Organizational and Operational Models for Certificate Management Entities as Part of the Connected Vehicle Program

REVISED WORKING PAPER (Task 2)

This document presents interim findings of ongoing work. Some contents, including assumptions about the CME models, have been superseded and do not represent the most recent analysis. This document is being released to DOT contractors and agreement holders who are working on connected vehicle security solutions to support further study.

August 8, 2012

FHWA-JPO-12-078

U.S. Department of Transportation
Research and Innovative Technology Administration

Produced by Booz Allen Hamilton for the
U.S. Department of Transportation
Research and Innovative Technology Administration
Intelligent Transportation Systems Joint Program Office

Notice

Technical Report Documentation Page

1. Report No. **FHWA-JPO-12-078**	**2. Government Accession No.**	**3. Recipient's Catalog No.**
4. Title and Subtitle Organizational and Operational Models for Certificate Management Entities as Part of the Connected Vehicle Program Revised Working Paper (Task 2)		**5. Report Date** 08/08/2012 **6. Performing Organization Code**
7. Author(s) Dominie Garcia, Andrea Kiernan, Blake Sheppard, Richard Walsh		**8. Performing Organization Report No.**
9. Performing Organization Name And Address Booz Allen Hamilton 8283 Greensboro Drive McLean, VA 22102		**10. Work Unit No. (TRAIS)** **11. Contract or Grant No.** **DTFH61-11-D-00019**
12. Sponsoring Agency Name and Address Research and Innovative Technology Administration Intelligent Transportation Systems, Joint Program Office 1200 New Jersey Ave SE Washington, DC 20590		**13. Type of Report and Period Covered** Formal Deliverable 05/21/2012 – 08/08/2012 **14. Sponsoring Agency Code**
15. Supplementary Notes		

16. Abstract

This report presents the analyses and conclusions from work focused on analyzing multiple options for organizational and operational models for certificate management entities (CMEs) within the connected vehicle system. The report discusses all functions and design alternatives for the CMEs. It also identifies technical and policy decisions that affect implementation of the CMEs, but have not yet been made by USDOT and its partners in this research. The report contains an extensive cost estimate and description of known functions and operations, and identifies multiple variables and factors that will influence the eventual implementation of the certificate management system. Finally, the report includes high level discussions of possible implementation plans and next steps, with suggestions for performance measures and metrics for the CMEs.

17. Key Words Connected vehicle system, connected vehicle program, Security Credential Management System, certificate management, vehicle-to-vehicle, vehicle-to-infrastructure	**18. Distribution Statement**		
19. Security Classif. (of this report)	**20. Security Classif. (of this page)**	**21. No. of Pages** 122	**22. Price**

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized**

Table of Contents

List of Tables

List of Figures

Executive Summary

Introduction

The U.S. Department of Transportation (USDOT) has established a multimodal program that is researching the potential for wireless communication among vehicles and with infrastructure to dramatically improve transportation safety, and to advance mobility and environmental goals. The connected vehicle program, as this research is known, is led by the Intelligent Transportation Systems Joint Program Office (ITS JPO) within the Research and Innovative Technology Administration (RITA) with support from four other modal agencies[1] within the USDOT. A critical requirement for a connected vehicle system (the term used to refer to future deployment of the system) is security, or a means to ensure that messages sent across the system are legitimate and have not been tampered with. At the same time, users want to have a reasonable assurance of appropriate privacy in the system. Research to date has indicated that use of a Public Key Infrastructure (PKI) security system, involving the exchange of digital certificates among trusted users, can support both the need for message security and for providing appropriate anonymity to users while in transit. Certificate Management Entities (CMEs) perform the back office functions required to administer a PKI security system, such as registering users and issuing and revoking certificates. The term Security Credential Management System (SCMS) is also used to refer to all CME organizations, or the certificate management system as a whole.

Security Credential Management System (SCMS) – The set of organizations that house the various functions and activities necessary for the certificate management process – the entirety of the PKI system for the connected vehicle system.

Certificate Management Entity (CME) – A legal and administratively independent organization that houses PKI functions.

The ITS JPO has contracted with Booz Allen Hamilton (Booz Allen) to analyze alternative approaches and models for CME organizations as part of an overall security network for connected vehicles. In order to be viable, a SCMS structure must be cost-effective, efficient, and scalable. The purpose of this project is to analyze alternative CME structures for the purpose of determining the extent to which each satisfies the objectives and functional requirements identified by USDOT.

Initially the team was focused on developing organizational models that could house the CME functions in order to identify how these different models might impact security and privacy. Over the course of the analysis the team determined that, fundamentally, the differences between

[1] Federal Highway Administration, Federal Motor Carrier Safety Administration, Federal Transit Administration, and National Highway Traffic Safety Administration.

organizational models are not anticipated to change levels of security or privacy protection due to the nature of the PKI system that will establish trust and direct how credentials are exchanged between users. The differences between organizational models are related to organizational policies and resulting efficiencies realized through various structures. At this point in the development of the connected vehicle system, it is premature to specify all organizational policies, but as progress towards implementation is made, further detail around how to stand up and operate the chosen structure should be specified, based on existing organizational best practices. We present alternative organizational models, with focus on one (Model G) that provides the highest levels of organizational and operational efficiency. The technical questions that came out of this analysis are outlined in the conclusion and will ultimately drive implementation plans.

CME Functions

PKI is the governing paradigm of communications security within the connected vehicle system. PKI is a certificate management system that features a central authority, known as the Certificate Authority (CA), that verifies that users in a system are trustworthy based on certain credentials. This allows users to trust one another and effectively interact, even if they have had no prior interaction, by virtue of their trust in the CA, which issues "certificates" that provide the indication of trust to all other users.

There are several functions that are central to PKI systems: the Registration Authority (RA), Certificate Authority (CA), and Misbehavior Detection and Management (MDM). USDOT and its industry partners have modified traditional PKI designs to meet the security needs of the connected vehicle system. Additions to PKI for this environment include organizational separation of RA and CA and the addition of two Linkage Authorities (LAs) – a new function that provides common values to a large set of certificates, and also provides an additional level of security because of the addition of another step in the process of creating certificates. The modification of the PKI system, by adding two LAs, enables the system to be able to maintain high levels of privacy and security. The result is a more complex process for creating and expanding public and private keys, which are cryptographically created protections added to encrypt and decrypt certificates in order to guard against the possibility of trip trackability (the ability of someone to track an individual vehicle through a portion or all of a trip).

Several of the PKI functions for the connected vehicle system have yet to be fully developed either theoretically or operationally. Gaps in the technical architecture include how the Misbehavior Detection and Management (MDM) function will operate.

As noted above, there are two terms that are used to refer to the entire PKI system of certificate management for the connected vehicle system (SCMS) and to those individual organizations within the system that will house, operate, maintain and be responsible for particular functions (CME). Figure 1 below demonstrates the relationship of CME to SCMS and also provides an illustration of the basic interactions between the SCMS functions. Additional detail about interactions between functions and the processes of certificate creation and management are included throughout the report.

Figure 1. Security Credential Management System

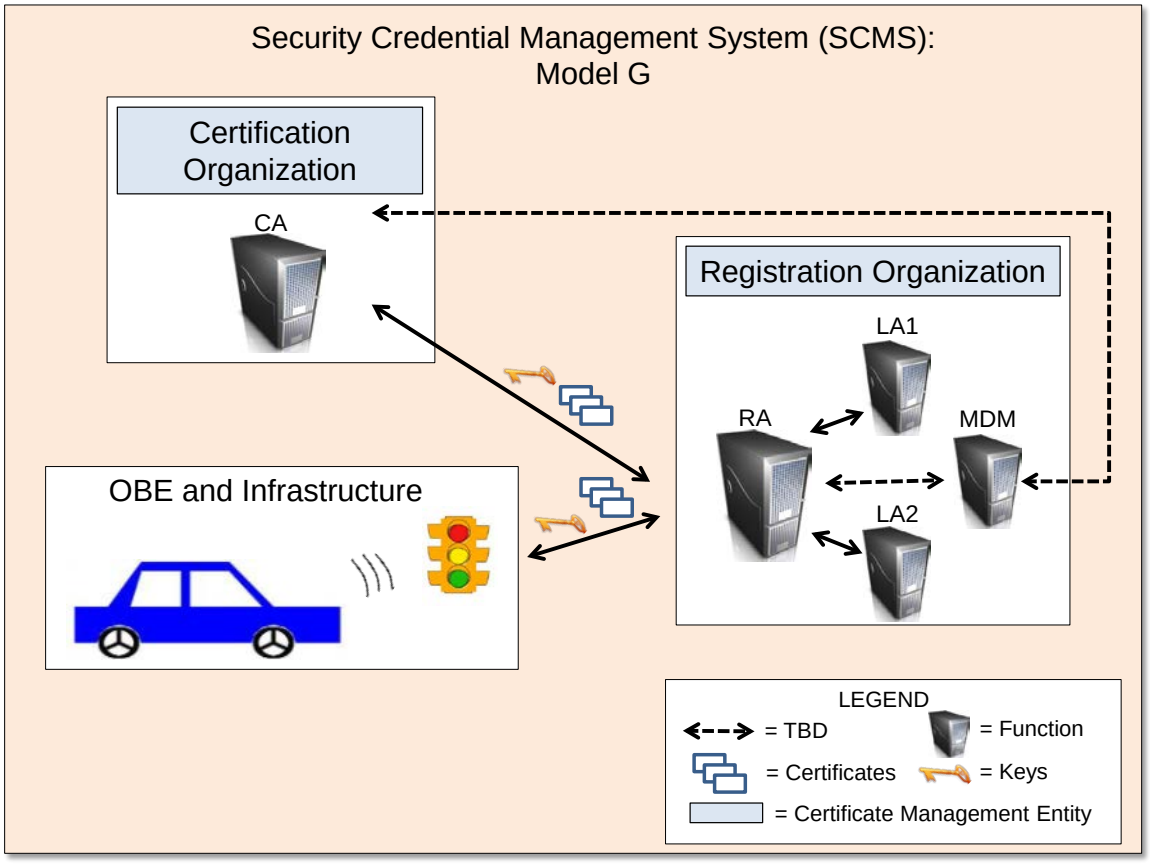

CME Models

The process of analyzing CME options began with an identification of all of the possible ways that the CME functions could be aligned within a model. A total of 15 models were identified, seven of which were acceptable in regards to ensuring privacy and security. Three models were then chosen for further exploration. The three models that are described in more detail include all responsibilities for each function in various combinations in order to realize cost and organizational efficiencies. The models combine various functions within separate entities in order to realize some operational and organizational efficiency, while maintaining security protections. No security or privacy differences exist between the models, as these are predicated on technical and policy decisions. Also included is a discussion of different ways of combining the LAs, both based on the team's analysis and stakeholder input. Development of sufficient technical, procedural, and physical controls is proposed as a way to maintain separation of data and functions while leveraging other efficiencies and operational connections. As noted above, the differences in the models are operational and organizational and don't change the analyses of security and privacy protections. Model G provides the most operational and organizational efficiencies, and so is the one used as an example when talking about the entire SCMS, as above in Figure 1.

Baselining

USDOT requested that Booz Allen identify a security baseline of acceptable vulnerability and risk by researching comparable systems in other industries. The team analyzed existing PKI systems in other industries and organizations and examined audit thresholds and protocols, when available. The team also examined industry protections against potential risks and vulnerabilities. Though unable to identify a specific security baseline relevant to the connected vehicle system, the team found that PKI as a choice of security system for the connected vehicle system provides as comprehensive a security protection as observed in other systems, though additional protections within the PKI environment can be applied in order to provide multiple, additional layers of security and protection against data sharing or unacceptable levels of functions sharing information. Other industries protect against threats to IT systems by implementing procedural, technical, and physical controls to hardware and software access. In addition, auditing procedures and protocols specify acceptable levels of security breaches for some industries, though exact numbers are not available.

Certificate policies within all industries specify how organizations are required to protect against hardware and software vulnerabilities, although to date almost all are based on X.509 certificates, while the connected vehicle system is designed to use the Institute of Electrical and Electronic Engineers (IEEE) 1609.2 certificates, for which new certificate policies and controls will need to be specified. Technical and policy direction about how to monitor, audit, and enforce standards will guide implementation of security standards within the SCMS PKI.

Personal Privacy Protection

USDOT has not yet made a determination about whether, and to what extent, the SCMS will need to collect users' personally identifiable information (PII) in order to permit users to participate in the connected vehicle system. For this reason, to ensure that the current policy analysis is comprehensive, we have considered all relevant collection options, including the option of keeping user credentials anonymous – an approach advocated by motor vehicle manufacturers and privacy advocates. If the CME collects PII for the purpose of being able to connect security credentials to individual users, vehicles or devices, there are various ways to do so – and the team has presented and evaluated each, as well as the option of collecting no PII, in terms of its impact on privacy protection and CME operations.

The options analyzed include:

- Collect no PII – This option ensures users' total anonymity and prevents the SCMS from being able to trace malfeasance or misbehavior back to a specific individual, vehicle or device, thus preventing the system owner(s) and/or government from taking legal action against hackers and others who might use the system for malicious or illegal purposes. Vehicle manufacturers and privacy advocates have taken the position that this is the only acceptable option -- the one way to ensure user acceptance without damaging the commercial interests of manufacturers, safety benefits of the technology, or privacy interests of consumers.

- Collection of PII during device activation. This can happen in a number of ways:

  - Leveraging existing systems that collect vehicle-based PII, such as the Vehicle Identification Number (VIN)
  - Collecting new PII
- Collection of PII within the pseudo system, which is the system that operates on an ongoing basis to verify, exchange, distribute, monitor, and accept certificates between vehicles and also between vehicles and Roadside Equipment (RSE).

The latter two collection options would provide a means for tracing malfeasance or misbehavior back to a specific individual, vehicle or device to enable the system owner(s) and/or government to identify and take action against those who attack the system for malicious or illegal purposes. Motor vehicle manufacturers have taken the position that collection of any PII by the CME will adversely affect user acceptance, consumer privacy, the safety benefits of V2V technology, and their commercial interests, as consumers will not buy new vehicles containing V2V devices due to a perception that the technology will enable the government to track their location. The collection of PII by the CME poses more of a risk to individual privacy than collection of no PII. However, it remains an open question whether a SCMS with no mechanism to find hackers and others who might use the system for malicious or illegal purposes also could undermine user confidence in the V2V system.[2]

This report discusses and compares against a set of USDOT-furnished criteria each of the methods of credentialing identified above for various users and communications environments (Vehicle-to-Vehicle [V2V], Vehicle-to-Infrastructure [V2I], Vehicle-to-Other Device [V2X]).

In the event that USDOT decides that the SCMS needs to collect users' PII, this report suggests that the most feasible collection option, in terms of technical viability, cost, and security protections, is integrating into SCMS processes the credentialing of users within existing government registration systems, such as state vehicle registration and USDOT's registration system for heavy commercial vehicles.

Expansion of Users to Infrastructure and Mobile Devices

As the connected vehicle system evolves and expands in scope, it is expected that additional applications and types of users will need to be authenticated and provided with the security credentials to participate in the system. Roadside infrastructure, such as traffic signals, will need to be authenticated as trusted members of the system as well. As requested by USDOT, this report discusses inclusion of nomadic and other non-vehicle-based devices in the trust network. The team's fundamental finding is that the CMEs will need to add levels of functional breadth to existing operations in order to accommodate such expansion. Additional specification of technical needs and types of applications will determine the kinds of authentication processes that may be needed for other equipment, but at this point the requirements for infrastructure and nomadic devices are not

[2] Rather than base its decision solely on the speculation of a select group of V2V stakeholders, the USDOT is considering conducting or sponsoring more research about user acceptance that specifically targets these issues, to better inform its upcoming regulatory decisions.

anticipated to be as complex as those envisioned for five-minute certificates on vehicles. Each certificate is valid for five and a half minutes, which means that a total of 105,120 pairs of keys are required for a one year supply of certificates (assuming five-minute certificates). Figure 2 represents the connected vehicle environments referenced in this report.

Figure 2. Connected Vehicle Environments

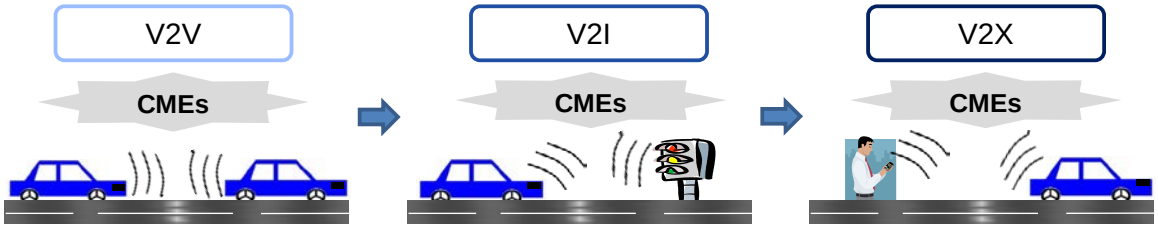

Misbehavior

Misbehavior detection and management are critical components of a SCMS. To date, USDOT and its research partners have conceptualized how these processes may work, but have yet to develop or specify the relevant mathematic or algorithmic operations. In this report, we have outlined some critical issues related to misbehavior processes and their implications to CME operations including:

- Technical malfunction and human malfeasance within the system – how the differences between these types of misbehavior are detected and what policies are in place to deal with consequences of each are yet to be specified.
- Certificate Revocation List (CRL) – exact technical specification of how the CRLs are constructed and distributed is yet to be specified.
- Regaining access to the system after placement on CRL – whether this happens through replacement of On Board Equipment (OBE) (as suggested by one stakeholder group), by reactivation of existing Certificate Signing Request (CSR), which is the certificate to authenticate a user's device, or by rekeying with a new CSR are all decisions that have yet to be made, and will impact how CMEs operate and communicate between each other.
- Suspension vs. revocation – decisions about what offenses would require suspension of certificates versus revocation are yet to be made and will also impact the previous point about how to regain access to the system once either suspension or revocation is reversed.

Technical Specifications

This report contains a detailed listing of technical needs for all the CMEs. Although there are several ways of potentially producing certificates, this report includes estimation of hardware security modules (HSMs) in combination with standard processors to produce, store, distribute, and manage certificates. There remain some gaps in being able to specify certain data and processing needs, as some of the functional operations are still under development – notably LA and MDM functions. We have used the current assumptions and our recommendations for how to produce certificates using HSMs to inform our cost estimates.

Costs

Based on estimates included in the technical specifications chapter, and research conducted within the PKI industry, the team was able to develop detailed cost estimates for most certificate management functions (with the exception of the MDM), including hardware, software, facilities, redundancy needs, and personnel. Estimates include initial, up front expenses associated with system development, facilities acquisition/build-out/construction, and hardware and software procurement. Annual estimates include operations and maintenance for software, hardware, and facilities, personnel costs, and a system refresh every four years. Up front and annual costs are specified in total and expressed as costs per OBE on vehicles. Our cost estimates include sensitivity analyses of roll out phases (penetration percentage), regional needs, and cost savings realizable per model. Highlights include:

- Total deployment build up and operational costs at full deployment (i.e., 250M vehicles) for a period of six years could range from $1.54-$2.11B, with an annualized cost per OBE of $1.03-$1.41.
- Large up-front costs can be somewhat offset by much lower maintenance costs annually.
- Renewal of hardware and software is estimated to happen every four years because of the large amount of data processing needed throughout the system.

Implementation Planning

This report includes initial high-level discussions of implementation planning and additional considerations that will be required for successful roll out of the CMEs, with particular focus on certain factors, including:

- Disaster recovery
- Technical and policy support for users and CMEs (Help Desk)
- Performance measurement and metrics

Conclusions

One of our primary conclusions about the SCMS design is related to the separation of the two LA functions. External stakeholder groups have suggested that these need to be within legally and administratively separate entities. This team's research and analyses suggest that this separation is not necessary, as it would increase costs and organizational complexity significantly, while improving security safeguards only negligibly. Our additional conclusions are set forth in a summary table in the report's Conclusions Chapter. The summary table identifies technical and policy decisions integral to implementation of any SCMS model that have yet to be made by USDOT and its stakeholders. These include:

- Phases of Roll Out – urban versus rural roll out, timing for implementation of CMEs across the country
- Misbehavior – how is it detected, how is malfunction differentiated from malfeasance (technical questions), and what are the consequences and enforcement policies (legal and policy questions)

- Credentialing – what, if any, PII will be collected at what point in the system and what are the rules governing access to data
- LAs – what the actual mathematical and algorithmic processes are within the LAs. Do the LAs need to be separated into legal/administratively separate entities
- CSR – what is the life span of the CSR and how it is rekeyed or renewed
- Back-Up Certificates – if they exist and how they are used
- Certificate Policy – what the policy will say regarding the roles and responsibilities, the rules governing obtaining certificates, the technical requirements for generation and protection of private keys and certificates, and the requirements for audit records and periodic compliance audits
- End of Life – how end of life is determined and defined and what the policies are that govern disposal of OBE and removal from the system
- Frequency of Certificate Download – currently assumed to be once a year, but may need to be revisited based on problematic size of downloads
- Number of RAs – determined by amount of hardware needed and communications delivery network decisions to communicate with OBE
- Number of CAs – determined by decisions about virtual or physical environment and how to distribute and produce redundancy around the system
- PKI Hierarchy – Root CA and hierarchy decision based on options and security determination
- Certificate Life Span – how would a longer certificate life span impact processing needs and trip trackability

Chapter 1 Introduction

The U.S. Department of Transportation (USDOT) has established a multimodal research program into wireless communication among vehicles and with infrastructure with the potential to dramatically improve transportation safety, and to advance mobility and environmental goals. The connected vehicle program, as this research is known, is led by the Intelligent Transportation Systems Joint Program Office (ITS JPO) within the Research and Innovative Technology Administration (RITA) with support from four other modal agencies[3] within the USDOT. The ITS JPO has contracted with Booz Allen Hamilton (Booz Allen) to analyze alternative approaches and models for Certificate Management Entities (CMEs) that could administer the security functions required to support the connected vehicle system. The CMEs must ensure the security of communications and protect the privacy of system users, with the goal of building user trust. To be viable, a CME structure must be cost-effective, efficient, and scalable. The purpose of this project is to analyze alternative CME organizational structures to determine how each satisfies system objectives and functional requirements identified by USDOT.

The connected vehicle program is focused on conducting foundational research to facilitate a secure communications system that could support Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications to enable safety, mobility, and environmental applications for the public good, and to further public acceptance of such a communications system. The security approach that is utilized was developed by a consortium of light vehicle manufacturers, the Crash Avoidance Metrics Partnership (CAMP), with the intent of satisfying the needs of balancing security and privacy.[4] Addressing these needs is critical to acceptance and use of the system. If users do not feel the system is trustworthy, they will not use the system and therefore the system will not work. The security approach was developed for research and testing purposes and does not necessarily reflect the end state system. The term Security Credential Management System (SCMS) is also used to refer to all CME organizations, or the certificate management system as a whole.

Security Credential Management System (SCMS) – The set of organizations that house the various functions and activities necessary for the certificate management process – the entirety of the PKI system for the connected vehicle system.

Certificate Management Entity (CME) – A legal and administratively independent organization that houses PKI functions.

For the connected vehicle system to work effectively, users of the network must be able to trust the validity of messages received from other system users. Establishing the basis of this trust network as

[3] Federal Highway Administration, Federal Motor Carrier Safety Administration, Federal Transit Administration, and National Highway Traffic Safety Administration.

[4] The system being tested and developed will be widely applicable (i.e., will be relevant for trucks, transit, and pedestrians).

well as other physical and software design considerations across the system are the key elements of a security design for the connected vehicle system. This project assumes use of a Public Key Infrastructure (PKI) scheme to achieve the security goals related to establishing trust among users, consistent with the security approach developed by CAMP. The use of a PKI in this system allows for the creation and management of digital certificates that bind an identity to its public key to certify the sources of messages, which enables users to trust the system.[5]

Certificate Management Project Goals

The goal of this project is to assess alternative organizational models for CMEs that balance the security of communications with system users' privacy. A total of 15 models were initially considered, although only seven were deemed acceptable based on privacy and security criteria identified by the USDOT (presented in December 2011). Of these seven models, three are presented in more detail in this report.

This report presents analyses and findings related to organizational models for CMEs, including the definition of the roles and responsibilities of these structures and significant background discussion of both technical and policy issues that impact these organizational models. The team has also identified initial cost estimates and resource requirements based on the assumptions and models presented herein.

This report includes detailed discussion of attempts to understand and develop a security baseline for the CMEs. Based on comparable industries and approaches to protecting security in PKI systems, the team has gathered evidence that may guide how the PKI system being designed for the SCMS may protect system and user security and privacy. It is critical to note that much of the technical foundation for the CMEs in development for the connected vehicle system is still theoretical in nature and has yet to be implemented or operationalized. The best estimates of security protection have been developed based on the theoretical underpinnings.

The team has conducted wide ranging research into industry examples of various other elements that impact the development of the CMEs at full deployment of the connected vehicle system. Investigation of audit practices, cost and resource estimates, privacy protections and standards, existing organizations and systems that may act as pieces of the CMEs' needs, and procedural and technical controls have all yielded important and useful examples and best practices that serve as starting points for design and implementation of CMEs.

Project Approach

The team has conducted the research and analyses of the development of CME models systematically and with the input of key stakeholder groups, both internal and external to USDOT. During the previous phase of work, the team conducted extensive stakeholder interviews and outreach and synthesized that input with current technical and policy thinking and approaches to

[5] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.

develop the full complement of models that were considered by USDOT. Based on feedback on a report and presentation given in December 2011, the team delved into greater detail around three of the seven models for full explication.

As part of this work, the team conducted deep analysis around technical architecture, specifications, and needs in order to fully understand the foundation upon which to base operational and organizational models for certificate management. Because of the quickly evolving and dynamic nature of the technical work, several assumptions are still in place to guide the discussions. The team will refer to these assumptions and their implications on the CMEs throughout this report.

One of the main goals of the current report, in addition to presenting different operational models for providing certificate management, is to highlight the decisions that are still to be made in both technical and policy arenas. The team has included options for many of the still-outstanding decisions and what the various implications are for different options.

The requirements of V2V safety applications form a critical underpinning of the work presented in this report. V2V messages are based on the notion that safety messages sent between vehicles provide the starting point for thinking about certificate management and for design of needed CMEs. Other components of the connected vehicle system, including V2I mobility and environment applications, are assumed to be part of the evolution of the system. The report includes a discussion of how additional needs and requirements for the CMEs may develop based on the incorporation of additional users and applications. Lastly, we outline next steps that are required for deployment and implementation of the CMEs envisioned herein.

Chapter 2 Definitions

Activation System	The Certificate Management Entity (CME) that activates or initializes the OBE. It includes a Certificate Authority for Activation (CA_{ACT}) and Registration Authority for Activation (RA_{ACT}) that are separate from the CA and RA of the pseudo system.
Basic Safety Message (BSM)	The outgoing message sent by a vehicle that communicates information and data about its current state to a set of neighboring vehicles. That information or data is used by Vehicle-to-Vehicle (V2V) safety applications in the neighboring vehicles to warn users of crash imminent situations.
Butterfly Keys	A set of public keys related to a single private key generated by the RA and CA. There are two – one for signing and one for encryption. The signing keys are used to validate BSMs signed by the On Board Equipment (OBE). The encryption keys are used to encrypt the certificates for transmission back to the OBE.
Caterpillar Keys	A pair of public and private key pairs generated by the OBE. There are two per set of OBE certificates requested. One pair is used for signing and one pair is used for encrypting. The public parts are sent to the RA where each is expanded into a set of keys that are sent to the CA as part of each certificate request.
Certificate Authority (CA)	The SCMS epicenter responsible for certificate production, certificate validation, and misbehavior detection management.
Certificate Identifier	A unique identifier in each certificate calculated from the linkage values specific to that certificate provided by the linkage authorities.
Certificate Management Entity (CME)	An organization that houses the certain functions and activities necessary for the certificate management process.
Certificate Policy	The document that describes the roles and responsibilities for implementing a PKI, the rules governing how certificates are obtained, the technical requirements for generation and protection of private keys and certificates, and the requirements for audit records and periodic compliance audits.
Certificate Revocation List (CRL)	A list of certificate identifiers that the Misbehavior Detection and Management (MDM) function identifies to be misbehaving due to technical error or human malfeasance.

Certificate Signing Request (CSR) Certificate	The certificate used to initialize the OBE in order to authenticate the device to be part of the CME and thus receive batches of five-minute certificates. Each certificate is valid for five and a half minutes, which means that a total of 105,120 pairs of keys are required for a one year supply of certificates (assuming five-minute certificates).
Cocoon Keys	A pair of public and private key sets generated by the RA from the caterpillar keys passed from the OBE. The purpose is to expand the caterpillar key into something the CA can use to return information that only the OBE can read.
Connected Vehicle Program	The U.S. Department of Transportation (USDOT) research program focused on the combination of applications, services, and systems necessary to provide safety, mobility, and environment data to users.
Connected Vehicle System	The deployed system of connected vehicle devices, infrastructure, and back end functions that will enable safety, mobility, and environment applications to be exchanged.
Credentialing	The process of linking certificates or On Board Equipment to individual credentials, such as PII.
Cryptography	The combination of mathematical algorithms and computer science to protect users, networks, and messages sent throughout a network by encrypting messages. Only authorized users of the network have the necessary information or credentials to access the data within the network.
Dedicated Short Range Communications (DSRC)	The one-way or two-way short-to-medium range wireless communication channels specifically designed for automotive use and a corresponding set of protocols and standards.
Elliptic Curve Cryptography (ECC)	A public key cryptography method that utilizes points found within a curve group to create keys. The point selected from the curve is multiplied by a random number numerous times.
Global Processing	The process of comparing a sample of random messages from the OBE against the behavioral models based on all message traffic to detect statistical anomalies that signal potential misbehavior.
Hardware Security Module (HSM)	A piece of hardware that exists as a layer of security that consistently protects communications, credentials, and requests by safeguarding and facilitating encoding, decoding, verification, and electronic signature procedures. Also accelerates cryptographic transactions per second.
Implicit Certificates	The five-minute certificates.
Linkage Authority	The CME entity responsible for generation and creation of linkage values

(LA)	at the request of the RA.
Malfeasance Notice	A report generated by the MDM that notifies the RA and other OBE(s) that a certain OBE(s) is being corrupted by a human intentional error(s).
Malfunction Notice	A report generated by the MDM that notifies an OBE(s) to be taken out of service and repaired due to a technical malfunction(s).
Misbehavior	The reference to technical errors and human intentional errors that have a negative impact on process operations within the pseudo system.
Misbehavior Detection and Management (MDM)	The CME function responsible for detection, tracking, and managing potential threats to the CME and connected vehicle system. Also responsible for CRL management.
On Board Equipment (OBE)	The user equipment that provides an interface to vehicular sensors for safety measures, as well as a wireless communication interface to the RA for CME processes.
Personally Identifiable Information (PII)	Any form of information that can be used to uniquely identify, contact, or locate an individual person, directly or indirectly.
Point Multiplication	The operation of successively adding a point along an elliptic curve to itself repeatedly. It is used in elliptic curve cryptography as a means of producing a key or encrypting an object.
Private Key	In a PKI, the key held secretly by the subject of a PKI certificate that contains a related public key. It is not made available to any other entity. The private key is mathematically related to the public key in such a way that what is encrypted with one is decrypted with the other.
Pseudo System	A system that operates on an ongoing basis to verify, exchange, distribute, monitor, and accept certificates between vehicles and also between vehicles and RSE.
Public Key Infrastructure (PKI)	A set of hardware, software, people, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates. PKI has been chosen as the mechanism to provide integrity and authentication within the connected vehicle system. This system creates and manages digital certificates that bind an identity to its public key to certify the sources of the messages.
Public Key	In a PKI, the key that is included in a PKI certificate. The public key is mathematically related to the associated private key in such a way that what is encrypted with one is decrypted with the other.
Registration	The CME entity responsible for certificate batching and issuance and

Authority (RA)	cocooned and decryption key generation. In many cases this function is an intermediary between the CA and other entities of the connected vehicle system and is the only CME entity that communicates with the OBE.
Roadside Equipment (RSE)	An infrastructure node that serves as an intermediary in Vehicle-to-Infrastructure (V2I) two-way communications between CMEs and vehicles, or sends out its own messages to OBE.
Root Certificate Authority (CA)	The master CA that provides the signatures on the certificates for its subsidiary CAs. The Root CA possesses a self-signed certificate that contains its own public key to differentiate itself from other CAs.
Security Credential Management System (SCMS)	The set of organizations that house the various functions and activities necessary for the certificate management process.
Server Farm	A collection of computer servers or processors maintained to accomplish computational needs associated with key generation, certificate production, signing, verification, and encryption.
Signal Phase and Timing (SPaT)	A message that is used to convey the current status of a signalized intersection. The receiver of this message is able to determine the current state of each phase and what the expected next phase is to occur.
Trip Trackability	The ability of someone to track an individual vehicle through either a portion of or an entire trip.
Trust Store Management	A process that provides procedures to import, edit, and remove certificates trusted by the system for validation of a digital signature and certificates. This process ensures that the issuing CA is self-signed and loaded directly into the trust store, or that the CA can trace the signature on its certificate though one or more CAs to a self-signed certificate loaded in the trust store.
Vehicle-to-Device (V2X)	The wireless communication exchange of messages and data between vehicles, infrastructure, and capable nomadic devices within the connected vehicle system.
Vehicle-to-Infrastructure (V2I)	The wireless exchange of critical safety and operational data between vehicles and highway infrastructure, intended primarily to avoid motor vehicle crashes but also to enable a wide range of other safety, mobility, and environmental benefits.
Vehicle-to-Vehicle (V2V)	A dynamic wireless exchange of data between nearby vehicles that offers the opportunity for significant safety improvements.
Virtual Machine	The software implementation of a machine (e.g., a processor) that executes software code program instructions for the functions of the CME.

X.509 Certificate In cryptography, X.509 is an International Telecommunications Union Telecommunication Standardization Sector (ITU-T) standard for public key certificates and attribute certificates. This international standard defines a framework for how certificates are formatted, revoked, and managed, among other things.[6]

[6] ITU-T, X.509 website: http://www.itu.int/ITU-T/recommendations/rec.aspx?rec=X.509.

Chapter 3 CME Functions

This chapter describes the functions necessary for an operational CME. The description of these functions is based on the team's review of documents and several working sessions with subject matter experts.[7] At the time of writing this report, detailed specifications of the technical design were largely under development and, thus, a set of assumptions about the technical architecture was needed as the basis for analysis. Assumptions about the CME functions, the authentication process of an OBE, the details of the CSR, and certificate life spans have been developed and vetted with USDOT and are described in this chapter and listed in Appendix C. With the evolving development of the connected vehicle system architecture and decisions regarding technical specifications, future changes to the current assumptions in this chapter may change the nature of the functions, which in turn, may impact the overall CME models and SCMS design.

This chapter also includes a detailed discussion of physical, procedural, and technical controls for PKI systems to demonstrate different methods of protecting and managing sensitive data and maintaining separation between functions.

Pseudo System Certificate Management Functions

As described in the introduction, PKI is the foundation upon which the proposed security system models and functions are based. For the PKI to meet the security needs of the connected vehicle system, four necessary functions were identified: Registration Authority (RA), Certificate Authority (CA), Linkage Authority (LA), and Misbehavior Detection and Management (MDM). The LA is a unique function introduced in the connected vehicle system to facilitate revocation of groups of certificates issued to a single OBE in the event that a malfunction or inappropriate use is detected. This is driven by the scale of the system and security needs. Because of the large number of five-minute certificates, there needs to be an efficient method of revocation. Additional explanation of the certificate revocation process is included in Chapter 8 of this report, with the caveat that the misbehavior detection processes and technical details have not yet been designed.

Together, these functions make up what the team refers to as the pseudo system. This name was coined by technical teams to describe the pseudonymous part of the SCMS that operates on an ongoing basis (day to day) to verify, exchange, distribute, monitor, and accept certificates between vehicles and also between vehicles and Roadside Equipment (RSE). This system is separate from the activation system (CA_{ACT}), which is the part of the SCMS that verifies that a given OBE should be in the system and activates its operation. Figure 3 below illustrates the

[7] Some of the documents reviewed include the Security Approach for V2V/V2I Communications Delivery System, the Core System Requirements Specifications (SyRS), Core System, System Architecture Document (SAD), and the Vehicle Safety Communications (VSC3) Interim Report. Additionally, guidance was provided by the technical teams that are researching the technical designs, namely the Crash Avoidance Metrics Partnership (CAMP) and escrypt, Inc. under direction from the USDOT.

basic relationships between the functions within the entire SCMS.

Figure 3. Security Credential Management System – Model G

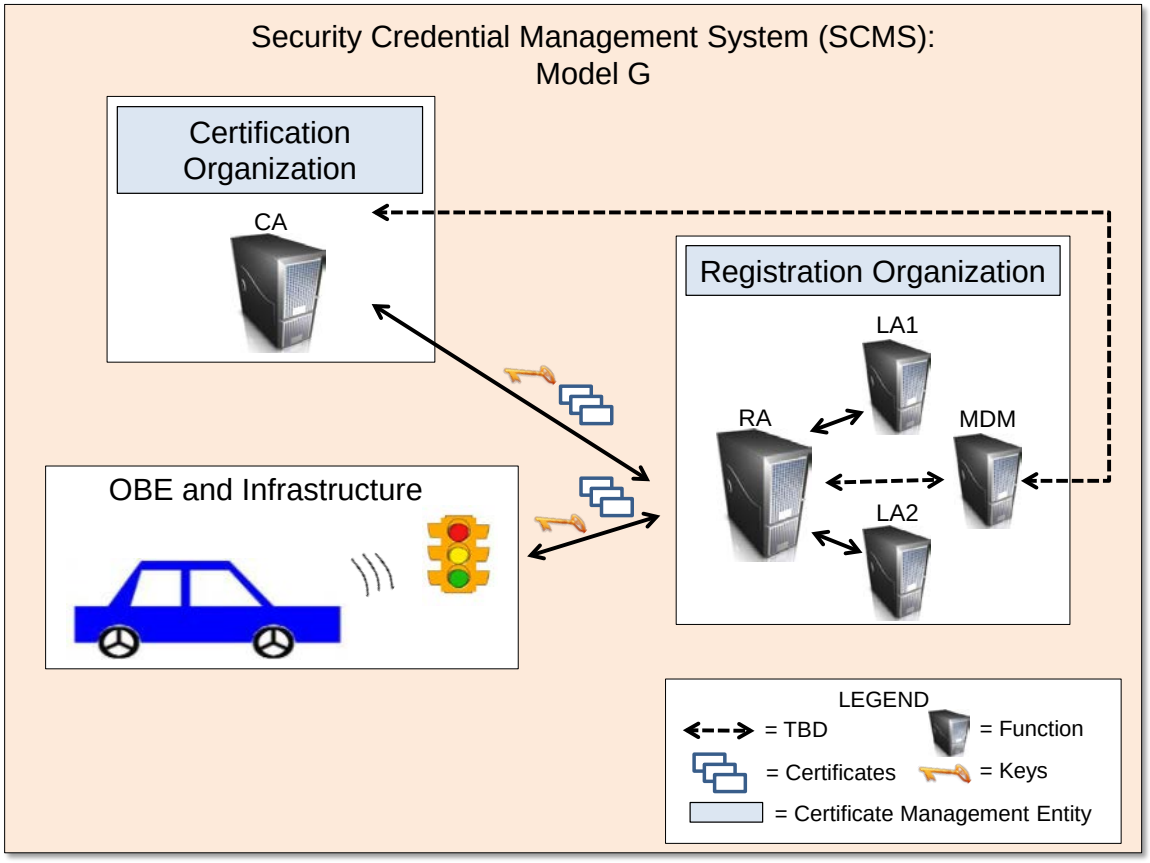

All descriptions below of processes and functions are based on the current technical design produced by CAMP. It should be noted that private and public keys are asymmetric cryptographic terms. Symmetric cryptography uses "shared secret keys." There are no "public keys" or "private keys" in a symmetric cryptographic system because if there are public keys, anyone can decrypt the information. On the reverse, if there is a private key then only one person could decrypt information.

On Board Equipment (OBE) generates two key pairs for a request - one for signing and one for encryption. The public keys for both (caterpillar keys) are sent to the RA which expands each into a set of cocooned keys. An encryption and a signature key are included in each request forwarded by the RA to the CA. The CA transforms each into a butterfly key. The signing public key is in the certificate. The encryption public key is used to encrypt the certificate.

Registration Authority (RA) communicates directly with the OBE and interfaces with each of the other CME functions. The RA receives OBE certificate requests, which include a signing and encryption caterpillar public key (explained in further detail under Figure 4). The RA expands each caterpillar key into cocoon sets of keys for each OBE. Each certificate is valid for five and a half minutes, which means that a total of 105,120 pairs of keys are required for a one year supply of certificates (assuming

five-minute certificates). The RA requests linkage values from each LA, and when it receives those values, it bundles them with the signing and encryption key. The RA collects sets of request data from multiple OBE and shuffles the requests to ensure that complete certificate requests are not sent to the CA in a sequentially identifiable order. The RA sends the certificate request to the CA for certificate issuance. The RA receives the OBE certificates back from the CA, batches them into groups which are encrypted, and forwards them to the OBE for use.

Linkage Authority (LA) communicates only with the RA and provides linkage values in response to a request by the RA. The linkage values provide the CA with a means to calculate a certificate ID and a mechanism to connect all 105,120 five-minute certificates for ease of revocation in the event of misbehavior. At this point in time, technical groups believe that at least two LAs are required to split formation of the certificate ID and improve the privacy of the system. In cases of user malfeasance or technical malfunction, the value produced by the LAs, known as the linkage value, is placed on the CRL. Part of the linkage value signifies the day that misbehavior occurred, which informs the RA to revoke all certificates on a misbehaving OBE from that day forward.

Certificate Authority (CA) issues the certificates used to ensure trust in the system by authenticating a device. The CA receives the certificate request from the RA. It does a final transformation of the cocoon keys, calculates a certificate serial number using the linkage values, and generates and signs the certificates. The CA encrypts each certificate with the encryption public key of the associated OBE, which is a public key related to the OBE private encryption key, and sends the encrypted data back to the RA for distribution to the OBE. In addition to certificate issuance, the CA collaborates with the LAs and RA to identify OBE values to place on the CRL if a malfeasance has been determined. If malfeasance is identified, the CA will place the linkage value on the CRL, which the CA generates, signs, and sends to the RA for distribution.

Misbehavior Detection and Management (MDM) receives misbehavior reports from the OBE and performs investigations or other processes to figure out levels of misbehavior. This is not an external law enforcement function, but rather a function that represents the internal CME work to detect when messages are not plausible or when there is potential malfunction or malfeasance within the system. Policy will control the extent to which any part of the CMEs perform misbehavior management outside of the system or work with law enforcement. To date, the MDM processes or algorithmic functions have not been defined.

Figure 4 represents the functions within the CMEs with a detailed listing of responsibilities for each function for the pseudo system.

Figure 4. Certificate Management Functions and Responsibilities

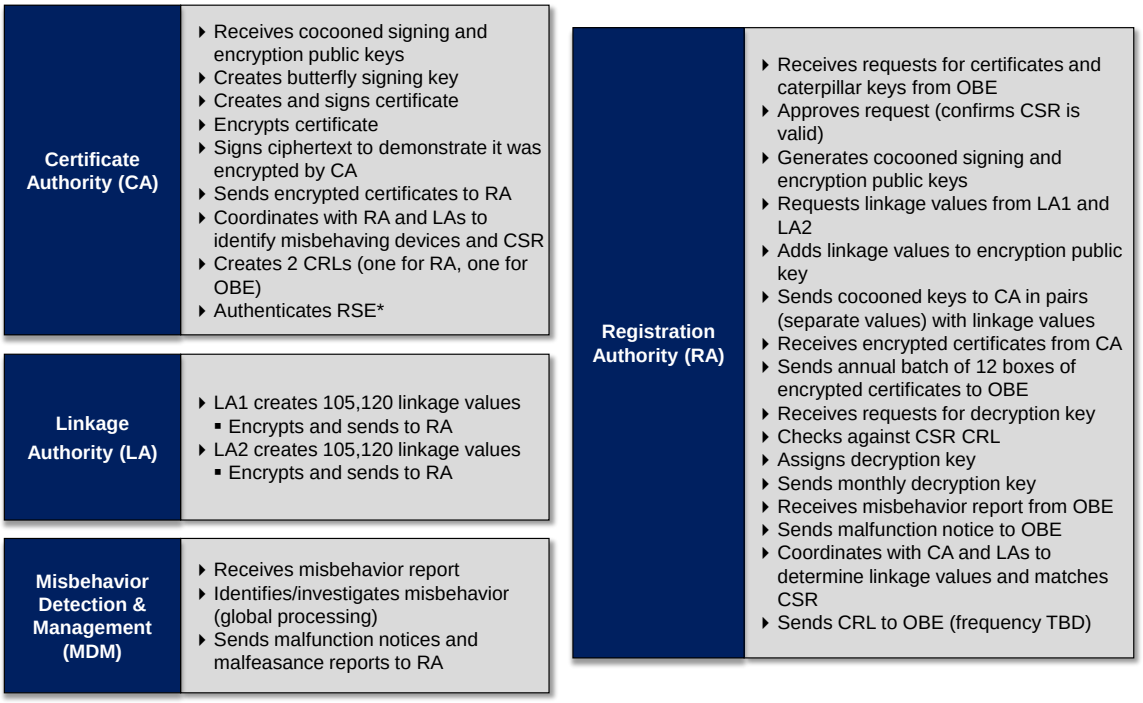

**Specifications for RSE certificate management have not been determined. Regardless, an authentication process for RSE will need to be designed.*

Each of these functions and its respective responsibilities within the CMEs plays a critical role in ensuring that certificates are created, sent, received, and monitored accurately and efficiently within the pseudo system for the connected vehicle system. A separate system, which is also described in the subsequent sections of this chapter, is responsible for the activation process. Each of these systems will be separate from each other and the OBE will interact only with the RA from the pseudo system.

Figure 5 illustrates the flow of activities and the process of performing the various responsibilities within each function, as described in the following sections of this chapter.

Figure 5. CME Process Flow

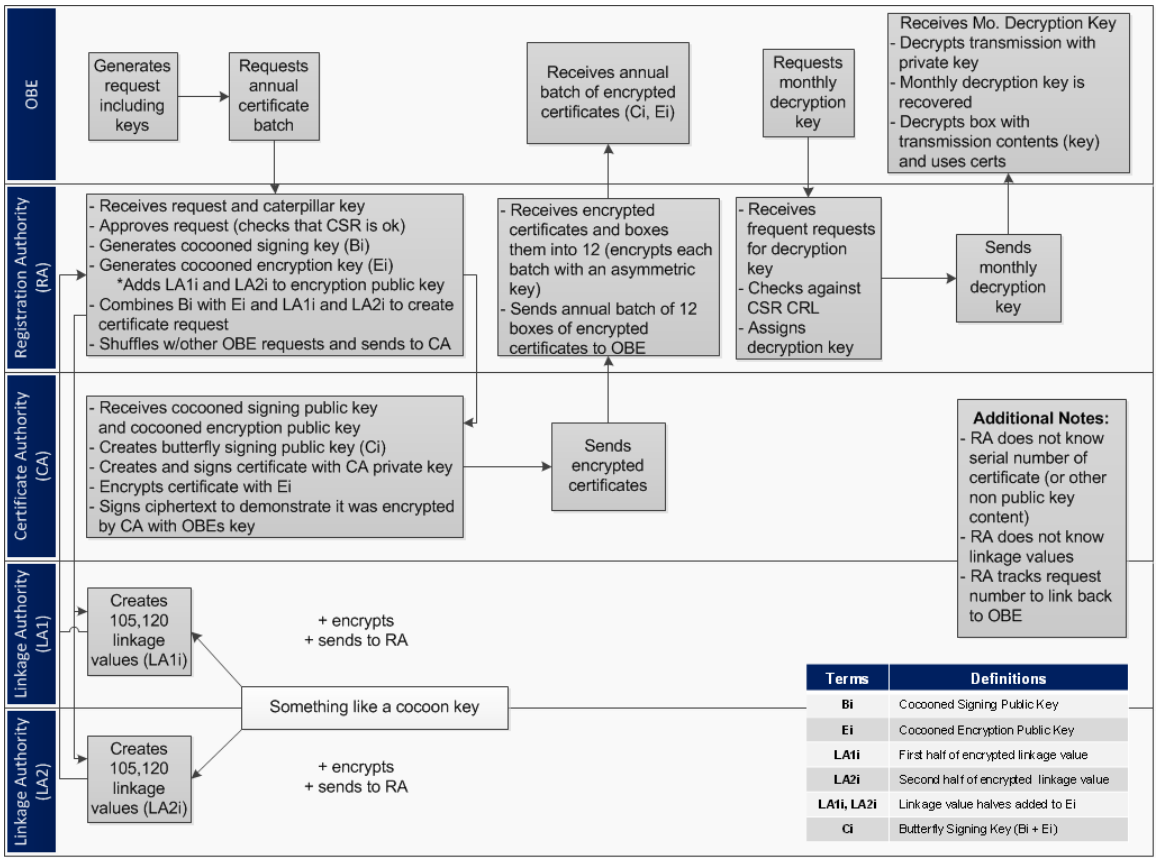

The CME processes described in Figure 5 can be summarized in the following points:

- The OBE will first create signing and encryption caterpillar key pairs. The OBE includes the public caterpillar keys in a certificate request for its batch of annual certificates, which it signs with its CSR certificate. The OBE then sends the request to the RA.
- Once the RA receives the request, it will first check against the CRL to ensure that the OBE CSR is valid.
- The RA expands the caterpillar keys into a set of 105,120 signing and encryption cocoon keys.
- The RA sends a request for linkage values to LA1 and LA2. *It is important to note that the RA in the pseudo system is the only CME function that interacts with the OBE.*
- The LAs will each produce a common identifier per batch of 105,120 as well as individual values for each certificate.
- The RA will add these encrypted linkage values to the signing and encryption public keys that it generates and will send them to the CA in sets so that each of the linkage values are represented for LA1 and LA2.
- The CA will create and issue certificates, encrypt each certificate with the corresponding encryption public key, and send them to the RA for distribution.

- The RA receives the encrypted certificates from the CA. Once all certificates for an OBE have been obtained, the RA will organize them into 12 "packages" that will each be encrypted with an asymmetric key.
- Each month the RA will receive a request from the OBE for the monthly decryption key to unlock that month's batch of certificates. Before the RA will assign the monthly decryption key it will check against the current CSR CRL to ensure that this CSR certificate is not revoked. If the CSR is not on the CRL, the RA will send the monthly decryption key to the OBE.
- This decryption key then allows the OBE to unlock the next month's batch of certificates to be used.

The MDM function was not included in this process flow since it has not been designed yet.

Linkage Authority

As mentioned previously, the LA function has been specifically developed for the CMEs in the connected vehicle system. This is driven by the scale of the system and security needs. Because of the large number of five-minute certificates (105,120), there needs to be an efficient method of revocation. An LA will produce a linkage value for each certificate with a common identifier that links batches of annual certificates, which will be the identifier placed on a CRL and will indicate revocation of the entire batch without having to place each individual certificate number on the CRL. Two LAs are necessary to satisfy the privacy concerns related to one LA being able to identify too much information about the OBE.

For each certificate set, the RA requests the LAs to provide linkage values. Each LA first generates a single value for each certificate set. It then calculates a value for each day and uses that value to encrypt the time period identifier for the day. This results in 105,120 unique values for five-minute certificates. Each LA provides this set of values to the RA for combination with the cocooned keys generated by the RA. The linkage values chain forward in time (i.e., the value of the "next" certificate linkage value is created using the previous value, but the process cannot be reversed).

When the CA receives the certificate request from the RA, it uses the pair of linkage values per certificate contained in the request to generate a certificate identifier for each issued certificate. The method used to create the certificate identifier and the large quantity of certificates being issued make it extremely difficult for an individual LA to identify which certificate used a specific linkage value.

In the event of a need to revoke a set of certificates, the RA, in combination with the CA and the LAs, identifies the set of values used to create the certificate. The resulting value is provided to the CA that will issue the CRL and place it on the CRL to allow the OBE to calculate which certificates are revoked.

Two Linkage Authorities

Because the set of linkage values forward chain, if the certificate identifier was created by a single LA, that entity would have the knowledge required to track a vehicle's location no matter how often the vehicle changed certificates. The suggested addition of another LA provides an additional security precaution to ensure that, as with the split between RA and CA, no one authority, function, or entity has the information needed to link multiple certificates to a specific vehicle for tracking or other security violation purposes.

CAMP has suggested that the two LAs should be housed in administratively and legally separate organizations in order to maintain the separation between them. This team believes it is possible that robust technical and procedural controls can be developed to separate the data and processes of the two LAs if they are collocated within one organization. Furthermore, having access to the two sets of values from the two LAs is not enough for one of the functions to unearth the individual certificate identifiers or OBE – additional information from the CA would be needed. A few highlights of the advantages and disadvantages associated with separating the LAs into separate organizations include the following critical points and are also highlighted in Table 1:

- The LAs are primarily automated functions which mean that their needed administration, personnel, and management are limited. If an entirely separate LA is included in the SCMS design, it will necessitate all of the associated administrative, legal, organizational, and operational overhead and structure to support it, all at additional costs.
- Although the notion of two LAs provides an added level of security by bringing in another value that has to be combined and transformed in combination with the keys produced by the CA and RA, the mere knowledge of both LAs would not be enough to recreate and identify an individual certificate (in the case of a breach of security between the two LAs in one organization).

Table 1. Organizational Considerations for Linkage Authorities

	Organizational Considerations for Linkage Authorities	
	One Single Entity for LAs	**Two Separate Entities for LAs**
Implications	▸ Increases need for internal controls	▸ Increases difficulty of collusion
	▸ Potentially decreases costs associated with organizational standup (e.g., PMO costs, staff, etc.)	▸ Potentially increases costs associated with organizational standup
	▸ Potentially decreases costs associated with building or leasing facilities	▸ Potentially increases costs associated with building or leasing facilities
	▸ May increase efficiency of communication with RA	▸ May lengthen process of communication with RA

Many very secure and critical systems include functions and data that need to be separate within one organization with the right technical, procedural, and physical controls to ensure that separation. Auditing and oversight of the controls provides the needed continuous view into compliance and enforcement of policies and procedures. We include several examples of successful technical and procedural controls that exist, upon which to base models for the CME.

Additional Technical Considerations

Following is a discussion of the current technical design of the SCMS. At this point, these are assumptions. If these assumptions change, it would not impact the models but could impact the system as a whole.

Certificate Life Span

Based on the need for short certificate life spans as a security measure, our analysis uses the following assumptions about the certificate life span. Future changes to these time spans will not necessarily change the CME structural and organizational models, but they should be noted to ensure that an accurate log of technical specifications is maintained and because they have significant impact on the overall cost and scalability of the system at full deployment.

- Short term certificates = five-minutes
  - Current technical design developed by CAMP and that is used as the basis for this CME analysis assumes five-minute certificates will be used in order to decrease the likelihood that a vehicle could be tracked through the monitoring of consecutive certificates.[8]

- The number of certificates required per OBE per year is the primary driver for scaling the CME processing requirements. Currently, the analysis assumes 105,120 certificates a year. If fewer certificates are required per year, that would require less processing time, which results in lower overall costs for the operation of the system.

- Overlap of short term certificates = 30 seconds
  - The overlap between certificates reduces the risk of not accepting a signed communication due to time synchronization issues.
- Back-up certificates = timeframe to be determined
  - These certificates provide a mechanism for vehicles that have not recently come into contact with an RA to send certificates when the OBE does not have any valid and decrypted five-minute certificates to use. The number of back-up certificates per vehicle has not yet been decided, nor have any policy decisions been made relating to back-up certificates.
- Batches of five-minute certificates are downloaded to an OBE once a year, totaling 105,120 certificates per yearly batch. This assumption limits the number of times an OBE must have high bandwidth connectivity with the CMEs and reduces the chances that an OBE will not have available certificates.

- Decryption keys are provided once a month to unlock monthly groups of the full yearly batch. This assumption limits the number of certificates available to an OBE that has been determined to be operating incorrectly.

[8] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.

- The life span of the CSR, which would represent the only connection to PII, is still undefined. Some have suggested a life span of one to two years, while others have suggested a life span upwards of 20 years.
  - A shorter life span of the CSR would require a mechanism for automated rekeying of the CSR in order to ensure no additional burden on the user.
  - Regardless of what life span is chosen, the timeframe needs to be such that the CSR will not expire when the yearly batch of certificates expires.
  - A longer life span would still require some form of rekey and may cause CRLs for the CSR certificates to grow very large due to vehicles removed from the fleet prior to expiration.

Certificate Signing Request (CSR) Certificate

In order to establish a procedure for a device to request annual batches of certificates, there must be a supporting authentication mechanism. CSR certificates are used for this purpose and allow the device to authenticate with its CSR certificate for both annual batches of certificates and monthly decryption keys. In order for the device to obtain an initial CSR certificate there must be a secure activation process. The options for how this process works are dependent on the type of device credentialing that is chosen, a full discussion of which is included in Chapter 6.

1. Under a scenario where there is no PII collected anywhere in the system, the CSR certificate can be installed on the device at the time of manufacture for built-in devices and at the time of installation for after-market devices. Fundamentally, once the device is operational, it will be authenticated to start receiving five-minute certificates and decryption keys. Not having PII ensures that user privacy is maintained, but also eliminates any accountability if there is malicious use of the certificates as there is no mechanism to track from a specific set of certificates back to the vehicle or the vehicle's owner.

2. Assuming that there is a desire for the ability to trace back from certificates to a specific vehicle or user, some minimum amount of PII will need to be collected and associated with the CSR. In order to limit any potential abuse of this information, it should be strictly maintained within the issuing CA environment and nothing directly related to the vehicle or owner would be included in any issued certificate. As the issuance and revocation mechanisms for CSR and five-minute certificates will be very different, and to provide maximum separation of PII from the certificate system, it is recommended that there be a separate CA responsible for the issuance of CSR certificates (we refer to this as the CA_{ACT}, see subsection below on activation System).

   - At this time, it is not clear if the CSR will be specific to the vehicle or the owner. There are several ways in which a device or a user could be authenticated, depending on the credentialing approach chosen. Some options include connection to VIN, collection of non-vehicle-based PII, or either of these options. These are described in Chapter 6.

   - Regardless of how the authentication happens under this scenario, there will likely be a separate storage location for the PII (either integrated into another system, such as the vehicle registration system, or in a separate, stand-alone system) that will isolate the PII from the actual CSR certificate that exists on the device. In this way, PII is not permanently stored on the device or on the certificate, nor can it be read by the RA function.

Once a device has been authenticated and it has a CSR certificate on it, subsequent requests of the RA for certificates and decryption keys can occur. The CSR acts as a verifier that the device should be in the system.

As previously mentioned, the CSR will periodically expire and need to be renewed. The requirement for periodic renewal of the CSR allows the OBE to be re-authenticated from time to time to ensure that only verified OBE are present in the system. It is envisioned that the renewal of the CSR will be automatic, so that users will not have to take specific actions for the renewal process to occur. Automatic renewal would ideally take place over the selected communications system (e.g., renewal through RSE) at the appropriate time to avoid expiration. Preventing the expiration of a CSR will ensure that a user's participation in the system is not interrupted.

In the event that a CSR did expire, if for instance a vehicle was not used for an extended period of time and the expiration date passed, a back-up certificate could be relied on to serve as a temporary verification of the OBE in order to send certificates. The technical specifications of back-up certificates have not been developed in detail in this report, as the decisions surrounding whether a back-up certificate can be maintained on an OBE have not yet been made. Presently, it is assumed that the CSR will be the primary method of verifying an OBE within the system, and it will periodically expire and require renewal.

Authentication System

Under a scenario where some PII is collected as part of the initialization of the device, an additional CME is proposed for authentication and assigning of the CSR certificate. The CSR will be used to authenticate the OBE for its one year batch of certificates from the pseudo system. For the purposes of ease of discussion the team will refer to the activation system CME as the CA_{ACT}. There is no need for the RA and CA within the activation system to be separated as in the pseudo system since the CSR is not used in location-based communications. Rather, it functions as an indicator to the pseudo RA that the device is a trusted member of the system and can receive its 105,120 five-minute certificates. Figure 6 outlines this process in more detail. It is worth noting that there are several options about how and where the activation system could operate.

- One option includes activation happening once a vehicle is sold (potentially by a dealer):
    - When the CA_{ACT} receives a request from the OBE, it will verify the information of the user and will assign the CSR to be distributed to the OBE.
    - This can work similarly for after-market devices.
- Another option is that the activation will happen at the time of manufacture:
    - When a vehicle is manufactured with an OBE, the manufacturer will request a CSR from the CA_{ACT} with a link to the VIN as the authorizing component. The CA_{ACT} will provide the CSR. This allows for the vehicle to always be active, even when stored on a vehicle lot, or being test driven.
    - After-market devices would have to have a separate process, similar to what was described previously (similar to the process of requesting a CSR once a vehicle has been delivered to a distributor or dealer).

Regardless of the method of activation, this would be the only place in the system that any type of PII might be collected. PII would not be stored on the CSR; it would be stored in a separate database to eliminate any PII from the CSR being accessible to the RA. It should be noted that this system and its functions would be separate from the system utilized to create, manage, distribute, monitor, and revoke five-minute certificates on a regular basis (refer to the Use Phase in Appendix A for further detail).

Figure 6. Authentication Process

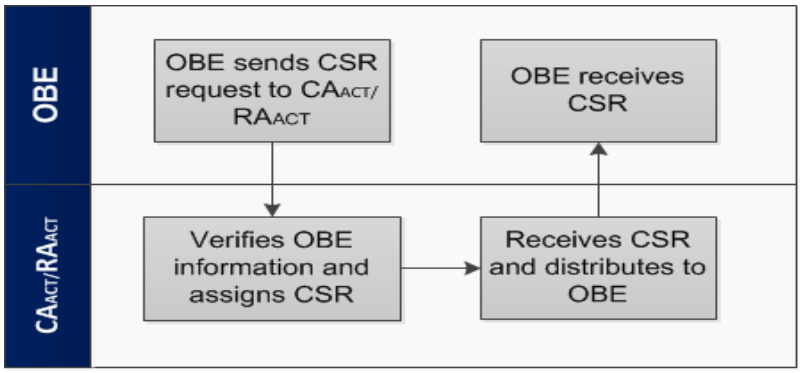

Under a scenario where no PII is collected, there would still have to be activation of the device. If there is no credentialing, initialization could happen at the time of manufacture, or installation of an after-market device. There would be no "check" that the user should be in the system if no PII is collected, and fundamentally, it would only require a pre-loaded CSR certificate that ensures that the device functions appropriately.

OBE Automation/Software

Much of the discussion to date about the CME functions, as described previously, has been based on an implicit assumption that as much automation as possible will be built into the OBE and its software. This includes programs that will automatically communicate with the RA for requests, reports, renewals of CSR and certificates, and other related activities. The current thinking around functions that should be automated within the OBE includes:

- Monthly requests for decryption keys
- Plausibility checks to ensure that a device is not misbehaving
- Plausibility checks on incoming messages and automatic rejection of messages coming from misbehaving devices
- Random selection of messages to put on a report to send to the RA for global processing
- Annual certificate batch requests
- CSR auto renewal
- CRL requests
- OBE and CRL processing

PKI Architecture and Hierarchy

The trust inherent in a PKI is based on the relying party (an OBE) knowing that a certificate issuer is one that is trusted. There are several possible architectures for implementing this trust relationship. There are advantages and disadvantages to each, and the need for backwards compatibility inherent in the connected vehicle system means that the decision on how to implement PKI trust must be established before the first production OBE are in development.

The basic premise is that just as the vehicles and infrastructure in the system need to be "trusted" by the use of certificates, the CME functions also need to be "trusted" by the vehicles or infrastructure receiving certificates from them. Before accepting a digitally signed message as valid, the OBE should look at the certificate that was issued and ensure:

- That it has not expired;
- That the CA that issued it is trusted; and
- That the certificate is not listed on a CRL.

The means to determine if the issuing CA and the RA and LAs are trusted is through a process called trust store management. This process ensures that the CA is self-signed and loaded directly into the trust store, or that the functions can trace the signature on their certificates through one or more CAs to a self-signed certificate loaded in the trust store. Because even root certificates expire at some point, there is a need to be able to update the trust store with new CA certificates and remove old ones as they expire or are no longer trusted. This is a way to ensure that the CA issuing certificates to the OBE, as well as the RA and LAs are authenticated as trusted parties and have a valid "CA (or RA or LA) certificate."

What follows is a description of various alternatives for establishing the structure of a trusted PKI architecture.

A single self-signed Root CA – In this structure, every relying party (e.g., OBE, RA, LA, RSE) only needs to know this one CA identity (that of the Root CA). That identity is loaded at the time the device is created. This is the simplest trust architecture but also has the least flexibility in the event of a catastrophic failure or a security breach of the Root CA. It also requires that this Root CA not expire or that there is a mechanism to update the CA within all relying party systems.

A hierarchical CA structure – In this structure, there is a single self-signed CA (Root) and one or more subordinate CAs (signing CA). This architecture provides the flexibility to add and remove signing CAs over time. The single Root is responsible for enforcing the security and technical policies on signing CAs. It requires that the relying parties have mechanisms to obtain signing CA certificates and, potentially, a new Root CA certificate. The signing CAs are those that would be communicating with RAs in the pseudo system.

Multiple hierarchical CA structures – This structure consists of two or more CA hierarchies. It is more flexible in that multiple independent organizations may provide CA services. It requires oversight and governance to ensure that each separate Root structure adheres to a minimum set of security and technical standards. This can be implemented using direct trust – where individual Root CAs are loaded into relying party systems or via cross trust agreements between PKIs (e.g., one Root is

trusted by the OBE and that Root issues cross certificates to the other Roots via a formalized process).

If the Root CA or CAs will expire, then the OBE design must incorporate mechanisms to do trust store management. Even in scenarios where there are long-lived certificates, at some point, some subset of devices will outlive even the longest lived certificate and require an update. There is a balance between the complexity of the PKI Hierarchy and the risk associated with a compromise of one of the CAs. A single Root CA is very simple but a catastrophic failure or compromise invalidates all of the certificates in the system. Using multiple Root CAs limits the damage that can be caused by any single security or catastrophic event.

Every PKI needs to have a specific policy or set of policies which govern its operations. In a traditional PKI such as the one that issues the certificates on the Federal Government Personal Identity Verification (PIV) cards, the policy is documented in what is referred to as a certificate policy (CP). This document describes the roles and responsibilities for implementing the PKI, the rules governing how certificates are obtained, the technical requirements for generation and protection of private keys and certificates, and the requirements for audit records and periodic compliance audits.

Organizational Boundaries

For the SCMS and the connected vehicle system to work, security and privacy need to be protected in both technical and procedural ways. Further discussion and examples of technical and procedural controls is included in the following section of this chapter. Thus far, we have described the functions independently, as each has its own requirements and specifications, as currently designed. In developing and analyzing various organizational configurations for operating these functions, certain threshold security measures have been built into the design of the system, including short-term certificates and the separation of the RA and CA functions. The addition of the LA is a mechanism to deal with the scale and security of the proposed system by providing an efficient way to revoke batches of certificates and to avoid any one entity having too much data that could allow someone to access information to track vehicles within the system. Another assumption, or rather working condition proposed, is that the best way to protect against attacks and unwarranted access to the system and vehicles is by creating distinct organizations, thus making it harder for people to share data across functions. The proposed models take the separation of the RA and CA, and the existence of two LAs, as working conditions. However, the models incorporate various combinations of functions into organizations or entities that support organizational efficiency.

Deciding on whether to use organizational boundaries or internal organizational procedural and technical controls to provide segmentation of data and protection against cross-functional collusion must be balanced by other considerations, such as cost, complexity, management, and oversight. The team addresses these issues in the discussion of CME models in Chapter 4.

Physical, Procedural, and Technical Controls

Physical, procedural, and technical controls are key features of a PKI system that should be considered when implementing the SCMS. The splitting of functions between legally separate organizations is one method of guarding against collusion within the CMEs, but it is not the only

method. Requiring functions to reside in administratively and legally separate organizations is a substitute for specifying the actual security controls needed to mitigate the risks to trip trackability based on the sharing of data by the CMEs. From a security perspective, administratively and legally separate organizations may provide separate security domains and would ensure that personnel are not shared across functions. However, those things can be provided inside a single organization if the policy requirements are clearly stated and the penalties for violation of the policies make it unlikely that any legal entity would not pay appropriate attention to the security requirements.

The principal avenues of attack that need to be addressed are:

- Physical Security. What are the "guards, gates, and guns" requirements to protect systems from unauthorized access? Physical access to a computer system makes many software attacks much harder to block. Do the CMEs need to be physically separated so that a successful physical attack on one does not provide physical access to them all? Physical separation is distinct from organizational separation.

- Logical Access. Any online system is subject to attack. The ability to exploit a vulnerability or compromise of an authenticated identity is always a possibility. For that reason, the CMEs must be engineered with high security requirements in mind and operated under a strict configuration control scheme that enables accurate and timely updates to mitigate security risks. The architecture of the system must also mitigate the ability of an intrusion into one CME allowing access to others. This would mean the CMEs need to be operated in separate security domains and the authentication mechanisms and privileges afforded via online access severely limited to no more than what is required for the specific task (e.g., there would be no ability to export "controlled" data based on an externally authenticated request).

- Insider Access. All computer systems are vulnerable to attacks by authorized insiders. Insiders, whether motivated by ideology or greed, provide an avenue to access systems no matter how good the physical and logical protections are. The insider threat is often considered the hardest to mitigate. Typically, systems that are to be operated at a high level of security have requirements related to the vetting of personnel who will operate the systems, separation of duties to keep any single individual from having too much access, and implementation of multi-party control on critical functions (such as configuration changes, access to controlled data, etc.). There can also be prohibitions from using personnel from one CME to support the operation of other CMEs.

Organizations can be designed with controls in place that allow multiple functions to operate within one structure while maintaining security. Regardless of the organizational structure chosen, there is a need for specific controls to ensure that inappropriate sharing of information does not occur. To better understand the physical, procedural, and technical controls in place in large PKIs that exist today, the team reviewed CPs for public entities, such as the Department of Defense and the Federal Bridge Certification Authority, as well as private entities, such as SAFE-BioPharma®[9] and CertiPath®.[10] The CP of a PKI is a document that outlines the policies for how certificates are created and used, along with details and guidelines about organizational design elements such as access to data and internal

[9] SAFE-BioPharma® is a registered trademark of SAFE-BioPharma, LLC.
[10] CertiPath® is a registered trademark of CertiPath, LLC.

controls. Every PKI must have a CP, and the X.509 standard for a CP is the template followed by PKI systems across industries throughout the world. Outlined in this section are examples of physical, procedural, and technical controls that can be applied to the CMEs.

PKI systems that follow the X.509 standard enumerate the physical and procedural controls in Section 5 of their CPs, in accordance with the X.509 CP template. Section 5 of these CPs is titled "Facility, Management, & Operational Controls." This section also describes personnel controls and audit logging procedures, but for the purposes of this discussion, they have all been included under the umbrella of procedural controls. Physical controls are intended to address the physical design elements of PKI equipment and the security of facilities and stored data. These controls are likely to involve the materials used to construct buildings or containers (e.g., steel, concrete), the types of locks necessary for different types of information, and the environmental conditions in which hardware and software should be stored (e.g., temperature of facility). Procedural controls provide direction for how processes are executed within the PKI. In addition to defining trusted roles and the responsibilities of staff, procedural controls also are used to specify the number of persons required per task, among other things. For example, the CP for the Federal Bridge Certification Authority specifies that when multiple parties are required for logical access to sensitive information, all parties must be in a trusted role, and at least one of the individuals present must be an Administrator. This type of control can be employed in the models reviewed in Chapter 4, where multiple functions are proposed to be collocated in the same organization.

Many controls are common across different PKI systems regardless of the information that is being protected. These controls are described in Table 2. Other controls will be tailored to the specific needs of the organization that manages the PKI. Some specific examples of physical and procedural controls tailored to specific organizational needs are listed in Table 3. All examples used in the subsequent tables are intended for illustrative purposes only.

Table 2. Common Physical and Procedural Controls for PKIs

Common Physical and Procedural Controls for Public Key Infrastructures	
Physical Control Examples	**Procedural Control Examples**
▸ Facilities for housing PKI equipment should be constructed using specified building materials (e.g., concrete walls and steel doors). ▸ When not in use, the CA equipment should be locked in containers that are appropriate for the classified information that the system is protecting, and should be stored separately from activation data. ▸ Environmental considerations such as air conditioning, water exposure, and fire prevention should be accounted for when designing a facility for the equipment. ▸ A security check to the facility housing the entity equipment should occur prior to leaving the facility unattended. Among other things, the check should verify that physical security systems (e.g., door locks, vent covers) are functioning properly.	"Trusted Roles" are clearly defined and responsibilities for each are outlined. ▸ Number of persons required for different tasks should be specified (e.g., multiple parties are often required to perform tasks associated with CA key generation at specific levels of assurance). ▸ System backups should be completed on a periodic schedule. ▸ Personnel controls should be implemented and encompass the qualifications and experience required to support the PKI system (e.g., background checks, security clearances, citizenship requirements, and/or trainings). ▸ CA operations should be administered by a person or body (e.g., Board of Directors). ▸ Audit log files should be generated for all events relating to the security of the PKI system.

Table 3. Specific Physical and Procedural Controls for PKIs

Specific Physical and Procedural Controls for Public Key Infrastructures	
Organization	**Physical/Procedural Control Example**
Department of Defense	▸ When classified government information is being protected by the system, the structure surrounding the equipment and any containers that hold equipment must be built to standards consistent with the classified information contained therein. ▸ Requires personnel to abide by strict qualifications.
Federal Bridge Certification Authority	▸ Specifies that executive branch agencies must follow policies for record archival consistent with the General Records Schedules established by the National Archives and Records Administration, or an agency-specific schedule.
Private Organizations (e.g., CertiPath, SAFE-BioPharma)	▸ Private organizations will often strive to adhere to the content of the X.509 certificate policy template, developed by the International Telecommunications Union.

Technical controls are used in conjunction with physical and procedural controls to ensure security of the PKI. PKIs following the X.509 standard outline technical controls in Section 6 of their CPs, in accordance with the X.509 CP template. Section 6 of these CPs is titled, "Technical Security Controls." Technical controls describe specific design aspects of the PKI hardware and software that ensure security of cryptographic material, especially in relation to the processes surrounding keys (e.g., generation, distribution, protection, and disposal). Table 4 describes technical controls that are

common among the PKI systems the team analyzed. The specific technical controls based on the unique design of the SCMS will need to be more fully specified as the CP for the connected vehicle system is authored and the organizations are stood up.

Table 4. Common Technical Controls for PKIs

Common Technical Controls for Public Key Infrastructures
Technical Control Examples
▸ States that key sizes are determined by algorithms scheduled to improve in efficiency over time. ▸ Establishes policies for private key protection, management, backup, and disposal. ▸ Specifies how activation data shall be used, protected, and controlled during the activation process. ▸ Lists specific computer security technical requirements, which differ between PKIs for public and private entities. Examples include the functionality of requiring authenticated logins and supporting recovery from key or system failure. ▸ Requires time stamping and synchronization of PKI entities with a time service such as the NIST Atomic Clock or the NIST Network Time Protocol (NTP) Service.

An aspect of CPs that is closely tied to technical controls is the different "levels of assurance." Levels of assurance are based on the Federal Information Processing Standard (FIPS) and are listed in a CP to define the amount of trust associated with a particular type of PKI user, as well as the security provided by the PKI itself. In this way, the level of assurance to which a participant belongs is a major part of defining the way that the user participates in the system. Different levels of assurance are associated with different technical controls. For example, a PKI system operating at a more advanced level of assurance might require that signing keys be generated in hardware cryptographic modules that meet higher FIPS standards than those PKI systems operating at a more basic level of assurance. Investigating the appropriate number of assurance levels for the connected vehicle system PKI will be an important part of the development of the CP. Table 5 lists the different assurance levels that are used among four industry PKIs.

Table 5. Levels of Assurance

Levels of Assurance Included in Select Public Key Infrastructures	
Organization	**Technical Control: Level of Assurance**
Department of Defense	▸ Specifies 9 levels of assurance for participants.
Federal Bridge Certification Authority	▸ Specifies 6 levels of assurance for participants.
CertiPath	▸ Specifies 9 levels of assurance for participants.
SAFE-BioPharma	▸ Specifies 3 levels of assurance for participants.

This section has discussed physical, procedural, and technical controls that should be considered for the CP for the connected vehicle system PKI, using existing X.509 standards as examples. It is

important also to consider elements of the system that are unprecedented and that may require new or unique controls. For instance, the physical separation of functions that have ”sharing” prohibitions (e.g., CA, RA, and LAs) should be considered, as it is reflected in the different models proposed. For example, housing the equipment that supports these (non-sharing) CME functions in different physical locations (e.g., rooms, buildings, etc.) could prevent the possibility for a physical security breach (e.g., breaking in to a facility) in one place to compromise entire sets of information from the different entities. These types of prohibitions and controls can be used to provide the needed security and separation of functions without placing them in entirely separate legal/administrative organizations.

This chapter reviewed how the key functions and responsibilities within the CME work together to process certificates and ensure the system remains secure and efficient. Understanding these processes allows for the analysis of how grouping specific functions together within the same entity can either create efficiencies or concerns for the SCMS, discussed in detail in the next Chapter (4), which also outlines how putting physical, technical, and procedural controls in place can be used to address security concerns, while realizing organizational and operational efficiencies.

Chapter 4 CME Models

The organizational models being evaluated are derived from the grouping of functions into different entities. To clarify – when functions are represented to be in one entity by a single box in a model, it implies that the operations, governance, administration, and legal definition/boundaries of an organization are distinct and independent of another organization or entity represented by a different box. As mentioned previously, there is an assumption that separating functions into distinct administrative and legal entities provides the highest levels of security and protection against data and information crossing functional lines. Individuals would have to collude across organizational boundaries in order to breach security protocols and controls in place. This assumption has not been tested, nor has it been agreed upon by all stakeholders, and the team maintains that it might not always be valid. Chapter 3 includes a detailed description of technical, procedural, and physical controls that exist in other industries, providing heightened security and separation of functions and data.

As the team examined all configurations of major functions, models that were eliminated consisted of models that housed the CA and RA, or CA and LA together. Due to the nature and scale of the connected vehicle system under a full deployment scenario, it is not viable to have the CA and RA together without strong (organizational) boundaries between them. Separation of these functions presents a basic level of privacy and communications security protection which acts as a baseline for this project.

Models with the combination of LA and CA within one organization were also eliminated. That configuration is assumed to include too high a risk of the CA being able to have access to the linkage values, therefore being able to identify any one particular vehicle. Note that the CAMP analysis proposes two LAs; it is this team's contention that neither should be collocated or in the same organization as the CA.

Following an initial presentation to USDOT and the public in December 2011, a few approaches to organizing CME functions in different operational and organizational models were chosen for further exploration. The differences between these models are operational and organizational and don't change the analyses of security and privacy protections. The team presents alternative organizational models, with focus on one (Model G) that provides the highest levels of organizational and operational efficiency. This paper also presents additional analyses on multiple topics that affect CME functions and thus the organizational models. Figure 7 includes the three high-level organizational models that are subsequently modified based on the additional analyses presented.

Figure 7. CME Models

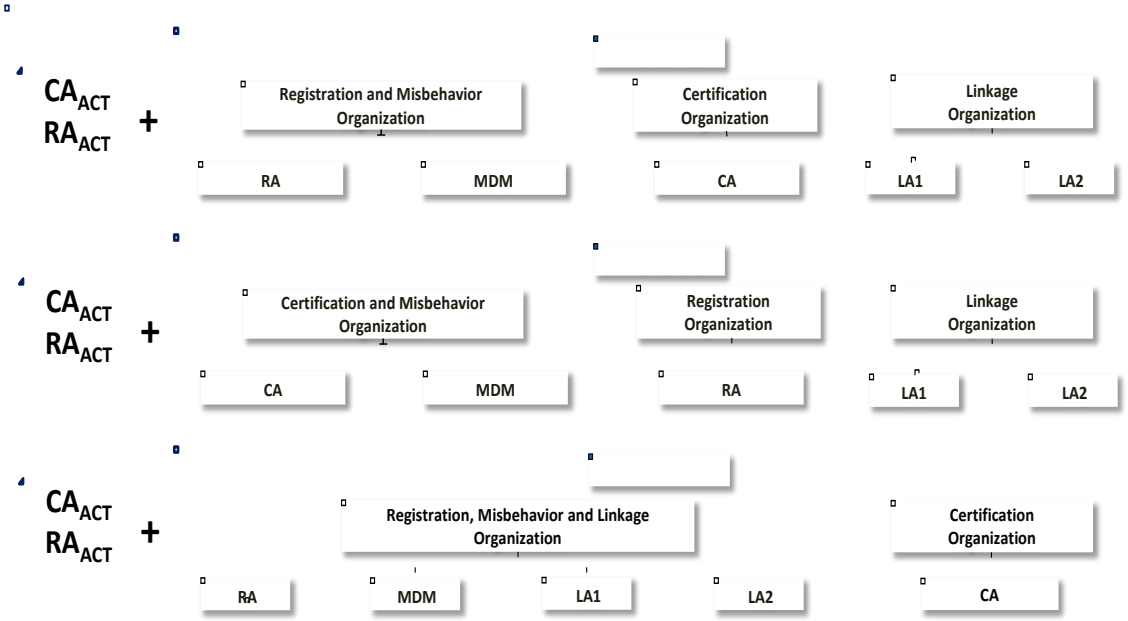

Model Differences

In this section of the report, each model is described, highlighting the differences and anticipated advantages and disadvantages of each. Note that there has been no technical specification of the MDM function, so its inclusion in the models at this point is based only on the conceptual understanding of what it will do, as described in Chapter 3.

Model C1

This model groups the RA and MDM functions together, with CA in a stand-alone entity and the two LAs in another stand-alone entity. This configuration adheres to the needed organizational separation of CA and RA, as well as CA and LAs. The main benefits of this model would come from the organizational and physical efficiencies provided by the RA and MDM functions being housed in one organization. Because the MDM would be communicating with the RA regularly to send misbehavior reports, there would be some savings from backhaul communications systems. In addition, because MDM is anticipated to be a primarily automated function with little human involvement, it is conceivable that maintenance, oversight, and management staff that perform certain RA functions would be able to cross-utilize their time with MDM activities as well. It should also be noted that although location strategies have not yet been determined, the team believes that there will have to be more geographically distributed RAs than CAs for the system (see discussion in Chapter 10). This may prove to be a disadvantage to this model depending on how the distribution of the MDM evolves, as there may not need to be as many MDMs as RAs.

The combination of the LAs into one organization provides some physical location efficiencies and savings, though, as noted previously, there will have to be strict technical, physical, and procedural

controls in place to ensure the separation of the two LAs. However, the team anticipates that these will also be primarily automated functions with little human involvement and collocating them under appropriate controls would make sense from a cost and location savings perspective.

Model D1

As with Model C1, this model leaves the LAs in their own stand-alone entity, though in this case the RA is by itself and the CA and MDM are grouped together. The advantages of model C1, vis-à-vis the LAs, are the same in this model. However, it may in fact prove to be more efficient, both operationally and financially, to group the CA with the MDM. Based on the concept of what the MDM will do and how the CA will eventually need to communicate and cooperate with the other functions in order to trace a misbehaving device back to an original CSR, *if that policy decision is made*, there will be communication and logistical efficiencies to be gained from the collocation and organizational combination of the MDM and CA. Furthermore, the team believes that it may prove beneficial for there to be one or two MDM locations, as this could more easily be built out with the one or two CA locations.

Model G

Model G presents the most efficient combination of functions – bringing RA, MDM, and both LAs into one organization. This would provide the most flexibility and opportunities for cost savings and operational efficiencies. Under robust controls (as with all models) that protect the integrity and needed separation of functions, the four functions could share physical locations, some oversight and management staff, and other non-security-sensitive resources. The communication between the LAs and the RA would be further facilitated as well. This model provides the greatest opportunities for efficiencies, logistics ease, and standup of a fewer number of entities, without gaps in security or additional risks, as can be evaluated at this nascent stage of the project and specification of functions.

In developing the models the team took input from stakeholder groups (namely the automobile industry through the Vehicle Infrastructure Integration Consortium [VIIC]) that suggested that the two LAs be separated into distinct legal entities. This flows from the assumption that legal separation provides more security than do procedural, technical, and physical controls. The team investigated the nature of different legal, technical, and physical controls available and in use in other industries in order to provide safeguards against corruption of data or crossing functional lines. Chapter 3 includes a more detailed discussion of these controls. The team believes that the LAs can be collocated with the appropriate controls and oversight in place, and there are organizational efficiencies and reduction of possible errors that would come from that colocation. For sake of completeness, Figure 8 shows additional modes that separate the two LAs. Other models were proposed that included the CA and RA together or CA and LA together in one entity, which was deemed unacceptable, so are not shown.

Figure 8. Additional Models Proposed by VIIC

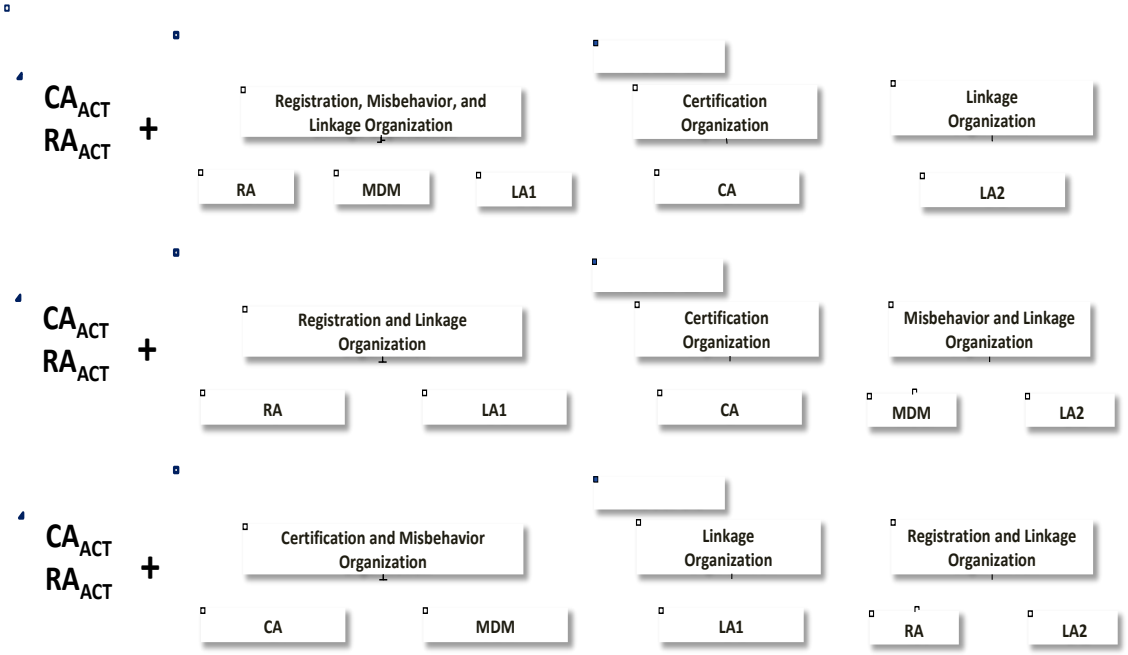

Note that the additional models include one of the LA functions placed within another entity that does not house its companion LA function. The cost chapter (Chapter 10) explores the differences in cost in more detail, as estimated by the operational and organizational efficiencies gained by combining functions into entities when possible. The main question arises from the total separation of an LA by itself (see models C2 and D2) and whether that additional expense of operating a small, mostly automated function in an entity of its own is justified by the assumption of the need for legal separation to ensure appropriate levels of security. This team believes that technical, procedural and physical controls are robust enough to obviate the need for the LAs to be in separate organizations. The advantages and disadvantages are described in Chapter 3.

Models C2, C3, and D2

These models, as noted, were suggested by the VIIC, based on the assumption that the LAs need to be housed in distinct organizations. This team believes that the security risk posed by grouping the LAs in one organization do not outweigh the costs and loss of efficiencies, as well as other resources sharing, and physical location sharing. The team refers to other industries and organizations that provide robust and highly secure controls to maintain separation of data and functions when needed. Model C2 poses additional cost and logistic challenges due to the separation of the LAs – it is questionable whether a function as automated as an LA would benefit from having to stand up its own organization, complete with all organizational and administrative oversight, management, legal, operational, physical, and other requirements.

The real differences between the models lie in implementation and cost implications. Grouping of large functional divisions allows for significant anticipated efficiencies and reduction of cost because of the ability to share personnel, some equipment, and physical locations. The team believes that data

and functional separation can be maintained in these cases through the use of effective procedural and technical controls, as described previously in Chapter 3.

An additional note is that the existence of an organization to house a particular function does not imply the number of physical locations or machines that may be required to administer its operations. For example, depending on processing needs and capabilities, and estimates of scale and possible geographical structure, there may need to be several locations across the nation that operate the functions within a legal and administrative entity, based on the model chosen. If Model G were selected, for example, there could be several locations or geographic centers for operating the Certificate Organization. The decisions about numbers and locations of physical entities will be predicated on scale of the system, as well as policy guidance.

Roles and Responsibilities

Each organization, housing various functions (or one function in some cases) will be responsible for the tasks and activities associated with each function, as previously described in Chapter 3. In addition, standard management and internal organizational governance activities and functions will be required of each organization, in order to ensure that compliance with policies, performance, internal practices, and resource management are all monitored. It will be necessary for each organization to communicate with the other organizations within the CMEs in order to perform the processes of creating, encrypting, batching, distributing, and revoking certificates. Policy guidance will be one of the driving forces in determining how functions and organizations communicate with each other and what kind of industry oversight will monitor the practices of all CMEs. As stated above, although it is premature to specify these policies now, it would be necessary to identify and outline these policies as progress continues towards the implementation.

Throughout this report, the team identified elements within the CMEs for which specific policies or standard operating procedures will need to be developed. The team believes it is premature to develop actual policies and rules of access as well as internal CME standard operating procedures without a full understanding of the technical needs of some of these functions or a decision of how the functions will be grouped. Included here are some initial high level descriptions of key areas in which identification of processes will need to happen with reference to where within the paper more detailed discussions exist.

- Managing user access and certificate issuance: Guiding policy about authorized system users will need to be developed. There are examples of other vehicle registration systems, such as state registration laws, that can serve as guidance in this area. Drivers with criminal records related to potential system attacks, for example, may be excluded from participation. In addition there is an assumption that in order to authorize users and conform to any individual user access guidelines, there will need to be a credentialing process. Details about options for credentialing users and including that process within existing registration or title systems is included in Chapter 6.
    - Based on the team's analyses, we believe the credentialing process will be a responsibility of the Activation CA that provides devices with initial certificates that allow them to be included in the certificate management system and processes. The CA will

be the entity tasked with issuing certificates to OBE and will be the function responsible for complying with any policy rules that govern certificate issuance.

- Misbehavior detection, certificate revocation, distribution of revocation information, and retrieval and storage of data for misbehavior detection purposes: All of these issues, while important to ongoing system integrity and building of user trust are parts of a technical design for misbehavior detection and management that have not yet been developed. The team has included in the organizational models a function, MDM that will ultimately be tasked with operating and conforming to any technical and policy decisions that guide misbehavior detection processes, revocation distribution, and retrieval and storage of data for these purposes.

- Database backup: In the team's discussion of disaster recovery plans, and in the development of cost models, we have accounted for estimated data backup and storage needs. Chapter 11 includes details of anticipated disaster recovery plans as included in the current core system design. In addition, the team built in a 30 percent hardware and software premium to account for backup of data.[11] The working assumption is that each function will manage its own backup and storage of critical data with appropriate technical, procedural and physical controls in place to guard against unauthorized access and potential data corruption or insider threat. A detailed discussion of these types of controls and examples of other industries that use them successfully can be found in Chapter 3.

- Security – physical and logical access: As with data backup and storage issues, physical, logical, and/or authorized access to data and physical locations must be tightly controlled. The team has included a detailed discussion of several technical, physical, and procedural control options that exist to ensure security of the system and all its functions (Chapter 3). The particular standard operating procedures and processes that guide levels of access will have to be developed in conjunction with policy guidance, but it is fundamental that the only people with access to data and information that could potentially lead to privacy or security violations be strictly controlled, monitored, and measured regularly. Consequences for violations of these processes and procedures must also be specified by guiding policy and enforcement mechanisms established. At a minimum, it is reasonable to assume that the only personnel with authorized access to any function's information would be those who must work with the data, and that no personnel from one function should have access to another function's data and processes, including during certain processes, like certificate revocation, that demand the sharing of information. It is premature to specify how those access rules should be implemented or enforced, but they can range from strict physical controls that require multiple users with different codes or authorization, to procedural and organizational controls. Please see Chapter 3 for more discussion.

- System performance metrics: It is critical that system performance metrics are included at all levels of an organization, for all CMEs. Levels include organizational, group, and individual.

[11] Research indicated that 10 percent is commonly used for back up of software systems. The team added and additional 20 percent based on the criticality of the system and the desire to produce conservative estimates at this point in time.

The metrics developed must relate directly back to goals and objectives of the system as a whole and the individual CME, as well as any group's role within that CME. Chapter 11 includes a description of a model that is commonly used to develop and evaluate performance metrics as well as initial concepts for how to measure various system needs.

- Audit policies and procedures: Auditing procedures and regulatory deterrence discourage malfeasance and provide consequences for violators. The team completed research in the health records, electronic voting, and payment card industries to reflect policies and procedures used and how they ensure compliance. Details for each industry are provided below.

    - Audits for HIPAA compliance are federally mandated and conducted by the Department of Health and Human Services (HHS). A majority of healthcare providers conduct routine audits of internal systems and technical controls to avoid fines and license revocation.

    - Electronic voting systems use real time audit logs to ensure accuracy of vote count by producing a printout of every individual vote without PII in the event a recount is needed. Individuals are punished for engaging in electronic voting misbehavior at different levels depending on the offense.

    - The Payment Card Industry (PCI) establishes security guidelines known as the Data Security Standards (PCI DSS™).[12] Common security measures include external network vulnerability scans, wireless intrusion detection and prevention systems (WIDS/WIPS). Merchants who do not comply face penalties that may include fines and participation revocation.

  Audit procedures also examine compliance with privacy guidelines and any standards followed for PII collection and storage, which will need to be considered. A further discussion regarding PII collection is provided in Chapter 6.

- Managing user privacy protection: Managing and ensuring user privacy throughout the system is a priority of all involved. There are several perceptions of what levels of privacy protection are needed or desired. For the purposes of a complete analysis, the team included several options for credentialing users, from not collecting any personal information, to collecting new and comprehensive information. A detailed discussion of these options and their implications for the CMEs is included in Chapter 6.

- Enforcement: As guiding policy for the connected vehicle system is decided, enforcement and outcomes for misbehavior will need to be specified. At this point, it is premature to speculate on what enforcement policies will be for total system policy. In addition, all CMEs will have to institute internal enforcement policies and standards to ensure compliance with organizational policies, processes, and controls, especially in cases where controls are in place to guard against potential security breaches. Appropriate incentive and reward systems within CMEs could be

[12] PCI DSS™ is a trademark of the PCI Security Standards Council.

implemented side by side with enforcement policies to encourage desired behaviors at all levels of the organizations.

- System administration and maintenance: Included in the description of each function as well as in the cost estimates for functions and CMEs are administration, management, and maintenance costs. It can be assumed that maintenance of hardware and software used in the system will conform to state-of-the-art standard operating procedures for the systems used, with continuous refresh and evaluation. Part of the performance management metrics will need to be evaluation of ongoing maintenance and administration, in addition to functional operations.

It is anticipated that the connected vehicle system will cross national borders as it is implemented. This implies that any organizational model or structure should be able to accommodate cross-border implications. No model presented in this paper has an advantage over another with respect to the ability to evolve into entities and a system that can coordinate and communicate across national borders. The needs and processes by which cross-border coordination should happen are driven by policy and technical decisions and guidance. Some of the considerations that will need to be built into the operations and structures of CMEs in order to accommodate cross-border operations and coordination include:

- Technical compatibility
- Communications system transcending political borders
- Fee and other revenue agreements
- Enforcement and policy adherence and agreements
- Reciprocal sharing of information when needed

Chapter 5 Baselining

In order to compare the various ways of configuring certificate management entities and functions, it is important to catalog a baseline of security. As mentioned in previous chapters, the primary goals of the CMEs are to provide a trusted system of secure communications and to protect users' privacy. The team's approach to determining a baseline against which the various CME models and other aspects of the system can be compared, was to determine the specific needs of the PKI for the connected vehicle system and to analyze how other organizations have addressed security baselining within PKI systems. This chapter describes the findings from this analysis.

Baselining in a traditional sense would allow for the specification of acceptable levels of the risks associated with any system. Once an acceptable level of risk or threat is identified, various approaches can be compared against this baseline to evaluate whether they meet the mark. For the CMEs, there is a significant challenge in defining the security baseline because of the novelty of the system and the anticipated full deployment scale. The team has gathered information about risks and security standards from various comparable industries and scenarios, but it is important to note that, as with all aspects of the connected vehicle program, there is no one-to-one comparison to date.

Potential Threats to the CMEs

Before delving into a discussion about security baselining, it is important to first understand the threats that are currently anticipated. Six significant risk categories associated with the CMEs and the connected vehicle system are based off of analyses by CAMP[13] and additional research by the team. These are outlined below.

- **Privacy violations**: Situations where user PII is accessed inappropriately (if PII exists in the system)
- **Risk to the entire system**: Involves attacks that can burden the system or create false messages. This risk category includes the scenario of a user being placed in an unsafe situation because the system failed to send a proper message, or no message was sent at all. Thus far, these possibilities are envisioned to include the following types of attacks:
    - **Software Manipulation**: Installing malicious software on a user's OBE to create false messages
    - **Sensor Manipulation**: Interfering with a vehicle's sensor output
    - **Denial of Service**: Preventing a user from receiving messages
        - **Denial of Computation**: Sending large amounts of bogus messages, overwhelming the OBE

[13] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.

  - **Denial of Communication**: Jamming the wireless band with a powerful signal, blocking vehicles from transmitting messages

- **Vehicle Tracking**: Opportunities or risks that allow someone (inside or outside of the system) to identify a particular vehicle and then track it
- **Message Linking**: Malicious use of certificates to identify or locate a vehicle
  - **Syntactic Linking**: Using static identifiers from certificates to establish that multiple messages come from the same vehicle
  - **Semantic Linking**: Using dynamic identifiers to "join the dots" and reconstruct a vehicle's trajectory
- **Device Cloning**: The act of reverse engineering, copying, or simulating the OBE, RSE, mobile device, or any other physical infrastructure that will participate in the connected vehicle system. Industry solutions to device cloning have been developed with technology that uses Physically Unclonable Functions (PUF), which is unpredictable even if an attacker has physical access in the CME. Although PUF has advanced in recent years and is intended to be essentially unclonable,[14] PUF can still be susceptible to certain attacks (e.g., reverse engineering).[15] The different performance measurement methodologies for PUF vary in their evaluation of how effective this technology truly is. Additionally, the method in which it could be incorporated into the OBE would need to be further investigated by technical design teams.
- **Disruption of Infrastructure Components:** Opportunities or risks to RSE hardware that could affect the system depending on RA and OBE interaction

Because a thorough understanding of the threats to the system is fundamental to defining a security baseline, it is important that new threats are documented, understood, and prevented as the system is implemented and rolled out in the future. Undoubtedly, as malicious technologies grow more advanced, the SCMS will need to respond and adjust its preventative measures over time.

Systems Security

As stated previously, several potential risks to the security of the communications within the system exist. The premise of communications security is that the exchanges of data between vehicles (and between vehicles and infrastructure or other devices) occur on a network and through a system that provides authentication, monitoring, privacy protection, and misbehavior detection and enforcement to all users in the system. The decision to use PKI as the foundation of that trusted system is due to the reliability of PKI in providing security of communications.

The concept of a security baseline for this system is complex. Two aspects of an overall baseline for security are: 1) PKI design baseline, and 2) a security vulnerability baseline. First we review what PKI

[14] Mudit Bhargava, et al., *Attack Resistant Sense Amplifier Based PUFs (SA-PUF) with Deterministic and Controllable Reliability of PUF Responses*.
[15] Mehrdad Majzoobi, et al., *Testing Techniques for Hardware Security.*

offers in terms of standardized communications security in order to provide a PKI design baseline against which to compare various CME configurations. Later in this chapter we take a deeper look at how vulnerability can be examined. Audit procedures are one method used to evaluate and guard against breaches to security of a system.

PKI Baseline

The prototype for test studies in the connected vehicle program is based on PKI for the security system. This is also the current assumption about what will be used for the security system at full deployment of the connected vehicle system. The reasons for this choice have been documented in several studies.[16] The PKI system under development allows users in an unsecure public network to securely and privately exchange data (via certificates) by using a public and private cryptographic key pair, the public key portion of which is distributed through a trusted authority.[17] This protocol offers the greatest levels of security of communications exchanges and protection from the tracking and identification of users within the system. For the purposes of the connected vehicle system, standard PKI systems have been altered in order to provide additional layers of security.

Two additions to standard PKI systems provide the foundation upon which additional changes are configured. The RA and CA are organizationally separated. They exchange limited information in ways that are designed to ensure that neither the RA nor the CA will have the needed information to track a vehicle. In addition, because of the scale of the system, there has been an addition of two LAs in order to provide an efficient way of batching certificates for placing them on the CRL when necessary.

Given these requirements, the current PKI baseline for CME configuration includes an RA, a CA, two LAs, and an MDM function. Table 6 summarizes the details of the PKI baseline.

Table 6. PKI Baselining Details

Baselining Topic	Details related to the CMEs
Three primary risk categories	▸ Privacy violations: PII being accessed ▸ Risk to the entire system: Attacks that can burden the system or create false messages ▸ Risk of tracking a vehicle: Opportunities or risks that allow someone (inside or outside of the system) to identify a particular vehicle and then track it
Differences between CMEs and standard PKI systems	▸ RA and CA are split ▸ Two LAs are added ▸ MDM is included
PKI design baseline	▸ A PKI system with the following functions: RA, CA, two LAs, and MDM

The goal of identifying a baseline for security was to understand how the proposed models address the needs of the connected vehicle system and what an acceptable level of vulnerability would be for

[16] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.
[17] Ibid.

the system. The team has conducted significant research into other PKIs in order to find this information. There is no consensus across industries or organizations about what constitutes allowable security breaches or malfunctions. Furthermore, as noted previously, it is difficult to assess the unique challenges the CMEs may face in regard to the monitoring, protection, and recovery from attacks, given the novelty of the technical architecture. Although documentation regarding vulnerabilities within other PKI systems in comparative industries is limited, the team was able to gather some industry examples.

Security Vulnerability Baseline

The nature of a PKI trust model with the CA serving as a single body that must sign all certificates has a goal of eliminating security vulnerabilities. Nonetheless, there may be some level of risk expected for end users or trusted agents, and any known vulnerability in the system used by the RA or CA needs to be prevented or mitigated as soon as possible. In attempting to understand how existing security systems approach the issue of vulnerability, the team researched several organizations and systems that use PKI to protect the security and privacy of users. Table 7 at the end of this section summarizes the examples the team reviewed.

Establishing Baseline Security Methods: ICAO and Machine Readable Travel Documents

The International Civil Aviation Organization (ICAO) standards for machine readable travel documents (a.k.a. ePassports) define a "Baseline Security Method" that countries must follow to be ICAO compliant. The Baseline Security method, in this case a process called Passive Authentication, represents the minimum that a country must to do protect the user's sensitive PII from being compromised. ICAO also defines a set of "Advanced Security Methods" that countries may elect to follow, such as Basic Access Control and Data Encryption, but that are not required to be ICAO compliant. This approach to security baselining for a PKI is essentially defined by the attacks that are sought to be prevented. It provides participants with a certain level of flexibility by allowing them to implement whatever additional security controls they feel are necessary, as long as they meet the basic standard. Although a level of acceptable vulnerability for a large scale PKI system like the CMEs is difficult to define, there are resources for the measurement and evaluation of vulnerabilities for *hardware and software*, which are included below.

Evaluating IT Vulnerabilities: Payment Card Industry Vulnerability Scans

Different tools are available for the assessment of IT vulnerabilities associated with the infrastructure used by the CMEs. The National Vulnerability Database (NVD) was developed by National Institute for Standards and Technology (NIST) to serve as a repository of different vulnerability measurement and evaluation tools. It is used by both public and private agencies to categorize vulnerabilities and communicate them to other entities in a common language. One of these tools is the Common Vulnerabilities and Exposures (CVE®)[18] listing. The CVE is essentially a dictionary of "different information security vulnerabilities and exposures that aims to provide common names for publicly known problems."[19] The CVE would be useful for specifying how a vulnerability is written (using XML

[18] CVE® is a registered trademark of The MITRE Corporation.

[19] The MITRE Corporation., CVE website: http://cve.mitre.org/about/faqs.html#a1.

specifications), which is important for cataloguing system vulnerabilities and identifying appropriate solutions.

A second tool included in the NVD that is applicable to the infrastructure used by the CMEs is the Common Vulnerability Scoring System (CVSS), Version 2.0. The CVSS, which is maintained by the Forum of Incident Response and Security Teams, uses a Base Score derived from an algorithm that measures a vulnerability's exploitability and impact to the system. The Base Score is measured on a 10-point scale, with the most serious vulnerabilities rated at 10.0. According to the NVD, the "CVSS is well suited as a standard measurement system for industries, organizations, and governments that need accurate and consistent vulnerability impact scores."[20] In fact, the PCI uses the CVSS to rate vulnerabilities discovered through the external system scans that are required as part of PCI DSS compliance. When Approved Scanning Vendors scan merchant computer systems, any vulnerabilities that are found are assigned a rating of high, medium, or low severity, based on the associated CVSS score. One aspect of achieving PCI DSS compliance is that merchants must not have any system vulnerabilities that are considered to be medium or high severity, which translates to a CVSS Base Score of 4.0 or higher.

The vulnerability tools included in the NVD can assist in the development of a security baseline for the CMEs and provide commonly accepted ways to define an acceptable level of vulnerability. It is important to understand that these tools are not comprehensive, however, because they are intended to evaluate only the IT vulnerabilities associated with the infrastructure. Tools such as the CVE and CVSS Version 2.0 do not address other vulnerabilities such as insider threats, side-channel monitoring, and physical access to the facilities housing equipment for the CMEs.

Identifying Vulnerabilities Through Compliance Auditing: The U.S. Department of Defense

The Department of Defense (DoD) uses PKI to achieve secure communications between Agencies or Services within DoD, and also between DoD Agencies or Services and specific external organizations or entities (e.g., contractors, foreign allies, etc.). The certificate policies of large scale PKIs like that of the DoD include brief sections that discuss vulnerability, but it is addressed primarily in terms of compliance audits, audit logging procedures, and other records that are to be checked by appropriate third parties. These non-specific requirements are in place because the DoD PKI is essentially a collection of smaller PKI systems, and instead of defining specific levels of acceptable vulnerability, the certificate policies simply outline technical and procedural controls that are required, and state that each participant must be in compliance with its certification practice statement. Instead of providing a clearly identified threshold for what is acceptable in relation to breaches, data spills, or other factors related to security of the system and protection against risks, these broad certificate policies rely on compliance auditing to identify vulnerabilities in the system.

Table 7 summarizes the approaches that various industries use to protect against threats that may exist within PKI systems.

[20] NIST, NVD website: http://nvd.nist.gov/cvss.cfm?version=2.

Table 7. Methods of Addressing Security Vulnerability

Organization – Topic	Method of addressing security vulnerability
International Civil Aviation Organization (ICAO) – ePassports	▸ Defines 'Baseline Security Method' that countries must follow for the processing of machine readable travel documents, and additional 'Advanced Security Methods' that countries may elect to follow. ▸ This approach names a specific security measure that participants must take instead of defining a numerable factor that is acceptable for breaches, identity theft crimes, etc.
Payment Card Industry Security Standards Council – PCI DSS	▸ Requires regular system scans to identify IT vulnerabilities in merchant computer systems. Any IT vulnerabilities with a CVSS score of 4.0 or above are unacceptable. ▸ The approach identifies a threshold above which existing IT vulnerabilities are not acceptable. However, it does not account for non-IT vulnerabilities, such as internal malfeasance or physical tampering.
Department of Defense, Policy Management Authority – PKI for Identity Management	▸ Specifies general requirements for compliance auditing and maintenance of audit logs for participating PKI systems. Audits are designed to evaluate adherence to each participant's certification practice statement. ▸ This approach relies on auditing to identify instances of non-compliance with the security measures that participating PKIs claim to follow. This general requirement for auditing is common among large scale PKI systems such as that of the Department of Defense and Federal PKI Policy Authority (FPKIPA) that are essentially collections of smaller PKI systems.

The development of a baseline for security of the CMEs resulted in an analysis of PKI design and system vulnerability. By outlining the current understanding of potential threats to the CMEs and how the organizations have been designed to address them, a baseline for PKI design was identified as the CA, RA, two LAs, and MDM functions. The goal of any PKI is to reduce and remove any vulnerabilities detected in the system through different tools, some of which have been reviewed in this section. Because the connected vehicle system has not yet been launched, it is impossible to know for certain which attacks will occur most frequently and what their impact will be on the system. Industry research indicates that vulnerability baselines have been defined in terms of the preventative measures used and the severity level of IT related vulnerabilities.

Chapter 6 Personal Privacy Protection

The USDOT has not yet made a determination about whether, and to what extent, CMEs will need to collect PII in order to permit users to participate in the connected vehicle system. For this reason, to ensure that the policy analysis is comprehensive, this chapter considers all relevant collection options, including the option of collecting no PII and keeping user credentials anonymous – an approach advocated by motor vehicle manufacturers and privacy advocates.

In essence, there are only two basic options relating to collection of PII within the SCMS. One option is to collect no information that could link an OBE to an individual or a particular vehicle. This would ensure complete anonymity throughout the system, and if assumed or implemented, will limit the risks to be guarded against only to system and trip trackability. This option protects personal privacy to the maximum extent possible, but creates no mechanisms for tracing or identifying attackers or others who use the system for malicious or unlawful purposes. A second option is to collect enough PII to be able to trace a misbehaving vehicle (either because of locally undetectable technical malfunction or human malfeasance) back to an individual in order to repair a technical problem or enforce policy, regulations, or laws against attacks on the system. USDOT requested an additional analysis of the option of PII being placed directly on the actual certificates, which is also included in this chapter. The analysis revealed that due to the significant privacy and security risks inherent in this type of system (i.e., increased hacking, trip trackability, etc.), it is likely not an attractive option to any stakeholder group.

The collection of PII by the CMEs poses more of a risk to individual privacy than collection of no PII. However, as discussed below, there are various ways to collect PII which differ in terms of how much information is collected and the processes by which this collection occurs. This report suggests that one way to minimize the risk to personal privacy inherent in collecting PII would be collection and storage of PII in an isolated part of the SCMS with articulation of specific technical protections and policy guidelines designed to limit access and prevent disclosure.

The subsequent sections of this chapter describe the differences in the relevant credentialing options and discuss the operational implications of each option. The options discussed include:

1. Collection of no PII (total anonymity)
2. Collection of PII in the activation system (SCMS)
    a. Using an existing registration system
    b. Using a new registration system
3. Collection of PII within the pseudo system

1. Collection of No PII

Proponents of a SCMS that collects and stores no PII maintain that the safety benefits gained from V2V technology outweigh any downside to not being able to track down malicious attackers or users, or even connect misbehaving devices to individuals. They insist that the CMEs should collect no

information that could be used to connect an OBE with an individual or particular vehicle. The security and privacy implications of this approach include:

- No ability to enforce legal or policy consequences on malicious users.
- No need for a separate activation system, though activation would still happen, most likely at time of manufacture. For after-market devices, a process of activation would need to be specified, but it would not include collection or linking to any forms of PII.
- Reduced concern about violations of individuals' privacy because there is no individual information collected or stored by the system.

It remains an open question whether the SCMSs inability to trace or identify hackers and others who might use the system for malicious or illegal purposes could undermine user confidence in, or acceptance of, the V2V system. One can argue that users would feel more comfortable knowing that there is no feasible way for their PII to be compromised. However, over time, the immunity for bad actors that the system effectively creates could reduce user trust and the safety benefits of V2V technology, if users withdraw from the system. Ultimately, decision makers will need to compare the privacy benefits of not collecting PII with the security weakness inherent in a system that lacks any mechanism for ensuring accountability for malicious or illegal actions. Rather than base its decision solely on the speculation of a select group of V2V stakeholders, the USDOT is considering conducting or sponsoring more research on user acceptance that specifically targets these issues, to better inform its upcoming regulatory decisions.

2. Collection of PII During Activation

One option for collection of PII is to collect limited amounts of information that can be used to identify individual vehicles or drivers during the activation phase of the certificate management life cycle (see previous discussion about activation in Chapter 3). The entities that manage the certificate authorization process for the activation system (CA_{ACT}) would be separate and isolated from the entities that manage, administer, and assign five-minute certificates, and track and manage misbehavior detection (pseudo system). The personal information collected would only be used for initial activation or in cases where misbehavior needs to be traced back to an individual in order to enforce policy and legal consequences for malicious behavior. Policy and organizational rules that govern access to personal information need to be set, and the technical separation and ability to access that information needs to be strongly controlled and minimized. One way to pursue this approach is for the PII collected to be no more, and possibly much less, than what currently exists across several registration and certification systems today. The consequences of this approach are as follows:

- No personal information is included in any certificates.
- The collection of PII only occurs within the activation phase or system; it is separate in all ways from the pseudo system.
- The only time a connection back to the activation system and PII is needed from the pseudo system is when a maliciously behaving user is to be identified for compliance and policy enforcement. Policy decisions must guide the previous point – when and how to connect back to users.
- Technical, policy, and administrative or legal controls will exist to provide separation of activation from the pseudo system.

- No ongoing connection to activation databases is needed for determining authorization.
- Rules of access for employees within the activation system must guide who has access to PII data.
- Reduced burden to CMEs in responding to discovery requests (the more PII data that is archived and retained, the more burden CMEs will face in responding to such requests).
- If PII is collected, current laws in many states and at the federal level may require notice and consent for data collection, retention, and transfer between entities.

Credentialing Approaches

There are several methods of providing the authentication credentials described in the previous section: leveraging existing PII-collection systems (and several considerations within this option), and collecting new PII within a new registration system. This report includes consideration and description of ways to integrate the CA_{ACT} within the existing vehicle title registration systems, as one of those options, in the following subsection. While it is evident that this approach is not in line with some stakeholder communities' desire to maintain total system anonymity, it has been developed in a way that provides anonymity throughout the pseudo system that the OBE will communicate with regularly. Regardless of the collection method or source of PII, the team's approach has been to analyze how PII could be used to verify a user's involvement in the system. Any personal data would be kept separate and protected from the rest of the CMEs, and not included on any device or certificates. Details surrounding where the sequestered PII would be stored have not been decided upon, but it is possible that the information could be maintained in a database by an organization separate from the CMEs. As mentioned previously, this information would only be accessed by appropriate parties in instances where misbehavior needs to be linked back to a specific OBE and user. Specification for rules of access and the management of the misbehavior process would be defined by policy and regulations.

2a. Leverage Existing PII Collection

Instead of overseeing a new system to collect, store, and manage user PII, the SCMS may be able to leverage a system that already collects PII. One approach to the collection of PII is to associate each OBE with the VIN of the vehicle to which it belongs, by integrating into existing vehicle registration and title systems. The implications of this are two-fold:

- First, each OBE must be linked to a VIN.
- Second, if decided by policy, the appropriate enforcement authority must be able to use the VIN to trace back to a specific user in the case of enforcing consequences for human malfeasance in the system.

The VIN, a piece of second-order PII, is the critical factor in this approach. First order PII refers to PII that is linked directly to the person (e.g., name, address). Second order PII (such as VIN) is one step removed as it does not identify an individual directly but may be used to link back to an individual. The VIN could be linked to the OBE when the vehicle is manufactured or when the first owner purchases the vehicle. A protected database of VIN and OBE pairings must be maintained by the CA_{ACT}, a third party contractor, or some other entity. In the event of misbehavior, the CMEs would be able to trace a malicious certificate back to a specific OBE per legal and policy protocols, and pass this information on to the appropriate enforcement authority. This authority would reference the VIN and OBE

database to match the OBE with a specific VIN, and then use the VIN to identify a driver through vehicle title information owned by state DMVs.

Although leveraging existing state registration and title systems would obviate the need for the CMEs to associate an OBE directly with an individual and store massive amounts of first order PII, it does beg the question of *how* law enforcement officials would communicate with the 50 separate state DMVs in the event of enforcing consequences of misbehavior. Forming a connection between the CA_{ACT} function and/or the appropriate law enforcement agency and each state's vehicle registration system is one solution, but this would require a great deal of coordination with disparate bodies, which could jeopardize the future progress of the connected vehicle system due to the scale of the system and organizational complexities implied.

Options for Vehicle Title and Registration Aggregators

Different systems exist in both the public and private sectors that aggregate vehicle title and registration data collected by the state Departments of Motor Vehicles (DMVs). The National Motor Vehicle Title Information System (NMVTIS) is a primary example of one that could be leveraged for the connected vehicle system. Overseen by the Department of Justice (DOJ) and managed by the American Association of Motor Vehicle Administrators (AAMVA), NMVTIS was established by the Anti-Car Theft Act of 1992. NMVTIS offers a national perspective on vehicle history by linking each state's vehicle registration information into one central location. By collecting daily data updates, some in real time, from state DMVs, NMVTIS allows states and individual drivers alike to inquire about a vehicle's history through submission of a VIN. These reports are intended to help protect consumers and states from title fraud. One category of data that states are required to submit for use in NMVTIS is specific titling information that includes the name of the individual or entity that has title to a vehicle.[21]

According to the NMVTIS Program Office, "NMVTIS is the only vehicle history database in the nation to which all states, insurance carriers, and junk and salvage yards are required by federal law to report."[22] As of July 2012, 88% of U.S. DMV data had been integrated into NMVTIS, and approximately 10 states and the District of Columbia are still working with DOJ and AAMVA to engage in the system as full participation is required by law.[23] Although there are aspects of NMVTIS that would need to be further explored (e.g., whether its scope could be expanded for the purposes of the connected vehicle system, the specific details of its privacy policy, and system capacity), it is a valuable tool that could potentially eliminate the need for the CA_{ACT} to maintain a substantial database of user PII or for law enforcement officials to coordinate VIN inquiries independently with each state.

Private companies that specialize in market research and information services represent other possibilities that could be leveraged for the purposes of the CA_{ACT}. Polk®[24] is a leader in the automotive industry when it comes to marketing information and data services. The company has "the most comprehensive automotive market and vehicle owner data available in the marketplace," including information for approximately 600 million unique vehicles dating back 20 years.[25] Unlike

[21] Department of Justice, NMVTIS website: http://www.nmvtis.gov/nmvtis_states.html.
[22] NMVTIS Program Office, *Don't Be Fooled* brochure: www.nmvtis.gov/NMVTIS_Consumer.pdf.
[23] Department of Justice, NMVTIS website: http://www.nmvtis.gov/nmvtis_states.html.
[24] Polk® is a registered trademark of R.L. Polk & Co.
[25] R.L. Polk & Co., Polk Approach website: https://www.polk.com/approach/data_and_technology.

NMVTIS, to which states and other entities are required to submit information, Polk gathers data from various sources based on commercial agreements. Presently, Polk has agreements with the DMVs of all 50 states, Washington, D.C., and Puerto Rico that allow it to access registration and title data for current drivers, including VIN data and user PII such as name and address. The company uses this information primarily to assist other businesses, such as vehicle manufacturers, with marketing efforts and business necessities like product recalls. It is feasible to consider integration of the CA_{ACT} into Polk's existing systems. The technical aspects and legal implications of any integration would need to be further explored, particularly in relation to the Drivers Privacy Protection Act, which controls the ways that Polk can access and utilize driver PII in its activities.

At this early stage, an important consideration in evaluating existing systems is whether each contains information that is comprehensive, current, and relevant to the purposes of the SCMS (i.e., identifying a bad actor in cases of identified misbehavior). Major tradeoffs between options include the fact that NMVTIS compliance is mandated by federal law, whereas Polk's agreements are privately negotiated with state DMVs and other sources. As of the writing of this report, all states have not reached full participation with NMVTIS, while Polk maintains consistently updated data from all states, and has done so for several years. If the decision is made that there needs to be integration with existing vehicle title and registration aggregators, considerations such as these will need to be taken into account in order to choose an appropriate system to assist or host the CA_{ACT} function.

Policy guidance for specific rules about participation in the CMEs has yet to be determined, though one imagines that there may be rules of access to the system in general that would prohibit certain potential users from inclusion. Examples include those individuals who are not allowed to register motor vehicles and those who have been found to participate in the kinds of attacks that are feared for the SCMS such as hacking into the system. Leveraging the existing state registration and title systems through NMVTIS or commercial vehicle title and registration aggregators has to be vetted from a technical standpoint, but as of this team's current research it seems to present a technically feasible option. The additional implications of this approach are as follows:

- No additional collection of PII than what is currently collected.
- No need for a new organization to operate the CA_{ACT} function – reducing both costs and organizational complexity.
- Need for cross jurisdictional coordination or a centralized system that operates across state borders. This is a function that may be able to be provided by third party contractors (e.g., AAMVA).

As mentioned above, additional research into the three primary options for integrating connected vehicle system credentialing into existing systems will need to include further analysis of the technical and legal implications, as well as any associated costs of integration.

Credentialing Different User Types Through an Existing System

Using an existing system to credential users may need to be adjusted based on the type of user within the connected vehicle system. Credentialing for each user type may imply a different process based on the most feasible method of connecting a device to its registered owner. Regardless of user type, any data collected under this approach would be stored in an existing database for use in investigating malfeasant users or technical malfunctions in accordance with policy or regulations. The following describes how different user types may be credentialed through existing systems:

- **For passenger motor vehicles (light vehicles)** with integrated OBE, the approach would be that at the time of manufacture, the OBE would be linked to the vehicle VIN (as in the previous discussion about leveraging existing systems, which would involve VINs). When a user registers the car and takes title, that VIN would also be associated with the user over existing registration systems. The CA_{ACT} would take part in authenticating the user at the same time that an individual registers the car – there would be a dual functionality to the existing registration function that also authenticates the user and provides a CSR certificate. The technical, procedural, and policy regulations that guard against release of personal data would have to be specified within this approach. Enforcement of the standards and regulations would fall to the industry oversight structure.
- **For commercial heavy vehicles**, the process would be similar to that for light vehicles but the link would instead be to the USDOT registration number. It is important to note that this is only applicable to trucks involved in interstate commerce, as they are the ones that use the USDOT registration number. For trucks that are involved in intrastate commerce and are less than 26,000 lbs., the registration system is through the normal state DMVs, and so the authentication process would be similar to that for light vehicles. An additional note is that the USDOT registration number registers the company, not the vehicle. It is printed on the vehicle to show the company is registered appropriately. Using this number would imply that any trace back to an individual would have to happen in coordination with the registered company.

  The overall implications for commercial vehicles would be that there is no additional system collecting PII. The activation of trucks would be embedded in the registration system, and protocols, policy, and procedural and technical controls and guidance would govern access to the PII.
- **For After-Market Safety Devices (ASDs) that are integrated into or connected to light vehicles**, the approach would be to include an activation link at the authorized installer that would communicate with the existing vehicle registration systems as for vehicles that have the device integrated through OBE. The connection between device and VIN would not be performed by the manufacturer, but instead by the installer, through the CA_{ACT} function that would be embedded in the existing vehicle registration systems.
- **For transit vehicles and public fleets**, a central difference compared with passenger vehicles is that they are owned and operated by municipalities instead of individuals. Therefore, a misbehavior investigation using a VIN would link back to the agency that owns the vehicle, not necessarily the individual driver. However, it is not uncommon for transit operators and drivers of public fleet vehicles (such as police officers, firefighters, and emergency responders) to log into a mobile data terminal featuring an automatic vehicle location device at the beginning of each shift. This process authenticates the identity of the driver and in most cases includes some collection of PII. This would allow an agency to track back to an individual at any point in time if misbehavior is reported and therefore no additional PII is needed for this user type.

2b. Create a New PII Collection System

An alternative approach to collecting PII and activating and registering devices would be to implement a system that would necessitate a separate registration of the device through a CA_{ACT} organization collecting individual PII that is not linked to VIN. The amount of PII collected in this scenario could be anything from name and address to social security number. This approach would provide the same benefits as the approach outlined previously, where the device is connected to VIN, but also includes several other risks and associated implications, as described here:

- Need to stand up a new organization for CA_{ACT} – additional costs and organizational complexity
- Creation of a system that collects individual information in addition to what is currently collected and stored
- New policies that guide collection, storage, and access to PII
- Potential for non-participation or increased resistance from the public due to additional privacy concerns

There is a potential additional benefit in this scenario that would apply to a system which includes non-vehicle devices such as mobile phones and tablets (V2X). As the connected vehicle system evolves and expands to include additional users on mobile devices, there may need to be a process or function that allows for the tracing of misbehavior back to these non-vehicle users. The creation of a new PII collection system could potentially allow for all user data to be in one database (instead of divided among existing systems), which may make the process of identifying malfeasant users or misbehaving devices much easier.

2a or 2b. Credentialing Through Existing System or Through a New System

Credentialing approaches do not have to be mutually exclusive. It is feasible that both leveraging existing registration systems *and* creating a new PII collection system could be implemented, based on the needs of different user types. As mentioned previously, non-vehicle-based users such as those on nomadic devices will need to become trusted users of the system, so that their messages are known to be authenticated and non-harmful, as with vehicle-based users. Ensuring that all users within the system are trusted, regardless of what kind of device they use, provides integrity and instills user confidence. This could be accomplished through a separate system that collects PII on an opt-in basis, depending on the applications anticipated to be used through the nomadic device. The implications of this approach (including either existing systems or creating a new system) are similar to the implications of each of the stand-alone approaches. However, there is an additional implication that stems from the fact that this would require additional resources and organizational complexity than would be required for either stand-alone system.

3. Direct Linking of Credentials to Certificates

As part of a thorough analysis, the team also examined one final option: that of directly linking all certificates (five-minute certificates and back-up certificates if they exist) to some form of user PII (first or second order). A potential benefit of this last credentialing option is that it would make the misbehavior enforcement process more efficient. A CA that detects malfeasance would not have to use the activation system to trace misbehavior to an individual in order to commence whatever enforcement action is required by policy. A security system in which certificates are directly linked to

PII still would be subject to procedural and technical controls on how and when PII could be accessed or shared with outside organizations (e.g., law enforcement agencies), in accordance with law, regulations, and system policies. However, the privacy and security risks inherent in this type of system, particularly vehicle tracking, would be significantly higher than with a credentialing system that either collects no PII or uses limited PII found only in the CA_{ACT} system. Such a system would present significant opportunities for collusion, hacking, or other methods of gaining access to PII. Based on the team's analysis and stakeholder perspectives, the team does not view this credentialing option as viable. The potential risks are high compared to other approaches presented previously, all of which would ensure the security levels required for the connected vehicle system, and trip anonymity.

Table 8 below reviews the implications to the CMEs based on the credentialing approaches for activation. Even under the approach of collecting limited PII during the activation phase, the model proposes a separation of that system from the pseudo system that manages and distributes the five-minute certificates. Those certificates would not contain any personal information nor would there be an active connection with the activation system that could possibly be breached or compromised to connect back to PII. The PII data would be sequestered and protected based on legal, technical, and/or policy controls and guidelines that would ensure that access to PII is limited to identification of, and enforcement against, misbehaving system participants, in accordance with applicable system policy.

Table 8. Credentialing Types

Type of Credentialing	Implications to SCMS
No Credentialing	
No PII is Collected	▸ Perceived to increase user acceptance and participation ▸ Unable to track back to prosecute malfeasance/bad actors
Credentialing within the Activation System	
Create New PII Collection System	▸ Increases costs and organizational complexity ▸ Duplicates information already collected by other systems ▸ Requires new policies and regulations for protection of PII ▸ May increase resistance to the system ▸ Requires the need for a separate database of PII to be maintained
Leverage Existing PII Collection System	▸ Reduces costs and organizational complexity since the activation system will not be needed ▸ Requires a centralized system that is used across jurisdictions ▸ Decreases ability to collect any other PII ▸ Potentially increases trust of system participants since no additional PII is collected
Credentialing within the Pseudo System	
Direct Linking of Credentials to Certificates	▸ Increases opportunity for collusion or hacking of PII since PII will be included in each certificate ▸ Eliminates the need for the activation system (CA_{ACT})

USDOT Criteria Review

In developing conceptual processes and approaches to different ways of credentialing, and how to credential various user groups, the team vetted each option against a list of criteria provided by USDOT. Each method of credentialing and each user group within the method of leveraging existing registration systems has been included in the table and measured against the criteria, shown in Table 9. Table 10 further explains the “Technical Feasibility” criterion highlighted in red in Table 9.

Table 9. USDOT Criteria

Criteria	Criteria Detail	Credentialing Options			
		1. No PII	2a. Thru VIN	2b. Non Veh PII	3. Direct Link
Technical Feasibility	Impact to security system (would this change the anticipated level of security in the system)	Yes	None	None	Yes (high)
	Back Office or CA_{ACT} policy choice (is a policy decision needed to implement this choice)	Yes	Yes	Yes	No
	Possible from technical standpoint	Yes	Yes	Yes	??
	Technical implications as they relate to additional applications beyond safety (V2I, V2X)	*See Table Below*			
	Technical implications that would make option desirable	Yes	Yes	No	No
Ability to Leverage Existing System	Ability to leverage existing motor vehicle or driver registration systems	No	Yes	No	Yes
	Reduce scale of CME functions (Act = activation system, P = pseudo system)	Yes for Act Decreases P	Yes for Act No impact to P	No for Act No impact to P	Yes for Act Increases P
Security	Pose any special risk to security of PKI design	Yes	No	No	No
	Which poses the least risk to security and why	All the same, no impact to security			
	Impact to extra-system enforcement as through law enforcement/state/Federal agencies	No way to take legal action against system attack	Guided by Policy (one more step)	Guided by Policy	Easier
Privacy	Does it have any impact on privacy	No	Yes, Low	Yes, Medium	Yes, High
	Which poses the least risk to privacy and why	Least risk to privacy	See no PII		
Scope of Data Collection	Scope of data collected and maintained by the RA (Act = activation system, Maintain = maintaining of data)	None	Act = 0, Maintain = Database pairing	Significant amount of data in both	Act = 0, Maintain = Increase data storage
Number of Transactions	Estimated number of transactions for the CA_{ACT} annually	None	Once Ever	Once Ever	Yearly
Ease of System Use/Implementation	Will the options have an impact on participation in V2V	Yes/Negative	No	Yes/Negative	Yes/Negative

Table 10. Technical Feasibility

	1. No PII	2a. Thru VIN	2b. Non Veh PII	3. Direct Link
V2V	No ability to trace to bad actors	Need to explore options for integrating into commercial vendors	Collection, storage and access to PII	Collection, storage, and access would need to expand size of certificates
V2I	No ability to trace to bad actors	Infrastructure does not have VIN	Collection, storage and access to PII	Collection, storage, and access would need to expand size of certificates
V2X	No ability to trace to bad actors	Need to explore options for connecting nomadic device through commercial vendors	Collection, storage and access to PII	Collection, storage, and access would need to expand size of certificates

Privacy Protections in Comparative Industries

As previously mentioned, if PII is to be collected by the SCMS it must be safeguarded to ensure that total system acceptance is not adversely affected. Individuals' PII is collected in many industries for different purposes, and several methodologies are employed to protect it. It is common for public and private organizations alike to rely on privacy regulations, industry guidelines, and technical security measures to protect PII that is collected as part of business processes. Each of these topics is described in this section.

Private organizations often look to federal laws and regulations as a basis, or a reference point, on which to build their own privacy policies. In the electronic tolling industry, the collection of driver PII and payment data expedites the process of moving vehicles through toll plazas. The E-ZPass®[26] Interagency Group (IAG) allows the different toll authorities that make up the E-ZPass network to institute their own policies for management of driver PII. However, to be in full compliance with the IAG standards, at a minimum, member toll authorities must comply with a privacy policy that is based largely on the requirements of the Drivers Privacy Protection Act. The Drivers Privacy Protection Act is a federal statute that governs the use and access to personal data by state Departments of Motor Vehicles, contained in Chapter 123 of Title 18 of the United States Code. In addition, each toll

[26] E-ZPass® is a registered trademark of the Port Authority of New York and New Jersey.

operator must comply with the state and local regulations that control PII collection, storage, and access.

For the healthcare industry, the Health Insurance Portability and Accountability Act (HIPAA) established the legal requirements for Electronic Health Records (EHR). This Act includes rules for the collection, storage, transmittal, privacy, security, access, and auditing of EHRs. All users of EHR must comply with HIPAA, which protects all information contained in health records, from PII to medical records and diagnoses.

The primary trade association of the payment card industry, the PCI Security Standards Council, is an example of a group of private organizations that has created a body to establish industry-specific privacy guidelines. Payment brand companies such as Visa®[27] and MasterCard®[28] worked together to establish the PCI DSS to ensure that merchants and service providers protect cardholder data to the greatest extent possible. The PCI DSS includes 12 requirements that touch on issues such as network security and physical access to data. Other industries, such as the electronic tolling industry, have incorporated the PCI DSS standard into their evaluation of privacy in payment transactions with customers.

Privacy laws and regulations, as well as voluntary self-regulation of privacy by private industry groups, are all central to the protection of PII. However, technical security measures are what make the protection a reality, as discussed in Chapter 3. The industries discussed thus far use technical protections such as a PKI and data encryption, antivirus protection, and external system scans in order to safeguard data. Electronic voting systems have become commonplace and Election Assistance Commissions at the state level assist with the implementation of E-voting security protections recommended by NIST. Some of these protections include smartcard requirements, audit logs, alert technology, and screening for unacceptable participants, such as ex-convict voters.

OnStar®[29] and SiriusXM™[30] Satellite Radio

At the recommendation of the USDOT, the team researched OnStar and SiriusXM Satellite Radio as services that could offer insights applicable to the CMEs. While these services are somewhat related to the CMEs in that they collect driver PII and communicate data to millions of users, Table 11 reflects their key differences.

[27] Visa® is a registered trademark of Visa International Services Association.
[28] MasterCard® is a registered trademark of MasterCard International, Inc.
[29] OnStar® is a registered trademark of OnStar, LLC.
[30] SiriusXM™ is a trademark of Sirius XM Radio, Inc.

Table 11. OnStar and SiriusXM Key Differences

	OnStar	Sirius	SCMS
Voluntary	Yes – users willingly submit first-order PII	Yes – users willingly submit first-order PII	If NHTSA mandates the use of connected vehicle technologies in new motor vehicles and it is determined that, for enforcement purposes, PII collection is necessary, both the service and the PII collection will be mandatory – neither voluntary nor opt-in
Fees	Monthly subscription	Monthly subscription	USDOT has stated that users should not be required to pay a subscription fee to use safety applications in a connected vehicle system
Trip Trackability	Uses GPS technology that is intended to track vehicle for purpose of sending assistance	N/A	A major priority for the SCMS is to prevent tracking of a system user's vehicle
Scale of System	6 million users	21 million users	250 million users at full deployment

In addition, SiriusXM relies on satellites to transmit radio signals to users' vehicles, whereas the use of satellite technology has been eliminated as a possibility for the communications platform of the CMEs.[31]

Conclusions

The protection of user PII is of utmost importance within the SCMS, and the way to ensure that zero privacy violations occur is to collect no PII. However, as detailed previously, if the security system collects no PII, there is no way of identifying or taking enforcement actions against misbehaving system participants or other bad actors. The lack of an effective misbehavior management and enforcement scheme could threaten user confidence in much the same way as the collection and misuse of PII.

There are credentialing options that balance the need to protect user privacy with the need to track down, identify, and take enforcement actions against bad actors adversely affecting the system. Some of these options leverage existing systems such as the NMVTIS. The team's analysis of comparative industries provide some examples of how PII is effectively collected, stored, and destroyed, and it is evident that some examples are more applicable than others to the CMEs. Ultimately, the team finds that the organization and functions of the SCMS will require a balance between the need for security and the need for some degree of user anonymity.

[31] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.

Chapter 7 Expansion of Users to Infrastructure and Mobile Devices

To date, most of the CME discussions, analyses, and models have involved a connected vehicle system that includes only vehicles communicating with each other – V2V communications. However, a fully realized connected vehicle system is envisioned to encompass communications and messages exchanged between vehicles, RSE, and other devices, such as cellular phones. Integrating V2I and V2X communications into the existing CME models will require additional research and considerations. As with other areas of investigation and development for the CMEs, a number of technical and policy decisions that have not yet been made will guide implementation of the additional levels of communications with infrastructure and other devices. Nevertheless, making use of viable assumptions and basic organizational design principles, the team can describe the additional functions and operations that V2V and V2X communications will present. Figure 9 is a pictorial view of how V2V, V2I, and V2X communications differ.

Figure 9. Connected Vehicle Environments

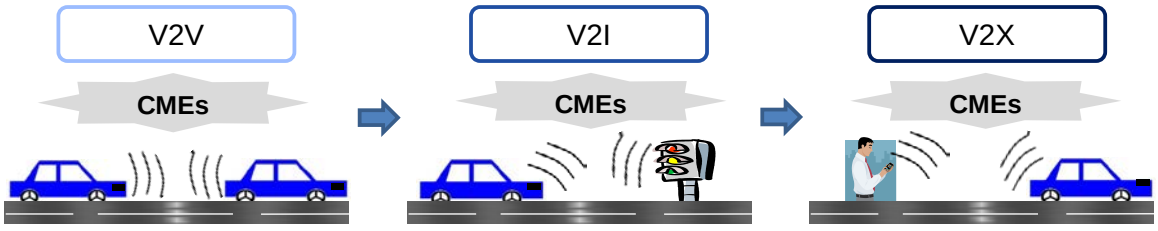

Vehicle-to-Vehicle (V2V)

V2V communications form the foundation of the need for CMEs. Facilitating safety messages sent between vehicles will be the initial function of the CMEs. Additional messages, in addition to safety messages, that may be exchanged between vehicles in the future will also need to rely on certificates for trusted communications between vehicles.

Vehicle-to-Infrastructure (V2I)

There are two roles that infrastructure or RSE will play in the connected vehicle system. One is as a conduit for communication between the CMEs and the OBE and the other is as an originator of messages directly to the OBE. The team discusses each in turn.

Infrastructure as a Communications Conduit Supporting Security

It is important to note that the security system will likely rely, in part, on communications via RSE, though the extent of that reliance has not yet been determined. Technical and policy decisions about whether vehicle communications go through RSE to CMEs have not yet been made. Some of the communications that might go through RSE are:

- Requests from OBE to CME for annual certificate batches and monthly decryption keys
- Sending of decryption keys from CME to OBE
- Sending of certificate batches from CME to OBE
- Distribution of the CRL

At the same time, transmission of some of these communications via RSE may pose technical difficulties due to the extended download time required for an annual batch of certificates (105,120 certificates per batch). The technical solutions to these issues are under development in a separate task dealing with the Communications Data Delivery System (CDDS) and are outside the scope of this report. Recommendations from that work will inform the decisions about how communications are delivered from CMEs to OBE and vice versa. As those decisions are made, the CMEs presented here and the model implemented will have to add the communications link from RA to RSE and then to OBE. The fundamental structures of the CMEs, in terms of configurations of functions, will not change based on those technical decisions, but the activities and ongoing management of the CMEs that produce and receive those messages through RSE will expand accordingly.

Infrastructure as Originator of Messages

There also are early stage discussions and development of applications for RSE to communicate with and send messages directly to the OBE. Messages such as Signal Phase and Timing (SPaT), which include safety, mobility, and environmental messages, are among those being developed. Some, if not all, V2I communications will require that the RSE be authenticated as a trusted member of the entire system. This authentication would likely occur in much the same way as an individual OBE is authenticated. Messages sent from an RSE to an OBE will need to contain certificates that serve as security credentials. At the same time, an RSE will also have to be able to read, process, and verify the validity of certificates contained in messages received from OBE.

Technical specifications of how this will happen have not yet been developed and are outside of the scope of this project. Nonetheless, the team recognizes the need for CMEs to include methods and policies for facilitating authenticated participation of RSE in the connected vehicle system.

Authentication Process of RSE

Since there is no sensitive PII collected during the RSE activation process, there should be no need for a separate activation system for these devices. The exchange of certificates from RSE to OBE would still take place within the same PKI environment, providing a trusted relationship for certificate exchanges. A proposed activation process for the RSE would likely include:

- The RSE sending a request, similar to the CSR of an OBE, to the pseudo RA to be activated.

- The RA would then send a signed certificate response allowing the RSE to operate as an intermediary between the OBE and itself.
- This certificate would grant the RSE functionality within the CME until the certificate expires.
- It is assumed that the certificate life span of the RSE could range from one year to the lifetime of the system, as it does not have to be the same life span as the certificates or the OBE. The main reason to have the expiration dates the same would be for ease of management.
- The certificate used to authenticate the RSE would be housed in a separate database so that one can track back if an RSE has malfunctioned.
- The life span of one certificate may be dependent on the updates scheduled for the CMEs, although a policy decision about the activation process of the RSE will need to be determined.

Vehicle-to-Other-Devices (V2X)

As the connected vehicle system develops and the number of applications making use of connected vehicle information increases, the security network will need to start incorporating communications between CMEs, vehicles, and other devices. This expansion implies that the CMEs will need to communicate with devices other than OBE in vehicles and RSE. Technical and policy decisions will be required to determine what kinds of messages and certificates cellular phones and other devices will receive and request. It is unclear at this stage whether the same certificate system and technical specifications will apply to other devices. It may be most effective, though not necessary from a policy or technical perspective, to apply the same standards and procedures and processes to communications with other devices. This approach would allow the CMEs built for V2V and V2I communications to expand their existing processes and standards to include new devices as applications for them become available.

Expanding the CMEs' activities and functions to include providing security for other devices raises a number of questions:

- Would new devices require different credentialing procedures (see discussion in Chapter 6 on Privacy)?
- If so, do new organizations need to be stood up, or are there ways of integrating the expansion of functions within existing CMEs?
- How do policy and technical specifications for the expanded communications impact the entire system and the needs of the various CMEs?

Analyses of these and related questions have yet to be completed and are outside the scope of this project. The CME models presented herein do provide a solid foundation for expansion and initial perspectives indicate that expansion of the operations and functions would be technically and procedurally feasible.

Authentication Process of Devices

A number of nomadic devices are envisioned for use in the connected vehicle system. It is likely that the process of authentication for each device would need to be unique. If the device is a cell phone, the authentication could be similar to the OBE authentication. Participation in the CMEs would only be

available for cellular accounts that are associated with a contract agreement between a consumer and a cellular company. The issue with change of ownership would exist but could be addressed with current cell phone industry policies. A CSR could be linked to a serial number on the phone that would be linked to the cellular account that belongs to the owner of the phone. A trusted party could activate a device within the system at the time of purchase, similar to the idea of activating the CSR in the V2V activation system. The difference would be that there is no need for separation of the activation and pseudo system since PII would already be protected by current privacy protection policies used by the cellular industry.

Another device to consider would be a GPS device. The proposed authentication process would be similar to that of the RSE in the V2I implementation. The device would not be linked to any PII and therefore would not require the separation of the activation and pseudo systems. The device could be authenticated and operated within the pseudo system, in a fashion similar to what was described previously for authentication of RSE.

By using either of these devices, short-lived certificates would still need to be utilized. Similar to V2V, there is the possibility of security and privacy attacks and vehicle tracking. The risks of vehicle tracking and software tampering are present in commercial consumer devices available on the GPS market today. Therefore, it is up to that industry to address these issues. Governing parties of the V2X environment may want to set standards regulating certificate life span to reduce the probability of tracking vehicles or individuals. To figure out the appropriate certificate life span for these devices, an analysis of the trade-offs involving transportation safety and functionality of the CME needs to be performed. Short-lived certificates guard against vehicle tracking, but require an exceptional amount of issuance and operation from the CMEs.

Because the use of nomadic devices would probably be considered voluntary, there may be less of a concern about the protection of PII. The protection of PII would be the responsibility of those parties (e.g., corporations such as cellular companies) that maintain contracts for nomadic devices with individual users. Table 12 below shows the implications to CMEs that need to be considered when expanding the system beyond V2V communications.

Table 12. Implications Beyond V2V

Type of Communications	Implications to SCMS
Vehicle -to-Infrastructure (V2I)	▸ Annual batch of certificates would take a long time to download onto vehicles if originated by RSE ▸ Expansion of activities and ongoing management of CMEs will change accordingly
Vehicle-to-Other Devices (V2X)	▸ May require different policy and technical controls compared to V2V and V2I ▸ May require additional organizations to be stood up in order to operate the expansion of CME functions if integration into existing systems is not a possibility ▸ May require certificate lifespans to increase from five minutes to a range (unspecified at this time)

All of the scenarios for the expansion and evolution of the connected vehicle system from V2V to V2I to V2X build on the basic V2V system and imply expansion of functionality and operations of CMEs, but not beyond their fundamental roles and responsibilities.

Chapter 8 Misbehavior

Standard PKI systems usually include checks on the system and its users in order to detect misbehavior. These security check processes are present across industries; from the payment card industry with its requirements of external system scans to healthcare providers, who undergo routine EHR compliance audits. Misbehavior can be categorized based on its origin: it either stems from technical (hardware or software) malfunction or from human malfeasance. As with other aspects of the CMEs for the connected vehicle system, the scope and scale are such that a customized and specified misbehavior function and process must be developed. In this chapter the team presents a synthesis of the literature and research on detection and management.

The critical issues that need to be addressed related to misbehavior detection and management – those that have a direct effect on the CMEs and their functions – are as follows:

- Technical process for misbehavior detection
- Differences in the processes for detecting technical malfunction versus human malfeasance
- Process by which bad messages are stopped to ensure they do not negatively affect the security and effectiveness of the system
- Rating of different kinds of misbehavior once they are detected – setting of different levels of consequences and reaction to different levels of threats or attacks
- Process by which CRL is distributed to the OBE
- Process by which a device or user can get back into the system once a CSR is revoked
- Process by which enforcement action, either internally or externally, is taken against misbehaving system participants and other bad actors

In the rest of this chapter, the team will catalogue the current understanding of how misbehavior detection and management may work in the CME system and also highlight the decisions (both technical and policy) that have yet to be made.

Misbehavior Detection and Management (MDM) Function

As described in the beginning of this report, the MDM function is a feature of the CMEs that is critical for maintaining the integrity of the system. Unlike more traditional PKI systems where misbehavior detection is fairly straightforward, the team has specified it separately from other functions within the PKI for the connected vehicle system. This separation highlights the importance of system security and user safety, and allows a focus to be placed on misbehavior management that may not be possible in other organizational design frameworks. Furthermore, as the processes and technical architecture by which misbehavior will be detected at local (in vehicle) and global (system-wide) levels have not been developed yet, it is not clear in which entity or other function MDM will reside.

Detecting Misbehaving Equipment

As noted previously, one kind of misbehavior can stem from malfunctioning equipment. Most of the research to date suggests that OBE will be able to detect misbehaving equipment messages, because the software will perform functionality checks and misbehavior detection processes at the local (vehicle) level.[32] This implies certain requirements for the OBE software, which have been noted previously in Chapter 3.

It is likely that OBE software would identify both malicious attacks and technical defects as misbehavior because it cannot determine the cause of the defect. Most technical defects would be detected by the on board diagnostics, and devices can be configured to stop sending messages when they detect that they are sending bad messages or having other technical issues. The remaining question is how human malfeasance which is not readily picked up by plausibility checks can be detected and managed.

Identifying Malfeasance

When bad messages are distributed because of human malfeasance, attacks, hacks, or other methods of adversely influencing the system, they may not be easily detected by the OBE. To address this issue, some researchers and stakeholders have suggested implementing a "global processing" function within the CMEs. As currently posited by the team developing system design,[33] global processing would require the OBE to collect a sample of random messages at regular intervals (the number of messages and frequency of intervals are to be determined by technical groups) and send these reports to the MDM. The MDM, having access to accurate information captured in real time, can compare information in these reports and gain a large-scale perspective of the connected vehicle system, and detect the origin of bad messages as statistical anomalies. The MDM would need to have the processing ability and software and technical sophistication to perform these kinds of checks. Even with the proposed software and hardware security measures in place, a recent estimate is that users would experience one false message every four days. Implications of this rate of bad message encounter have not been made clear by technical teams analyzing security impacts, nor have there been more details proposed about how this global processing function will find the statistical anomalies that indicate malfeasance.

Once malfeasance or other bad messages have been detected, the system must revoke or flag certificates from that device. There are divergent views on how this revocation or flagging might occur, summarized in the following:

1. The CMEs somehow (exact method TBD) "back-link" the individual certificate identifier to its batch in order to put the common linkage value on the CRL so that all certificates from that batch are ignored by receiving devices for the rest of the current month. The linkage value that is placed on the CRL revokes all certificates issued from that point forward. In addition, when the OBE requests a new decryption key for the subsequent month's batch of certificates, the RA does not send this key along.

[32] CAMP and Volpe, *Security Approach for V2V/V2I Communications Delivery System*.
[33] Ibid.

- There is no ability to remove certificates from a device, so the onus is on the receiving OBE to check messages against the CRL to determine which messages should be ignored or accepted.
- The CA sends the CRL to the RA for distribution and the RA distributes the CRL regularly (on a periodic basis, perhaps daily) to all OBE. This also implies that the RA maintains a database of CSRs that were distributed to OBE, so that it can realize when not to distribute to misbehaving OBE. The details surrounding CRL distribution are still being determined and may be affected by policy decisions, as well as technical design.

2. There are two CRLs. One CRL, called the CSR CRL, would be maintained by the RA alone and not distributed to OBE. This CRL is necessary for the RA to recognize who should not receive a decryption key. A second CRL would be distributed to OBE. This CRL would include the linkage values used to represent batches of five-minute certificates that have been flagged as originating from misbehaving devices. In this process, the CSR certificate of the misbehaving device is put on the RA CRL but not distributed to the devices. This is the mechanism by which the RA knows which requests for decryption keys it must deny.

3. There has also been discussion by some groups of an alternative process that does not involve the use of a CRL. This approach proposes that once a device is recognized as emitting bad messages, there would be a mechanism to remove the remaining certificates from that device. This would effectively remove any kind of bad message or CRL detection from the receiving devices, but it is unclear how this would be accomplished from a technical standpoint since the process by which this would happen is unknown. It is important to note that it may not be a realistic expectation that a misbehaving device would comply with self-removal of its certificates.

Consequences for Malfeasance

A key policy question that has not yet been addressed involves enforcement against misbehaving system participants or other bad actors. What should be the consequences for intentionally trying to influence or negatively affect the system? Will enforcement actions take place solely within the CMEs or will enforcement involve external legal action, either civil or criminal? Discussions later in this chapter review consequences in comparative industries such as user's being fined, jailed, and permanently revoked from the system. It should also be noted, as the comparative industry examples outline, well-defined consequences are necessary to deter malfeasance and malicious attacks from occurring.

As noted previously in Chapter 6, there are ways for the CMEs to collect minimal identifiable information about users and to implement adequate controls, both technical and procedural, to protect that PII appropriately. For example, access to user information could be limited to instances in which an administrative or judicial order requires disclosure of such information, in connection with an enforcement action. Or, the enforcement policies governing the system might limit access to user information to a specific function or role within the CME once clear evidence of malfeasance has been identified.

If no PII is collected anywhere in the system, there is no way to link an OBE device back to a vehicle or individual. For this reason, even when there is ample evidence of malicious, widespread damage

to the system or hacking caused by a specific device, there would be no way to identify or take action against the bad actor who caused the damage – in effect, the system would be unable to manage the malfeasance in an effective manner. In such cases, the only enforcement option available to those administering the CMEs would be to revoke the authentication of the device from which the malfeasant behavior emanated from the system.

The Certificate Revocation List (CRL)

The CRL is the master list that contains records of all unexpired misbehaving certificates. The distribution of the CRL allows devices to protect themselves from malicious messages. Once a device has sent a bad message, it is reported to the MDM function that reports the bad actor to the RA. The RA will get the information from the CA and linkage authorities in order to bind the certificate ID and its associated linkage value to a batch. The batch consists of the remaining certificates on the OBE that were previously authorized to be used for the month. This linkage value is placed on the CRL, which is distributed to all OBE daily. Placement on the CRL signals the OBE to reject messages from misbehaving devices. Additional technical assumptions about the CRL include:

- The linkage value from the LA allows for efficient revocation of all certificates in a batch.
- The current thinking is that the CRL will be distributed daily to OBE through RSE. Because the CRL could be large in size, the need for CRL distribution must be balanced against the capability of the communication system chosen. The specific technical details of the communication system are still being decided.
- There is also a discussion about possible updates to CRLs that might include only changes since the last distribution (i.e., new certificates that should be ignored or certificates that have been reinstated as trusted). The design and implications of this option are still being developed by technical teams.
- Each OBE holds a dynamic list of revoked certificates based on the most recent CRL downloaded.
- There will be at least two CRLs – one for the CSR and one for the pseudo system. Technical and policy specifications about the connection between pseudo system CRL and CSR CRL have yet to be determined.

Regaining Access to the System

An area of uncertainty is what process and policies will permit users to get back into the connected vehicle system once the issue that led to their placement on the CRL has been resolved. There have been a few options mentioned, none of which have been vetted for technical feasibility. The team mentions a few here with some discussion of implications to the CME system and the users.

- One proposal is that the user would need to replace the device. This puts a large burden on the user and is likely to result in reduction of participation due to costs, inconvenience, and potential shortages of OBE.
- Providing a new CSR certificate to reactivate the device is an option, and the process by which that can happen is to be determined through technical and policy decisions. If a new CSR certificate has to be provided, then a new batch of five-minute

certificates will also have to be downloaded. The idea is that whichever process is followed for the initial download of those certificates will be replicated.

- The CMEs will have to update all systems and CRLs to reflect the removal of a device from the CRL as well as the removal of CSR.

Suspension vs. Revocation

While the CRL will be a first line of defense in helping to protect devices against malicious messages, it may also be possible to identify the specific misbehaving device and halt its ability to send the malicious messages in the first place. Suspension and revocation are methods of capturing potentially disruptive messages and other cryptographic material to prevent system disruption while the CMEs evaluate the OBE in question. Suspension involves the temporary removal of material such as certificates from the communication system with the intent of restoring it when the issue of misbehavior or technical malfunction has been resolved. Revocation, on the other hand, implies a permanent removal of material from the system. The intent of revocation is to destroy any connection with a device that has been deemed unfit for the system.

Because suspension and revocation have different implications for the system, clear definitions of the circumstances requiring each action need to be developed. One example of an instance where suspension could occur would be if a hacker manipulated OBE software so that nefarious messages were sent to other vehicles. After detecting the misbehavior, the five-minute certificates from the OBE could be suspended indefinitely until the software issue was resolved. Although some five-minute certificates would expire during the remediation process, those certificates that were still valid could be used once suspension was lifted. The more serious act of revocation would be appropriate for situations where egregious misbehavior were to occur, such as Sybil attacks or vehicle theft. In these instances, the CSR associated with the OBE could be revoked so that the compromised device would not be able to re-enter the system. It is important to note that the levels of severity with regard to misbehavior (i.e., what constitutes egregious misbehavior that requires revocation) require policy decisions that are still under discussion by the USDOT. Misbehavior implications are listed in Table 13.

Table 13. Misbehavior Implications

Misbehavior Issue	Implications to CMEs
Technical malfunction and malfeasance within the system	▸ OBE software must have the capability to detect both malicious attacks and technical defects. ▸ OBE and MDM must collaborate to support global processing.
Consequences for Malfeasance	▸ Penalties for intentional attacks on the system or disruptions of communications must be defined.
Certificate Revocation List (CRL)	▸ If option with one CRL is pursued, the CA must be able to "back-link" from an individual certificate to its batch. ▸ If the option with two CRLs is pursued, the RA must be able to store a separate CRL that it uses to document misbehaving OBE. ▸ If the option with no CRL is chosen, the technical process for removal of certificates would need to be determined.
Regaining access to the system after placement on the CRL	▸ Replacement of the OBE is one method, but it has serious implications on user inconvenience. ▸ Providing a new CSR and batch of five minute certificates is another option, and the technical process would need to be further developed.
Suspension vs. Revocation	▸ The offenses that would require suspension of certificates must be differentiated from those that would require revocation. The mechanisms and rules for suspension still need to be developed.

Industry Approaches to Addressing Misbehavior

Misbehavior in one form or another is an issue across industries. The processes used by industries to address misbehavior are where differences can be observed. Generally, both laws (federal and state) and industry guidelines (e.g., best practices or standards developed by a trade association) are developed in order to deter potential malfeasant users from launching an attack on the system. These methods provide different approaches to how an organization or system can or should involve enforcement external to the system, either criminal or civil.

As noted previously in the PCI, merchants and service providers must agree to comply with the PCI DSS, a set of guidelines designed to ensure that systems are secure against attackers. If a merchant is found to be in violation of the PCI DSS, a merchant's compliance status can be revoked; the act of penalizing merchants or service providers is dictated by the voluntary agreement that the merchant has with the specific payment card brand(s) with which it is under contract (e.g., Visa, MasterCard). External law enforcement officials would not be involved unless a merchant's lack of compliance led to misbehavior in the form of criminal activity, such as identity theft or credit card fraud.

Within the healthcare industry, examples of misbehavior can be seen when nefarious internal actors tamper with patient data, or when external attackers steal sensitive patient PII. Users of EHR in this industry are subject to the HIPAA's legal requirements, which mandate strict standards for the handling of information collected about patients. To prevent unauthorized access to information or data breaches, EHR users must meet system specifications and undergo compliance audits from HHS. Preventative measures such as these are intended to prevent misbehavior from occurring in the first place. HIPPA violations can lead to external law enforcement actions and result in administrative or criminal sanctions (e.g., fines, license revocation, and imprisonment).

Additional Implications for CMEs

It is also important to note that other decisions related to misbehavior will have direct implications on the CMEs. The way in which the OBE hardware and software are configured, the policies and technical procedures that could be put in place to collect and protect PII, and the policies that are developed to deal with attacks on the system or users will impact how the CMEs function on a daily basis. Another item to consider would be the size and distribution frequency of the CRL, which is a core function of the CA (creating and updating the CRL) and the RA (distributing it to OBE), and the determination of the revocation process and how the CRL is created and distributed could affect how those entities perform their work.

While the security of the system and its communications, as well as the need for privacy protections, are at the foundation of all CME models and designs, the methods by which to ensure privacy and security are still under discussion. It is uncertain whether in fact the MDM will need to be a separate entity in order to maintain protections of privacy and communications security. Other options imply that there may be sufficient technical and procedural controls and methods available to protect data and security to an acceptable level. Chapter 3 of this report includes more information about technical and procedural controls, as well as examples of effective use of these types of control in other industries.

Chapter 9 Technical Specifications

In order to understand the hardware and software needs of the system, the team used current estimates for the functions involved in generating and distributing certificates. Most of the technical specifications included in this chapter come from another team's technical design and analysis (namely CAMP, noted throughout the chapter). Because the design of the certificates and the processes by which they will be generated and encrypted are all new and still under development, many of the numbers included here are estimates based on technical understanding to date, and may change in the future. All numbers and calculations are presented in the present day's numbers, per existing technology.

This chapter provides estimates of the types and numbers of hardware needed to produce certificates for full deployment of the connected vehicle system. Different levels of system needs can then be calculated based on various deployment and user penetration scenarios as they evolve.

Cryptographic Operations

The current technical design being used for this analysis, developed by CAMP, specifies the use of elliptic curve cryptography (ECC) for the encryption of certificates and keys in the connected vehicle system. The current specification being used is for 256-bit ECC keys. The required cryptographic horsepower of each function, as it pertains to ECC, is described in terms of point multiplication (PM). PM is an exercise used in the three main areas of operation within ECC which include key generation, signing, and verification of certificates. This calculation is used to figure out how many hardware security modules (HSMs) will be needed to produce the certificates as well as the data storage and sending and receiving needs of each function. HSMs are utilized to perform fast cryptographic transactions and provide private key protection. The cryptographic operations of each function within the CME require the following PMs per certificate:[34]

- 8 per CA
- 2 per RA
- 3 per LA1
- 3 per LA2

The total number of required PMs per certificate is 16. Research indicates the most secure hardware currently available for use in these cryptographic operations would be HSMs, which could be used in each function to perform the needed cryptographic operations. One annual batch of certificates would require 840,960 PMs per CA, 210,240 PMs per RA, and 315,360 PMs per LA. Information about the most efficient HSM on the market indicates that it has the maximum capability to execute 1,100 Elliptical Curve Digital Signature Algorithm (ECDSA) cryptographic operations per second with ECC

[34] Point multiplication estimates provided by escrypt, Inc.

256-bit keys.[35] However, given that systems generally are not able to operate at maximum/highest performance at all times, and to account for different needs of the systems, we have used half that amount of cryptographic operations per second (550) to estimate the total number of HSMs needed to produce annual batches of certificates for full deployment. Table 14 shows how many HSMs are needed in each function at full deployment to support the cryptographic operations, based on current design specification provided by the technical design team (CAMP and its partners). The table assumes an even distribution of production of certificates, which may not be the most appropriate way to implement the system, but provides an estimate and understanding of the extent of hardware and software needed.

Table 14. HSMs per Cryptographic Operation

Vehicles/OBEs	Scale	Certs/OBE/year	Total Certs/year
250,000,000	Full Deployment (100%)	105,120	26,280,000,000,000
CA Load/Sec		**HSM Max**	
Point Multiplication	8	Crypto/Second	550
Total Crypto Operations/Sec	6,666,667	Required	12,121
LA1 Load/Sec		**HSM Max**	
Point Multiplication	3	Crypto/Second	550
Total Crypto Operations/Sec	2,500,000	Required	4546
LA2 Load/Sec		**HSM Max**	
Point Multiplication	3	Crypto/Second	550
Total Crypto Operations/Sec	2,500,000	Required	4546
RA Load/Sec		**HSM Max**	
Point Multiplication	2	Crypto/Second	550
Total Crypto Operations/Sec	1,666,667	Required	3030
Total			**24,243**

[35] Based on the team's research, the SafeGuard CryptoServer Se-Series by Utimaco Safeware®, specifically the Safeguard Se50 PCIe, Se400 PCIe, or Se1000 PCIe products, present the fastest ECC processing times. All SafeGuard Products are registered trademarks of Utimaco Software AG.

Data Sizes of the CME Functions

Understanding certificate and key sizes is important in estimating the data storage needs for the system. The current assumption is that all keys associated with encryption in the following functions are compressed 96 bytes, and keys used for Hashing are estimated to be 256-bit Standard Hash Algorithm (SHA); 256-bit SHA is a function of SHA-2, a set of standard cryptographic hash functions, designed as a novel hash function computed with 32-bit words[36].

Certificate production begins as the OBE sends a certificate request to the RA, which will contain the asymmetric "caterpillar" key of the butterfly key expansion process. Escrypt, Inc., the developers of this production process, have indicated the size of this certificate to be approximately 189 bytes. The receipt of the caterpillar key triggers the request from the RA to the LAs for production of 105,120 linkage values by each LA. The linkage values are later used with the cocoon key in the butterfly expansion process.

Once the linkage values are delivered from the LAs to the RA, they are incorporated with the encryption public key. Prior to the use of linkage values, the RA generates the "cocoon" signing public key, as an expansion function from the caterpillar key. It also constructs a "cocoon" encryption public key to which it attaches the linkage values. The RA then sends these keys to the CA for final production of the implicit five-minute certificates. (See Chapter 3 for the detailed process description of the key expansion and creation functions).

The final step in the creation of the five-minute certificates is performed by the CA. The CA receives the cocoon signing public key and cocoon encryption public key from the RA. The CA uses these values to expand the cocoon signing key into a butterfly key. The CA also creates and signs each individual certificate with the CA private key and encrypts it with the encryption public key. The CA function uses eight PMs of ECC. Data load sizes as discussed here are included in Table 15 and subsequently used to estimate the hardware needs for certificate storage at each function.

Table 15. Certificate and Key Sizes

Operation Type	Data Size
Certificate request size	189 Bytes
Linkage value size and encryption size (256-bit ECC):	40 Bytes
Cocooned signing public key size:	96 Bytes
Cocooned encryption public key size:	96 Bytes
Butterfly signing public key size:	96 Bytes
Butterfly encryption public key size:	96 Bytes
Certificate size (Butterfly signing public key + Butterfly encryption public key + Linkage values (Linkage Authority 1 + Linkage Authority 2)):	272 bytes

[36] RITA ITS JPO, *Security Credential Management System Design*.

Industry standard understanding of 8 bits per byte

Non-Cryptographic Operations

Current estimates are that approximately 75 percent of the operations needed by the CME functions, as calculated by the needed PMs, can be performed by the HSMs, as they are primarily cryptographic operations. Each HSM is paired with a CPU for operations, so there are as many CPUs for cryptographic operations as there are HSMs needed. In addition, the other 25 percent of the operations (that are non-cryptographic operations) will need to be performed by standard CPU processors. Table 16 includes estimates of the need for *additional* CPU processors to account for the additional 25 percent of operations across the system and per function.

Table 16. Processor Needs for Non-cryptographic Operations

	HSMs needed (75%)	Additional CPUs needed (25%)
CA	12,121	4040
RA	4546	1515
LA1	3030	1010
LA2	3030	1010
Total	**22,727**	**7575**

There will also need to be servers for data storage at each function. Based on estimates of certficate and key sizes provided by CAMP, the team estimated multiple years worth of data storage to account for those hardware needs. As discusssed in Chapter 10, additional standard servers will be used for data storage throughout the SCMS.

Certificate Revocation

It is yet to be determined how certificate revocation or certificate suspension will address misbehavior. In the instance that a certificate is revoked, the CSR could become inactive for the remaining life span of the vehicle or OBE In the instance of certificate suspension, the CSR could be placed on a list similar to the CRL until certain performance criteria are met again, as determined by system policies. Policy and technical discussions regarding this issue need to be conducted to identify the best approach. Processing and storage, as well as any hardware and software needs for sending of CRLs and receiving misbehavior reports from OBE, will need to be added to the full complement of hardware and software discussed here once the technical design has been established.

Server Software Platforms

Another technical aspect that needs to be considered while standing up the CMEs is the software on which the encrypted operations function. Despite the fact that there are no available off the shelf software products that would be able to perform the needed functions, the team has analyzed potential options for software needs. These estimates include software development and customization, as well as data management.

At the time of this paper, very few Commercial Off the Shelf (COTS) server products exist that can support IEEE 1609.2 certificates. Security Innovation®[37] sells the Aerolink™[38] product, which is available as a Linux®[39] package or Windows®[40] library. Escrypt, Inc. has the CycurV2X®[41] product, which comes in a software form, but is primarily marketed for use in embedded systems such as vehicle OBEs. Additionally, there are open source libraries such as OpenSSL™[42] and Bouncy Castle[43], which can be modified to support IEEE 1609.2 certificates. Although a wide variety of data exists surrounding existing CA products for X.509 PKI certificates, most of the data is irrelevant since existing CA COTS products do not support IEEE 1609.2 certificates. Additionally, proxy products that exist to interface between end users and the CA for enrollment and certificate issuance are not designed to support the level of privacy and complexity associated with the current design.

The development of the SCMS components will entail a significant research and development (R&D) effort. Although a prototype system containing CA, LA1, LA2, and RA is currently being used by the model deployment team, development will still need to occur to transition the proof of concept system to an optimized system capable of handling annual certificate issuance of trillions of certificates per year at full deployment. Additional R&D will also need to occur to optimize the cryptographic processing either by leveraging HSMs or clustering CPU cores. The software products will need to be able to utilize a significant number of HSMs or CPU cores, which may entail additional R&D or integration efforts. The other challenge with the sheer volume of the system is managing the data across all of the distributed system components. Database management for this system will require planning and integration. The RA(s) will need to track certificate request data for 250 million vehicles (at full deployment) and ensure that the request data and status is maintained in a database accessible to all RAs.

Backward Compatibility of the System

Since total connected vehicle system deployment and implementation of the SCMS could potentially take 20+ years, it is imperative that the technology and business choices used for the initial roll out be able to evolve and adjust to future technologies at all levels (within the CMEs and for the

[37] Security Innovation® is a registered trademark of Security Innovation, Inc.
[38] Aerolink™ is a trademark of Security Innovation, Inc.
[39] Linux® is a registered trademark of Linus Torvalds.
[40] Windows® is a registered trademark of Microsoft Corporation in the United States and other countries.
[41] CycurV2X® is a registered trademark of escrypt, Inc.
[42] openSSL™ is a trademark of The Open SSL Project.
[43] Bouncy Castle free cryptography software is available from The Legion of the Bouncy Castle website.

communications network). Consideration will need to be given regarding how to adjust to new technologies in mobile data hardware, software and services, and providers to ensure capabilities don't become incapable of support at some point in the future.

Chapter 10 Costs

Cost Considerations

Since the focus of this effort is to assess possible organizational structures to support the use of PKI through CMEs, one must consider the functions of the organizations, hardware and software requirements, personnel needs, and associated facilities when estimating costs. To establish a reasonable estimate, current PKI systems were examined. To that end, several key points must be addressed to provide a high-level estimate of the costs associated with the connected vehicle system:

- Types of software and volume of licenses that will be required for system development and operation
- Types of hardware, and volume, to meet system requirements
- System roll out possibilities and personnel costs, in terms of skill set, level of effort, and salary, necessary to develop the system and maintain it into the future
- Facilities necessary to house hardware and personnel, including number of facilities, space requirements, and potential construction costs

This chapter includes estimated costs for many of these elements in order to provide ranges of total system costs at various levels of deployment, which are based on underlying assumptions about needed functions and operational models under evaluation.[44]

PKI Industry Findings

The PKI that will be implemented for the connected vehicle system is unique and includes elements that do not currently exist in other PKIs, implying that brand new organizations will have to be stood up. This will require heavy customization at an enterprise level. Software, hardware, licensing, and development are some of the primary costs that must be considered when standing up a PKI system. The system will require heavy startup costs, annual operation and maintenance, and the cost of auditing security and privacy procedures. VeriSign™, a leader in the PKI industry, estimates initial and ongoing costs will include[45]:

[44] This sample estimate is made available to the government for independent evaluation of the associated direct costs of implementing a PKI system of this scale. This is not intended to provide financial or investment advice, and should not be relied on as such. The information presented is only to highlight issues for your consideration. Strict assumptions are adhered to and some scenarios, where information is lacking, are hypothetical and for illustrative purposes only. Deployment/investment decisions should not be based upon this sample cost estimate alone. There are no representations or warranties of any kind, either express or implied.

[45] VeriSign, Inc. *Reducing Complexity and Total Cost of Ownership with VeriSign® Managed PKI.*

Capital Expenditures:

- Software acquisition and development
- Hardware and networking infrastructure
- Highly available validation infrastructure (CRL and Online Certificate Status Protocol [OCSP])
- Secure facilities
- Backup and disaster recovery
- Scalability to support user and application growth

Annual Operating Expenses:

- Annual operations and maintenance for systems and facilities
- Creation and auditing of policies and procedures
- Management of the certificate lifecycle
- End user support
- IT training
- Electric power

As part of the team's research efforts to identify the cost elements of modern PKI systems, we held discussions with several vendors within the PKI community. Though there is no precedent for the scope and scale of the SCMS under the connected vehicle system, the team can isolate the elements of existing PKI systems, identify their technical components, and define the personnel required to develop and maintain the network. At full market penetration, over 250 million vehicles would be equipped with OBE, utilizing over 26 trillion certificates per year. Research and industry contacts revealed that there are basic elements that drive the costs of systems with high-volume certificate issuance, presented in Table 17.

Table 17. Cost Drivers of Systems with High-Volume Certificate Issuance

Thematic Areas	Cost Drivers of Systems with High Volume Certificate Issuance
Certificate Volume	▸ The Department of Defense (DoD) maintains the largest PKI system in the U.S.: • 4.5 million active subscribers • 6 million certificates issued per year (1.5 million cards are issued to users per year with each card containing 4 certificates). This figure excludes certificates for devices such as SSL, domain controllers, and network devices, amounting to at least 50,000 other miscellaneous certificates (code signing, trusted agents, etc.). • 60 million certificates have been issued over the life of the PKI (~10 years active deployment). • Roughly 30 issuing CAs handle the load, all using HSMs for the encryption needs, excluding redundant equipment needs. • Lifetime of the CAs is 6 years (3 years in active issuance, 3 years issuing CRLs only)
Cost Tendencies	▸ Costs are most extreme during system development and implementation (including hardware, software, consulting, and system auditing). ▸ The volume and frequency of cryptographic operations necessary to keep pace with certificate issuance is the main cost driver relating to the amount of hardware, software, and horsepower necessary to meet the system requirements. ▸ Software platform costs are driven by the number of licenses necessary for development. ▸ Each user (OBE) will bear a software cost associated with each unique piece of equipment. ▸ Existing commercial models charge fees on either a per certificate or per user basis. ▸ PKI operates in three distinct and separate environments: development, operation and maintenance, and third party verification.

Thematic Areas	Cost Drivers of Systems with High Volume Certificate Issuance
Organizational Needs	▸ Redundant RAs and CAs must be established to accommodate high volume needs as well as system failures. ▸ It is possible, and common, for RAs and CAs to operate as entirely separate organizations. ▸ CA components can largely be virtual. ▸ Scale will be determined based on how many certificates can be produced per CA and the processing and distribution speed per RA. ▸ Facilities often associated with PKI systems include registration sites, data centers, and usage centers.
Staffing Characteristics	▸ Control policy requirements will drive staffing needs. ▸ Development staffing may potentially require 10-20 system architects to develop the platform and write code. ▸ CA staffing needs will be driven by the security demands of the system, the volume of certificates being produced, and the need to distribute the certificates. ▸ Regardless of certificate volume, PKIs require at least a two person access control for physical and logical access to the CA for administrative operations.
Impact of Technology	▸ ECC provides much faster key generation and signing operations than RSA (another type of algorithm for public key cryptography). (RSA stands for Ron Rivest, Adi Shamir, and Leonard Adleman, who first described it publicly.) ▸ Servers may process cryptographic operations faster than HSMs but there will be a tradeoff between speed and security. ▸ Software and hardware must be replaced/refreshed every few years. • The lifespan of a CA is inversely proportional to the amount of certificates being issued through it. This relates directly to the lifespan of hardware and software. Hardware that is exposed to heavy use will only last for a few years before it simply wears out. Software typically becomes obsolete or vulnerable after several years of use and must be replaced for these reasons.

Costing Assumptions

The following assumptions are built into the cost estimation:

- Estimates are provided for an initial period of six years.
- System benefits are not calculated.
- A discount rate of 7 percent is used to calculate Net Present Value (NPV), based on industry standards and government averages over the last 10 years.
- 105,120 five-minute certificates will be issued per year to a unit of OBE.
- IEEE 1609.2 certificates are assumed to be the certificate type, which determine estimates for the number of servers and processors required to support data loads and cryptographic processes.
- Cryptographic standards at 256-bit ECC are used to estimate the number of HSMs required. Maximum performance of this hardware is 1100 cryptographic operations per second[46]; however, in line with best practices, performance should not be estimated at the maximum potential. This estimate assumes performance to be 550 cryptographic operations per second.
- As a rule of thumb, software and hardware supporting cryptographic operations account for 75 percent of the system costs. The remaining 25 percent of system

[46] SafeGuard® CryptoServer Se-Series Benchmarks, http://hsm.utimaco.com/nc/en/products/se-series.

costs support other administrative functions and additional shuffling and bundling of certificates.

- Personnel costs are estimated using the average rate across a team of individuals supporting one particular sub-function, with 2088 hours in a year.
- Several functions are assumed to require around the clock staffing (e.g., CA, RA, LA), based on PKI industry practices. As such, staff at these facilities are assumed to work in eight-hour shifts, with three crews supporting a facility on a daily basis.
- Software estimates are provided on a per license basis for the software platform and database software. The platform will likely support multiple servers under one license but is assumed to be limited to a point. Database software is assumed to support one entire physical location per license.
- Hardware and software will be fully refreshed in the fourth year of use, resulting in cost surges that year. Costs in this refresh period are assumed to be 100 percent of the original investment for hardware and 50 percent of the original investment for software. Moore's Law is currently not being applied to this estimate.
- Total facilities costs are estimated using the average cost per square foot to furnish a facility to support PKI. Leasing costs account for the potential for a variety of geographically dispersed locations by using the average cost between leasing a General Services Administration (GSA) facility and commercial information technology office listings[47].
- Current estimates provide for construction of CA, RA, and LA facilities and leasing of space for other functions.
- The number of locations is highly uncertain; however, due to the system requirements for redundancy and continuity of operations, a minimum of two facilities should be accounted for each function that requires heavy data processing.
- Space requirements of 1.5 square feet per server are assumed to calculate space needs for data centers or facilities that house servers. This figure factors in the need for each facility to accommodate generators, extensive cooling systems, fire suppression systems, redundant communications, security, and administrative space, and is based on general industry information about average space per server across several large server farms and data centers.
- Developmental costs for data centers are assumed to be in the range of $600-$1500 per square foot, based on several construction projects undertaken in the past five years by Fortune 500 companies[48].
- The cost of the OBE is considered to be a sunk cost, necessary for system implementation. Therefore, it is excluded from this estimation because it has no bearing on the organizational evaluation. All other costs reviewed subsequently are affected by different organizational models.
- PKI typically includes three environments: development, operation and maintenance, and third party verification. Each environment requires independent equipment, software, and personnel. For this estimate, only the operation and maintenance environment is considered. Costs associated with the other environments may or may not match the operation and maintenance environment. At this point, the

[47] Commercial listings obtained via LoopNet®, a leading commercial listing service.
[48] Richard Miller, "Details of Google's The Dalles Site Now Public."

uncertainty surrounding the needs of the other environments is too great to include as part of this cost estimate.

Estimation Method

Due to the uncertainty inherent in the design and the groundbreaking, unique nature of the proposed system, as well as dependent factors currently being developed in other projects, namely the analysis of wireless communications systems to be used in supporting the connected vehicle system, the "analogy estimating" method was primarily used for this effort. To accomplish this, data were captured from existing PKI standards, software, hardware, and staffing methods (thus providing an analogous system and set of data from which to estimate). From this starting point, the assumed scope and scale of the proposed system were applied across the cost elements in the model to arrive at a range of costs that may be associated with this system.

Because the team is examining the possibility of several different organizational alignments, costs were first estimated for the primary functions, independent of organizational alignment – RA, CA, LA, activation, Root CA, and program management and oversight. Then the team accounted for potential efficiencies gained and cost savings associated with the unique organizational models. Under the models, the Misbehavior Division is always aligned with another entity. As such, the MDM would be able to leverage equipment, facilities, and resources of either the RA or the CA depending on the entity with which the MDM is aligned. Because the functionality and processes of the MDM function are largely undefined at this point, the team believes this is the most appropriate approach at this time, though we have not yet estimated costs for the MDM. As the processes and functions of the MDM are further specified by the technical teams designing the system, the team will adjust cost estimates accordingly. The major variable remaining for the MDM then becomes personnel costs. Calculations in the following sections should be considered rough and are based on technology available at the time of the estimate. The final subsection of this chapter includes detail about the anticipated cost differences between organizational models.

Resources and Staffing Considerations

Standing up the system will require a highly diverse staff with backgrounds in the areas of systems engineering, PKI, and IT consulting. The bulk of the personnel costs may be borne during the first years of the system, during design, development, and implementation. Due to the nature of a highly automated system, long term operation and maintenance costs can be relatively lean, from a personnel perspective. As research and vendor interviews have shown, monitoring and maintenance of equipment drives long term staffing needs. Existing PKI systems tend to operate with a staff of two technical PKI experts per CA, working in eight hour shifts, potentially 24 hours a day. Therefore, it may be feasible to assume a staff of at least six per CA, and potentially a total of ten for additional support for a system with higher transaction volume and complexity than traditional PKI. Long term costs can be kept low through system virtualization, in which CA, RA, and LA activities can be processed through a few highly dispersed server farms and virtual machines.

Staffing the program will also depend on the organizational model. Because one of the assumptions is the separation of the RA and CA, software, hardware, and infrastructure cannot be leveraged between those components; therefore, personnel cannot be leveraged across the two components. The inherent automation in the system, through the use of algorithms and technology, will likely

eliminate the need for a large staff of PKI experts. Rather, staffing through the operation and maintenance phase will be driven by the amount of hardware that needs to be monitored and the number of physical locations, which is uncertain at this point.

A help desk function will also need to be staffed. In this system, due to the relationship between the RA and the OBE, it may be feasible to position a help desk function within the RA. In existing systems supporting over 700,000 users, a daily staff of seven may be sufficient to handle call volume – two system engineers handling technical matters and five coordinators handling business matters with varying levels of effort due to PKI automated service support. For this estimate, the team assumed that a help desk component will accompany an RA at each physical location.

Ultimately, decisions about the number of personnel needed to monitor server activity should include consideration of the size and locations of data centers, as well as the data load being processed in a given location. The need to monitor equipment and the flow of information through the system is a major driver of the size of the workforce.

Cost of PKI for Certificate Management Entities

Costs for CMEs are estimated according to function (i.e., CA, RA, LA, activation, Root CA, and Program Management and Oversight), for a base period of six years, accounting for hardware and software refresh in year four. All data is appended to this report in the Excel workbook, "Cost Model for CMEs." Tables presented on the following pages were derived through the use of the model, which is designed to be open and is linked to technical specifications that drive large portions of the costs. When technical specifications are changed in the model, costs react accordingly across all certificate management functions.

To determine personnel costs, the team selected likely job functions that would be required to develop and support the system and captured information from three websites, www.indeed.com, www.glassdoor.com, and www.simplyhired.com. Based on the salary averages returned from these three sites for rates paid in the Washington, DC, area, an average salary was established per job title and per functional team and was then multiplied by the number of times the function would need to be repeated. Given the established average salaries, the team could then estimate either hourly rates or yearly commercial rates, factoring in a rate premium multiple of 2.2 times salary to yield a burdened salary expense. All personnel costs mentioned in this section were calculated in this fashion.

To account for facilities costs associated with the system, the team examined costs associated with the construction of large data centers and server farms erected by such companies as Google™[49] and Microsoft®[50]. Based on public information, the size of a data center depends on the amount of equipment it houses, computing requirements, data load balancing requirements, and power requirements, to name a few. For example, Google operates a server farm out of Oregon which includes several facilities of roughly 70,000 square feet each; housing around 45,000 servers each. This equates to roughly 1.5 square feet of space per server, including the necessary generators, cooling systems, wiring, and space for administrative functions. Developmental costs for facilities

[49] Google® is a registered trademark of Google, Inc.
[50] Microsoft® is a registered trademark of Microsoft Corporation.

such as these are estimated to range from $600-$1500 per square foot. On the other hand, it may be possible to leverage existing data centers, build-out the space to suit the needs of the system, and lease the facility for any length of time. Possible leasing rates were determined by researching GSA facilities in different regions of the country and commercial listings for IT office space through LoopNet®[51] (www.loopnet.com), a popular commercial real estate listing service. In addition, the following table provides a representation of PKI facility costs for one facility, regardless of size and exclusive of the technical requirements of the connected vehicle system.

Based on the assumptions and technical specifications, Table 18 reflects the total cost of the system at full deployment for 250 million vehicles. Costs are estimated separately per PKI function (i.e., CA, RA, LA, activation, root CA, and program management) over a six year period. Costs reflected in the table account for software, hardware, personnel, and facilities. The costs are shown in two different configurations, one that breaks them down by function and one that breaks them down by type of cost (i.e., software, hardware, facilities, etc.). Note that the Standard Deviation column in both tables provides some insight into the rates of uncertainty built into the cost estimates, as described above in this chapter.

Table 18. Total Estimated Cost of Certificate Management Entities (in Millions)

Functions	YR 1	Std Dev	YR 2	Std Dev	YR 3	Std Dev	YR 4	Std Dev	YR 5	Std Dev	YR 6	Std Dev	Total	Std Dev	NPV
CA Functions	$ 327.1	$ 58.4	$ 54.4	$ 8.1	$ 54.4	$ 8.1	$ 291.3	$ 47.9	$ 54.4	$ 8.1	$ 54.4	$ 8.1	**$ 836.2**	$ 138.6	**$ 695**
RA Functions	$ 119.4	$ 20.7	$ 29.6	$ 2.2	$ 29.6	$ 2.2	$ 92	$ 12.4	$ 29.6	$ 2.2	$ 29.6	$ 2.2	**$ 329.7**	$ 42.1	**$ 272.5**
LA Functions	$ 235.6	$ 41.1	$ 44.2	$ 6.2	$ 44.2	$ 6.2	$ 220.2	$ 35.9	$ 44.2	$ 6.2	$ 44.2	$ 6.2	**$ 632.6**	$ 101.9	**$ 523.9**
Activation	$ 1.2	$ 0.2	$ 0.8	$ 0.1	$ 0.8	$ 0.1	$ 0.8	$ 0.1	$ 0.8	$ 0.1	$ 0.8	$ 0.1	**$ 5.1**	$ 0.6	**$ 4.1**
Root CA	$ 1.7	$ 0.2	$ 1.4	$ 0.1	$ 1.4	$ 0.1	$ 1.4	$ 0.1	$ 1.4	$ 0.1	$ 1.4	$ 0.1	**$ 8.7**	$ 0.7	**$ 7**
Program Management and Oversight	$ 2.5	$ 0.2	$ 2.6	$ 0.1	$ 2.6	$ 0.1	$ 2.6	$ 0.1	$ 2.6	$ 0.1	$ 2.6	$ 0.1	**$ 15.5**	$ 0.7	**$ 12.3**
Total	**$ 687.5**	$ 120.9	**$ 133**	$ 16.8	**$ 133**	$ 16.8	**$ 608.3**	$ 96.4	**$ 133**	$ 16.8	**$ 133**	$ 16.8	**$ 1,827.8**	$ 284.6	**$ 1,514.8**

Cost Category	YR 1	Std Dev	YR 2	Std Dev	YR 3	Std Dev	YR 4	Std Dev	YR 5	Std Dev	YR 6	Std Dev	Total	Std Dev
Software	$ 1.7	$ 0.3	$ --	$ --	$ --	$ --	$ 0.5	$ 0.1	$ --	$ --	$ --	$ --	**$ 2.2**	$ 0.4
Software, O&M	$ 0.2	$ --	$ 0.3	$ 0.1	$ 0.3	$ 0.1	$ 0.3	$ 0.1	$ 0.3	$ 0.1	$ 0.3	$ 0.1	**$ 1.7**	$ 0.3
Hardware	$ 474.8	$ 79.5	$ --	$ --	$ --	$ --	$ 474.8	$ 79.5	$ --	$ --	$ --	$ --	**$ 949.7**	$ 159
Hardware, O&M	$ 66.5	$ 11.1	$ 95	$ 15.9	$ 95	$ 15.9	$ 95	$ 15.9	$ 95	$ 15.9	$ 95	$ 15.9	**$ 541.3**	$ 90.6
FTEs: Initial Cost	$ 0.2	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	**$ 0.2**	$ --
FTEs: Annual	$ 23.5	$ --	$ 33.6	$ --	$ 33.6	$ --	$ 33.6	$ --	$ 33.6	$ --	$ 33.6	$ --	**$ 191.3**	$ --
Facilities: Initial Costs	$ 117.7	$ 29.2	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	$ --	**$ 117.7**	$ 29.2
Facilities: Annual	$ 2.9	$ 0.6	$ 4.2	$ 0.9	$ 4.2	$ 0.9	$ 4.2	$ 0.9	$ 4.2	$ 0.9	$ 4.2	$ 0.9	**$ 23.7**	$ 4.9
Total	**$ 687.5**	$ 120.9	**$ 133**	$ 16.8	**$ 133**	$ 16.8	**$ 608.3**	$ 96.4	**$ 133**	$ 16.8	**$ 133**	$ 16.8	**$ 1,827.8**	$ 284.6

To account for a degree of uncertainty, Monte Carlo simulations were conducted for each independent cost element. To simulate costs, the Monte Carlo method requires the capture of inputs for each cost

[51] LoopNet® is a registered trademark of LoopNet, Inc.

element, including low, mean, mode, and high costs, in order to map costs against a unique probability distribution for each element. Following the calculation of a unique value for each element, based on a probability distribution combined with the use of a random number to account for uncertainty, the mean, standard deviation, minimum value, maximum value, 25th, 50th, and 75th percentiles were determined through the course of 500 simulations per cost element. Cost inputs were determined based on publicly available information via web searches and discussions with technical Subject Matter Experts (SMEs).

As with any build out of technology and facilities, a large percentage of the costs are incurred at the outset, represented under year one, and again when hardware and software near the end of their effective life and must be replaced, represented under year four. Because requirements and the organization must still be defined, development costs are associated with a diverse 10 person team working on the software development of the entire system at once. The 10 person team includes one principal systems architect, one senior information assurance analyst or engineer, two junior systems engineers, one senior developer, two junior developers, one senior IT consultant, and two junior IT consultants. Based on the salary averages for rates paid in the Washington, DC area, average salary for a development team was estimated to be approximately $102,000 per year. Initial development and customization of software may only take 160 hours because the platform would be based on existing PKI standards and practices. Given this hypothesis, billings for development may not exceed $200,000.

Costs associated with the CA are outlined in Table 19 following the description of CA costs. One of the first elements that must be considered is the server software platform. In this particular instance, the Red Hat®[52] Linux®[53] Server is chosen as the operating platform. While other software options may be more suited for the processing of IEEE 1609.2 certificates (such as Aerolink™ by Security Innovation), Red Hat Linux Server software was chosen for the purpose of the base estimate. Under this option, Oracle®[54] database software would need to be purchased as well. While each software platform can support up to 400 quad core servers, for ease of calculation and comparison against other software licenses, costs are estimated on a per license basis, rather than a per platform basis. Each license can support 13.33 servers; calculations are provided with these parameters.

Based on technical specifications of the processing power required to support the system at full deployment, 4,040 quad core servers are needed to support ancillary CA functions, including shuffling and bundling certificates, at a total cost of around $5.5-$10.1 million. The need to support high volumes of cryptographic operations translates to the procurement of 12,121 HSMs at an estimated cost between $124.5 million and $165.5 million. The HSMs would be accompanied by 12,121 quad core servers at a cost of $16.5-$30.1 million. The amount of data moving through the CA is 461 bytes per certificate, which will require 18,173 terabytes of storage to secure the data for a period of 1.5 years. Storage costs are relatively inexpensive when compared to other components, ranging from $1.8-$2 million. As a rule of thumb for IT implementations, annual Operations and Maintenance (O&M) costs for software can be estimated at 18 percent of the purchase price while annual O&M costs for hardware can be estimated at 20 percent of the purchase price.

[52] RED HAT® is a registered trademark of Red Hat, Inc.
[53] Linux® is a registered trademark of Linus Torvalds in the U.S. and other countries.
[54] Oracle® is a registered trademark of Oracle International Corporation.

The CA function could be supported by four person teams, working eight hour shifts, around the clock, yielding a need for eight staff members during the day with two working the night shift, for a total of 10 staff at each CA location. Based on research and discussions with PKI experts, a typical four person team may be comprised of one senior systems engineer, one associate systems analyst or engineer, one senior developer, and one junior developer. Based on the salary averages for rates paid in the Washington, DC, area, a CA team average salary was established at roughly $96,000 per year. Given this average salary, the team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $212,000 per year per staff member and $2.12 million per 10 person team at a given location.

To account for facilities costs associated with the CA, the team examined the possibility of constructing brand new data centers to suit the needs of the system. New construction was chosen instead of leasing because of the unique security requirements of the system, as well as the scale, and the need to protect user privacy. Based on the requirement for 1.5 square feet of space per server, mentioned above, and taking into account the servers and HSMs associated with the CA, over 49,000 square feet of space is required to support CA functions. It is estimated that three facilities would need to be constructed to satisfy the requirements for continuity of operations, redundancy, and data load distribution. If the equipment, as well as backup units, is distributed equally, one could assume roughly 16,400 square feet of space per facility at a cost of $12.9-$21.5 million each.

Table 19. CA Functional Costs

Cost Categories	Avg Cost Per Unit	Std Dev from Avg Cost	Number of Units	Total	Total Std Dev
Software					
Red Hat Enterprise Linux Server (Cost per license)	$ 533	$ 104	1,212	**$ 646.6K**	$ 126.5K
Database Software (Oracle Real Application Clusters)	$ 266	$ 51	3	**$ 797**	$ 153
Backup Software (for server software)			30% of total	**$ 194K**	$ 38K
Total				**$ 841.4K**	$ 164.6K
O&M (as a percentage of total cost)	18%			**$ 151.5K**	$ 29.6K
Hardware					
Servers with Quad Core Processors	$ 1,923	$ 562	4,040	**$ 7.8M**	$ 2.3M
Hardware Security Modules	$ 11,965	$ 1,693	12,121	**$ 145M**	$ 20.5M
Servers Required for Coupling with HSMs	$ 1,923	$ 562	12,121	**$ 23.3M**	$ 6.8M
Memory (64 GB of RAM in 1 card)	$ 1,894	$ 394	2,020	**$ 3.8M**	$ 795K
Storage (1 TB Drives)	$ 104	$ 6	18,173	**$ 1.9M**	$ 110K
Monitors	$ 457	$ 81	60	**$ 27K**	$ 5K
Keyboard, Mouse Combinations	$ 85	$ 20	30	**$ 2.5K**	$ 604
Personal Computers	$ 1,191	$ 229	30	**$ 36K**	$ 7K
Backup Hardware (servers, HSMs, memory, and storage)			30% of total	**$ 54.5M**	$ 9.2M
Total				**$ 236.4M**	$ 39.7M
O&M (as a percentage of total cost)	20%			**$ 47.3M**	$ 7.9M

Costs associated with the RA are outlined in Table 20, following the description of RA costs. As with the CA, one of the first elements that must be considered is the server software platform. Red Hat Linux Server software is also being estimated for the RA as the operating platform, and servers with quad core processors are used as the server option under the hardware category. While the RA handles data load of 690 bytes per certificate, fewer cryptographic operations are required to process this information, leading to a need for fewer HSMs and servers than the CA but a greater amount of data storage space. Based on technical specifications of the processing power required to support the system at full deployment, 1,010 quad core servers are needed to support ancillary RA functions, including shuffling and bundling certificates, at a total cost of around $1.3-$2.6 million. The volume of cryptographic operations in the RA translates to the procurement of 3,030 HSMs at an estimated cost between $31.2 million and $41.4 million. The HSMs would be accompanied by 3.030 quad core servers at a cost of $4.1-$7.5 million. Storage costs for housing the aforementioned larger data loads per certificate would require 27,200 terabytes of space to store the data for 1.5 years at a cost ranging from $2.6-$3 million. Information technology implementation rules of thumb also apply for the RA, yielding the same percentage of annual O&M costs for software and hardware as the CA, 18 percent and 20 percent, respectively.

The RA function could be supported by four person teams, working eight hour shifts, around the clock, yielding a need for eight staff members during the day with two working the night shift, for a total of 10 staff at each RA location. Based on research and discussions with PKI experts, a typical four person team may be comprised of one senior systems engineer, one associate systems analyst or engineer, one senior developer, and one junior developer. Based on the salary averages for rates paid in the Washington, DC, area, an RA team average salary was established at roughly $96,000 per year. Given this average salary, the team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $212,000 per year per staff member and $2.12 million per 10 person team at a given location.

In addition to the staff supporting operations and maintenance of the RA, the RA may be a logical location for a help desk function, given its frequent interaction with the OBE and user base. As such, the help desk function could be supported by a total of seven people during any given day at each RA location. Based on research and discussions with IT related help desks, a typical seven person team may require the skillset of a range of individuals, including systems engineers and general help desk coordinators handling business matters. Based on the salary averages for rates paid in the Washington, DC, area, a help desk team average salary was established at roughly $70,000 per year. Given this average salary, the team could then estimate hourly rates that would be charged to clients, including a commercial rate premium multiple of 2.2 times salary, yielding a total cost of around $154,000 per year per staff member and $1.08 million per 7 person team.

To account for facilities costs associated with the RA, the team examined the possibility of constructing brand new data centers to suit the needs of the system. New construction was chosen instead of leasing because of the unique security requirements of the system, as well as the scale, and the need to protect user privacy. Based on the requirement for 1.5 square feet of space per server, mentioned above, and taking into account the servers and HSMs associated with the RA, over 34,000 square feet of space is required to support RA functions. It is estimated that five facilities may need to be constructed to satisfy the requirements for continuity of operations, redundancy, and load distribution. If the equipment, as well as backup units, is distributed equally, one could assume about 6,900 square feet of space per facility at a cost of $5.4-$9 million each.

Table 20. RA Functional Costs

Software								
Red Hat Enterprise Linux Server (Cost per license)	$	533	$	104	303	**$**	**161.6K**	$ 31.6K
Database Software (Oracle Real Application Clusters)	$	266	$	51	5	**$**	**1.3K**	$ 255
Backup Software (for server software)					30% of total	**$**	**48.5K**	$ 9.5K
						$	**211.5K**	$ 41.4K
O&M (as a percentage of total cost)		18%				**$**	**38K**	$ 7.5K
Hardware								
Servers with Quad Core Processors	$	1,923	$	562	1,010	**$**	**1.9M**	$ 568K
Hardware Security Modules	$	11,965	$	1,693	3,030	**$**	**36.3M**	$ 5.1M
Servers Required for Coupling with HSMs	$	1,923	$	562	3,030	**$**	**5.8M**	$ 1.7M
Memory (64 GB of RAM in 1 card)	$	1,894	$	394	505	**$**	**956K**	$ 199K
Storage (1 TB Drives)	$	104	$	6	27,200	**$**	**2.8M**	$ 165K
Monitors	$	457	$	81	170	**$**	**78K**	$ 14K
Keyboard, Mouse Combinations	$	85	$	20	85	**$**	**7K**	$ 2K
Personal Computers	$	1,191	$	229	85	**$**	**101K**	$ 20K
Backup Hardware (servers, HSMs, memory, and storage)					30% of total	**$**	**14.3M**	$ 2.3M
						$	**62.3M**	$ 10.1M
O&M (as a percentage of total cost)		20%				**$**	**12.5M**	$ 2M

The concept of an LA is challenging to estimate through analogy because it is an evolving function that does not currently exist in traditional PKI systems. Because the primary function of a LA is to execute cryptographic operations, the bulk of the cost is associated with HSMs and their supporting servers, as outlined in Table 21, following the description of LA costs. The base case for LAs assumes that each LA, (LA1 and LA2) will be separate, both organizationally and physically. However, for the purposes of this estimate, costs associated with the LAs will be reported in combination, rather than breaking out identical costs for LA1 and LA2. Red Hat Linux Server software is also being estimated for the LAs as the operating platform, and servers with quad core processors are used as the server option under the hardware category. While the LA handles very small data loads of 40 bytes per certificate, a high volume of cryptographic operations are required to process this information, leading to a need for a lesser amount of data storage space but more than double the number of HSMs and servers than the RA. Based on technical specifications of the processing power required to support the system at full deployment, 3,030 quad core servers are needed to support ancillary LA functions, at a total cost of around $4.1-$7.5 million. The volume of cryptographic operations in the LAs translates to the procurement of 9,091 HSMs at an estimated cost between $93.4 million and $124.2 million. The HSMs would be accompanied by 9,091 quad core servers at a cost of $12.4-$22.6 million. Storage costs for housing the data loads per certificate would require 3,154 terabytes of space to store the data for 1.5 years at a cost ranging from $308-$336 thousand. Information technology implementation rules of thumb also apply for the LAs, yielding the same percentage of annual O&M costs for software and hardware as the CA and RA, 18 percent and 20 percent, respectively.

The LA function could be supported by two person teams, working eight hour shifts, around the clock, yielding a need for four staff members during the day with two working the night shift, for a total of six staff at each LA location. Since the LA function is unique to the connected vehicle system, it's challenging to anticipate the skill set required to support the function; however, for the purpose of arriving at an estimate, the team may assume that the two person team could be comprised of one associate systems analyst or engineer and one junior developer. Based on the salary averages for rates paid in the Washington, DC, area, an RA team average salary was established at roughly $96,000 per year. Given this average salary, the team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $212,000 per year per staff member and $1.27 million per 6 person team at a given location.

To account for facilities costs associated with the LA, the team examined the possibility of constructing brand new data centers to suit the needs of the system. New construction was chosen instead of leasing because of the unique security requirements of the system, as well as the scale, the need to protect user privacy, and the uniqueness of the LA function. Based on the requirement for 1.5 square feet of space per server, mentioned above, and taking into account the HSMs associated with the LAs, almost 27,000 square feet of space is required to support LA functions. It is estimated that as many as six facilities may need to be constructed to satisfy the requirements for continuity of operations, redundancy, and load distribution, as well as the physical separation of each LA. Construction costs may be $3.5-$5.8 million per facility.

Table 21. LA Functional Costs

Cost Categories	Avg Cost Per Unit	Std Dev from Avg Cost	Number of Units	Total	Total Std Dev
Software					
Red Hat Enterprise Linux Server (Cost per license)	$ 533	$ 104	909	**$ 485K**	$ 94.9K
Database Software (Oracle Real Application Clusters)	$ 266	$ 51	6	**$ 1.6K**	$ 306
Backup Software (for server software)			30% of total	**$ 145.5K**	$ 28.5K
Total				**$ 632K**	$ 123.7K
O&M (as a percentage of total cost)	18%			**$ 113.7K**	$ 22.3K
Hardware					
Servers with Quad Core Processors	$ 1,923	$ 562	3,030	**$ 5.8M**	$ 1.7M
Hardware Security Modules	$ 11,965	$ 1,693	9,091	**$ 108.8M**	$ 15.4M
Servers Required for Coupling with HSMs	$ 1,923	$ 562	9,091	**$ 17.5M**	$ 5.1M
Memory (64 GB of RAM in 1 card)	$ 1,894	$ 394	1,515	**$ 2.8M**	$ 596K
Storage (1 TB drives)	$ 104	$ 6	3,154	**$ 327K**	$ 19K
Monitors	$ 457	$ 81	72	**$ 33K**	$ 6K
Keyboard, Mouse Combinations	$ 85	$ 20	36	**$ 3K**	$ 724
Personal Computers	$ 1,191	$ 229	36	**$ 43K**	$ 8K
Backup Hardware (servers, HSMs, memory, and storage)			30% of total	**$ 40.6M**	$ 6.8M
Total				**$ 175.9M**	$ 29.6M
O&M (as a percentage of total cost)	20%			**$ 35.2M**	$ 5.9M

Costs associated with the activation function are primarily tied to personnel and facilities. Though the activation function would be required to utilize compatible software and hardware, minimal cryptographic requirements, as well as minimal data storage requirements, suggest that hardware costs would range from $33-$47 thousand, represented largely by two servers, one HSM, and the associated backup requirements.

The activation function could be supported by a two person team. It is not anticipated that this function would require 24 hour support. Based on research and discussions with PKI experts, a typical two person team may be comprised of one information assurance analyst or engineer (mid) and one junior developer. Based on the salary averages for rates paid in the Washington, DC, area, an activation team average salary was established at roughly $96,000 per year. Given this average salary, the team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $212,000 per year per staff member and $425,000 per two person team.

To account for facilities costs associated with activation, the team is relying on leasing space. New construction would not be necessary for a function that is somewhat standard for existing PKI systems. Because of the small amount of hardware required, as well as minimal staffing requirements, requirements for space would be minimal. If one, independent, facility of roughly 5,000 square feet was required to support this function, initial costs to build out a space could be $500-$760 thousand; annual costs could be $250-$400 thousand. These calculations rely on leasing costs using an average price per square foot per month of $1.38-$2.38.

Costs associated with the Root CA function are primarily tied to personnel and facilities as well. Though the Root CA function would be required to utilize compatible software and hardware, it is estimated that two servers, one HSM, and the associated backup requirements would support the function. Costs for hardware are estimated at $38-$55 thousand.

The Root CA function could be supported by a five person team. It is not anticipated that this function would require 24 hour support. Based on research and discussions with PKI experts, a typical five person team may be comprised of one senior systems engineer, two junior developers, and two junior IT consultants. Based on the salary averages for rates paid in the Washington, DC, area, a Root CA team average salary was established at roughly $87,000 per year. Given this average salary, team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $191,500 per year per staff member and $957,500 for the five person staff.

To account for facilities costs associated with the Root CA, the team is relying on leasing space. New construction would not be necessary for a function that is somewhat standard for existing PKI systems. Because of the small amount of hardware required, as well as minimal staffing requirements, requirements for space would be minimal. If one, independent, facility of roughly 10,000 square feet was required to support this function; initial costs to build out a space could be $500-$760 thousand annual costs could be $335-$540 thousand. These calculations rely on leasing costs using an average price per square foot per month of $1.38-$2.38.

Costs associated with the Program Management and Oversight function are primarily tied to personnel and facilities as well. Though the Program Management and Oversight function would be required to utilize compatible software and hardware, it is estimated that only two servers and zero

HSMs would be required to support the function, because the Program Management and Oversight function would not have a need to perform cryptographic operations. Hardware costs may be between $30-$45 thousand, spent largely on equipment for the support staff.

The Program Management and Oversight function could be supported by a 14 person team. It is not anticipated that this function would require 24 hour support. This management team may be comprised of one deputy director of operations, one operations support staff member, three administrative assistants, one general counsel, two associate staff attorneys, one director or administrator of strategy, one deputy director of strategy, one deputy director of policy, one chief financial officer, one deputy chief financial officer, and one financial analyst. Based on the salary averages for rates paid in the Washington, DC, area, a Program Management and Oversight team average salary was established at roughly $96,000 per year. Given this average salary, the team estimated a commercial rate premium multiple of 2.2 times salary, which would account for contractor or third party services being utilized, and yielding a total cost of around $153,700 per year per staff member and $2.15 million for the 14 person staff.

To account for facilities costs associated with Program Management and Oversight, the team is relying on leasing space. New construction would not be necessary for a function with standard space needs and technical requirements. Because of the small amount of hardware required, requirements for space would be minimal. If one, independent, facility of roughly 10,000 square feet was required to support this function, initial costs to build out a space could be $500-$760 thousand; annual costs could be $335-$540 thousand. These calculations rely on leasing costs using an average price per square foot per month of $1.38-$2.38.

The physical locations of the organizational components will be a critical cost driver. Savings could be realized through an effective location strategy that aims to leverage facilities, equipment, and personnel across multiple organizational components and PKI functions. Segmenting the system, while providing an effective means of increasing security, may be cost prohibitive when viewed through the lens of facilities costs. Evaluating vehicle registration data may be one manner of executing a location strategy. Even though the system is largely virtual, physical locations and hardware housed within could be distributed closely to major populations of registered vehicles. Table 22 illustrates how vehicle registrations[55] can be used to calculate the number of certificates required in a particular state, and ultimately the number of HSMs and associated elements required to process that certificate load. This data could theoretically be used to dispatch hardware and personnel to specific locations and determine the kinds of facilities necessary to support those locations.

[55] FHWA, Highway Statistics Series website: www.fhwa.dot.gov/policyinformation/statistics/2010/mv1.cfm.

Table 22. Impact of Vehicle Registration Volume on Location Strategy

State	Motor Vehicle Registrations	Total Certificates Required	HSMs Required for RA to Process Certificates
California	31,014,128	3,260,205,135,360	188
Florida	14,372,807	1,510,869,471,840	87
Montana	925,854	97,325,772,480	6
Ohio	9,800,933	1,030,274,076,960	59
Virginia	6,148,794	646,361,225,280	37

With proper attention to load balancing and bandwidth requirements, and a connected data delivery system that utilizes a truly national network, nothing would preclude RAs from assuming a centralized structure. In this situation, RAs would only utilize enough locations to support certificate loads while satisfying the need for redundancy and continuity of operations.

Realizing Efficiencies and Cost Savings

Estimating costs for the PKI system is a basic problem that is driven primarily by determining certificate issuance requirements and the kind of computing power and resources necessary to support those requirements. Certain software, hardware, facilities, and resource costs cannot be avoided, regardless of organizational model, but they can be affected by sharing some resources and combining certain functions. Table 23 illustrates the impact of various cost elements on the potential for savings by showing cost elements in year one as a percentage of the total. As shown previously, hardware is a high cost driver that is challenging to minimize due to the heavy computing requirements and need for separate systems for several functions.

Table 23. Impact of Cost Elements on Potential Savings

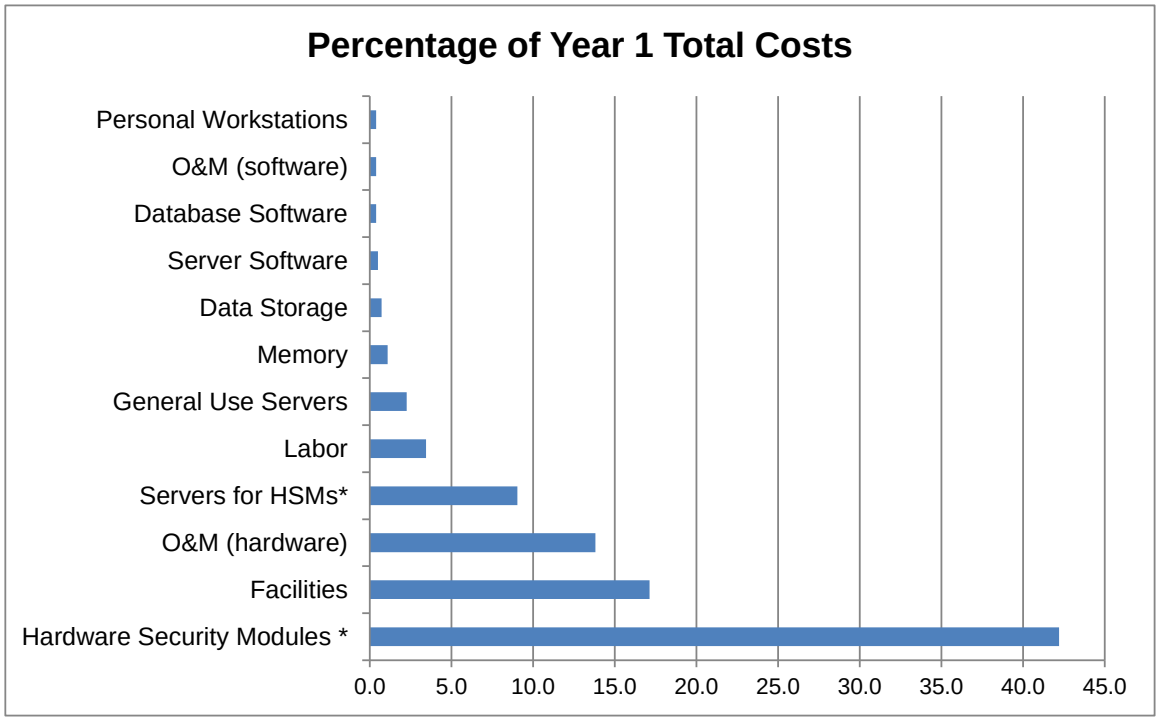

* Elements do not include backup or redundancy costs.

The cost of servers, HSMs, and memory is difficult to minimize under any organizational scenario. Regardless of organizational structure, processing and encryption requirements will remain the same. The CA currently bears the greatest cost burden, primarily due to the higher number of HSMs required to perform CA functions. However, the CA may be able to realize cost savings by operating virtual environments because it is not necessary to maintain proximity to the users and the OBE. Thus, the CA could save money on labor and facilities by managing only a few physical locations for the entire system. The RA, on the other hand, may have to operate more facilities than the CA because of its more demanding communications requirements with the OBE, though technical direction on how many RAs will be needed and what the communications system will be has not yet been finalized and is the topic of parallel projects. While the CA spends more money on HSMs, the RA will likely spend more money on facilities. Similarly, if the LAs remain physically separate, they will incur higher facilities costs than they might otherwise.

Any time functions are combined, select facilities, equipment, and resources may be shared. This could even lead to the cross training of personnel across functions, as long as they adhere to the necessary controls and policies. We have estimated projected cost savings for each model at multiple levels of deployment, based on anticipated organizational, resource, logistical, and operational efficiencies that can be gained from the combination of functions. Table 24 below deducts these cost savings over the base case of costs for all stand-alone functions.

Models C1, C2, and C3

In models C1 and C3, the RA and MDM are bundled together, leaving the CA and LAs to operate separate facilities. In C3, one of the LAs is bundled with RA and MDM. Though this might imply an additional cost savings, we anticipate it not realizing as much cost savings as C1 due to the separation of one of the LAs. At best, its savings would be equal to those of C1. In these models, the RA and MDM could share some equipment and facilities costs, and potentially employ personnel who are responsible for multiple roles. However, the CA and LAs would not realize any cost savings due to their need to operate in separate facilities with separate staff.

Model C3 bundles the RA and one LA and another LA and the MDM, with the CA being on its own. Estimated savings for this model are lower than they are for the other C models because of the need to stand up and operate, maintain, govern, legally administer and manage an entire entity for two automated functions (LA and MDM). The savings that come from combining automated functions with more staff-intensive functions (like RA and MDM), are not as great in C3.

Models D1 and D2

In organizational Models D1 and D2, the CA and MDM are combined while the RA and LAs remain separate in the former and are bundled in the latter (one LA with the RA, one LA on its own). As with the model described above, C2, that has an LA on its own, the cost saving are less for this iteration of Model D. For both Models D1 and D2, the CA and MDM would be able to leverage shared equipment and potentially personnel. In addition, the combination of the MDM with the CA is anticipated to realize higher operational and cost efficiencies because of the tight coordination that will need to be done between these two functions (depending on final MDM design and specs).

Table 24 below illustrates the cost savings, and makes clear that the differences between the C and D Models are minimal, in terms of estimated efficiencies or savings.

Model G

Organizational Model G combines the RA, MDM, and LAs under one organization, leading to the potential for these entities to share servers, facilities, and possibly staff, potentially saving $146.2 million over a six year period. As mentioned previously, under this model the CA may be able to operate in a largely virtual environment, covering the system from a few distributed locations with a lean staff. In this virtual environment, facilities and personnel costs of the CA could be reduced by at least one third.

Table 24. Total Costs for each Model over 6 Years

Models	Base	C1/D1	C2/D2	C3	G
Savings	None	5%	4%	3%	8%
Vehicle Penetration Level:					
25%	$ 514.7 M	$ 488.9 M	$ 494 M	$ 499.2 M	$ 473.5 M
50%	$ 934 M	$ 887.3 M	$ 896.7 M	$ 906 M	$ 859.3 M
75%	$ 1,355.9 M	$1,288.1 M	$1,301.7 M	$1,315.2 M	$1,247.4 M
100%	$ 1,827.8 M	$1,736.4 M	$1,754.7 M	$1,772.9 M	$1,681.6 M

Another way that savings can be realized is through the power of purchasing in bulk. The preceding estimates on the price of servers are given on a per unit basis without the application of a wholesale discount. It's not unreasonable to think that a discount of at least 25 percent on the cost of hardware could be achieved. Because the cost of hardware is a significant burden in this system, wholesale discounts could lead to savings of around $370 million over six years.

Estimating the total cost of ownership for CMEs as part of the connected vehicle system can be done by examining the organizational functions required to execute the program and the IT horsepower that is necessary to support those functions. With enough users in a system, costs tend toward low dollar amounts on a per user basis. In the base case, as shown in Table 25, with 250 million vehicles in the system at full deployment, annual cost associated with the functions (CA, RA, LAs, activation, root CA, and program management) per vehicle over a six year period ranges from $1.03-$1.41 taking into account one standard deviation from the mean ($1.22).

Table 25. Total Annual SCMS Cost per OBE

Cost Categories	Base Totals	Std Dev	C1/D1	C2/D2	C3	G
Total OBEs	250M					
Average Annual Cost	$ 304.6M	$ 47.4M	$ 289.4M	$ 292.5M	$ 295.5M	$ 280.3M
Total Annual Cost Per OBE	$ 1.22	$ 0.19	$ 1.16	$ 1.17	$ 1.18	$ 1.12

It is challenging to estimate other cost elements at this point, such as the number of facilities required to operate the system, the cost of unique facility construction, as well as when new facilities should be constructed versus leasing, the number of staff required to support the system, and total payroll. Altering these variables in any number of ways can easily shift costs above or below the thresholds displayed above, per OBE per year. The establishment of such a unique system will require the procurement of large amounts of equipment, large capital investments in facilities, regardless of whether a facility is leased or constructed, and the retention of a specialized, dedicated staff.

While the estimates provided in this report are preliminary and are based on current technical design specification, many of which have not been detailed, they are informative in thinking about decisions related to policy and stand up of SCMS organizations. The main cost drivers are hardware and software needs and are in large part determined by the numbers of certificates in the system. Additional analyses and scenarios that take into account alternative assumptions about technical and organizational/operational designs and decisions will be conducted moving forward and may yield different conclusions.

Industry Comparison

The connected vehicle system will ultimately reach approximately 250 million users in-vehicle, and likely more with the inclusion of nomadic mobile devices. To serve such a large group of users effectively and efficiently, an appropriate level of administrative functionality, network infrastructure, and customer service must be achieved. While no existing PKI system can compare with the size and amount of data transactions that will be supported through the connected vehicle system, the wireless telecommunications (telecom) industry may be an appropriate comparison for gaining a sense of total system costs, based only on scale, infrastructure needs, and user volume.

The team examined the U.S. wireless telecom industry in general and looked in more depth at three of the industry leaders: Verizon®[56] Wireless (Verizon), AT&T®[57] Inc. (AT&T), and Sprint Nextel (Sprint®[58]), to develop a summary of industry financials. The wireless telecom industry realized $198.7 billion in revenue for the year 2011 from 322.9 million subscribers[59]. This translates to an average of $615 in revenue per user.

Three companies currently account for just over 80 percent of the market share: Verizon, AT&T, and Sprint.

- Verizon represents roughly 38 percent of the market. For the five years ending with 2011, Verizon averaged $57.4 billion in revenue, with an average subscriber base of 88.8 million, yielding average annual revenue per subscriber of $649.69 over that period.[60]
    - Capital expenditures for the company have ranged from $7.2 billion to $9 billion per year over the last three years.[61]
- AT&T is a very large and diversified carrier, representing 30 percent of the market. For the five years ending with 2011, AT&T averaged $48.4 billion in revenue, with an average subscriber base of 86.2 million, yielding average annual revenue per subscriber of $561.71 over that period.[62]
    - Capital expenditures for the company have ranged from $7.9 billion to $9.6 billion per year over the last three years.[63]
- Sprint has the smallest market share of the three, with 14 percent. For the five years ending with 2011, Sprint averaged $27.9 billion in revenue, with an average subscriber base of 43.9 million, yielding average annual revenue per subscriber of $635.85 over that period.[64]
    - Capital expenditures for the company have ranged from $1.6 billion to $6.3 annually over the last five years.[65]

---

[56] Verizon® is a registered trademark of Verizon Trademark Services, LLC.
[57] AT&T® is a registered trademark of AT&T Intellectual Property, Inc.
[58] Sprint® is a registered trademark of Sprint Communications Company L.P. U.S. Telecom, Inc.
[59] Dale Schmidt, *IBISWorld Industry Report 51332: Wireless Telecommunications Carriers in the US.*
[60] Ibid.
[61] Verizon Wireless. 2011 10-K Report. Retrieved July 3, 2012 from the SEC online Edgar database.
[62] Dale Schmidt, IBISWorld Industry Report.
[63] AT&T Inc. 2011 Annual Report. Retrieved July 3, 2012 from the SEC online Edgar database.
[64] Dale Schmidt, IBISWorld Industry Report.
[65] Sprint Nextel. 2011 10-K Report. Retrieved July 3, 2012 from the SEC online Edgar database.

As detailed in previous sections in this chapter, current estimates of both total system capital costs and costs per user do not fall close to the ranges seen by the wireless industry. Although the systems, their intended usage, and their operations are not completely analogous, this comparison can yield some insights into the willingness of users to pay for valued services, and the extent of capital, operating, marketing, and other costs for a nationwide infrastructure and user services. Additional research into this, and other potentially comparative industries, to understand cost categories and expenses, needs to be conducted to aid in any commercial feasibility study for the connected vehicle system.

Scenario Modeling

As part of the ongoing evaluation of the SCMS, five scenarios will be subsequently modeled to assess the effects on costs of varying assumptions including: rollout schedules, system ownership, level of organizational centralization, system component performance, and facility location. Prior to developing these scenarios, the base model will be augmented with the following elements that can be simulated for all additional scenarios:

- Facility Acquisition and Development – the choice to lease or build will be an option for all functions.
- Impact of Technological Progress – refresh rates and historical trends for rates of improvement in technology (e.g., Moore's Law) will be applied.

Using the updated model developed for the base case, the following scenarios will be constructed, presented, and evaluated:

1) **Incremental Rollout and Adoption of the SCMS.** To examine the effects on costs of rolling out the system nationwide on an incremental basis, the following scenarios will be modeled:
    a. A realistic rollout schedule based on NHTSA parameters about new vehicle penetration/adoption rates;
    b. A more accelerated rollout schedule that also considers increased rates for after-market adoption.

2) **SCMS Ownership.** To examine the effects on cost of different system ownership arrangements, the following scenarios will be modeled:
    a. An entirely privately owned system;
    b. A hybrid system of private ownership and minority government ownership.

3) **SCMS Centralization.** To examine the effects of different organizational/institutional arrangements, the following scenarios will be modeled:
    a. A structure in which the organizations that run the certificate management functions are centralized (i.e., one "CA" organization operating out of two or three physical locations);
    b. A decentralized structure in which there would be several organizations that do the same functions within a CME organization (i.e., several signing "CA" organizations). Different levels of decentralization will be modeled to examine a distributed system across several regions and/or states.

4) **Quality of SCMS Components**. To examine the effects of using system components of varying levels of performance and quality, the following scenarios will be modeled:
    a. High-performing components capable of faster processing speeds and higher cryptographic operations per second;
    b. Mid-level-performing components capable of enough processing power and cryptographic operations to meet system requirements.

5) **SCMS Facility Location.** To examine the effects of a variety of geographic locations for facilities, the following considerations will be modeled:

    a. Energy needs, labor costs, network and bandwidth capacity, and facilities costs will be evaluated across a sampling of locations. High-density, urban areas, mid-market communities, and lesser populated areas will be considered. However, at a minimum, an area must have the power capacity and fiber optic bandwidth to merit evaluation.

Chapter 11 Implementation Planning

Implementation of the CMEs involves many interrelated pieces regardless of the final organizational model. The team analyzed the impact of critical implementation elements including integrating CME functions into existing systems (e.g., VIN or Title), multiple technical specifications such as processing times and capacity of the software, updates that will need to be made to the software and hardware over time, and creation of a help desk to assist with answering questions related to technical errors and issues. This chapter includes summaries of key topics that may be relevant to an implementation plan.

Implementation of the connected vehicle system is likely to take a phased approach. Since the scale of the system is so large, it would be impossible for all vehicles to be active in the system on day one. A thorough analysis of locations, infrastructure, timing, business requirements, technology and costs is needed to determine the appropriate roll out approach to ensure a successful implementation of the CMEs. The team has referenced throughout this paper where and how these issues affect implementation of the CMEs.

Measuring CME Performance

Following the establishment of strategic goals and objectives for CMEs, and during the early stages of implementation planning, a performance measurement approach can be defined to create incentives for staff, track organizational and system performance, and determine how to address identified performance gaps. It will be critical that such an approach focus not only on system performance but also on the abilities of CME staff to address deficiencies in the system and meet unit and individual goals and objectives. Relevant measures will provide management and leadership with a line of sight into support and functional operations. There are many means of approaching performance measurement, but all approaches tend to employ a general cascading top-down or bottom-up method to target how program management, critical system functions, and individual support staff should be assessed against the ability to meet goals. The following characteristics are common to organizations that practice effective performance measurement:

- Strong leadership that leads by example and inspires a performance focused culture
- Clear mission, vision, and goals that address the needs of the users
- Well-developed plans supported by budgets and resources
- Comprehensive balanced measures (and how to use them) to link strategic, organizational, and individual level plans and enable management of performance versus plan
- Well-defined organizational structure aligned to users and business requirements to promote accountability
- Efficient processes that focus on what the user values
- Appropriate technologies that support business needs

An overarching framework of performance management is needed to support a comprehensive measurement system that aligns with organizational and system goals. One such framework that

could be leveraged for the SCMS is the Goal-Question-Metric (GQM) model,[66] which was developed specifically for use in software development and systems design. This model, if used for CMEs, would begin at the conceptual level by defining goals for CMEs and the system. These goals should relate to the dimensions of people, technology, processes, and organizational structure. The model would then delve into the operational level by asking questions to characterize the way a dimension should be addressed in order to achieve the goal. The third level of focus for the model targets the quantitative metrics and measures that can be utilized to answer every question. These measures can be both objective and subjective. For example, one may objectively measure the staff hours spent on a certain task but subjectively measure the level of user satisfaction that is derived as a result of the hours spent on the task. As an example of how the GQM model can be employed to identify measures that support CME goals, Table 26 summarizes an initial draft of how the results of stakeholder interviews, a critical component in the development of performance measures, can be translated into measures that could be used to test the performance of the CMEs in the connected vehicle system, based on the GQM model. Target ranges for metrics should be considered; however, they are contingent upon an organization's tolerance for risk and therefore have not been included in this discussion.

Table 26. Performance Measurement Framework

Goal	Question	Metric/Measure
Convey general statements of success as a desired condition or outcome	**Address issues that , when answered, are a means of determining if goals and objectives have been attained or progress is being made**	**Provide information necessary to answer the questions**
Vehicles and linked infrastructure must establish mutual trust to exchange information	▸ How is trust determined? ▸ How is trust lost or gained? ▸ What makes the CME reliable? ▸ What is the acceptable error rate for misbehavior?	▸ Data accuracy ▸ Security audit fail rate ▸ Efficient signing and encrypting of certificates ▸ Disaster recovery testing and response ▸ Quality control of PII ▸ CRL suspensions, revocations, and reinstatements ▸ Amount of time before a misbehaving device receives a CRL
Vehicles and linked infrastructure should exchange meaningful data to facilitate safety, traffic, and environmental messages	▸ Is data useful to the user? ▸ What are the consequences of a lack of data?	▸ Percentage of accidents avoided by vehicles equipped with on board units ▸ Percentage of data used to relieve traffic congestion ▸ Fast and efficient information distribution ▸ User satisfaction
System operation, maintenance, and updates should not have a discernible impact on system performance	▸ How long do standard updates take? ▸ How often is the system tested, evaluated, and updated?	▸ Average server and HSM downtime and repair time ▸ Percentage of staff time devoted to O&M ▸ OBE and CSR CRL content ▸ Frequency of technology upgrades
The CMEs should be cost effective without sacrificing security	▸ What are the explicit and implicit costs of each CME function? ▸ How do costs relate to security?	▸ ROI (Return on Investment) ▸ Percentage of expenses levied for security

Based on best practices and industry comparisons, we have compiled a list of initial metric examples and ideas that may provide direction and understanding of the kinds of metrics that should be

[66] Victor Basili, et al., *The Goal Question Metric Approach.*

considered within the overarching framework. Full performance measurement systems, setup, measurement tools, and tracking systems involve extensive effort, and should be a part of the overall implementation and ongoing oversight and management of all CMEs. The metrics included in Table 27 are a foundation and starting point for consideration and to inform policy and technical decision making. The following metrics could be used to track errors through regular reports.

Table 27: Select Potential Metrics for System Design and Development

Potential Metrics for System Design and Development		
Metric	**Description**	**Unit of Measurement**
Project effort	Total project team time spent on project related activities during the life cycle of the project.	Hours
Project effort: Direct delivered team hours	Total time spent directly contributing to defining or creating outputs (e.g., software code, HSM configuration, OBE user manuals) that are delivered as part of the system.	Hours
Project effort: Direct non-delivered team hours	Total time resulting in production of outputs (e.g., requirements definition, change request log) that are not delivered as part of the final system.	Hours
Project effort: Support hours	Total time expended on work that does not directly define or create products but assists those who do.	Hours
Productivity	The ratio of system size over project effort where size is based on function points, use cases, and units of hardware configured.	System size unit per hour
Project duration	A measure of the length of the project.	Work days
Schedule adherence	A measure of how much the original duration estimate differs from the actual duration.	Work days
Requirements completion ratio	The extent to which planned functional requirements were satisfied in the final system implementation.	Number of requirements completed out of total requirements
Problem resolution rate	Amount of time required to resolve issues/errors once discovered.	Hours
Post release defect density	Defects discovered (in OBE, HSMs, etc.) at various time intervals after deployment.	Number of unique defects

The performance of the system can also be evaluated using PKI metrics.[67] Examples of some of these metrics are included in Table 28.

[67] Randall Gumke, *Navy/Marine Corps Intranet Information Assurance Operational Services Performance Measures*.

Table 28: Potential PKI Metrics

Certificate revocation	The time it takes to update and/or distribute the CRL.	Time (seconds, minutes, hours)
Registration time	The time it takes to activate the OBE.	Time (seconds, minutes, hours)
Interoperability	The system will have to perform across different localities, using a variety of technologies and networks by leveraging different service agreements. Many metrics could be established to measure the effectiveness of system interoperability.	Time for coordination between functions; accuracy of data exchange between functions
HSM and server throughput	A measure of the amount of data processed through hardware units (data load balancers should automatically ensure even distribution of data).	Bytes per time unit
HSM and server utilization rate	A measure of the capacity being utilized by the system (e.g., the number of linkage values produced and stored).	Ratio of actual use to maximum capacity
Staff efficiency	A measure of alignment of resources against system demands. System targets should be set at the minimal number of personnel to meet the demands of the system and handle typical outages.	Response time (minutes, hours, days)
Help desk efficacy	A measure of the effectiveness of the help desk.	Number of calls taken; response time to calls; issue resolution rate

For the most effective performance measurement plan, the CMEs will need to consider broader security and safety goals such as the their ability to support meaningful data exchange as well as more specific operational goals such as identifying and monitoring functional activities.

As a supplement to the GQM model previously mentioned, an initial estimated timeline for the priority of measure implementation with respect to the roll out phases in the CME process is essential for overall strategic guidance. Table 29 indicates the preliminary considerations of measurement priorities versus the vehicle penetration rate in the CME. Low or medium priority is relative to the costs and organizational needs to setting up comprehensive performance measurement systems. Rather, as the connected vehicle system and the CMEs are implemented and expanded to meet increasing customer demands, more robust and detailed measures will need to be developed.

Table 29. Measure Implementation Priority

Metric/Measure Category	Low Vehicle Penetration in CME (10%)	Medium Vehicle Penetration in CME (50%)	High Vehicle Penetration in CME (100%)
Vehicle and Infrastructure Trust	High Priority	High Priority	High Priority
Meaningful Data	Low Priority	Medium Priority	High Priority
System Updates	Low Priority	Medium Priority	High Priority
Cost	Low Priority	Medium Priority	High Priority

Each of the CME organizational functions (i.e., CA, RA, LA, and MDM) should also have a subset of goals and measures that can roll up to the larger goals of the CME. These metrics and goals will need to relate back to dimensions of people, technology, processes, and organizational structure on an ongoing basis. To measure performance from another angle, system audits and vulnerability assessments should be considered as potential methods of discovering underlying causes of performance issues. As discussed in Chapter 5, vulnerability assessments are also useful in identifying potential weaknesses that could cause security breaches. Ensuring that CME performance goals align with security processes can help to support the effectiveness of the system as a whole.

Disaster Recovery Plan

For large scale mission critical systems such as the CMEs, a disaster recovery plan is necessary to define the actions to be taken in the event of a crisis. A sound disaster recovery plan enables an organization to respond to emergency situations rapidly by establishing the priority of response activities and providing guidance to complete those activities. The goal is to restore mission critical systems to normal operating levels as soon as possible. At a high level, planning for disaster recovery involves preventative measures, detective measures, and corrective measures:

- **Preventative measures** are put in place to avoid disasters and are important for the ongoing operation of the system. Examples of preventative measures include backing up data and critical files on a regular basis and designing infrastructure, processors, and other parts of the system so that they can effectively perform their duties without being overwhelmed.
- **Detective measures** can be employed to discover and identify potential issues that may have circumvented preventative measures or that were overlooked by personnel. A common detective measure is running various tests in parallel with the system during a startup boot or a necessary reboot process to identify any performance problems in the system.
- **Corrective measures** are aimed at restoring the system after an emergency event occurs. The specific corrective measures for an organization must be customized

> based on the disaster itself and on the parts of the system that are considered to be mission critical, critical, essential, and non-critical.[68]

A disaster recovery plan for the CMEs needs to include specifications for backup systems for all system components, alternative power sources, and specialized personnel for operations of disaster recovery implementation, among other preparatory measures. NIST recommends that a cost-benefit analysis be conducted during the planning process to identify a contingency strategy that is most appropriate for the organization.[69] While the team's current cost estimates and models include additional costs for backup systems, there has not been a specific disaster recovery plan suggested by the technical team. With more information and development of such specifications, the cost model may need to be adjusted.

One additional topic to consider when discussing disaster recovery is dissimilar redundancy. It is common practice in some industries that one area serve as a backup for the other. For example, if the locations are split into regions and the northeast region fails, the southeast region would serve as the backup system so that the northeast region continues to operate as normal until the problem can be corrected. Having these plans in place allows for critical systems to continue operating as usual which means no disruption to users of the system.

Help Desk

Requirements for a help desk include interoperable CME hardware and software and access to the telecommunications and wireless networks. Variable costs for the CMEs can have a significant effect on total cost due to its number of users, employee capacity, and phone and internet use. Certain elements will need to be calculated to understand expected telecommunication costs and efficiencies when the program is implemented: industry help desk call cost averages, economies of scale achieved from user base, call duration, call types, and the number of employees available to support calls.

[68] Maartens, *Plan Against Disaster.*
[69] Swanson, et al., *Contingency Planning Guide for Federal Information Systems.*

Chapter 12 Conclusion

Due to the unprecedented nature of the connected vehicle system, there are several elements that are still under development by technical teams. Analysis of possible organizational models for the SCMS revealed several outstanding technical and policy decisions that need to be made prior to implementation. Table 30 includes a list of these issues and how they impact the CMEs.

Table 30. CME Issues and Implications

Topics	CME Issues	CME Implications
Misbehavior	▸ Penalties/level of law enforcement ▸ How do you identify malfeasance (global processing) ▸ What behavior requires suspension vs. revocation	▸ Oversight structure ▸ Pseudo system process flow ▸ PII and security policies
Credentialing	▸ No PII ▸ PII during activation (new or existing system) ▸ PII connected to certificates	▸ Existence of separate CA_{ACT} system ▸ Ability to trace back to users in cases of misbehavior
LAs	▸ Two LAs in one entity or two entities ▸ Computational requirements of LA ▸ Linkage value production process	▸ Level of vehicle trip trackability ▸ Organizational structure and costs ▸ Certificate distribution
Certificate Life Span	▸ Life span of certificates (five-minutes is currently being used)	▸ Longer life span reduces processing needs but may increase trip trackability
CSR	▸ ~2-5 year CSR ▸ ~20 year CSR	▸ Frequency of certificate renewals, size of CRL, infrastructure transition requirements
Back-Up Certificates	▸ If used, how many per year per vehicle ▸ If used, how long will they last	▸ Additional burden on CME functions for creation, distribution, monitoring of certificates
Certificate Policy	▸ What the policy will say regarding the rules for obtaining certificates, the technical requirements for generation and protection of private keys and certificates, and the requirements for audit records and periodic compliance audits	▸ PKI structure compliance for operations

Topics	CME Issues	CME Implications
Size of CRL (2 CRLs)	▸ CSR CRL – depends on life span of CSR ▸ OBE CRL – depends on how often it's distributed	▸ RA and OBE processing and storage needs
End of Life	▸ How it is defined ▸ How to ensure CSRs are taken off CRLs and RA is informed	▸ Sets guidelines for transfer of ownership ▸ Returned devices can allow for reuse ▸ Promotes better misbehavior detection
Frequency of Certificate Download	▸ Once a year or multiple times a year	▸ Affects resources required to manage, track, authenticate, sign, and revoke certificates
Number of RAs	▸ Computational requirements ▸ Distribution of RAs by physical or network geography	▸ Certificate distribution ▸ Organizational structure and costs
Number of CAs	▸ Processing needs ▸ Data loads ▸ Structure of Root CA ▸ Structure of assigning CA	▸ Certificate distribution ▸ Organizational structure and costs
Phases of Roll Out	▸ Timing for policies ▸ Number of OBE ▸ Number of RSE	▸ Stakeholder buy-in and communications ▸ Training and staffing plans ▸ Performance and outcome measurement

Development of additional technical design specifications and analyses to inform policy decisions in all these areas are evolving rapidly. This report uses current assumptions about technical designs to analyze various options for operational and organizational models. Using current industry examples when appropriate provides additional insight into the identification of possible impacts of various organizational choices. Initial cost estimates also provide some grounding in the needs of the system, though they are at a very early stage and will change as technical designs evolve and become more detailed. Ensuring that the SCMS and the individual CMEs are implemented in ways that maintain the balance between security and privacy throughout the system is the goal of all conceptualizations of future organizations.

APPENDIX A Certificate Management and On Board Equipment Life Cycle

Describing the entire life cycle of both the OBE and its connection to CME illustrates that human interaction and prompts are present at each stage. Leading up to the December webinar, most of the focus was placed on the Activation and Use Phases. During this next phase, the team has outlined in further detail the activities that occur at all six stages of the life cycle. Figure 10 provides high level phases and functions of the life cycle.

Figure 10. Certificate Management Life Cycle

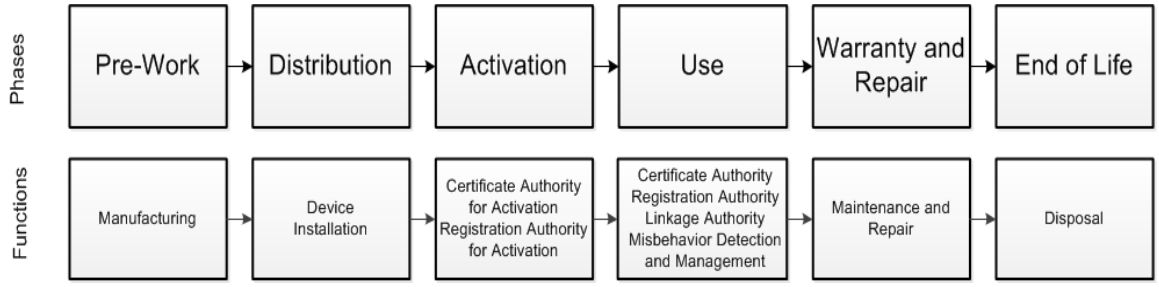

Pre-Work

The Pre-Work Phase involves the initial activities in the development and production of OBE hardware and software. Developing the software programming and hardware design will allow for all new vehicles to be manufactured with the OBE built in. Until that time, retro-fitted devices will begin to be manufactured and will be available for distribution and installation within the next two years by authorized dealers and installers. Once these devices have been installed, plans for any software updates will be implemented and rolled out over time.

Distribution

During the Distribution Phase, users of the system will gain access to their devices, either by purchasing new vehicles or an after-market device. As stated previously, coordination of the installation of the device will depend on when or how the device is purchased. For all new vehicles purchased from the dealer lot, either the dealer or the OEM will be responsible for activating the device in the vehicle, depending on how the device is credentialed. If a user purchases an after-market device from a retailer like Best Buy for example, the user would then bring the device to the

authorized dealer to have it installed and activated. Policies will need to be determined in order to encourage maximum user participation in the system.

Activation

The Activation Phase includes the CA for Activation (CA_{ACT}) and RA for Activation (RA_{ACT}) functions, which are related to the request and assigning of CSR certificates that activate the OBE. The CA_{ACT} and RA_{ACT} functions will ensure that the OBE is activated and receives its CSR so it can begin sending and receiving certificates. The CSR certificate will be the only place in the system where PII will be collected and stored. There is no need for RA_{ACT} and CA_{ACT} to be separated as in the pseudo system because the CSR is different from the 105,120 certificates issued by the CA and RA. Unlike the yearly batch of certificates, the CSR is not intended to provide an additional layer of security and privacy, and it is not used in location-based communications. The CA_{ACT} and RA_{ACT} will be in an entirely separate entity than all other CME functions that are part of the Use Phase (the pseudo system). Figure 6 previously in Chapter 3 includes a process representation of the activities involved in the Activation Phase.

Use

The Use Phase is the phase where the pseudo system performs all of the ongoing activities of the functions (i.e., CA, RA, LA, MDM) on a daily basis. At this stage, the OBE will request an annual batch of certificates from the RA. The RA will send a request to LA1 and LA2 for linkage values. The linkage values provided by the LAs will be sent to the RA. The RA receives the linkage values, shuffles the certificates, and sends a request to the CA that will assign a single linkage value to each certificate. Once the RA receives the certificates from the CA, it will box them into 12 one-month batches to distribute these encrypted certificates to each OBE. These functions happen once a year for certificate requests and distribution and once a month for decryption key requests and distribution.

Current thinking about MDM is that an OBE that receives a “bad” message will then report it to the CMEs (exact location to be determined). The report will include both the message and the linkage value. The MDM will investigate whether the misbehavior is due to system or technical malfunction or human malfeasance (see Figure 11). It will then send the report to the RA who will interact with the CA to determine, through a mathematical process still under development, which linkage value matches to the batch and will place that value on the CRL. The frequency of the CRL distribution is a decision that has yet to be made. Once the RA receives the CRL, it can begin to alert each user’s system of messages from other OBE that should be ignored. In addition, once a certificate and its accompanying OBE are placed on the CRL, it will be unable to acquire a decryption key to unlock the new batch of certificates for the subsequent month. Without the new certificates, the OBE will be unable to function within the system. As discussed, the process a user must follow to regain access to the system after being placed on the CRL either through suspension or revocation is a policy decision that still has to be made. Figure 5 in Chapter 3 includes a detailed process representation of the four major functions and activities that occur during the Use Phase.

Figure 11. Certificate Revocation

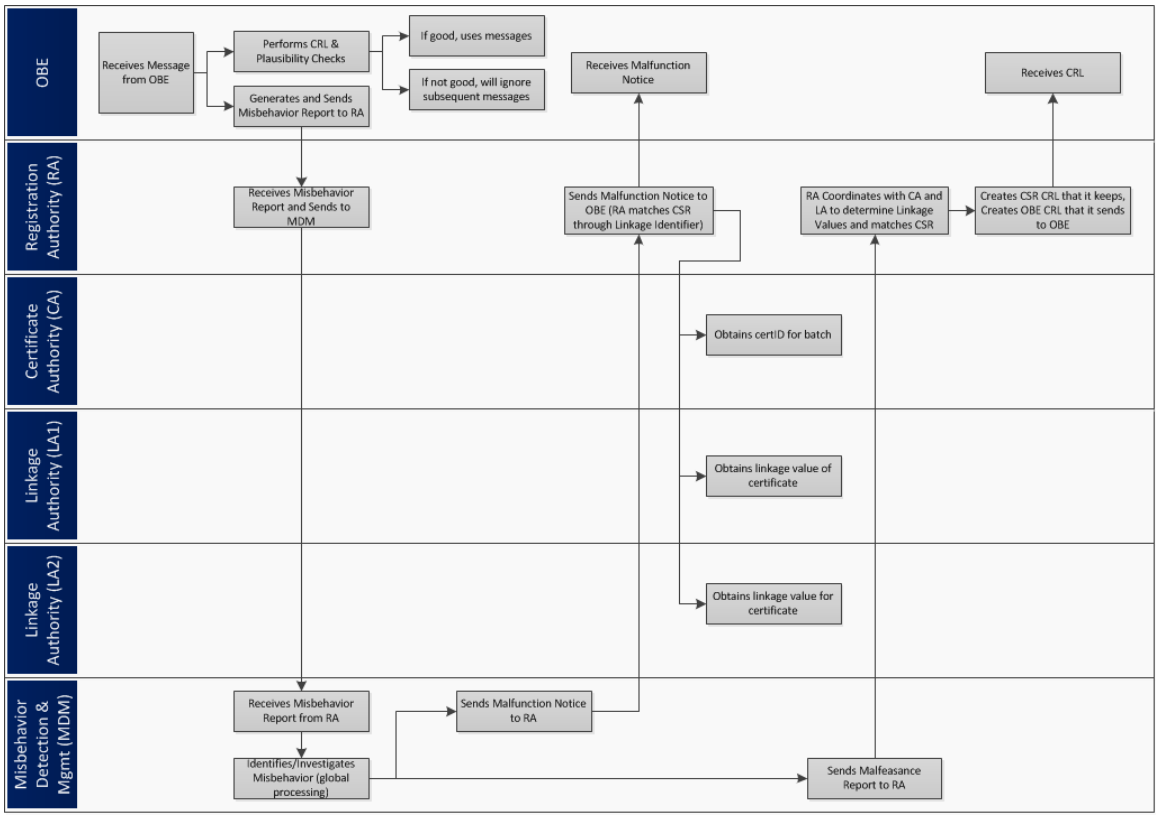

Warranty and Repair

The Warranty and Repair Phase encompass activities associated with the upkeep, repair, and replacement of OBE. The policies and standards for repairing or reactivating devices when they are damaged or have been placed on the CRL are under development but a few options have been identified. One option for a damaged device would be to have the user take the device to their local car dealership to have the device serviced. The user could also take the device to a specified location, such as the local vehicle inspection station to have the device repaired.

As stated previously, decisions on how the user would get a device back into the system once it has been placed on the CRL have not been made. One option is the user replacing the device. Another option includes providing the device with a new CSR and downloading a new batch of certificates. This would be similar to obtaining a new device and would require the user to go to a specified location to have this process performed. Both of these options would likely decrease participation in the system.

The process by which to distribute and download regular updates and maintenance activities is part of the technical architecture and guidance that is being developed.

End of Life

During the End of Life Phase, devices are either removed from the system and destroyed or are transferred to another user if the vehicle is sold. For devices which are removed from the system, some type of incentive should be considered to motivate users to return their devices to an appropriate location. One option would be to have users return their devices to the local dealership for the make of their vehicle (e.g., Honda, Toyota, etc.). The dealership would remove and destroy all unused certificates from the device and destroy the device as well through a recycling process.

Another option includes users returning their devices to a specified location such as the DMV or an authorized electronics retailer (e.g., Best Buy) where the device would be destroyed by removing unused certificates and recycling the device.

In the event that the user sells the car to a new owner, the device is also transferred to the new owner. If no PII is used in the system, it is unlikely that the new owner will have to take any action in regard to the device. A system that does include PII may require the new owner to take actions to register the device, but this process could also occur automatically through existing state title and registration systems as outlined in Chapter 6 of this report.

APPENDIX B CME Acronyms

BSMs	Basic Safety Messages
CA	Certificate Authority
CA_{ACT}	Certificate Authority Activation
CAMP	Crash Avoidance Metrics Partnership
CDDS	Communications Data Delivery System
CMEs	Certificate Management Entities
CRL	Certificate Revocation List
CSR	Certificate Signing Request
DOT	Department of Transportation
DSRC/WAVE	Dedicated Short Range Communications/Wireless Access in Vehicular Environments
ECC	Elliptic Curve Cryptography
ECDSA	Elliptical Curve Digital Signature Algorithm
HSM	Hardware Security Module
IEEE	Institute of Electrical and Electronic Engineers
ITS	Intelligent Transportation Systems
JPO	Joint Program Office
LA	Linkage Authority
MDM	Misbehavior Detection and Management
OBE	On Board Equipment
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
RA	Registration Authority
RA_{ACT}	Registration Authority Activation
RITA	Research and Innovative Technology Administration
RSE	Roadside Equipment
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-Device
VIIC	Vehicle Infrastructure Integration Consortium
VIN	Vehicle Identification Number

APPENDIX C Glossary of Assumptions

Vehicle-to-Vehicle Security Assumptions:[70]

- NHTSA intends to make an agency decision in 2013 about the future of connected vehicle technology. There are several scenarios that are possible:
  - One scenario is that data will clearly demonstrate that the V2V technologies are ready and sufficiently beneficial to warrant pursuit of future regulatory action for all new vehicles;
  - Or the data could demonstrate that the technologies are ready, but benefits of V2V cannot justify a NHTSA regulatory action – but instead recommend it become part of the New Car Assessment Program (NCAP) program where auto manufacturers would voluntarily equip to receive improved safety ratings;
  - Or the data could show that V2V technology is not sufficiently beneficial, or ready, to warrant either approach. In this case research would continue.
- To the extent possible, the technical V2V crash avoidance security system design should support a balanced approach to safety, security, and privacy – but the primary focus is on safety
- A trust relationship must exist among all participants in the V2V crash avoidance security system
- The V2V safety system must be interoperable:
  - Between vehicles from different manufacturers and other devices;
  - Throughout all geographic regions of the US.
- Interoperability requires accepted technical standards.
- Uniform rules of operation throughout the U.S. also need to be in place to support interoperability.
- 5.9 GHz DSRC will be used to send basic safety messages among vehicles and other devices.
- DSRC and/or other communications technologies may be used to provide necessary communications between vehicles and off-board security functions.
- The system must be able to withstand attacks and effectively recover from the effects of attacks.
- There will be no consumer subscription fees for any mandatory safety applications:
  - This does not preclude mandatory universally applicable taxes or fees to finance the system[71];
  - Subscription or other fees for non-mandatory, opt-in applications are possible.

[70] CAMP, *V2V Security Assumptions (working document).*

[71] Subscription fees refer to ongoing fees that a consumer voluntarily chooses to pay for a service. Mandatory universally applicable fees differ in that they are not voluntary and are therefore likely to either be collected by government agencies (such as in conjunction with vehicle registration) or included in the purchase price of the vehicle or equipment.

- The structure of the SCMS should be planned to allow gradual and seamless evolution from initial deployment to full deployment system with backward compatibility.
- The size of the SCMS is expected to grow with the number of equipped vehicles. However, all of the SCMS functions and interfaces necessary for full deployment should exist for initial deployment.
- On board security functions should be planned for the long lifetime of vehicles (i.e., greater than 10 years) from initial deployment.
- Interaction between the on board security and the SCMS should:
  - Be supported with limited DSRC security communication at strategic locations for initial deployment;
  - Support (opt-in) OEM-provided alternate communications to supplement increased level of interaction between the on board security and the SCMS.

Technical Design Assumptions:

- The prototype for test studies in the connected vehicle program is based on Public Key Infrastructure (PKI) for the security system
- Certificates have a lifespan of five minutes.
- The overlap between short-term certificates will be 30 seconds, reducing the risk of not accepting a signed communication due to time synchronization issues.
- Batches of five minute certificates are downloaded to an OBE once a year.
- Decryption keys are provided once a month to unlock monthly groups of the full yearly batch.
- The CSR will be the primary method of verifying an OBE with the system, and it will periodically expire and require renewal.
- As much automation as possible will be built into the OBE and its software. This includes programs that will automatically communicate with the RA for requests, reports, renewals of CSR and certificates, and other related activities. The current thinking around functions that should be automated within the OBE includes:
  - Monthly requests for decryption keys
  - Plausibility checks to ensure that a device is not misbehaving
  - Plausibility checks on incoming messages and automatic rejection of messages coming from misbehaving devices
  - Random selection of messages to put on a report to send to the RA for global processing
  - Annual certificate batch requests
  - CSR auto renewal
  - CRL requests
  - OBE and CRL processing

Certificate Revocation List Assumptions:

- The linkage value from the LA allows for efficient revocation of all certificates in a batch.
- The CRL will be distributed daily to OBE through RSE. Because the CRL could be large in size, the need for CRL distribution must be balanced against the capability of the communication system chosen. The specific technical details of the communication system are still being decided.
- Each OBE holds a dynamic list of revoked certificates based on the most recent CRL downloaded.

- There will be at least two CRLs – one for the CSR and one for the pseudo system. Technical and policy specifications about the connection between pseudo system CRL and CSR CRL have yet to be determined.
- Even distribution of production of certificates is assumed, which may not be the most appropriate way to implement the system, but provides an estimate and understanding of the extent of hardware and software needed.
- All keys associated with encryption in the following functions are compressed 96 Bytes, and keys used for hashing are estimated to be 256 bit SHA.

Organizational Assumptions:

- PII will not be permanently stored on the device (OBE) or on the certificate, nor can it be read by the RA function.
- PII would not be stored on the CSR; it would be stored in a separate database to eliminate any PII from the CSR being accessible to the RA.
- The best way to protect against attacks and unwarranted access to the system and vehicles is by creating distinct organizations, thus making it harder for people to share data across functions.
- Models with the LA and CA within one organization pose too high a risk of the CA being able to have access to the linkage values, therefore being able to identify any one particular vehicle.
- In order to authorize users and conform to any individual user access guidelines, there will need to be a credentialing process.
- Each function will manage its own backup and storage of critical data with appropriate technical, procedural and physical controls in place to guard against unauthorized access and potential data corruption or insider threat.
- The only personnel with authorized access to any function's information would be those who must input or work with the data. No personnel from one function should have access to another function's data and processes, including during certain processes like certificate revocation which demand the sharing of information.
- Maintenance of hardware and software used in the system will conform to state-of-the-art SOPs for the systems used, with continuous refresh and evaluation.
- V2V communications form the foundation of the need for CMEs. Facilitating safety messages sent between vehicles will be the initial function of the CMEs. Additional messages, in addition to safety messages, that may be exchanged between vehicles in the future will also need to rely on certificates for trusted communications between vehicles.
- The certificate lifespan of the RSE could range from one year to the lifetime of the system as it does not have to be the same lifespan as the certificates or the OBE.

Cost Assumptions:

- Estimates are provided for an initial period of six years.
- System benefits are not calculated.
- Base case estimates assume the separation of all PKI functional elements – CA, RA, LAs, activation, root CA, and program management and oversight. Cost saving efficiencies may be gained through the organizational models under discussion in this report.
- A discount rate of 7 percent is used, based on industry standards and government averages over the last 10 years.
- 105,120 five-minute certificates will be issued per year to a unit of OBE.

- IEEE 1609.2 certificates are assumed to be the certificate type, which determine estimates for the number of processors and HSMs required to support data loads and cryptographic processes.
- Cryptographic standards at 256 bit ECC are used to estimate the number of HSMs required. Maximum performance of this hardware is 1100 cryptographic operations per second[72]; however, in line with best practices, performance is assumed to be 550 cryptographic operations per second.
- As a rule of thumb, software and hardware supporting cryptographic operations accounts for 75 percent of the system costs. The remaining 25 percent of system costs support other administrative functions and additional shuffling and bundling of certificates.
- Personnel costs are estimated using the average rate across a team of individuals supporting one particular sub-function, with 2088 hours in a year.
- Several functions are assumed to require around the clock staffing (e.g., CA, RA, LA), based on PKI industry practices. As such, staffs at these facilities are assumed to work in 8 hour shifts, with three crews supporting a facility on a daily basis.
- Software, hardware, infrastructure, and personnel cannot be leveraged between the CA and RA due to organizational model constraints.
- Software estimates are provided on a per license basis for the software platform and database software. The platform will likely support multiple servers under one license but is assumed to be limited to a point. Database software is assumed to support one entire physical location per license.
- Hardware and software will be fully refreshed in the fourth year of use, resulting in cost surges that year. Costs in this refresh period are assumed to be 100 percent of the original investment for hardware and 50 percent of the original investment for software. Moore's Law is currently not being applied to this estimate.
- Total facilities costs are estimated using the average cost per square foot to furnish a facility to support PKI. Lease costs account for the potential for a variety of geographically dispersed locations by using the average cost between leasing a General Services Administration (GSA) facility and commercial information technology office listings[73].
- Current estimates provide for construction of CA, RA, and LA facilities and leasing of space for other functions.
- The number of locations is highly uncertain; however, due to the system requirements for redundancy and continuity of operations, a minimum of two facilities should be accounted for each function that requires heavy data processing.
- Space requirements of 1.5 square feet per server are assumed to calculate space needs for data centers, or facilities that house servers. This figure factors in the need for each facility to accommodate generators, extensive cooling systems, fire suppression systems, redundant communications, security, and administrative space.

[72] SafeGuard® CryptoServer Se-Series Benchmarks, http://hsm.utimaco.com/nc/en/products/se-series.
[73] Commercial listings obtained via LoopNet®, a leading commercial listing service.

- Developmental costs for data centers are assumed to be in the range of $600-$1500 per square foot, based on several construction projects undertaken in the past five years by Fortune 500 companies[74].
- The cost of the OBE is considered to be a sunk cost, necessary for system implementation. Therefore, it is excluded from this estimation because it has no bearing on the organizational evaluation. All other costs reviewed subsequently are affected by different organizational models.
- PKI typically includes three environments: development, operation and maintenance, and third party verification. Each environment requires independent equipment, software, and personnel. For this estimate, only the operation and maintenance environment is considered. Costs associated with the other environments may or may not match the operation and maintenance environment. At this point, the uncertainty surrounding the needs of the other environments is too great to include as part of this cost estimate.
- A help desk component will accompany an RA at each physical location.

[74] Richard Miller, "Details of Google's The Dalles Site Now Public."

APPENDIX D References

"About EAC: The U.S. Election Assistance Commission (EAC)." *The U.S. Election Assistance Commission (EAC): Help America Vote Act, National Voter Registration Act*. Web. 06 Jan. 2012. <http://www.eac.gov/about_the_eac/default.aspx>.

Alterman, Peter. *The U.S. Federal PKI and the Federal Bridge Certification Authority*. Publication. Federal PKI Steering Committee and Federal Bridge Certification Authority, 13 May 2001. Web. 22 Feb. 2012. <www.cendi.gov/presentations/alterman_pki_05-13-01.ppt>.

AT&T Inc. 2011 Annual Report. Retrieved July 3, 2012 from the SEC online Edgar database.

Barr, James. “HIPAA Records Management.” *Faulkner Information Services*. Web. 9 Jan. 2012.

Basili, Victor, et al. *The Goal Question Metric Approach*. University of Maryland. 1994. Web. 6 June 2012. <www.cs.umd.edu/~mvz/handouts/gqm.pdf>.

Boutin, Chad. "NIST Delivers Updated Draft Standards for Electronic Voting Machines." *National Institute of Standards and Technology*. 02 Jun. 2009. Web. 6 Jan. 2012. <http://www.nist.gov/itl/csd/voting_060209.cfm>.

Bowker, Mark, and Brian Garrett. *Virtualizing SQL Server Workloads with Microsoft Hyper-V R2*. Publication. Enterprise Strategy Group, 2011. Web. 9 Mar. 2012.

Brand, Adam. “QSA’s Don’t Assess PCI Compliance.” Web log post. *PCIfun Blog*. WordPress, 6 Dec. 2010. Web. 22 Feb. 2012. <http://pcifun.com/2010/12/16/qsas-dont-assess-pci-compliance/#comments>.

CertiPath. *CertiPath X.509 Certificate Policy*. Version 3.15. 2011. Web. 13 Feb. 2012. <https://www.certipath.com/library/policy-management-authority/policy-documents>.

Cormier, Bob. *Total Economic Impact of McAfee’s Network Security Platform’s Intrusion Prevention System: Single Company Analysis*. Publication. Cambridge: Forrester Research, Inc., 2009. <http://www.mcafee.com/elqNow/elqRedir.htm?ref=http://www.mcafee.com/us/resources/reports/rp-forrester-economic-impact-ips.pdf >.

"Disaster Recovery Pressures Rise Due to Cost of Downtime and More Stringent RTOs." *Symantec - AntiVirus, Anti-Spyware, Endpoint Security, Backup, Storage Solutions*. Symantec Corporation, 30 Jun. 2009. Web. 19 Jan. 2012. <http://www.symantec.com/about/news/release/article.jsp?prid=20090630_01>.

Dispatch Magazine On-Line. *Mobile Data for Public Safety*. Allen Media, 2011. Web. 1 Aug. 2012. <http://www.911dispatch.com/info/mobiledata.html>.

Erika McCallister, Tim Grance, and Karen Scarfone. *Guide to Protecting the Confidentiality of Personally Identifiable Information (PII).*Publication no. 800-122. Publication. National Institute of Standards and Technology, 2010. Print.

Evensen, Knut. “Slide 19: IEEE Security Services.” *CALM: Continuous Communications for Vehicles*. PowerPoint Presentation. SeVe COM Workshop, 1-2 Feb. 2006.

E-ZPass Interagency Group. "Welcome to the E-ZPass Interagency Group." *E-ZPass Group Home Page*. E-ZPass Interagency Group, 2011. Web. 7 Dec. 2011. <http://www.e-zpassiag.com>.

Forum of Incident Response and Security Teams. “Common Vulnerability Scoring System (CVSS-SIG)” FIRST.org, Inc., 2007. Web. 22 Feb. 2012. <http://www.first.org/cvss>.

Gumke, Randall. *Navy/Marine Corps Intranet Information Assurance Operational Services Performance Measures.* Publication no. NSN 7540-01-280-550. Thesis, U.S. Naval Post Graduate School. Monterey: June 2001. Web. 6 June 2012 <www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA396135>.

Health Information Technology for Economic and Clinical Health (HITECH) Act. Title XIII of Division A and Title IV of Division B of the American Recovery and Reinvestment Act of 2009 (ARRA). Pub. L. no. 111-5 (Feb. 17, 2009). Codified at 42 U.S.C. 300 et seq.; 17901 et seq.

"HIMSS Annual Security Survey Results." *HIMSS: Transforming Healthcare Through IT*. Publication. Health Information and Management Systems Society, 2 Nov. 2011. Web. 10 Jan. 2012.

HIMSS. *Patient Identity Integrity White Paper*. Publication. Health Information and Management Systems Society, 2009. Web. 9 Jan. 2012. <http://www.himss.org/asp/ContentRedirector.asp?ContentID=76274>.

Hite, Randolph C. *Electronic Voting Offers Opportunities and Presents Challenges.* Tech. no. GAO-04-766T. United States General Accounting Office, 2004. Print.

International Civil Aviation Organization. *Machine Readable Travel Documents Part 1: Machine Readable Passports.* Publication no. 9303. 6th ed. Vol. 1. Publication. ICAO Secretary General, 2006. Web. 22 Feb. 2012. <www.icao.int/publications/Documents/9303_p1_v1_cons_en.pdf>.

International Telecommunications Union, *ITU-T Recommendation X.509, ISO/IEC 9594-8*. 13 Nov. 2008. Web. 20 Jan. 2012. <http://www.itu.int/ITU-T/recommendations/rec.aspx?rec=X.509>.

Jacobson, Leslie. *Vehicle Infrastructure Integration Privacy Policies Framework*. 0.2 ed. Vol. 1. Publication. Institutional Issues Subcommittee of the National VII Coalition, 2007. Print.

Jafri, Amir and Leung, June. *PKI Deployment Business Issues- Resource Planning*. Publication. OASIS Member Section, 9 Aug. 2005. Web. 24 Feb. 2012.

Kasunic, Mark. *A Data Specification for Software Project Performance Measures: Results of a Collaboration on Performance Measurement.* Technical Report no. CMU/SEI-2008-TR-012. Software Engineering Institute, Carnegie Mellon. July 2008. Web. 07 June 2012. <http://www.sei.cmu.edu/library/abstracts/reports/08tr012.cfm>.

Keston, Geoff. "Electronic Medical Record: Trends". *Faulkner Information Services*. Web. 9 Jan. 2012.

Maartens, Hendrik. "Plan Against Disaster." *Business Brief* (2010): LexisNexis. Web. 13 June 2012.

"Meaningful Use of Electronic Health Records." *Health Policy Brief*. Health Affairs, 24 Aug. 2010. Web. 6 Jan. 2012. <http://www.healthaffairs.org/healthpolicybriefs/brief.php?brief_id=24>.

Mehrdad Majzoobi, et al.,(2008). *Testing Techniques for Hardware Security*, from *International Test Conference*, 2008, paper 31.3. IEEE International.

Mell, Peter, Karen Scarfone, and Sasha Romanosky. *CVSS: A Complete Guide to the Common Vulnerability Scoring System*. Vol.2. Publication. FIRST org., Inc., 2007. Web. 22 Feb. 2012 <http://www.first.org/cvss/cvss-guide.html>.

Miller, Richard. "Details of Google's The Dalles Site Now Public." *Data Center Knowledge*. 2008. Web. 4 June 2012. <http://www.datacenterknowledge.com/archives/2008/02/18/details-of-googles-the-dalles-site-now-public/>.

MITRE Corporation. "What Is CVE?" CVE – About CVE. MITRE, 2012. Web. 26 Jan. 2012. <http://cve.mitre.org/about/faqs.html>.

Mudit Bhargava, et al. (2010). *Attack Resistant Sense Amplifier Based PUFs (SA-PUF) with Deterministic and Controllable Reliability of PUF Responses*, from *International Symposium on Hardware-Oriented Security and Trust (HOST)*, 2010, 106-111. IEEE International.

National Institutes of Health National Center for Research Resources. *Electronic Health Records Overview.* Publication. McLean: MITRE Corporation, Apr. 2006. Web. 9 Jan. 2012.

Oltsik, Jon. *The True Costs of E-mail Encryption: Trend Mcro IBE (Identity-based) vs. PKI Encryption*. Publication. Enterprise Strategy Group, Inc., 2010. Web. 24 Feb. 2012.

Parker, Doug J. *Transit Cooperative Research Program Synthesis 73 – AVL Systems for Bus Transit: Update*. Publication. Transportation Research Board, 2008. Web. 1 Aug. 2012.

PCI Security Standards Council. *Payment Card Industry Approved Scanning Vendors*. Publication. PCI Security Standards Council, LLC., 2010. PCI Online Document Library.

PCI Security Standards Council. *Payment Card Industry (PCI) Security Standard: Requirements and Security Assessment Procedures*. Vol. 2. Publication. Wakefield: PCI Security Standards Council LLC, 2010. Web. 18 Jan. 2012.

PCI Security Standards Council. "What is the PCI Security Standards Council?" *PCI Standards & Documents*. PCI Security Standards Council, LLC. 2006. Web. 21 Dec. 2011. <https://www.pcisecuritystandards.org/security_standards/role_of_pci_council.php>.

Polk. "Approach – Data & Technology: Information is Only the Beginning." R.L. Polk & Co, 2012. Web. 31 July 2012. <https://www.polk.com/approach/data_and_technology>.

Practice Fusion. *Data and System Security White Paper*. Publication. Practice Fusion, 5 Oct. 2010. Web. 6 Jan. 2012. <http://www.practicefusion.com/resources/whitepaper/practice-fusion-emr-security.pdf>.

Regenscheid, Andrew, and Geoff Beier. *Security Best Practices for the Electronic Transmission of UOCAVA Election Materials*. Publication. Gaithersburg: Information Technology Laboratory. 2011. Print.

SAFE-BioPharma Association. *SAFE Certificate Policy*. Version 2.4. 2009. 13 Feb. 2012. <www.safe-biopharma.org/cp-pdf >.

Schmidt, Dale. “Wireless Telecommunications Carriers in the US: Industry Report.” *IBISWorld*. Report No. 51332. 2012.

Schneider, David. “Under the Hood at Google and Facebook.” *IEEE Spectrum: Inside Technology*. IEEE. 2011. Web. 4 June 2012. <http://spectrum.ieee.org/telecom/internet/under-the-hood-at-google-and-facebook/0>.

"Smartmatic: InfoCenter - Successful Case Studies." *Smartmatic: Electronic Voting Systems - Identity Management Solutions - Technology Consulting - Public Safety - Public Transportation*. Smartmatic, 2011. Web. 5 Jan. 2012. <http://www.smartmatic.com/case-studies/case-studies-home>.

Sprint Nextel. 2011 10-K Report. Retrieved July 3, 2012 from the SEC online Edgar database.

"State Felon Voting Laws." *Felon Voting*. ProCon.Org, 20 Dec. 2011. Web. 8 Jan. 2012. <http://felonvoting.procon.org/view.resource.php?resourceID=286>.

Swanson, Marianne, Pauline Bowen, Amy Phillips, Dean Gallup, and David Lynes. *Contingency Planning Guide for Federal Information Systems*. 800-34th ed. Vol. 1. Publication. National Institute of Standards and Technology, 2010. Print.

Swedberg, Claire. “PUF Technology Catches Clones.” *RFID Journal*. RFID Journal LLC., 4 Sep. 2008. Web. 22 Feb. 2012. <http://www.rfidjournal.com/article/view/4304>.

United States Department of Commerce, National Institute of Standards and Technology. “National Vulnerability Database Version 2.2.” *NIST Computer Resource Center*. 2012. Web. 8 Feb. 2012. <http://nvd.nist.gov/>.

United States Department of Commerce, National Institute of Standards and Technology. “NVD Common Vulnerability Scoring System Support v2.” *NIST Computer Resource Center*. 2012. Web. 15 Feb. 2012. <http://nvd.nist.gov/cvss.cfm?version=2>.

United States Department of Defense. *United States Department of Defense X.509 Certificate Policy*. Version 10.1. 2010. Print.

United States Department of Health and Human Services, Office of the National Coordinate of Health Information Technology. "The Nationwide Health Information Network, Direct Project, and CONNECT Software." *The Office of the National Coordinate of Health Information Technology*. 2011. Web. 9 Jan. 2012. <http://healthit.hhs.gov/portal/server.pt?open=512&mode=2&objID=3340>.

United States Department of Justice, Bureau of Justice Assistance. *Welcome to the National Motor Vehicle Title Information System (NMVTIS): Reporting Entity Webinar*. Publication. Bureau of Justice Assistance, 31 Oct. 2011. Web. 24 Feb. 2012. <www.vehiclehistory.gov/Webinar10_31_11.pdf>.

United States Department of Justice, Bureau of Justice Assistance and American Association of Motor Vehicle Administration. “Approaches for states to participate in NMVTIS.” *American Association of Motor Vehicle Administration*. 2012. Web. 28 Feb. 2012.

<http://www.aamva.org/KnowledgeCenter/Vehicle/NMVTIS/ApproachesforstatestoparticipateinNMVTIS.htm>.

United States Department of Justice, Bureau of Justice Assistance. "Don't Be Fooled: Are You About to Buy a Rebuilt Wreck or a Cloned Car?" Online brochure. National Motor Vehicle Title Information System Program Office. 2009. Web. 25 May 2012. <www.nmvtis.gov/NMVTIS_Consumer.pdf>.

United States Department of State, Bureau of Consular Affairs. "The U.S. Electronic Passport Frequently Asked Questions." *Bureau of Consular Affairs.* Web*.* 22 Feb. 2012. <http://travel.state.gov/passport/passport_2788.html#Five>.

United States Department of Transportation, Federal Motor Carrier Safety Administrations. *Safety Measurement System (SMS) Methodology*. 2nd ed. Vol. 2. Publication. Cambridge: The VOLPE Center, Jan 2012. Print.

United States Department of Transportation, Federal Highway Administration. "State Motor-Vehicle Registrations 2010." Highway Statistics Series 2010. 2011. Web. 15 Mar. 2012. <www.fhwa.dot.gov/policyinformation/statistics/2010/mv1.cfm>.

United States Department of Transportation, Research and Innovative Technology Administration. *Security Approach for V2V/V2I Communications Delivery System*. Publication. Washington, D.C.: Crash Avoidance Metrics Partnership and John A. Volpe National Transportation Systems Center, Aug. 2011. Print.

United States Department of Transportation, Research and Innovative Technology Administration. *Security Credential Management System: Security System Design for Cooperative Vehicle-to-Vehicle Crash Avoidance Applications Using 5.9 GHz Dedicated Short Range Communications (DSRC) Wireless Communications*. Intelligent Transportation Systems Joint Program Office, Apr. 2012. Print.

United States Federal Public Key Infrastructure Policy Authority. *X.509 Certificate Policy For The Federal Bridge Certification Authority*. Version 2.25. 2011. Web. 17 Feb. 2012. <www.idmanagement.gov/fpkipa/documents/FBCA_CP_RFC3647.pdf>.

United States General Services Administration. Public Affairs. "E-Government Moves Forward as Federal Bridge Certification Authority Conducts First Cross-Certifications." *Public Affairs*. 2002. Web. 22 Feb. 2012. <http://www.gsa.gov/portal/content/100338>.

United States Government Accountability Office. *GAO Cost Estimating and Assessment Guide: Best Practices for Developing and Managing Capital Program Costs*. Publication no.GAO-09-3SP. Publication. Washington, DC: *United States Government Accountability Office*, Mar. 2009. Print.

United States Office of Management and Budget. "Circular A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs, Appendix C." *Office of Management and Budget*. Washington, DC: Dec. 2011. Web. 26 Jan. 2012.

Utimaco. "Se-Series". *Utimaco Products*. 2012. Web. 30 May 2012. <http://hsm.utimaco.com/nc/en/products/se-series>.

Vehicle Infrastructure Integration Consortium. *VIIC Key Policy Issue – Security and Privacy*. Publication. VIIC, 6 Oct. 2011. Print.

VeriSign, Inc. *Reducing Complexity and Total Cost Of Ownership With VeriSign Managed PKI.* Symantec Corporation. 2011. Web. 27 Jan. 2012. <https://www4.symantec.com/.../whitepaper-cost-effective-pki.pdf>.

Verizon Wireless. 2011 10-K Report. Retrieved July 3, 2012 from the SEC online Edgar database.

Visa International Operating Regulations. Publication. Visa, 10 Oct. 2011. Web. 22 Feb. 2012. <http://corporate.visa.com/_media/visa-international-operating-regulations-oct-2011.pdf>.

Weimerskirch, André. *Security and Privacy in V2X: Current Approaches for Deployment*. PowerPoint Presentation. Escrypt Inc.: Embedded Security, Ann Arbor. Jan. 2012.

.

U.S. Department of Transportation
ITS Joint Program Office-HOIT
1200 New Jersey Avenue, SE
Washington, DC 20590

Toll-Free “Help Line” 866-367-7487
www.its.dot.gov

FHWA-JPO-12-078