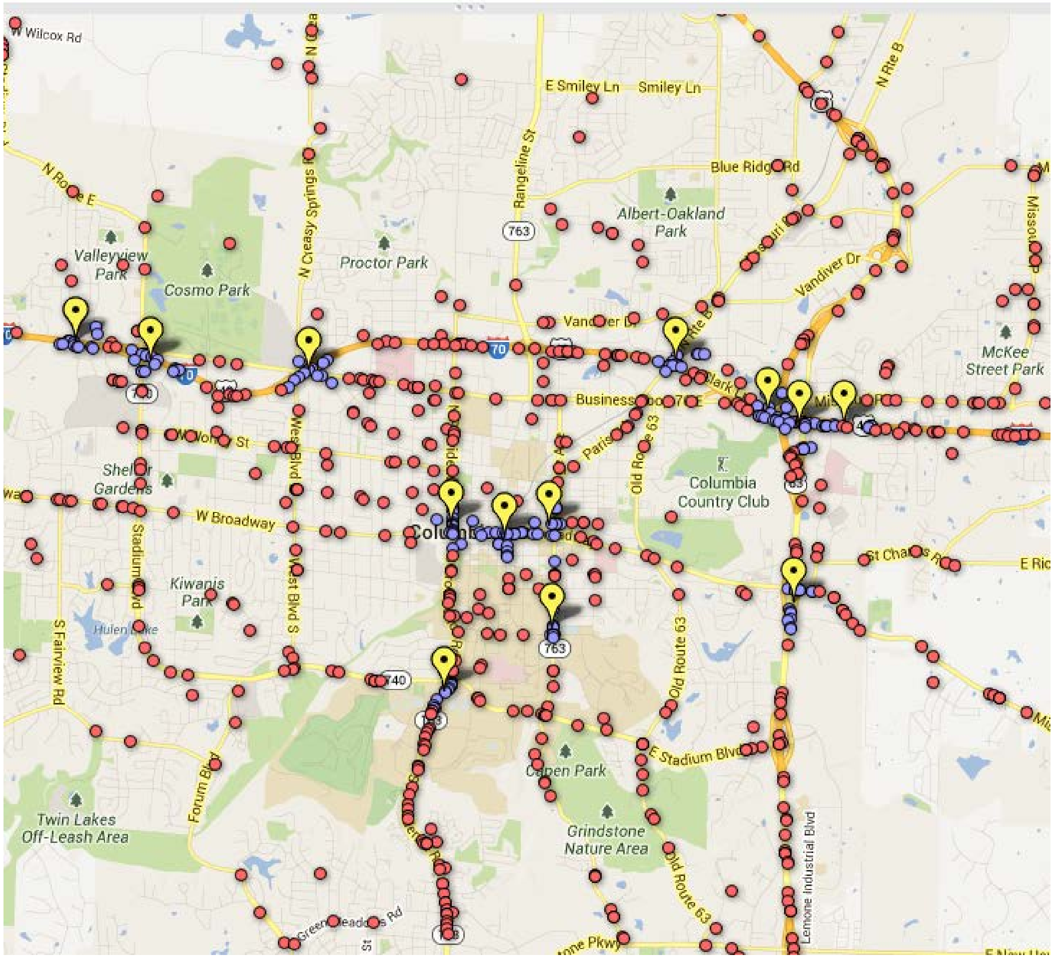

Yellow markers are calculated RSE locations. Blue dots are historical crash locations within 300m range of RSE. Red dots are historical crash locations outside RSE range.

The map is designed simply to be illustrative of a potential deployment. It is important to keep in mind that these RSE placements only represent the calculated optimal locations for addressing severe crashes. The geographic model did not consider real-world factors such as the presence or absence of electrical power, telecom, or other ITS equipment; nor did it consider non-safety V2I applications or the potential need for RSE to support network security functions. Also, while most of the calculated RSE locations appear to be along roadways, the specific locations were not constrained to be located within the highway right-of-way and thus may not be practical.

Safety benefits for each RSE buildout scenario were calculated as follows:

- Safety benefits for a given year = Total V2I safety benefit potential for that year * Share of crashes geographically covered by RSE buildout size * Application effectiveness rate
- Total safety benefits = Sum of present value of yearly estimates

For these calculations, the total V2I safety benefit potential is as defined and calculated above in Section 3.1, taking account of the level of OBE penetration and changes in future VMT and crash rates. The share of crashes covered is taken from the geographic relationships taken from the Missouri crash data, as described earlier in this section. The effectiveness rate use a range of 4 assumed rates, from 15% to 100%.

The table below provides a summary of these calculations. Note that estimated benefits rise with larger RSE buildouts, but at a decreasing rate. In other words, a 500,000-unit deployment provides greater benefits than a 250,000-unit deployment, but not *twice* as much; there are diminishing marginal returns to RSE investment due to the clustering of crash locations. Conversely, benefits do rise linearly with the safety effectiveness rate.

Table 2. Estimated incremental V2I safety benefits ($ billion, total PV at 7%, 2020-2040) over V2V-only baseline, by deployment scenario and application effectiveness rate

	15%	25%	40%	100%
50,000 units	$3.2	$5.3	$8.4	$21.0
250,000 units	$8.1	$13.4	$21.5	$53.7
500,000 units	$11.6	$19.4	$31.0	$77.4
2.3 million	$22.2	$37.1	$59.3	$148.2

3.3 RSE Costs

The overall cost-effectiveness of RSE deployments depends not only the geographic coverage of crashes that they can provide, as analyzed above, but also on the upfront and ongoing costs of those RSE units. Because Connected Vehicle has only been deployed on a few smaller-scale test beds, information on RSE costs is incomplete and is continuing to evolve. More reliable information on RSE costs is expected to become available as more testing and more real-world installations are completed and as final decisions are made about communications protocols. More widespread deployment would also be expected to yield greater production volumes and potential economies of scale.

For the purposes of this analysis, the Volpe Center team drew on informal estimates produced by a contractor working on the Connected Vehicle program. These cost estimates were as follows: $8,839 per RSE in upfront costs for site preparation; $22,719 per RSE for installation costs, and for periodic replacement of the unit on a 15-year lifecycle; and $7,482 per RSE per year for recurring costs, including maintenance and telecommunications.

Using these cost figures, the total lifecycle costs for RSE buildouts of various sizes were estimated.

Table 3. Estimated Lifecycle Costs by RSE Deployment Size

Size of National RSE Deployment	Estimated NPV (at 7%) of Lifecycle Costs, 2020-2040
50,000 units	$5.0 billion
250,000 units	$25.1 billion
500,000 units	$50.3 billion
2.3 million units	$232.5 billion

Based on informal discussions, it appears that the recurring costs (rather than upfront equipment costs) have the greatest uncertainty associated with them, because they depend on estimates of future telecommunications costs. This is somewhat unfortunate because the recurring costs represent the larger share – just over two-thirds – of overall lifecycle costs. Thus, even small differences in the estimated annual costs of an RSE unit can change the overall cost significantly.

As a sensitivity test, the Volpe Center team also developed a scenario in which RSE upfront costs were the same as estimated above, but recurring costs were significantly lower, at 20% of the initial estimate. This scenario highlights the potential impacts of a major reduction in telecommunications costs – something that is not implausible given the tendency of telecom costs per gigabyte to decline over time, as well as the potential purchasing power that could be exercised by state DOTs (or other Connected Vehicle entities) who would be purchasing wireless data in bulk.

A draft version of the AASHTO footprint analysis presents cost estimates that are somewhat higher than those used by the Volpe team. For example, although AASHTO's estimated RSE installation costs are in the same general range (very roughly $20,000 per unit), their analysis also envisions $3,000 to $40,000 in backhaul upgrades per site. AASHTO also assumes a 5-10 year replacement cycle for equipment rather than 15 years, which would imply substantially higher lifecycle costs. AASHTO's analysis does not yet present figures for ongoing telecommunications costs.

3.4 Overall Benefit-Cost Analysis

Total benefits and costs for the four RSE buildout scenarios were estimating by combining the analyses discussed above. The table below shows the ratio of benefits to costs over the period 2020 to 2040, discounted at 7%, for the combination of four different buildout scenarios and four safety effectiveness rates. As can be seen in the table, the benefit-cost ratio rises with effectiveness, since each RSE deployed is that much more effective in actually reducing crashes. For any given effectiveness level, the benefit-cost ratio declines as the deployment size increases, because of the nonlinear relationship between RSE coverage and safety benefit potential.

Scenario 1: RSE costs as above (recurring costs $7,482 per unit per year).
Benefit-cost ratio uses lifecycle benefits and costs for period 2020-2040, discounted at 7%.

	15% Safety Effectiveness	25% Effectiveness	40% Safety Effectiveness	100% Safety Effectiveness
50K	0.63	**1.04**	**1.67**	**4.18**
250K	0.32	0.53	0.85	**2.14**
500K	0.23	0.38	0.62	**1.54**
2.3M	0.10	0.16	0.25	0.64

Scenario 2: Significant reduction in RSE annual recurring costs ($1,496 per unit per year).
Benefit-cost ratio uses lifecycle benefits and costs for period 2020-2040, discounted at 7%.

	15% Safety Effectiveness	25% Effectiveness	40% Safety Effectiveness	100% Safety Effectiveness
50K	**1.36**	**2.27**	**3.64**	**9.10**
250K	0.70	**1.16**	**1.86**	**4.65**
500K	0.50	0.84	**1.34**	**3.35**
2.3M	0.21	0.35	0.55	**1.39**

3.5 Benefit-Cost Discussion

The benefit-cost analysis presented above is based on a number of assumptions that will need to be revisited over time as the Connected Vehicle program evolves and as additional data become available. It should be viewed as a preliminary indicator rather than a definitive set of findings. With those limitations in mind, the analysis does suggest several key conclusions and implications for further research and analysis.

- First, V2I applications have a very large safety benefit potential, even when viewed as an incremental add-on to V2V safety systems. Even though V2V can potentially address many of the same crashes, there are other crash types for which V2I is more suitable; moreover, V2I provides additional benefits during the years when OBE penetration is low because it can be available when only one vehicle (rather than both) is OBE-equipped.
- V2I's incremental benefit potential is projected to peak relatively early -- about 6 years into the Connected Vehicle deployment, due to the influence of rising OBE fleet penetration and increasing V2V interactions. Therefore, there is at least a plausible argument for facilitating more rapid RSE deployment in order to capture as much of this benefit as possible.
- Based on current estimates, the costs – both upfront and ongoing – associated with RSE infrastructure and backhaul telecommunications are substantial. As such, the overall benefit-

cost proposition for an incremental V2I system and the "optimal" RSE deployment size both depend strongly on these costs. Telecommunications costs in particular appear to be subject to uncertainty due to the difficulty in forecasting costs in this fast-changing sector.

- There is also relatively little information available on the other two key factors that influence the benefit-cost ratio for V2I: effectiveness rates of the applications themselves, and the actual geographic distribution of crashes across the road network. More information on effectiveness rates may come from the Safety Pilot, and more geocoded crash data (beyond the Missouri data used here) may become available over time. Until then, this analysis relies the relationships identified in the Missouri data and sensitivity analysis for other variables.
- ***Overall, using the RSE cost estimates and other assumptions as detailed above, and with realistic (but as yet unconfirmed) effectiveness rates of 25% to 40% for V2I safety applications, V2I appears to be most cost-effective as a "tailored" deployment with roughly 50,000 RSE units nationwide, deployed at locations optimized to capture the largest share of addressable crashes. Based on the sensitivity analysis, a major breakthrough in telecommunication costs would make a somewhat larger deployment cost-effective.*** Beyond that, larger deployments would still prevent additional crashes, but they would also entail the costs of siting RSE at lower-frequency crash locations, and would not yield net benefits. Again, however, this balance would change to the extent that the safety effectiveness rates of the applications may turn out to be higher than estimated here, for example with an eventual move to partial automation that reduces the need for human intervention.
- The optimal RSE deployment size could also be much larger to the extent that RSE are needed for non-safety applications and/or network security.
- Although the analysis was done using OMB's recommended discount rate of 7% , some sensitivity testing was also conducting using the 3% alternative rate. In general, net benefits are higher when using the 3% rate because of the way RSE and OBE costs are front-loaded compared to the future stream of benefits. Future analyses may need to consider issues related to the choice of an appropriate discount rate.

4. Next Steps

Going forward, there are three main areas in which this analysis can be enhanced to make it more useful for the Connected Vehicle program. The first is with basic technical updates to the model itself. The second is expanding the benefit and cost estimates to encompass areas that were previously excluded, including mobility and environmental applications, as well as the possibility of onboard equipment for heavy commercial vehicles and buses. The third is coordination with other related efforts and stakeholders, particularly AASHTO's deployment analysis.

With regard to technical improvements, the Volpe Center team plans to update the current assumptions and forecasts for key variables such as VMT, crash rates per VMT, light-duty fleet turnover, and V2V- and V2I-relevant crash totals and severity rates. These updates are not expected to result in major changes to the model. However, some of the underlying data stretches back to 2007 or earlier, and does not reflect some of the trends and developments of the intervening years, including a noticeable decline in fatal crash rates and a rebound in new vehicle sales. These updates are thus important to making the model as credible and defensible as possible.

In addition, findings from the Safety Pilot and other ongoing research should be used to update assumptions about the applicability of V2V/V2I to specific crash scenarios and the range of plausible safety effectiveness rates for applications. OBE deployment scenarios should be developed for trucks and buses to gauge the impact of expanding Connected Vehicle deployment beyond light-duty vehicles. To the extent possible based on research from the AERIS and DMA programs, Connected Vehicle Pilots and other efforts, the model should also be expanded to include the benefits and costs of non-safety applications, such as those designed to reduce congestion and emissions.

Coordination with stakeholders will likely focus on aligning key assumptions for the RSE deployment scenario, in terms of the number of units, their deployment timeframe, logic for the placement of those units, and upfront and ongoing costs for RSE. This coordination should include AASHTO and its deployment analysis effort, but ideally also with the local (county and municipal) transportation agencies who will ultimately own and operate much of the roadside infrastructure. This is particularly true for issues such as RSE placement, where real-world constraints such as the presence of power and telecommunications may strongly influence deployment. This coordination helps to ensure that the technical insights from the analysis are matched by practical relevance and awareness of developments in the Connected Vehicle world. This coordination element would also include aligning assumptions about OBE deployment with any new scenarios from NHTSA and/or the automotive industry.