National Security Credential Management System (SCMS) Deployment Support

SCMS Baseline Summary Report

www.its.dot.gov/index.htm
Final Report – January 12, 2018
FHWA-JPO-18-688

U.S. Department of Transportation

Produced by Booz Allen Hamilton
U.S. Department of Transportation

Notice

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
FHWA-JPO-18-688		

4. Title and Subtitle	5. Report Date
National Security Credential Management System (SCMS) Deployment Support: SCMS Baseline Summary Report	January 12, 2018
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Joshua Kolleda, Tyler Poling, David Fitzpatrick, Scott Andrews, James Marousek, Lawrence Frank	

9. Performing Organization Name And Address	10. Work Unit No. (TRAIS)
Booz Allen Hamilton 8283 Greensboro Dr McLean, VA 22102	11. Contract or Grant No.

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
	Final draft
	14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

This report provides a baseline of information on the Security Credential Management System (SCMS) to ensure a shared understanding of the SCMS ecosystem and how it will ensure trusted and private connected vehicle (CV) communications. A shared understanding will help the United States Department of Transportation (USDOT) and SCMS stakeholders develop potential SCMS ownership and governance models to deploy a National SCMS. This report reviews factors impacting ownership and governance at a high-level, as well as tradeoffs among technical SCMS design options to help inform future design and implementation. This report also includes an explanation of the PKI policies necessary to ensure a functioning and secure National SCMS. Finally, the report provides a description of the current SCMS Proof of Concept (PoC), which the USDOT is developing for use during the CV pilots and other federally-funded CV initiatives.

17. Key Words	18. Distribution Statement
Security Credential Management System (SCMS), Proof of Concept, Connected Vehicles, Pilots	

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
		84	

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

Table of Contents

List of Tables

List of Figures

Executive Summary

This report provides a baseline of information on the Security Credential Management System (SCMS) to ensure a shared understanding of the SCMS ecosystem and how it will ensure trusted and private connected vehicle (CV) communications. A shared understanding will help the United States Department of Transportation (USDOT) and SCMS stakeholders develop potential SCMS ownership and governance models to deploy a National SCMS. This report reviews factors that impact ownership and governance at a high-level, as well as tradeoffs among technical SCMS design options to help inform future design and implementation. This report also includes an explanation of the PKI policies necessary to ensure a functioning and secure National SCMS. Finally, the report provides a description of the current SCMS Proof of Concept (PoC), which the USDOT is developing for use during the CV pilots and other federally-funded CV initiatives.

The National SCMS Deployment Support project is intended to help identify and explore potential strategies for the establishment and governance of a National SCMS ecosystem through thoughtful engagement with stakeholders to seek guidance and potentially gain consensus on these strategies. Ideally, the outcome will also produce next steps and milestones to implement the consensus strategy or strategies.

Efforts by the National Highway Traffic Safety Administration (NHTSA) and industry around vehicle-to-vehicle (V2V) communications have presented a need to establish a National SCMS. Over the past several years NHTSA has been progressing with the necessary activities to implement V2V communications on a national scale. In August 2014, NHTSA issued an Advance Notice of Proposed Rule Making (ANPRM), as well as a report assessing the readiness of V2V technologies. The report included an analysis of the USDOT's research findings in several key areas including technical feasibility, privacy and security, and preliminary estimates on costs and safety benefits. The ANPRM sought public input on findings to support the regulatory work to require V2V devices in new light vehicles.

In December 2016, NHTSA decided to proceed with its regulatory and research authority by issuing a Notice of Proposed Rule Making (NPRM). The NPRM would require that all future light vehicles (model year 2022 and beyond) be equipped with V2V technology capable of transmitting basic safety messages (BSMs).[1] More recent indications from NHTSA question whether this NPRM will proceed but the need for a National SCMS to ensure trusted communications remains regardless of what decision is ultimately made. This is because a National SCMS is needed to protect CV communications that fall outside of the NHTSA's V2V rulemaking, as well as those that would fall within it. The CV ecosystem does not only rely on V2V communications. Other necessary CV operations require communications that are enabled by either dedicated short range communications (DSRC) or other communications technologies, such as cellular, Wi-Fi, or satellite, which will be included within this environment. The National SCMS needs to consider how various technologies and communications services will interact and operate within the anticipated CV environment, supporting safety and other types of applications and messages. Regardless of the communications medium, ensuring authentic

[1] BSMs contains information about vehicle position, heading, speed, and more relating to vehicle state and predicted path. The BSM contains no personally identifying information (PII) and is broadcast in a very limited geographical range. The BSM is defined in the SAE J2735 standard.

communications is critical for vehicle operators to trust the vehicle safety applications that rely on CV messages such as BSMs.

To deploy and oversee the multifaceted SCMS, there must be an ownership and governance model or models to ensure effective governance and continued operations. Without establishing these models now, the SCMS could organically grow into a non-sustainable system characterized by varying levels of security and enrollment of vehicle-to-everything (V2X) devices that do not meet standard requirements. Ownership is a key factor to ensure there is adequate funding for initial deployment, and to support sustainable operations. Essentially, there must be an SCMS Manager, which will serve as the governing body for the SCMS ecosystem. The SCMS Manager will also coordinate and monitor the operations of SCMS functions. The owner or owners of the SCMS Manager and SCMS functions will also greatly influence the level and type of industry governance, as well as stakeholder input for the development of governing policies.

There are three basic options to deploy an industry ownership and governance model: (1) public; (2) public-private partnership (P3); and (3) private. There are many potential SCMS Manager and broader SCMS ecosystem ownership and governance models based on the desired (and potentially necessary) public and private involvement. Models can range from completely public to completely private based on the objectives of the organization, government mandates, market needs, and many other factors. Each model will have its own strengths and weaknesses, along with specific implementation challenges. The deploying entity must balance fulfillment of public interest objectives with considerations such as cost, deployment schedule, risk, and desired government authority. It is also important to understand that the model does not have to be a static selection. For example, it could evolve from an initially completely government owned and operated model to a version where the government still has oversight and authority but the SCMS is primarily operated by private entities. Along with the development of ownership and governance model(s), a strategy for deployment and implementation of that model must be developed. The implementation plan is as important as the selected ownership and governance model in making the National SCMS a reality. Depending on the selected model, an implementation plan would contain differing activities and milestones.

The ownership and governance models must consider the entire “SCMS ecosystem,” which includes the SCMS itself and the peripheral industry participants that play a role in developing, provisioning, operating, and maintaining the equipment and systems necessary to support the security functions identified for the overall CV enterprise. The SCMS ecosystem includes the entities responsible for originating CV equipment and applications (including services provided to the vehicle/user), entities responsible for certifying that this equipment and these applications conform to specified requirements and standards, entities responsible for selling and provisioning the equipment and/or the applications, entities responsible for maintaining and servicing the equipment and/or the applications, end users such as vehicle owners/drivers, and state and local agencies that may implement applications using vehicle based and/or roadside equipment.

The SCMS itself encompasses all Public Key Infrastructure (PKI) functions necessary to establish and maintain privacy and security within the V2X ecosystem. It provides the various functional elements (described in detail in Section 2.3) that will perform these security management functions over the equipment and/or application lifecycle. This includes various levels of certificate authority; functions to detect, identify, and remove misbehaving devices from the system; and functions to facilitate the operation of the SCMS without compromising the privacy of the system users.

At the core of the SCMS ecosystem are the root certificate authority (CA) or authorities, trust anchor management function, the SCMS Manager, and its associated policies and regulations. The SCMS Manager provides the core policy and governance foundation for the SCMS ecosystem in general, and the SCMS specific functions in particular. The SCMS Manager’s authority, responsibilities, ownership, and organizational

structure has yet to be determined, but it is likely that it will serve as the motivating force to establish the SCMS functions through policy and regulation. The SCMS Manager will also likely serve in an ongoing capacity as the core of a governance body, coordinating and monitoring operations among the various functions and certificate management entities (i.e., entities that group multiple SCMS functions/components under a single owner/operator). It is also expected that the SCMS Manager will collaborate with entities and organizations outside of the immediate SCMS such as certification and testing shops, state and local transportation organizations (e.g., state departments of transportation and divisions of motor vehicles), vehicle inspection facilities, automotive repair shops, and automotive or device dealerships (described in detail in Section 2.4). The SCMS Manager and/or its governance board will need to interface with other governance bodies, such as those overseeing credential management systems in Canada and Mexico.

Other than the hierarchy created by the certificate trust chain, the specific ownership, operational, and governance models of the SCMS Manager and SCMS components (e.g., certificate authorities) is not yet defined. In developing these models, the USDOT and V2X ecosystem stakeholders need to consider the following questions and many more to ensure a functional and secure system:

- Who operates each CA and other SCMS components and under what governance model?
- Is everything operated by the same organization?
- Are there multiple PKIs (i.e., multiple roots)?

Depending on the activities carried out by each function (specifically those related to identifiable information), the function may need to be handled in a centralized manner separate from other functions. For example, the certificate management entity (CME) owning and operating the MA function cannot own and operate any other function, and thus it would need to be considered as a centralized function, isolated from the other CMEs. Other functions may or may not be isolated in this way. PKI policies developed by the SCMS Manager will determine the necessary separation of functions based on the final PKI design and root structure. Refer to Sections 2.3.2 and 2.6 for an analysis of component separation and potential CME groupings.

Within early discussion of potential ownership and governance models for this project, the team identified high-level models ranging from completely publicly owned, governed, and operated to completely private. More specifically, the team begins to explore the potential deploy a National SCMS though five potential high-level models (summarized in Section 3.1):

1. Completely public
2. Government-led P3
3. P3 concession
4. Industry-led P3
5. Completely private.

There are several factors that will influence the development and deployment of ownership and governance models. Throughout the early stages of the National SCMS Deployment Support project, the team has identified public interest objectives that must be addressed and fulfilled by the selected ownership and governance model. Current public interest objectives consist of secure communications, privacy, availability (e.g., interoperability, redundancy, flexibility), stakeholder representation, affordability, and performance. The team also identified evaluation criteria or considerations that the selected model will greatly influence and, at the very least, must be thoroughly discussed during model development. Current evaluation criteria consist of ownership, funding, policy creation and approval, oversight and auditing, trust anchor management, legislation and regulation, competition, and overall risk. Many of these objectives and evaluation criteria overlap or influence each other. Refer to Section 3.1 for descriptions and high-level tradeoff analysis of the objectives and

criteria. In addition to the public interest objectives and evaluation criteria, there are other technical and governance areas of interest identified through the CV pilots that should not have a significant impact on the ownership or governance structure, but will need to be addressed within the PKI Certificate Policy (CP). These areas of interest are:

- Trust anchor and certificate authority management (Section 3.2)
  - Single versus multiple root certificate authorities
  - Single versus pseudonym certificate authorities
  - Certificate authority retirement
  - Elector establishment and management
- Policy compliance and enforcement (Section 3.3)
  - Audits
  - Device certification (and potential re-enrollment) criteria for end entities and back office systems
- Communications options for providing SCMS services (Section 3.4)
  - Impact of commercial communication services
  - SCMS-provided certificate usage.

A comprehensive PKI CP is necessary no matter the structure of the final technical SCMS build, the SCMS Manager owner, CME owner/operator, or governance model. The PKI policy should be structured similarly across the board according to industry best practices (i.e., RFC 3647). However, the content and guidance within the policies will differ based on technical build, ownership model, and governance model. Vice versa, these models may be developed to support desired policies identified by stakeholders during the model development process. A PKI CP describes the operational and security requirements that will be implemented within the PKI. The CP is typically made publicly available so that any interested party can examine the requirements under which a specific implementation of the policy is operated. Having an overarching CP is especially important in a distributed environment, as is anticipated for the SCMS PKI. As currently envisioned, portions of the PKI must be operated by separate organizations while still meeting the overall functional, security, and privacy requirements. Without an overarching CP to which all elements of the PKI align, it will not be possible for the SCMS Manager to provide appropriate guidance and oversight to the implementing organizations.

Currently, the USDOT is leading the build and deployment of the SCMS PoC to support the CV pilots and other federally-funded V2X related efforts. While the National SCMS will look substantially different from the SCMS PoC, the government and industry can make use of SCMS PoC practices, policies, lessons learned, and potentially even SCMS PoC infrastructure when deploying the National SCMS ecosystem. As the SCMS PoC effort continues to advance, the National SCMS Deployment Support team will stay engaged to implement new information and lessons learned during the development of potential ownership and governance models, as appropriate.

Chapter 1: Introduction to a National SCMS

This chapter provides additional information on the scope of the National SCMS Deployment Support project, the need for a National SCMS, and the need for specific ownership, governance, and operational models for the National SCMS. Additionally, this chapter provides context on the necessary activities to implement the selected National SCMS model or models. For background information on the CV concept and Public Key Infrastructure, refer to Appendix A.

1.1 Project Scope

The National SCMS Deployment Support project is intended to help identify and explore potential strategies for the establishment and governance of a National SCMS ecosystem through thoughtful engagement with stakeholders to seek guidance and potentially gain consensus on these strategies. Ideally, the outcome will also produce next steps and milestones to implement the consensus strategy or strategies. The strategies will include guidance and plans around:

- Establishment of an SCMS Governance Board (or similar oversight entity), including definitions of functions, roles, and responsibilities
- Establishment of an overall SCMS Manager (or similar system management entity), along with definitions of functions, roles, and responsibilities for managing ongoing operations and executing any functions deemed to be "inherently central"
- Establishment of management entities that will be part of the larger SCMS delivery system (and whose authority is directly dependent on and linked to the SCMS Manager)
- High-level policies and procedures that define and guide interactions among the various entities that make up the SCMS
- Roles and responsibilities of other entities that are not directly part of the SCMS but who may play a supportive, authorization, administrative, or other indirect role (such as the federal government, state governments, industry associations, etc.)
- Business and financial options for initial deployment and sustainable operations.

This report is an integral part of this project and serves as a baseline of SCMS information to create a shared understanding among stakeholders. It provides background, assumptions, design trade-offs, feasibility considerations, and other design or operational issues that may impact the ownership, governance, and operations of National SCMS entities, elements, and functions.

1.2 A Need for a National SCMS

The USDOT has been supporting CV research, development, testing, and deployment for more than a decade. As discussed in Appendix A, vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications, and other supporting technologies, are creating a new environment that allows vehicles and transportation infrastructure to communicate with each other. These communications systems are composed

of devices, installed in vehicles and on roadside units (RSUs), that exchange messages over wireless communications media and protocols. For instance, V2V and V2I devices can use information from other vehicles to determine if a crash-avoidance warning to the vehicle's driver is needed.

CV technologies have the potential to significantly prevent or reduce the impact of millions of accidents every year, increase transportation mobility, and reduce emissions. Despite the potential benefits, CV operations introduce challenges that are not present in existing vehicle systems. Trust and the appropriate protection of individual privacy are critical to ensuring successful messaging among vehicles and infrastructure.

To address the need for trust and privacy concerns around connected vehicles, the USDOT has partnered with the Crash Avoidance Metrics Partnership (CAMP) to design and develop an SCMS Proof-of-Concept (PoC). CAMP is an automotive coalition that was formed to accelerate the implementation of crash avoidance countermeasures in passenger cars to improve traffic safety. The SCMS PoC provides early CV deployers and other research initiatives with the mechanism for CV devices to exchange information in a trustworthy and private manner using digital certificates.

Recently, the USDOT has transitioned their research efforts to activities around adoption and eventual deployment of CV systems. In preparation for deployment, the concepts tested in the SCMS PoC must transition to support CVs on a national level. Therefore, a National SCMS must be established to govern and manage the security credentials needed for the gradual national roll-out of CVs.

1.2.1 A National SCMS is Needed to Address the V2V Rulemaking

NHTSA's efforts around V2V have also presented a need for establishing a National SCMS. Over the past several years NHTSA has been progressing with the activities needed to implement V2V communications on a national scale. In August of 2014, NHTSA issued an ANPRM, as well as a report assessing the readiness of V2V technologies. The report included an analysis of the USDOT's research findings in several key areas including technical feasibility, privacy and security, and preliminary estimates on costs and safety benefits. The ANPRM sought public input on findings to support the regulatory work to require V2V devices in new light vehicles.

In December of 2016, NHTSA decided to proceed with its regulatory and research authority by issuing a NPRM. The NPRM would require that all future light vehicles (model year 2022 and beyond) are equipped with V2V technology capable of transmitting BSMs.[2] More recent indications from NHTSA question whether this NPRM will proceed. But the need for a National SCMS to ensure trusted communications will remain regardless of what decision is made.

As described in Appendix A, the messages must be authentic and trusted for an effective national deployment of V2V technology. At the same time, privacy (or anonymity), from the perspective of private sector vehicle operations, is equally as important. The National SCMS is a critical element of this approach and will need to be established prior to vehicles equipped with V2V communications technology to ensure that both V2V (specifically BSMs) and V2I communications are both secure and appropriately anonymous.[3]

[2] Basic safety messages (BSMs) contains information about vehicle position, heading, speed, and more relating to vehicle state and predicted path. The BSM contains no personally identifying information (PII) and is broadcast in a very limited geographical range. The BSM is defined in the SAE J2735 standard.
[3] Reference(s): NPRM and ANPRM Federal Motor Vehicle Safety Standards; V2V Communications

1.2.2 A National SCMS is Needed to Address the Vehicle-to-Everything (V2X) Environment

A National SCMS is also needed to protect CV communications that fall outside of the NHTSA's V2V rulemaking. The CV ecosystem does not just rely on V2V communications. Other necessary CV operations require communications that are enabled by either DSRC or other communications technologies, such as cellular, Wi-Fi, or satellite, which will be included within this environment**.** The National SCMS needs to consider how various technologies and communications services will interact and operate within the anticipated CV environment, supporting safety and other types of applications and messages. Regardless of the communications medium, ensuring authentic communications is critical for vehicle operators to trust the vehicle safety applications that rely on CV messages such as BSMs. Otherwise, no action will be taken based on the new CV technology and exchanged CV information. Basically, there is no benefit unless there is trust in the system. Instead, an unsecure or untrustworthy system could produce a multitude of harmful effects on the transportation system.

1.3 SCMS Ownership and Governance Model(s) Development

In addition to NHTSA mandates on V2V technology, it is expected that there will be substantial growth in ubiquitous CV communications, and its security and trust must be protected. As these developments draw new suppliers into the market and address new use cases, these suppliers must have clear, consistent guidance from a formalized SCMS Manager and Certificate Management Entities (CMEs), entities that own and/or operate one or more SCMS functions, that explain how devices will be granted certificates to allow them to plan for deployment in volume. Great strides have been made in establishing and operating the technical SCMS Proof of Concept (PoC), and the National SCMS Deployment Support project will address the last missing pieces – ownership, governance, management, policy, and oversight for a National model. However, the structure and policies suitable to operate the significantly smaller-scale PoC will not be sufficient to govern the security credential needs of a full national deployment of V2X devices. The SCMS may be considered an entirely new public service and, as such, will require ongoing and relevant policies, practices, auditing, oversight, and compliance to ensure efficient and effective operations.

To deploy and oversee the multifaceted SCMS, there must be an ownership and governance model or models to ensure effective governance and continued operations. Without establishing these models now, the SCMS could organically grow into a non-sustainable system characterized by varying levels of security and enrollment of V2X devices that do not meet standard requirements. For example, without a feasible ownership and funding model, there would likely be a lack of transparent ownership of SCMS functions which would also lead to a lack of accountability. There may also be various, possibly inconsistent funding streams that could lead to issues in availability and inconsistent services. Without a governance model and accompany policies and processes, there could be varying security, privacy, and device standards across components and/or geographical areas. This could result in interoperability concerns and lack of confidence in the system. Of course, a lack of consistent PKI policies could also result in exploitable system vulnerabilities that could cripple the entire CV system. Without considering the worst effects, this would at least render the system useless.

Ownership is a key factor to ensure there is adequate funding for initial deployment, and to support sustainable operations. Essentially, there must be an SCMS Manager which will serve as the governing body for the SCMS ecosystem. The SCMS Manager will also coordinate and monitor the operations of SCMS functions. Refer to Chapter 2 for additional information on the various SCMS functions. The owner or owners of the SCMS Manager and SCMS functions will also greatly influence the level and type of industry governance, and stakeholder input in the development of governing policies. An important question to answer, which the team

will explore further within this project, is how the authority to govern the SCMS public service will be bestowed upon the SCMS Manager.

There are three basic options to deploy an industry ownership and governance model: public, P3, and private. There are many potential SCMS Manager and broader SCMS ecosystem ownership and governance models based on the desired (and potentially necessary) public and private involvement. Models can range from completely public to completely private based on the objectives of the organization, government mandates, market need, and many other factors. Each model will have its own strengths and weaknesses, along with specific implementation challenges. The deploying entity must balance fulfillment of public interest objectives with considerations such as cost, deployment schedule, risk, and desired government authority. It is also important to understand that the model does not have to be a static selection. For example, it could evolve from an initially completely government owned and operated model to a version where the government still has oversight and authority but the SCMS is primarily operated by private entities.

While this report discusses some of the factors that impact the development of ownership and governance models, those factors will be fully explored within a follow-on task focusing on model development.

1.4 SCMS Deployment and Implementation

Along with the development of an ownership and governance model(s), a strategy for deployment and implementation of that model must be developed. The implementation plan is as important as the selected ownership and governance model to making the National SCMS a reality. Depending on the selected model, an implementation plan would contain differing activities and milestones. For example, a variation of a P3 model would likely include facilitation of industry working sessions, development of industry consortia, and establishment of official agreements among key stakeholders.

The strategy would minimally include a transition plan to move from model planning to initial deployment. Implementation would include the following activities and artifacts:

- **Establishment of the National SCMS implementation workgroup.** Following the structure of the model planning and development process, the transition plan begins by setting the foundations for an implementation workgroup, industry consortium, and/or task force committee as necessary. These groups will could be comprised of government officials and industry stakeholders needed for the selected governance model and must have a guiding organizational charter.
- **Roles and responsibilities document.** Many entities will be involved in the implementation of a National SCMS. To ensure all necessary entities have a role and that the relevant skill sets are covered, the transition plan will include a roles and responsibilities document outlining this information.
- **Communications plan.** A successful transition requires open and designated lines of communication among participating parties. The communications plan will detail the key individuals who will interact between the planning and implementation teams to ensure the proper levels of information sharing and transparency.
- **Project plan and timeline.** Another crucial element of the transition plan is the project plan and timeline. This will turn the “next steps” for implementation into actionable tasks for the implementation team as well as a timeline for completing each task. The project plan will ensure a seamless transition from planning into deployment. Activities within the project plan could consist of the following, with subtasks, with owners.
    - Establish the SCMS Manager with internal departments, including a technical operations oversight function

  - Establish PKI policies including those for all types of certificate authorities (CAs), registration authorities (RAs), and linkage authorities (LAs), as well as the communications between these components
  - Establish policies for certification labs and authorize at least one certification lab to evaluate and approve components
  - Set up initial set of electors (or other trust anchor management mechanism) with one logical misbehavior authority (MA) with a certificate revocation list (CRL) store and at least one root CA
- **Evaluation and feedback plan.** The implementation team will monitor the progress of standing up the selected ownership and governance model.

1.5 SCMS PoC Relation to a National SCMS

Federally-funded CV pilots and other initiatives will use the SCMS PoC both to secure communications for CV operations and as a means to provide valuable insight into SCMS operations. Essentially, the SCMS PoC is a prototype for a National SCMS. The SCMS PoC will test and refine operational concepts that will help to extract the policies, procedures, and lessons learned from the design, deployment, and operations – which will, in turn, be used to help establish a National SCMS.[4] The National SCMS Deployment Support project will review the products and lessons learned from the SCMS PoC and pilot initiatives – along with input from industry, government experts, and stakeholders – to inform recommendations to deploy an ownership, governance, and operational model for the National SCMS. Chapter 5 of this document provides more detail on the SCMS PoC.

[4] Reference(s): Fundamental Principles and Research of the SCMS PoC Presentation

Chapter 2: SCMS Overview and Functionality

This chapter provides an overview of the SCMS functions and ecosystem. It discusses how the SCMS works, and where it will fit in relation to the National Intelligent Transportation Systems (ITS) Architecture. It also introduces the SCMS ecosystem. This includes descriptions of the functions and interdependencies within the SCMS; SCMS entities and functions that are external to the SCMS; and end entities (i.e., devices) that will connect to the SCMS. Furthermore, this chapter discusses the assumptions pertaining to SCMS design and operation.

2.1 High-Level SCMS Overview

As discussed in Appendix A, the SCMS is the back-end system intended to support secure communications and protect privacy in V2V and V2I security-based communications within the connected vehicle environment (CVE). The SCMS provides digitally signed certificates that can be used as part of the process for signing and encrypting messages. It uses a PKI-based approach that employs highly innovative methods of encryption and certificate management to facilitate trusted and private communication. PKI systems use asymmetric key systems consisting of two separate but mathematically related keys. The private key is kept secret by its owner, while the public key may be distributed to anyone (hence the name public key). Knowledge of the public key does not enable anyone to derive the private key without impractically extensive mathematical processing. Use of a public key system simplifies issues of key management and distribution, because public keys require no security. However, an infrastructure device will be required to generate and manage its private and public keys (i.e., a PKI).

PKI systems support a variety of secure communication functions, as described in Appendix A. For example, users can protect data intended for a particular recipient by encrypting that data using the recipient's public key. The data can only be unencrypted by someone who possesses the corresponding private key (i.e., the intended recipient). Messages can also be digitally signed by computing a digest of the message (a mathematically computed hash) and using the sender's private key as input to the digital signature algorithm (e.g., RSA, ECDSA). Recipients will use the message digest (computed independently from message body), the sender's public key and the digital signature attached to the message as inputs to the signature verification function. The signature verification function will compare two separate mathematical values to verify the authenticity of the signature. If the mathematical values match, then the message must have been sent by the claimed sender, as only they have their private key, and the message was not altered during transmission (because if it had been changed, the hashes would not match).

The USDOT's Architecture Reference for Cooperative and Intelligent Transportation (ARC-IT) merges the USDOT's National ITS Architecture and its Connected Vehicle Reference Implementation Architecture (CVRIA). ARC-IT has been developed to further facilitate deployment support and provides a common framework for planning, defining, and integrating ITS services.

ARC-IT is defined through approximately 100 ITS service packages; each containing a reference architecture; which can be used and uniquely adapted by ITS projects deploying one or more of the ITS services defined by the respective ITS service packages. Collectively, the ARC-IT Security and Credentials Management[5], Privacy Protection[6], and Core Authorization[7] service packages provide a reference architecture to address the need to facilitate secure ITS communications for most of the other ITS service packages. Accordingly, as with the SCMS POC, the National SCMS will be aligned with these two reference architectures.

It is noted that these three ARC-IT service packages leverage findings from other analyses, including the FHWA's V2I cybersecurity tasks, the CV Pilot Deployment Program, and the outputs of the US-EU Steering Group's Harmonization Task Group (HTG) 6 ("Candidate Harmonized Policies for Cooperative ITS Security Implementation") and AU-EU-JP-US Steering Group's HTG 7 ("Standards Selection, Gap Analysis, and Identifiers for Cooperative Intelligent Transportation Systems"). The result is a finer level of security objective assessment. All information flows defined in ARC-IT have been assessed for their confidentiality, integrity, and availability (CIA) objectives, and those assessments justified. Assessment of the security objectives related to information flows allows the architecture to derive security objectives for physical objects. This has resulted in the creation of device classes, or groupings of device security classifications, organized to ease manufacture and procurement.[8]

The European Commission (EC) is taking a similar security approach. The Cooperative ITS (C-ITS) has security-related requirements owing to its need to establish and maintain trust between disconnected entities with no prior relationship. While any communications system needs to provide a mechanism to allow communicating partners to trust each other, the environmental and performance characteristics of the system have an impact on what kinds of technologies might work. For C-ITS, a PKI has been developed to support the needs unique to the wireless vehicle environment. This PKI may be applicable to other systems in C-ITS. To better understand the policies surrounding applicability of PKI and other security mechanisms, HTG 6 produced an analysis of the systems necessary to operate this PKI and how they might interact (in case there are multiple credential management systems).

The foundation of the reference architecture provided for with the three ARC-IT services, listed above, is the Cooperative ITS Credentials Management System (CCMS). The CCMS is a high-level aggregate representation of the interconnected systems that enable trusted communications between mobile devices and other mobile devices, roadside devices, and centers and protect unauthorized access of data. It will provide the framework for the National SCMS, and accordingly the SCMS Manager will need to provide relevant guidance for the core CCMS functions, which are listed in Table 1 below.

[5] For more information on the ARC-IT Security and Credentials Management service package, visit: http://local.iteris.com/arc-it/html/servicepackages/sp63.html#tab-3.
[6] For more information on ARC-IT Privacy and Protection service package, visit: http://local.iteris.com/arc-it/html/servicepackages/sp119.html#tab-3.
[7] For more information on ARC-IT Core Authorization service package, visit: http://local.iteris.com/arc-it/html/servicepackages/sp12.html#tab-3.
[8] Reference: ARC-IT: https://local.iteris.com/arc-it/index.html.

Table 1: CCMS Functions[9]

Functionality	Description
Authorization	Provides authorization credentials (e.g., pseudonym certificates) to end entities. The end entity applies for and obtains authorization credentials, enabling the end entity to enter the "Operational" state. This function requires an interactive dialog, including at minimum a Certificate Request from the end entity desiring certificates. This request will be checked for validity, with the embedded enrollment certificate checked against an internal blacklist. If all checks are passed, it will distribute a bundle of linked pseudonym certificates suitable for use by the requesting end entity, with the characteristics and usage rules of those certificates dependent on the operational policies of the CCMS. It also provides the secure provisioning of a given object's Decryption Key in response to an authorized request from that object. The retrieved Decryption Key will be used by the receiving object to decrypt the "next valid" batch within the set of previously retrieved Security Credential batches.
Enrollment	Provides enrollment credentials to end entities. The end entity applies for and obtains enrollment credentials that can be used to communicate with other CCMS components, entering the "Unauthorized" state. CCMS Enrollment components also participate in de-registration processes through interaction with CCMS Revocation components.
Misbehavior Reporting and Action	Processes misbehavior reports from end entities. Misbehavior reports are analyzed and investigated if warranted. Investigated misbehavior reports are correlated with end entities and systemic issues are identified. If revocation is warranted, this component provides information to Authorization or Revocation components to initiate revocation and/or blacklisting, as appropriate.
Provisioning	Provides the end entity with material that allows it to enter the 'Unenrolled' state. This consists of root certificates and the crypto material that allows it to communicate securely with the Enrollment components. This function ensures the requesting entity meets requirements for provisioning and provides the certificates and relevant policy information to entities that meet the requirements.
Revocation	Generates the internal blacklist and CRL and distribute them to other CCMS components and end entities. Once placed on the CRL, an end entity is in the Unauthorized state. Once placed on the blacklist, an end entity is in the Unenrolled state.

[9] Reference(s): Cooperative ITS Credentials Management System functionality section of ARC-IT: https://local.iteris.com/arc-it/html/physobjects/physobj86.html#tab-1,

2.2 SCMS Ecosystem and Boundaries

The SCMS ecosystem includes the SCMS itself and the peripheral industry participants that play a role in developing, provisioning, operating, and maintaining the equipment and systems necessary to support the security functions identified for the overall CV enterprise.

As illustrated in Figure 1 below, the SCMS ecosystem includes the entities responsible for originating CV equipment and applications (including services provided to the vehicle/user); entities responsible for certifying that this equipment and these applications conform to specified requirements and standards; entities responsible for selling and provisioning the equipment and/or the applications; entities responsible for maintaining and servicing the equipment and/or the applications; end users such as vehicle owners/drivers; and state and local agencies that may implement applications using vehicle based and/or roadside equipment. These entities will interact in some way with the SCMS functions over the lifecycle of any given application and any given equipment implementation (i.e., on-board equipment [OBE], aftermarket safety device [ASD], or road side equipment [RSE]).

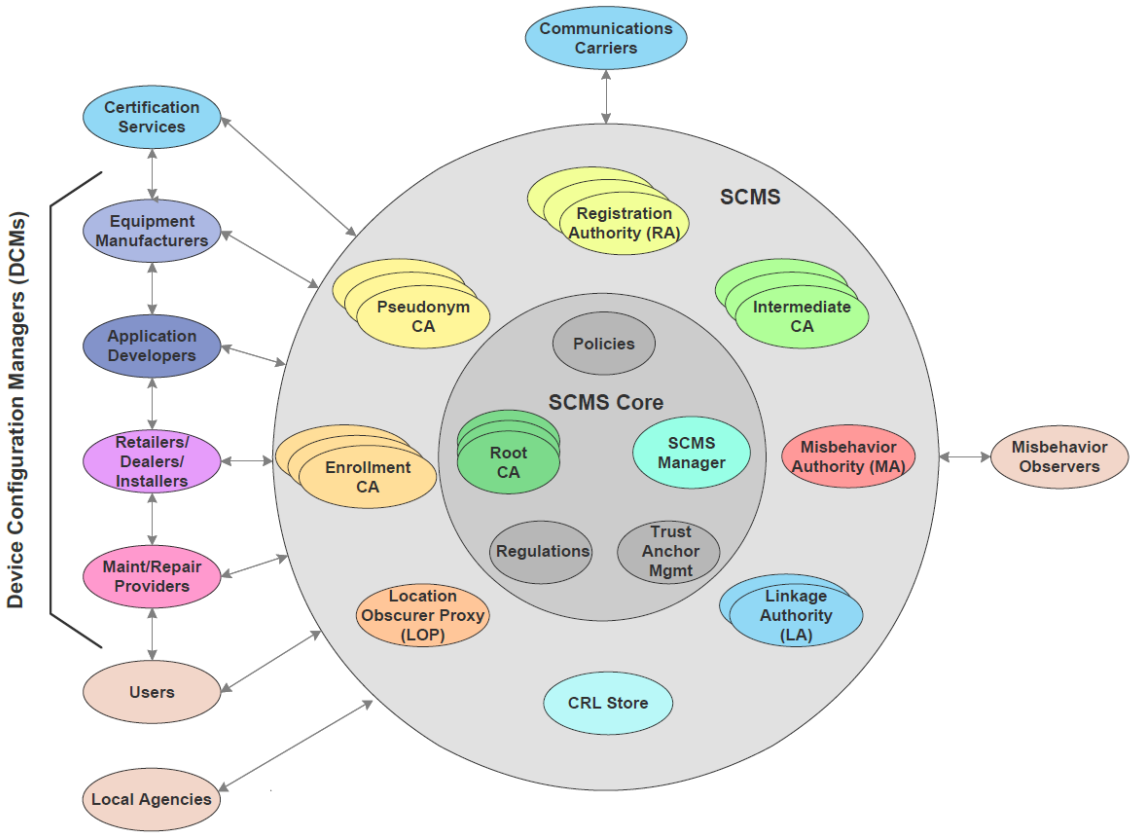

Figure 1: SCMS Ecosystem

The SCMS itself encompasses all PKI functions necessary to establish and maintain privacy and security within the V2X ecosystem. It provides the various functional elements (described in greater detail in the following section of this report) that will perform these security management functions over the equipment and/or application lifecycle. This includes various levels of certificate authority; functions to detect, identify, and

remove misbehaving devices from the system; and functions to facilitate the operation of the SCMS without compromising the privacy of the system users.

At the core of the SCMS ecosystem are the root CA(s), the trust anchor management function, the SCMS Manager, and its associated policies and regulations. The SCMS Manager provides the core policy and governance foundation for the SCMS ecosystem in general, and the SCMS specific functions in particular. The SCMS Manager's authority, responsibilities, ownership, and organizational structure has yet to be determined, but it is likely that it will serve as the motivating force to establish the SCMS functions through policy and regulation. The SCMS Manager will also likely serve in an ongoing capacity as the core of a governance body, to coordinate and monitor operations among the various SCMS CMEs and functions. It is also expected that the SCMS Manager will collaborate with entities and organizations outside of the immediate SCMS, such as certification and testing shops; state and local transportation organizations (e.g., state departments of transportation and divisions of motor vehicles); vehicle inspection facilities; automotive repair shops; and automotive or device dealerships. The SCMS Manager and/or its governance board will need to interface with other governance bodies, such as those overseeing credential management systems in Canada and Mexico.

2.3 SCMS Functions and Potential Functional Configurations

Other than the hierarchy created by the certificate trust chain, the specific ownership, operational, and governance models of the CA is not yet defined. In developing these models, the USDOT and V2X ecosystem stakeholders need to consider the following questions and many more to ensure a functional and secure system:

- Who operates each CA and other SCMS components and under what governance model?
- Is everything operated by the same organization?
- Are there multiple PKIs (i.e., multiple roots)?

Depending on the activities carried out by each function (specifically those related to identifiable information), the function may need to be handled in a centralized manner separate from other functions. For example, the CME owning and operating the MA function cannot own and operate any other function, and thus it would need to be considered as a centralized function, isolated from the other CMEs. Other functions may or may not be isolated in this way. PKI policies developed by the SCMS Manager will determine the necessary separation of functions based on the final PKI design and root structure. This chapter's later sections explore potential SCMS structural configurations and separation of components in greater detail. Refer to Chapter 4 for a description of the types of PKI policies necessary to ensure the functionality and security of the SCMS.

2.3.1 SCMS Function Descriptions

Table 2 describes the components of the overall system structure as identified in Figure 1 and hence contains all components for a full-featured full deployment. Some components (e.g., the ICA) may not be necessary for initial deployment.

Table 2: SCMS Functions[10]

Function Name	Activities
MA	The MA performs multiple functions to manage risk in the SCMS. It receives misbehavior reports from EEs, investigates potential misbehavior, and blacklists or revokes other components in the system. The MA sends out an updated CRL to the devices in the field. This is an intrinsically-central function except for the CRL Generator (CRLG), which is one of the following subcomponents of the MA: • Global Detection (GD): This entity collects the misbehavior reports and decides on revocation of certificates • CRLG: This entity compiles and signs the CRL which contains linkage information that all receiving devices can use to identify the non-trustworthy device
Root CA	The root CA provides system wide trust through certificates issued to all CMEs. It represents the basis of trust for the system. This is not an intrinsically-central function. The system design supports multiple root CAs via the elector mechanism, but there could be as few as one (such as during initial deployment). The root CA certificate is different from all other types of certificates in several ways: 1. It is the end of the trust chain, i.e. verification of any certificate in the system ends at verifying this certificate. 2. The signature on the root CA certificate does not have any cryptographic value, is there only for namesake, as the signature is by the root CA itself, and therefore the trust in a root CA certificate is established through a process external to the SCMS certificate chain. 3. Usually the root CA certificate has a very long lifetime, as changing a root CA certificate requires updating all EEs in the system. Depending on the system and root structure, this is time consuming, difficult to ensure, and likely very expensive. 4. Assuming the Elector model is employed as the trust anchor management method, only a quorum of Electors can issue root management messages and add them to a CRL to revoke a root CA certificate. Refer to Section 3.2 for additional information on trust anchor management. The root CA certificate does not have an encryption key, as the root CA is mostly offline and does not accept any incoming messages, whether encrypted or not. The root CA certificate needs to be made available to everyone in the system. Also, for the reason explained in (2) above, integrity of root CA certificate must be ensured by other means (other than the cryptography used in generating the certificate itself), such as tamper-proof hardware. For the same reason, provisioning and/or update of root CA certificate is done through out-of-band means. Root CA certificates can be revoked, and new root CA certificates can be added by using the elector model. Trust anchor management methods other than the elector model could also be used to add and revoke root CAs

[10] Reference(s): CAMP SCMS PoC Implementation EE Requirements and Specifications Supporting SCMS Software Release 1.2.1; CAMP Technical Design of the Proof-of-Concept Security Credential Management System for V2X Communications

Function Name	Activities
Elector	The Electors are responsible for managing the Certificate Trust List (CTL). The CTL is a list of root CA certificates that are to be trusted by actors within the system. There are multiple electors to avoid the single point of failure risks that come from single root certificates. The electors manage the CTL by signing trust management messages that instruct the receiver to carry out an action (Add or Remove trust) on a certificate (elector or root CA). The recipient does not act on a trust management message until it has received the same instruction from a threshold number, or quorum, of electors. System parameters govern how many electors there are and the value of the quorum.
Pseudonym Certificate Authority (i.e., End Entity CA)	The Pseudonym Certificate Authority (PCA) is an intrinsically non-central component of the SCMS. It issues pseudonym, identification, and application certificates for end entities (EEs). There may be multiple PCAs in the SCMS. Individual PCAs may for example be limited to a particular geographic region or to a particular manufacturer or type of device. Each PCA is associated with a single RA and a pair of LAs to perform its core functions. The PCA responds to requests from the MA to investigate potential misbehavior. The National SCMS will also provide enrollment and application certificates to roadside units (RSUs) and other non-vehicular end entities. Application certificates are required for end entities to digitally sign messages, such as Traveler Information Message (TIM), Signal Phase and Timing (SPAT), and MAP messages. These certificates are distinct from the pseudonym certificates issued to vehicles because privacy is not a requirement for roadside units as they are typically owned by a public agency or toll authority.
RA	The RA is an entity authorized to validate, process and forward certificate requests. The RA receives and responds to requests for certificates from EEs via the Location Obscurer Proxy (LOP) (which masks the source IP address and route of the end entity (EE) from the RA). The RA will only accept requests from EEs that have enrollment certificates from ECAs that are authorized to use the RA. The RA can initiate certificate requests to a PCA to generate certificates for a requesting EE. Each PCA is also associated with a pair of LAs (LA1 and LA2) (that generate pre-linkage values for pseudonym certificates), and to the MA. • The RA initiates requests to both LAs to obtain pre-linkage values. • The RA must respond to requests from the central MA to add EEs to its internal blacklist and to support misbehavior investigation. The SCMS may include multiple active RAs at any given time. Although multiple RAs may exist, a given device may access only one RA (as seen from the device). However, the RA that the device accesses may employ load balancing (i.e., the deploying of multiple instances of its internal parts to improve computational performance). Load balancing will happen behind the sole interface that the device interacts with and thus be invisible to the device. All communications between devices and the RA is protected via Transportation Layer Security (TLS). The device authenticates the RA by using the RA's TLS certificate (X.509), and the optional use of the Online Certificate Status Protocol (OCSP). The device authenticates to the RA by using its IEEE 1609.2 enrollment certificate, which is

Function Name	Activities
	validated at the application layer. This is a supplement of the one-way TLS authentication, to provide two-way authentication with a TLS/1609.2 hybrid scheme.
Intermediate Certificate Authority (ICA)	This entity is a CA, the certificate of which was issued by a different root CA or ICA. Its value is that it shields the root CA from traffic and attacks. It may also allow for greater granularity in permission granting: for example, ICAs (and the CAs below them) may be limited to a particular geographic region or to a particular manufacturer or type of device, to make auditing simpler. ICAs may not be needed for initial deployment. The ICA authorizes all other non-central components including ECAs, PCAs, RAs, LAs, or additional ICAs. Similar to a root CA, an ICA is intended to be an offline component meaning that it should be configured with no direct network access or address. A local ICA Manager operates the ICA manually. The specific details of how the operator presents messages to the ICA is implementation specific and subject to review by a certification procedure approved by the SCMS Manager. The ICA is not an intrinsically-central function.[11]
Enrollment Certificate Authority (ECA)	The Enrollment Certificate Authority (ECA) is an SCMS back-end component that signs and issues enrollment certificates for EE devices. The ECA receives and responds to requests from one or more device configuration managers (DCMs). This is not an intrinsically-central function. Individual ECAs may for example be limited to a particular geographic region or to a particular manufacturer or type of device. The process for obtaining an enrollment certificate is developed in such a way that no single organization has sufficient information to re-identify a device. It will take the cooperation of two entities, e.g., in response to a court order, to re-identify a device.
Location Obscurer Proxy (LOP)	The Location Obscurer Proxy (LOP) obscures the locations of requesting EEs (e.g., OBEs requesting certificates) from SCMS functions such as the RA. This is intended to mitigate the possibility that the EE's location and/or route could be determined from requests made to the RA. In the simplest sense, one might think of this device as a router performing network address translation. The LOP is not an intrinsically-central function.
Linkage Authority (LA)	The LA generates linkage values for a given EE based on a request from the RA. The certificates for a given device make use of linkage values (LVs) from two LAs, referred to as LA1 and LA2. The splitting is done to make tracking difficult. When necessary, the linkage values are used by the MA to determine if multiple misbehavior reports, potentially associated with multiple pseudonym certificates are actually from the same EE, and allows the revocation of all of the EE pseudonym certs via the CRL. This is not an intrinsically-central function.

[11] The use of an ICA does have implications on the message size and the validation of the message. The message size is increased as the PCA that signs the certificates is not directly signed off by the root CA but by the ICA which is in turn signed off by the root CA. For the same reason, the computational load on verifying the message is increased. However, there are simplifications to reduce that impact. At a certain point in time, all vehicles will have verified at least one message and thus can trust the ICA. From that point, its certificate does not have to be attached to every message. The same set of circumstances eliminates the need to continue verifying once one reaches a CA with previously verified trustworthiness.

Function Name	Activities
CRL Store	The CRL Store is a repository that contains the most up to date certificate revocation lists generated by the MA. The CRL Store is accessible by all CV equipment and SCMS entities so that they may obtain the most up to date certificate revocation information. This function is not designated as a central or non-central function. This is a simple pass-through function since the CRLs are signed by the CRLG.
CRL Broadcast	The CRL Broadcast is the entity that makes the current CRL available on a broadcast basis, e.g., these may be RSEs or the satellite radio system. This function is not designated as a central or non-central function. This is a simple pass-through function since the CRLs are signed by the CRLG.

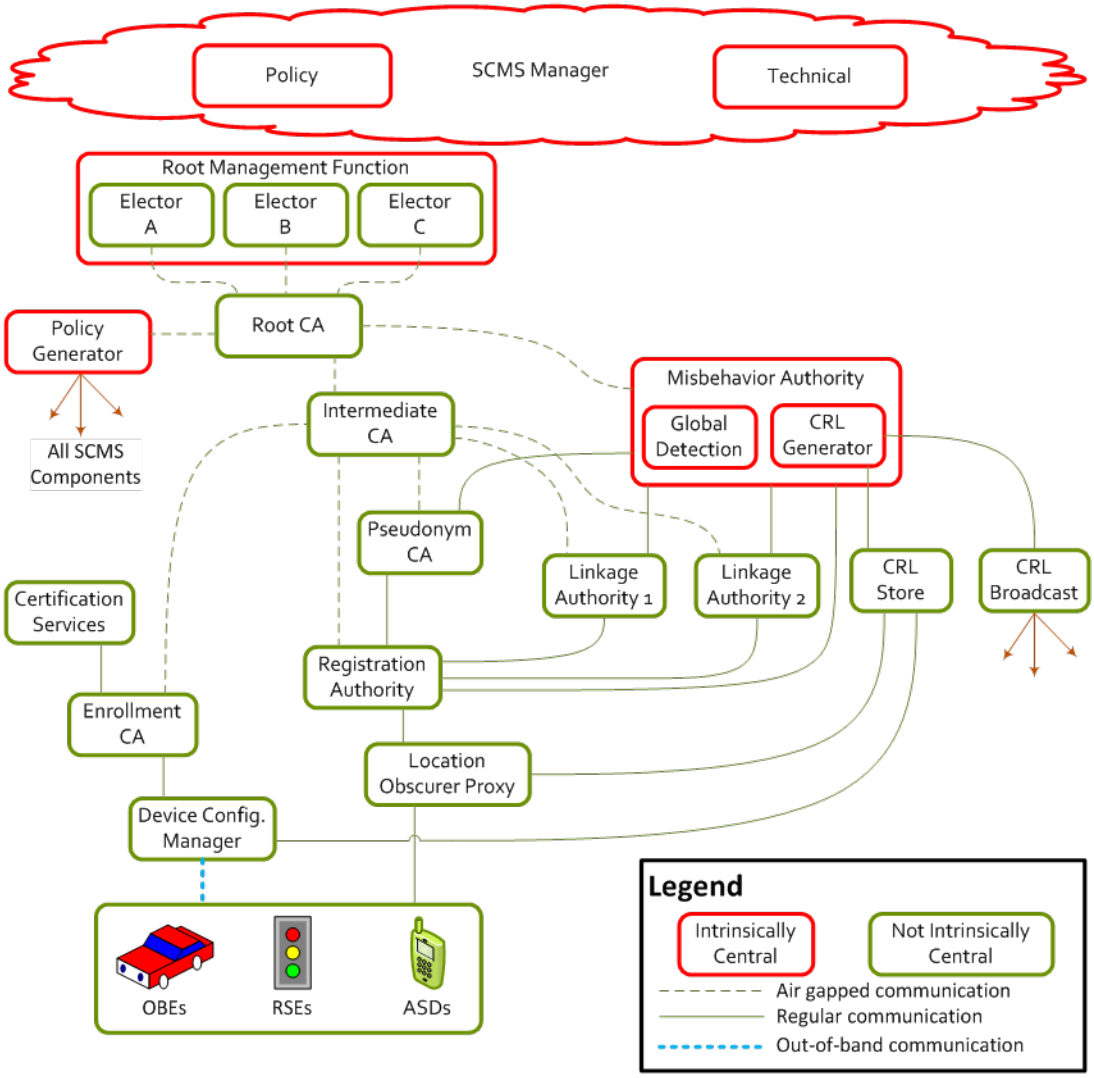

Figure 2: Overall SCMS Architecture
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

All communications between components in the SCMS are assumed to be authenticated and encrypted. This does not require all communications to be public-key-encrypted and signed. Instead, the assumption is that components in the SCMS will, in general, periodically establish and re-establish a symmetric session key which they use for encrypting and authenticating their communications. An exception to this "in general" assumption is communications that require non-repudiation; these communications will be signed rather than authenticated with symmetric cryptography. The overall assumption is that the symmetric key negotiation will be carried out with public keys and certificates belonging to the relevant SCMS components. Therefore, we note that all SCMS components will, in general, need a certificate, even if this certificate is not used by the device. It makes sense for these certificates to be issued by a root or an ICA. Typical mechanisms for this task are TLS and Virtual Private Network (VPN).

2.3.2 Reasons for Centralizing Select SCMS Components[12]

This section reviews advantages and disadvantages of selected components and, as applicable, also provides reasons as to why a component is considered intrinsically-central.

Reasons for Intrinsically-central Components:

- SCMS Manager: Sets policies and rules for the system
  - Single authority for establishing policies and rules ensures consistency
- MA: Detects and possibly revokes misbehaving devices
  - Both activities benefit from a central clearing house function
  - All operators depend on the MA to make the cooperative system work
  - The MA would use data from other entities, e.g., LAs, RAs and PCAs which might be run by different OEMs. Answering requests to an intrinsically-central and thus OEM-independent authority provides better overview and is more likely to be accepted by the OEMs.
  - Misbehavior reports come from devices of different manufacturers. An intrinsically-central component eliminates the need to send reports to different entities as well as the need for coordination between them.
  - A cross-jurisdictional MA (as opposed to, for example, having a series of national MAs) would better cover all vehicles because vehicles can drive across borders. A common definition of misbehavior across jurisdictions is assumed, i.e., that misbehavior is a technical rather than a policy issue.

Reasons for and Against Centralized Components:

Reasons for centralizing an element are identified by a plus sign (+) while reasons against are identified by a minus sign (-).

- Root CA
  - If there is a single root CA, there is a greater potential for an organizational error which could result in a potentially catastrophic compromise of that root CA potentially; having multiple root CAs avoids that single point of failure (-). However, whether there is a single root or multiple, all must conform to the centrally managed policy which specifies security
  - If there is a single root CA, the business relationships relevant to access to that root CA may be hard to run in an impartial way. Incumbents (either product suppliers or credential suppliers) who have a say in the operation of that root CA may use it to create barriers to entry for new competitors. (-)

[12] Reference: CAMP Technical Design of the PoC SCMS for V2X Communications

  - If there is a single root CA, it may be easier to create, communicate, and enforce certification policies such as choice of algorithm, physical security requirements, lifecycle management for devices, etc., as these policies can be developed once, centrally, rather than potentially being developed by multiple different roots (+). It is the SCMS Manager's responsibility to set and enforce the certification policy
  - Easier interoperability and less overhead (+)
  - In practice, it is not possible to avoid changing the root CA. The certificate may expire, the operator may be compromised, the business relationship may change so that a different service provider provides the root CA service, or the current root CA certificate may need to be replaced for any one of many other reasons. Since the system needs to be architected to support unplanned change of the root CA, and root CA rollover with overlap, it requires very little additional functionality to have multiple root CAs valid at the same time rather than only when the system is in transitional states. As such, architecting for multiple root CAs makes sense even if there is an expectation that the actual number of root CAs will be relatively small.

- LOP
  - Centralizing this component gives equal treatment to all devices (+)
  - Equal treatment at this point seems crucial for privacy, although this potentially could be accomplished by policies determined (and enforced) by the SCMS Manager (+)
- PCA
  - A multitude of PCAs results in potential tracking issues as the issuing PCA can be identified from the BSM's certificate (+)
  - In case of a compromise, the certificates issued from the compromised PCA need to be renewed. It is beneficial if that affects only a subset of the devices. (-)
  - A non-central PCA gives operators the option to opt-in on running one (-).

2.4 Roles and Functions of Entities that Interact with the SCMS

As illustrated above in Figure 1, the SCMS ecosystem includes entities and functions that are external to the SCMS itself. The CV device production, certification, and lifecycle management activities of these entities will be governed by policies and regulations that originate with the SCMS Manager, but they are owned and operated independently. As part of the SCMS ecosystem they will interact with each other, and with elements of the SCMS throughout the CV device lifecycle to assure that CV devices (e.g., on-board units [OBUs], ASDs, and RSUs) operate within the policies and requirements of the overall system, and specifically within the policies and requirements associated with the security and integrity of the CV system. Table 3 briefly describes these entities.

Table 3: Roles and Functions of Entities Outside the SCMS

Function Name	Activities
Device Configuration Managers (DCM)	Device Configuration Managers (DCMs) are the entities responsible for provisioning CV equipment (e.g., OBUs, RSUs, and ASDs) so that it can successfully interact with the SCMS and obtain the security credentials appropriate to its operation. While the DCM is external to the majority of the

Function Name	Activities
	SCMS, it plays a critical part of the enrollment certificate mechanism. DCMs must be subject to the same audit and oversight as other parts. The DCM is used during bootstrap to provide essential information to a bootstrapped device, and to relay information between a device and the ECA. DCMs will coordinate initial trust distribution with CV equipment so that it may then request, and successfully receive certificates from the Registration Authority (RA). The communication link between a bootstrapped device and DCM is out-of-band, e.g., a non-cryptographically protected communication in a secure environment. The SCMS Manager will need to establish (and enforce) the minimum level of security required for such out-of-band communications in order to maintain the integrity of the overall system. The DCM is not an intrinsically-central function; for example, different out-of-band communications methods could be used by different OEMs. Because these entities are not necessarily constrained to any specific structure or scope of operations, a variety of different industry entities may perform the role of the DCM. DCMs may include, for example, equipment manufacturers (i.e., device manufacturers, vehicle manufacturers, systems integrators), application developers (to the extent that applications running on a CV device are developed and installed separately from the device itself), retailers, dealers and installers (to the extent that the necessary configuration activities are implemented at the point of sale or point of delivery), and maintenance and repair facilities (to the extent that CV equipment must be re-configured following repair or upgrade operations).
Certification Services	Certification services are responsible for evaluating devices to assure that they perform to specified operational requirements, and that they conform to the security policies specified by the SCMS Manager. The SCMS Manager is responsible in this example for accrediting certification labs. This is not an intrinsically-central function. The certification lab can exist in two variations. In one, the OEM performs self-certification and the certification lab acts as a proxy between the OEM's internal lab and the SCMS. A second variation would be to use a test lab which performs intensive tests for a given device type. It is generally assumed that such evaluations and tests will be performed on a "type" basis rather than on every manufactured device. The Certification Service will then provide verifiable documentation that communicates to the Enrollment CA that units of that particular type are eligible for enrollment certificates. Note: since there is no way to assure that any individual device cannot be tampered with, the certification process and documentation will need to provide some means for verifying that a device presented to the ECA is in fact a bona fide example of the certified "type" of device[13].

[13] This verification step is not defined in the current SCMS design, but presumably it would take the form of some form of signature over the state of the device itself. For example, a signature over a hash of the binary code base (or suitable

Function Name	Activities
Communication Carriers	Communications carriers are external to the SCMS, and may or may not operate under policies and regulations associated with the SCMS. For example, cellular carrier may provide an IP link between a device and the SCMS, but the packets carried in that link would be secured independently from the carrier's operations.
Users	Users are vehicle owners and operators, and thus are responsible for the operation condition and maintenance of the CV equipment (i.e., OBUs or ASDs) associated with their vehicles. Users are expected to have very limited, if any, interaction with the SCMS unless there is a problem with their vehicle. For example, in cases where a vehicle is misbehaving (either because it has been tempered with, or because it is malfunctioning, the vehicle owner would be compelled, in some way. To have the vehicle examined and/or repaired by a service technician.
Local Agencies	Local agencies are responsible for public vehicles using CV equipment, and for roadside equipment under their jurisdiction. Local Agencies are expected to have somewhat limited, interaction with the SCMS, for example, they may need to specify the permissions for certain vehicles operating in their fleet, and will be responsible for approving the applications and permissions associated with roadside equipment, which will then be reflected in the certificates associated with that equipment.
Misbehavior Observers	The misbehavior system has not been defined sufficiently to determine exactly what operations it will entail. However, it is reasonable to assume that any form of misbehavior will need to be observed and reported by some entity on the roadway. This could be other CV equipment or roadside equipment operating in a special role to detect and report observed misbehaving CV equipment.

2.5 The Boundary Between End Entities and the SCMS

End entities (EE) are devices that will connect to the SCMS and receive certificates and the other files associated with certificates. These devices include OBUs, RSUs, and back office systems, such as those found in a traffic management centers. It is important to understand the device itself is not part of the SCMS. However, the SCMS Manager will impose requirements upon devices before they can be enrolled and interact with the SCMS to receive pseudonym certificates and other information necessary for a functional and trusted CV communications system.

There are many functions that the EE must perform before being able to communicate with the SCMS and receive the operational certificates necessary for CV communications. Based on the current design, these functions include:

portion thereof) of the certified device type, or a blockchain system wherein the history of the device from its manufacture to its initial enrollment is encrypted in to a chain that only the RA can decrypt. This way the RA can prove to itself that this instance of the device has not been altered from the type that was initially certified by the certification service.

- Generating public-private key pairs. The public key is used by the SCMS to generate certificates for the EE, including enrollment certificates and pseudonym certificates.
- Creating the various requests for the SCMS. Examples include the enrollment certificate request, pseudonym certificate request, and application certificate request. These requests are all formatted in Abstract Syntax Notation.One (ASN.1) format and can be found at the SCMS EE Requirements Wiki page.
- Creating HTTPS connections with the SCMS. All communications with the SCMS utilize HTTPS TLS connections, which will include the ASN.1 formatted requests in the payload.[14]
- Verifying the chain of trust.

It is the responsibility of the vendors and CV device deployers to ensure that these functions are integrated into their devices and functional prior to connecting to the SCMS or device certification.

Based on the current design, EEs must encrypt misbehavior reports to be sent to the MA. Therefore, all EEs will need the current MA certificate, which they obtain during enrollment from the DCM or during operation from their assigned RA.[15]

2.6 Interdependencies Among SCMS Components and CME Configuration Considerations

This section describes the interdependencies among SCMS components or functions and how these functions should be kept separate or may be combined with other functions within one CME. This analysis starts to provide more details on how SCMS components could be potentially owned and/or operated by various entities (e.g., the U.S. Government, vehicle manufacturers, PKI companies). Although, more analysis will need to be conducted during actual PKI policy development.

2.6.1 Separation of Specific SCMS Components or Functions[16]

This section describes the reasons why certain SCMS components (e.g., RA, PCA) allow an insider to track a vehicle, if run by the same organization. Attention is first directed to non-central components. Each combination of components listed below should not be contained within a single CME, where a single entity owns, operates, and manages the activities of that component.

- PCA and RA: The PCA and RA partition the knowledge regarding the enrollment certificate as to which requested a certificate (RA) and the content of the requested certificate (PCA). Separating these components avoids linking the enrollment certificate and the short-term certificate.
- PCA and LA1, PCA and LA2: If an organization ran the PCA and either the LA1, or LA2, it could link certificates to each other, i.e., determine whether two certificates belong to the same enrollment certificate or not. To show how this is possible, assume that an organization runs both the PCA and LA1. Two certificates are presented to the PCA. For both, the PCA looks up the unencrypted LVs and finds plv1 by plv1 = lv XOR plv2. It looks up Enc(plv1) for each certificate and presents the result to LA1. LA1

[14] Reference(s): SCMS PoC Government Management ConOps
[15] Reference(s): SCMS PoC Implementation EE Requirements and Specifications Supporting SCMS Software Release 1.2.1
[16] Component combination analysis was sourced from the CAMP Technical Design of the PoC SCMS for V2X Communications.

checks whether they belong to the same hashing chain and by this can determine whether they ultimately originate from the same device. This is done using a database lookup. Without the PCA, LA1 could not infer plv1 from lv. It is important to understand that both the enrollment certificate used for the certificate batch request as well as the device cannot be identified by this approach.

- LA1 and LA2: If an organization ran both LA1 and LA2 together, the organization could run these LAs to track vehicles by building up the forward-hashing chains of LA1 and LA2 and consecutively XORing them. Then the organization could build lists of consecutive certificate IDs of a given device.
- LOP and RA: The RA knows about the certificate batch request and the LOP can in general infer the location from which the certificate batch request came. (Note: The request is assumed to come over an IP network and that the location of the requester can be determined by means of the requester's IP address. The Internet Service Provider (ISP) can at least be identified and generally, there are databases, which can be utilized to get further details on the location.)

It is desirable that every central element is run by an independent organization. However, this may be impractical. To address such circumstances, it may be possible for a single operational entity to run more than one SCMS component provided that proper policies are established to ensure the level of organizational separation required to maintain security. There are constraints on which central / non-central elements may not be run by the same organization. The SCMS Manager will need to include these types of policies as part of the overall PKI policies. Chapter 4 describes the necessary PKI policies for the SCMS.

- LOP and MA: Some safeguards to prevent the MA from trivially being able to determine the location from where reports were filed (e.g., by analyzing the source address from the IP header) would help preserve privacy for vehicles reporting misbehavior. The latter functionality is obtained by the network address translator feature of the LOP.
- MA (plus CRL broadcast and CRL store): If combined with the RA, LA, or PCA, the MA could circumvent protocols, which would be in place when the MA requests information on a suspect certificate. These protocols are supposed to ensure that the MA shows proof of following a case that requires the requested information. The MA could use both LA and PCA to obtain linkage information on a case it is currently building. As an example, assume that the MA was combined with the PCA. The MA could acquire linkage information from the LA through official means. If the case was not yet strong enough to obtain information from the PCA through official means, the MA could use its ties to the PCA to gain this information thereby circumventing privacy-by-design. In another example, assume the MA is combined with the RA. The MA could build a case and gather all the required linkage information from RA and PCA. It could then use these ties between the MA and the RA to learn the enrollment certificate of the corresponding device as the RA can map the (partial) linkage information from the LAs to the enrollment certificate. Note: The enrollment certificate is information, which the MA does not need to accomplish its duty.
- Root CA: Running the root CA with a very trustworthy and well-monitored organization would help secure the root. However, there is no function or security need for organizational separation between the root and issuing or intermediate CAs. All the CAs in a specific trust chain could be operated by a single organization.
- SCMS Manager: As this component issues policies and rules defining the behavior of all the components of the SCMS, it would be advisable to run it completely independently, i.e., by an organization running no other components. Although running operations presents a potential conflict of interest with oversight and governance, there is no formal security requirement reason the SCMS Manager cannot run any of the components if the required organizational separation is enforced - e.g., the SCMS Manager could run electors, CAs, and MA but not the RAs or LAs.

2.6.2 Functional Relationships and Groupings of Certificate Authorities[17]

The architecture described in this document allows for multiple instances of many SCMS components, including the RA, LAs, PCA, and ECA. This subsection looks at relationships between those components. For example, the existence of lots of RAs and PCAs does not mean that each RA talks to each PCA. Put differently, even though there are lots of ECAs and lots of RAs, there might be advantages to having each given RA talk to a single ECA.

This section further discusses possible models for organization of these components and pairwise relationships. For pairwise relationships, there are four possible relationship types, those being all possible combinations of cases where component 1 or 2 has a relationship with only one, or more than one, instance of the other component. In general, if one component has a relationship with only one instance of the other component (for example, one RA -> one PCA), the advantage is reduced system complexity and the disadvantage is reduced flexibility; if one component has a relationship with multiple instances of the other components, the advantage is increased flexibility, but at the cost of increased system complexity. The following information only explicitly addresses cases where there are other advantages or disadvantages to be noted.

- RA/ECA:
  - ECA -> One RA: An ECA produces certificates that are only trusted by a single RA, though the RA may trust multiple ECAs
  - ECA -> Many RAs: An ECA produces lists of enrollment certificates that are trusted by RAs. Note that the lists must be disjoint, i.e., each enrollment certificate is trusted by exactly one RA only.
- PCA/RA:
  - One RA <-> One PCA: In this case, the signer ID in the pseudonym certificate reveals the RA used, which may in turn reveal information such as the vehicle OEM
  - Many RAs <-> One PCA: This paired relationship is better for privacy because PCA identity does not reveal which RA was used
  - One/Many RAs <-> Many PCAs: This case is also better for privacy as long as the PCA used to fulfill a given request is randomly chosen (i.e., the relationship is truly one RA to many PCAs rather than a series of different one to one relationships)
- RA/LA:
  - One/Many RAs <-> One LA: If an LA is used by multiple RAs, the LA must ensure that it does not give duplicate values to more than one RA
  - One/Many RAs <-> Many non-exclusive pairs of LAs: There is no reason why pairs of LAs should be fixed, e.g., one RA could use LA1 and LA2 and another RA could use LA1 and LA3
- PCA/LA:
  - If an LA is used by multiple PCAs, the LA must ensure that it does not give duplicate values to two PCAs

[17] Analysis was sourced from the CAMP Technical Design of the PoC SCMS for V2X Communications.

2.6.3 Potential High-Level Deployment Models[18]

The architecture shown in Figure 2 shows the most flexible, full-featured possible system. An initial deployment of the system may be made simpler (in terms of number of deployed components and complexity of relationships) by omitting some of the flexibility of the full model. The two deployment models, initial and full, represented within this section are based on existing research by CAMP and not firmly set. These deployment models are indicative of what the team believes the actual structure of SCMS might look like at the two deployment stages.

2.6.3.1 Initial Deployment Model

In the initial deployment model, the load on the SCMS is expected to be moderate, so some of the components that are not intrinsically-central are made central-by-choice (i.e., they have only one distinct instance), such as root CA, ICA, ECA, Certification Services, etc. The central-by-choice components are listed below accompanied with brief reasoning behind the choice.

- Root CA: A central-by-choice root CA facilitates easier interoperability and has less overhead
- ICA: A central-by-choice ICA facilities easier interoperability and has less overhead.
- LOP: A central component guarantees equal treatment of all devices, which seems crucial at this point for privacy, although this could potentially be accomplished by policies determined (and enforced) by the SCMS Manager
- CRLG / CRL Store / CRL Broadcast: A central CRLG seems sufficient as at most only one CRL (with one approved format) will be released to the devices.
- ECA / Certification Services: These components have minimum interaction with the device (once per device for ECA and once per device-type for certification services), so centralizing these components seems sufficient for initial deployment. The ECA could even potentially offer certification lab services. Combining these two components under one roof affords the ECA better control over the quality of the certification services that it relies on. However, this depends on the ubiquity of end entities available during initial deployment. A single, central certification service many not be sufficient to support all certification needs. For example, vehicle manufacturers are already rolling out vehicles capable of V2X communications and more manufacturers will likely follow soon.

[18] Analysis was sourced from the CAMP Technical Design of the PoC SCMS for V2X Communications.

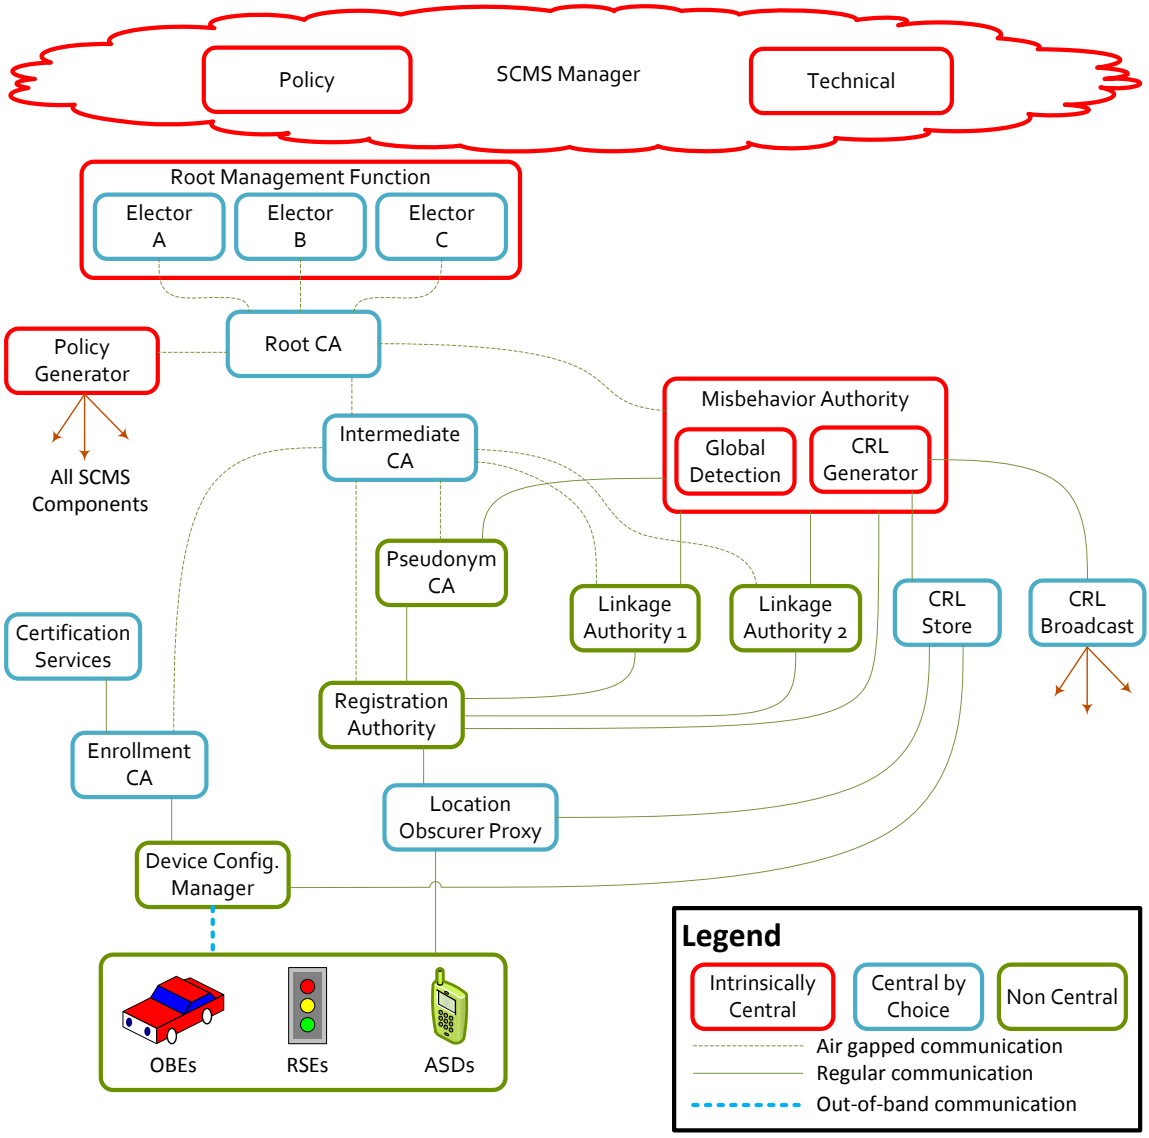

Figure 3: Potential Initial Deployment Architecture for a National SCMS
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

2.6.3.2 Full Deployment Model

In the full deployment model, all components that are not intrinsically-central except the CRLG are made non-central (i.e., have multiple distinct instances) to have maximum flexibility and to support a much larger load compared to the initial deployment model. The CRLG remains central-by-choice as it does perhaps the least amount of work (collecting and signing a list of revoked devices) among all the SCMS components. Along with allowing multiple root CAs, the full deployment model makes use of the electors. Again, this is only a potential model and additional deployment models will be explored during the course of this effort.

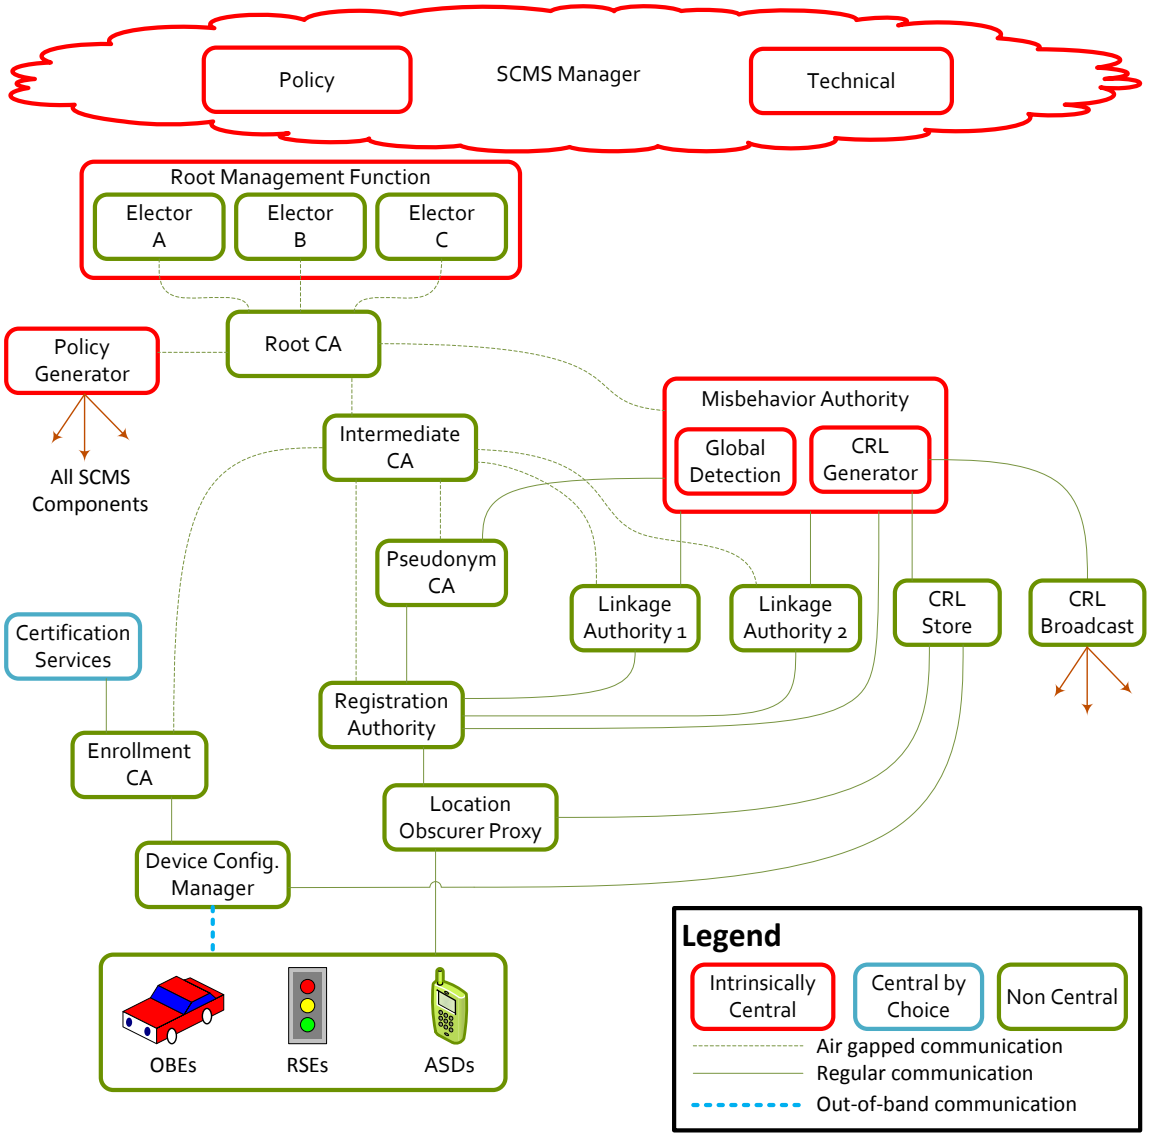

Figure 4: Potential Initial Deployment Architecture for a National SCMS
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

2.7 National SCMS (Internal) Security Considerations

One of the SCMS Manager's primary focus areas will be the effective governance and oversight of all SCMS assets and operations as they fulfill their mission, described in Section 2.3 of enabling secure communications between the community of ITS end entities. As with any critical IT system, securing the National SCMS assets and operations is a paramount concern, and the SCMS Manager and its functions will be key in mitigating vulnerabilities and threats through effective governance and operational oversight; and by providing appropriate internal security requirements and policies. Key areas of considerations include:

1. **Secure Communications** – Securing the information transmitted between National SCMS components through policies: assuring that the information was sent by the originator, and at the time indicated; assuring that the originator is authorized to send the information; that the information has not been accessible and/or altered by external entities while in transit.
2. **Cybersecurity** – Mitigating potential compromise of National SCMS assets through cyber-attacks with policies regarding network access controls (e.g., firewalls, intrusion detection, demilitarized zones), platform and application hardening (e.g., denial of service / distributed denial of service (DOS/DDOS) prevention, security modules, failure state processing, Common Criteria evaluations, FIPS 140 compliance), penetration testing requirements, and restrictions.
3. **Access Control** – Protecting National SCMS assets from unintended or inappropriate access to and usage of National SCMS assets through policies regarding identification, authentication and authorization (e.g., smart cards, biometrics, emanation controls). This includes system access and data access controls,
4. **Operations Security** – Ensuring that personnel do not inadvertently or maliciously cause harm to National SCMS assets and have the proper training necessary to properly operate these assets and to respond to security related incidents. This also includes that personnel operating the National SCMS do not intentionally or unintentionally subvert or compromise the integrity of the system. Auditing and monitoring of system usage.
5. **Operations Continuity –** Systems performance monitoring and control, proactive systems maintenance, failover and restoration; redundancy and sparing policies, reliable backup and recovery procedures, staffing and escalation procedures.
6. **Physical Security** – Protecting National SCMS assets (e.g., facilities, servers, datastores) from physical attacks.
7. **Operations Continuity Planning** – Contingency planning and disaster recovery policies.

Chapter 3: Factors Impacting Governance, Ownership, and/or Deployment Strategies for a National SCMS

There are many factors that must be considered when selecting an industry ownership and governance model, and planning for the subsequent deployment of that model. Chapter 3 discusses public interest objectives and evaluation criteria, and delves deeper into areas of interest identified during the CV pilots and development of the SCMS PoC. A follow-on task will focus on developing detailed potential ownership and governance models that address these factors and analyze trade-offs among the models.

3.1 Factors that Influence the Development and Deployment of Ownership and Governance Models

There are several factors that will influence the development and deployment of ownership and governance models. Throughout the early stages of the National SCMS Deployment Support project, the team has identified public interest objectives that must be addressed and fulfilled by the selected ownership and governance model. The team also identified evaluation criteria or considerations that the selected model will greatly influence and, at the very least, must be thoroughly discussed during model development. Many of these objectives and evaluation factors overlap or influence each other. Table 5 and Table 6 list these objectives and factors with brief descriptions. The following subsections analyze topics that the CV pilots and SCMS PoC participating entities have identified as areas of interest to ensure a functioning and secure National SCMS and how those may influence an ownership and governance model, and vice versa.

Within early discussion of potential ownership and governance models for this project, the team identified high-level models ranging from completely publicly owned, governed, and operated to completely private (see Table 4). The team will continue to build out these models and variations of these models throughout the course of this project.

Table 4: High-Level SCMS Manager and CME Deployment Models Based on Ownership and Initial Funding

Model A: Completely Public	Model B: Government-led P3	Model C: P3 Concession	Model D: Industry-led P3	Model E: Completely Private
• Standup new government office to serve as the SCMS Manager • Develops all policies with input from key stakeholders identified through this project • Stands up electors, root, and other SCMS functions • There must be separation of CMEs (per requirements specified in Chapter 2) so the government cannot operate all functions • May contract out operations but government maintains overall control • Initial funding for National SCMS standup comes from department budget • Sustainment funding through department budget. Need legislation for OBU fees or other funding mechanism	• Standup new government office or team to provide oversight • Team develops all initial policies with input from key stakeholders and potential CME owner/operators identified in this project • Team stands up electors, root, and other SCMS functions to then be auctioned off through RFP and run based on MOU from OST • Team develops new market place for additional CMEs to work through the SCMS Manager for validation to own/ operate • Initial funding for National SCMS standup comes from department budget • Sustainment funding is the responsibility of the new owner-operators	• Government team to serve as the facilitating agent and governor • Team develops initial policies with input from key stakeholders and potential CME owner/operators • SCMS Manager is run as a concession (Government oversees policies and operations, but concessionaire performs operations for a fee from CMEs and participants) • Government releases Cooperative Agreement RFP for implementation and operation (20% fed funded/ 80% performer funding split) • Awardee takes lead on standing up electors, root, and other SCMS functions under oversight by government • USDOT chairs the governance board to ensure public interest objectives are met • Government funding is to assist with deployment and operate governance/oversight office	• Government team to serve as the facilitating agent and oversight • Government to facilitate charter development, organization of initial consortium/ consortia, planning sessions • Team develops initial policies with input from key stakeholders and potential CME owner/operators identified in this project • Industry takes lead on standing up electors, root, and other SCMS functions • USDOT remains on the SCMS Manager governance board to ensure public interest objectives are met • Only government funding is to assist within initial facilitation	• Industry leaders form their own consortia facilitated by the deployment support project • Industry-led SCMS Manager develops all policies • Industry funds governance and PKI implementation • USDOT becomes a stakeholder and potential member (e.g., seat on an executive/governance board and/or advisory board) of the completely private SCMS ecosystem

Table 5: Public Interest Objectives

Public Interest Objective	Description	High-level Tradeoffs
Security	Security is dependent on technical design and policies, which must ensure security of the system and data regardless of the ownership and governance structure. USDOT will be challenged to provide the necessary oversight in a completely private model and a completely public model may not be appropriate to rapidly respond and evolve based on identified vulnerabilities, threats, and technology advances.	No matter the ownership and deployment model, the PKI policy (refer to Chapter 4 for more information) must detail the certificate policy to ensure security both within the National SCMS itself and across the National SCMS ecosystem and enforce this policy through audits and accredited device certification labs.
Privacy	Privacy is dependent on technical design and policies, which must ensure an appropriate level of vehicle and operator data privacy regardless of ownership and governance structure. Based upon SCMS Manager and CME ownership there may be increased privacy levels depending on government and private sector involvement. The government would likely need to focus involvement on maintaining security, privacy, and adequate stakeholder representation.	A completely public model has the potential for increased privacy levels because of the government focus on privacy. In the government-led P3, the transfer of select responsibilities to industry and increase in external CME owner/operators could increase risk if the policies do not maintain adequate separation of SCMS components and protection of data. In the more private models (industry-led P3 and completely private), there is likely decreased overall privacy and potential to move away from privacy policies as initially designed without USDOT or government representation.
Availability (i.e., interoperability, redundancy, flexibility)	Valid certificates issued by the SCMS must be available to end entities to ensure a functioning V2X communication system that provides safety benefits. The root structure and trust anchor management method will greatly impact system availability, interoperability, redundancy, and flexibility, as well as determine the specific information required within PKI policies (Refer to Chapter 4 for more information on necessary policies). Based on the technical design structure, the SCMS Manager will need to develop the appropriate, detailed policies to ensure that the system, no matter the root and trust anchor structure, is readily available to enable trust among end entities.	Section 3.2.1 analyzes tradeoffs in availability based on the root structure. When considering high-level ownership and governance models on the completely public to completely private scale, a public model will likely have less redundancy and flexibility than models with more private involvement and competition provided that the P3 and completely private models enforce policies for efficient trust anchor management. The team will complete additional analysis on the effects of various combinations of SCMS component ownership and operations on system availability

Public Interest Objective	Description	High-level Tradeoffs
Stakeholder Representation	In initial and current phases of CV pilots and SCMS design, stakeholder engagement and representation is an important enabler in the development of a SCMS PoC that considers public interest and goals. Stakeholder representation during the National SCMS implementation and deployment process, and in the SCMS Manager governance and operational oversight activities will help ensure transparency and trust in the system itself among the government, the private sector, and the general public.	There is likely to be increased stakeholder representation and transparency where government is a leader. There is likely to be less representation and transparency in models where private entities are the owners and operators. Stakeholders who want to be adequately represented usually need to "buy in" with membership fees.
Affordability	The technical design (e.g., initial single root with plan to introduce other roots), ownership (e.g., P3 non-profit SCMS Manager), and governance models will greatly impact affordability of the system. Deployment and implementation plans for the National SCMS must consider initial funding sources, sustainment funding sources, and how internal organizational and external industry governance affects efficiency.	A completely public model is likely the least efficient/ affordable. A government owned and operated model will have high overhead and lengthy processes. A government-led P3 would have initial low efficiency and high cost for gov't to facilitate initial consortium and root implementation with decreasing cost as industry takes the lead. In a P3 concession, the primary burden is on industry with initial modest support from government. A completely private model is likely the most efficient and affordable if it is a non-profit. However, costs may be less transparent, especially if there is no government representative on the board.
Performance	Performance can be viewed from an SCMS technical and functional perspective, as well as an organizational and governance perspective. The final SCMS design and PKI policies will determine the technical and functional performance of everyday National SCMS operations. Ownership and whether the SCMS ecosystem is based on profit, non-profit, or potentially a combination of features will influence organizational and governance performance within the industry.	A completely public model will likely have less than optimal organizational performance because of competing USDOT priorities and resources. A government-led P3 will likely have an initial period of low performance during standup. A P3 concession would generally have high levels of performance due to the competitive nature of concession. The industry-led P3 and completely private models will likely have increased internal organizational performance due to streamlining of operations to cut costs and competition.

Table 6: Evaluation Criteria

Evaluation Criteria	Description	High-level Tradeoffs
Ownership	Chapter 1 briefly discussed various ownership models. It is important to understand that these ownership models may evolve based on the needs of the system and the appropriate level of government oversight. There could also be different ownership models to various functions within the SCMS ecosystem. For example, the SCMS Manager could be a federal government owned and operated entity while select CMEs are owned and operated by private entities.	A completely public model would be owned by the federal government, and potentially by the USDOT, with a government-led P3 initially owned by the government with potential sale or transfer to industry. The government could maintain "ownership" and authorize vendors to operate. In a P3 concession, the National SCMS is owned by the government but funded and operated by private industry. In the industry-led P3 and completely private models, an industry consortium would likely own the SCMS Manager with various private entities owning and operating the SCMS components.
Funding	As mentioned in the Affordability objective, the National SCMS deployment and implementation plan will need to address initial stand-up funding and sustainment funding. Initial stand-up funding will be largely determined by ownership. For example, a completely private model may fund initial deployment through an implementation fund provided by consortium members. Sustainment funding could, and most likely will, be generated by similar methods no matter the ownership model (e.g., fee automatically included within the purchase of a new vehicle). However, the way in which the sustainment funding flows to the SCMS Manager and CME will depend on the root CA structure and governance model.	In the completely public and government-led P3 models, initially deployment funding would wholly come from the government. Initial deployment funding is a 20/80 cost share between the government and the concessionaire in the P3 concession. In the industry-led P3 and completely private models, industry would need to completely fund deployment. There are many options for sustainment funding including the automatic fee included with the vehicle price and annual fees such as an excise tax or included within vehicle registration. Other potential revenue sources include membership fees, accreditation and services fees, and auditing fees.
Policy Creation and Approval	The entity that takes the lead on initial National SCMS Manager and CME stand-up will likely lead the initial PKI policy development. The SCMS Manager will develop policies with the approval of a governance board. Chartering the SCMS Manager with initial policies already developed may help accelerate the stand-up of CMEs. The team recommends that these policies follow the structure outlined	For a completely pilot model, the government-led SCMS Manager develops policies. A governance board may not be necessary; however, it is recommended that policies are opened for public comment. In the industry-led P3, industry gradually takes the lead and updates policies with approval by an executive/governance board (USDOT has a seat). In the P3 concession, the government's role is to assure that

Evaluation Criteria	Description	High-level Tradeoffs
	in Request for Comments (RFC) 3647 which is the PKI industry standard. The personnel make-up and structure of the SCMS Manager and the governance board will depend on SCMS Manager and CME ownership.	concessionaire implements and enforces SCMS Manager policies. In the industry-led P3, the SCMS Manager develops policies facilitated by government with approval by an executive/governance board. The government has a seat on the board. This is the same for the completely private model except that the government could request a seat on the board, but not guaranteed.
Oversight and Auditing	The type of ownership will determine the type and level of National SCMS oversight. If there is specific legislation or regulation that provides authority to a National SCMS Manager is some way, or specifies use of a specific root for example, these actions would need to specify the entity providing oversight for the National SCMS Manager and larger SCMS ecosystem (e.g., Federal Communications Commission, NHTSA).	Congress would oversee a completely government owned and operated SCMS. A department or administration within the federal government would likely oversee any P3. A completely private entity would likely not have significant government oversight, other than the Department of Justice. For auditing, the PKI industry standard is to contract out a third party to conduct audits, potentially with intermediate government inspections based on the level of oversight.
Trust Anchor Management	The National SCMS must have an effective method to manage trust anchors no matter the technical design, ownership model, or governance model. The current default trust anchor management method is the elector concept, which is further described in Section 3.2.4. The SCMS Manager must develop policies and procedures for trust anchor management to ensure security within the selected root structure and technical design.	The trust anchor management method does not necessarily effect the high-level ownership and governance model. Any of the models could accommodate the current elector model as proposed. However, it is important to consider that the trust anchor management function is a core function and would ideally be separate from the SCMS Manager. The more important questions are how many electors are necessary without becoming cost prohibitive.
Legislation and Regulation	Depending on the ownership and governance model, the federal government may need to enact new legislation and/or regulation such as granting authority to new government entities and/or the SCMS Manager, or levy new taxes and fees.	The type and content necessary within legislation and regulation needs more analysis. However, unless there is consensus for a completely private model, some legislation and or regulation will be required to bestow the appropriate authority to stand up a National SCMS Manager which may or may not also provide oversight for the entire ecosystem. Regulation may also be necessary to create a technical mechanism to ensure

Evaluation Criteria	Description	High-level Tradeoffs
		early deployments are carefully managed and vetted until the system is mature (e.g., a single specific root is considered valid for initial deployment).
Competition	The ownership and governance model will greatly impact competition within the new SCMS ecosystem. Depending on the final goals and objectives of the National SCMS, the industry and government may not initially want competition to ensure that the nascent system is under tight oversight and control. The SCMS Manager and governance board could gradually introduce the ability for external entities to offer CME services as along as these entities conform to the National SCMS PKI policies and requirements.	In the completely public model and government-led P3, the only competition would be established through federal contracting practices to potentially operate SCMS functions over a specific period. The government would need to ensure the contract requirements are performance-based. However, this would limit flexibility and lock in vendors for a period. The P3 concession also locks in vendors, but the vendor is motivated to increase efficiency to increase their return on investment. The industry-led P3 and the completely private models will offer the most competition within the ecosystem to potentially increase performance and decrease costs. However, this will complicate governance, oversight, and auditing which will increase the workload for the National SCMS Manager and the aligned oversight entity, if one exists.
Overall Risk	Risk within the National SCMS ownership, governance, and operational models will take many forms. For example, there will be financial risk for the entities that stand up and own the SCMS Manager or CMEs. There is also operational risk – what is the impact of a specific governance model and certificate authority structure on the ability of the National SCMS to provide meet the public interest objectives?	The team needs to conduct considerable more analysis on potential models before determining the overall risk for each model. All risk falls on the government in the completely public model and gradually transfers to industry in the completely private model. For operational risk, the lowest overall risk model to ensure efficient operations while maintaining the necessary levels of security likely needs to include the government in some capacity, even if it is only in a minor oversight role. The completely public and completely private models are probably too risky for varying reasons which will be explored further in later project tasks.

3.2 Trust Anchor and Certificate Authority Management

The approach for trust anchor management and general management of certificate authorities is an ongoing discussion. There are tradeoffs in operating a single root PKI versus a multiple root PKI. Based on the PKI root approach, there are multiple ways to employ trust anchors for those roots. No matter which root structure is employed, there is a need for an effective trust anchor management solution to revoke and add roots as necessary. The current trust anchor management method favored for deployment is the elector model. Inevitably, a root will eventually retire and there must be processes and procedures for retiring the root and adding a new root to the system. The final design decisions will influence the overall cost and affordability of the SCMS, as well as impact performance and availability. When the system is implemented it will be difficult, if not impossible, to change the trust anchor management method without disrupting overall trust within the SCMS and V2X ecosystem.

3.2.1 Single vs. Multiple Root Certificate Authorities

The operational and public interest objectives of the SCMS system, along with the desired ownership and governance model will drive the decision whether to base the system on a single root or multiple root model. The design of the SCMS allows the use of more than one root. Both models have their advantages and disadvantages which vary based on the specific implementation scheme and accompanying PKI policies. In short, while a single root would be less expensive to stand up and operate, a root compromise could have far-reaching negative impacts. However, a single root structure could ease initial deployment and provide a regulatory lever to mandate use of that root. This could be a way to grant the necessary authority to the SCMS Manager to govern the industry and oversee operations.

A multiple root structure, while more expensive to stand up and manage, would provide additional redundancy and interoperability. A multiple root structure allows for more flexibility in expanding and decentralizing operations if necessary. However, the multiple root structure would require more effort and funding to effectively govern and maintain security and privacy. Looking at the options from a long-term perspective, a multiple root structure would likely increase competition and CME supply which would decrease operational costs. Although, governance and oversight costs would likely increase. As described in Section 2.6.3, the initial deployment could be done under a single root with the intent of expanding to multiple roots in the future which would reduce the short-term governance and oversight costs. The SCMS Manager would need to serve as the gatekeeper to add new root CAs to the system based on established policies and requirements.

A decision to operate with a single root limits future flexibility. If OBUs are designed with the assumption of a single root, it may not be possible to transition to a multiple root architecture in the future should it be determined that would be desirable.

Section 4.2.4 provides additional detail on the Certificate Policy implications of having a single or multiple root PKI model, along with examples of those policies and organizations that manage governance and operations. No matter the selected structure, EEs must be capable of operating under multiple roots to ensure interoperability.

While there is theoretically no limit to the number of root CAs in the system, the following aspects are relevant considerations for adding root CAs:

- At least one root CA is required.

- If a root CA is compromised, i.e., an attacker is able to issue a valid PCA, RA, or ICA certificate or to extract the corresponding private root CA certificate key, all devices are potentially affected. Therefore, each root CA certificate needs high security standards, defined by SCMS Manager, regardless of the root CA's size.[19]
- A large number of root CAs increases the risk of a security breach of a root CA. This risk is mitigated by applying high security standards for each root CA.
- Distributing the new root CA certificate may be expensive, it may take time for all devices in the system to receive that certificate, and therefore to be able to trust messages that chain back to that root CA. In contrast, if a new ICA is used, devices can trust messages from that ICA as soon as the ICA's certificate is received. This may be attached to a signed message from an EE device or distributed by some other means. Note that in the former case the new ICA's certificate would be included in the chain of trust (i.e., signed off by a trusted root CA).

In considering operational models and the number of eventual roots, stakeholders should understand that the number of root certificates presents several potential costs.

1. Each OBE must have all the root certificates (which are x bytes). More roots would add to the memory requirement.
2. Each root will have a CRL – hopefully small, but there will be memory cost associated with each.
3. A lot of roots may necessitate more issuing CAs (e.g., every brand has its own root and issuing CA).
4. Potential processing difference – all certs chain through some issuing CA to a root but root and issuing CA validity only need to be checked once each time CRLs are received – a lot of different chains would mean that there is a need to process each one and cache. This could be a processing versus memory discussion – process "all" the chains at start-up (or when new CRLs are obtained (costs memory to store) or process in "real time" as needed to validate a certificate (processing time/power).

3.2.2 Single vs. Multiple Pseudonym Certificate Authorities[20]

If there is a single PCA, there may be scaling issues. This makes the use of multiple PCAs attractive. However, if there are multiple PCAs, and each RA uses a single PCA (meaning that each vehicle that uses that RA is sent to the same PCA), each vehicle will be identifiable to the granularity of its PCA (because the PCA certificate is identified in the signer_id field in the certificate). This may significantly reduce privacy against an eavesdropper.

If there are multiple PCAs, and the RAs use multiple PCAs, but the RA routes all individual certificate requests from a given vehicle to a single PCA, that vehicle may be easier to trace by an eavesdropper. This can be exploited by an insider at the RA by routing only one set of certificate batch requests from a region to a particular PCA, meaning that there will be only one vehicle in that region with certificates from that PCA. It also increases the chance that an insider at the PCA can use information about certificate request time to track.

Therefore, if there are multiple PCAs, the best use of them from a privacy point of view is for RAs to use multiple PCAs, for each PCA to be used by multiple RAs, and for RAs to share individual certificate requests from a single vehicle between the different PCAs at random.

This is somewhat auditable by the vehicles, which may check (a) that their certificates come from a range of PCAs and (b) that averaged across time periods, they encounter the expected percentage of vehicles that use certificates issued by the same PCA in any given time period.

[19] The number of EE certificates that chain back to a CA is referred here as the CA size.
[20] Analysis was sourced from the CAMP Technical Design of the PoC SCMS for V2X Communications.

An alternative approach is that OEMs could run their own PCAs. This is under consideration by OEMs as the availability of the PCA has a direct impact on their customers. However, it affects privacy against eavesdroppers as the certificate of a vehicle will give away its OEM (or reveal that it comes from one of a group of smaller OEMs). It might be better to address the availability concern by reducing the requirements for availability on the PCA.

Since the PCA is not an intrinsically central component, an arbitrary number (a positive integer) of logical instances can be introduced in the system. If a single instance existed in the system, all vehicles are threaded uniformly throughout the system, e.g., the same authority would have signed all pseudonym certificates. Although this approach can be most beneficial for privacy, it induces several difficulties for the implementation and overall system performance.

First, opting for a central PCA creates a single point of failure in the system. Should the PCA be compromised or should it suffer some technical malfunction, all certificates would have to be revoked and reissued, assuming that the threat had been removed and the PCA is back online. While technical issues are prone to relatively quick fixes, a bad organization can inflict severe damage to users' privacy. Furthermore, servicing all end devices in the system could result in a major communications burden for the interface of the central PCA. Implementation and maintenance of an interface as such can be very costly.

Problems with reliability, flexibility, and component complexity can be reduced by introducing multiple PCAs in the system. This way the load can be distributed, and in the case that any PCA goes offline, its peers can ensure uninterrupted service. Moreover, whole trust would not have been put into a single party.

Unfortunately, the latter approach can cause privacy issues. Namely, in a system with multiple PCAs it is reasonable to assign an area of ordinance to each one of them. Now suppose that a user originating from the East Coast decides to drive to California. It is very likely that the number of vehicles with the certificates signed by the East Coast PCA will rapidly drop, proportionally to the length of the journey to the West. Therefore, the user could be tracked based on the signature of the PCA and not the content of the pseudonym certificate itself.

To prevent this sort of attack, the RA could introduce an additional dimension of randomness. Namely, whenever an RA requests a batch of pseudonym certificates it could randomly choose to which PCA it would submit the individual certificate request. Unfortunately, this approach would lead to an increased communications requirement in the LA-RA-PCA trio.

As an alternative, a hybrid scheme can be implemented. More specifically, segregation of the authority of the individual PCAs can be made based on the OEM brands, not geographic locations. While on the production line, all vehicles would be supplied with a three-year lifecycle of certificates, signed by the PCA of the OEM that produces them. Afterwards all vehicles in the system would send certificate top-off requests to a central PCA in the system. Pre-stored certificates would ensure the uninterrupted service in a three-year period, and consequently enough time to allow uniform distribution of load.

3.2.3 Certificate Authority Retirement[21]

Certificate lifetimes affect the security of PKIs. The longer a public/private key pair is in use, the greater the chances are that the keys can be compromised. As computing power increases and technologies improve

[21] Majority of this analysis was sourced from the CAMP Technical Design of the PoC SCMS for V2X Communications.

over time, cryptanalysis becomes a risk. For these reasons, excessively long-lived CA certificate lifetimes are undesirable. Figure 5 illustrates how the minimum lifetime of a typical CA certificate is calculated.

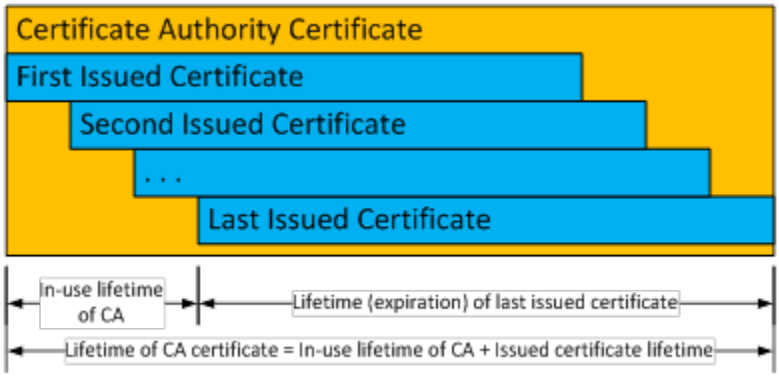

Figure 5: Calculating In-Use Lifetime of a CA Certificate
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

Some certificate authorities may issue certificates that are not valid until a significant time in the future. Examples of this within the SCMS are pseudonym certificates and rollover enrollment certificates. At the time of the writing of this report, the validity lag for these certificates can be up to 3 years. For example, a pseudonym certificate generated (issued) today may have a "Valid from" date that is up to 3 years from now. Figure 6 illustrates the impact of the validity lag on the lifetime of the issuing CA certificate.

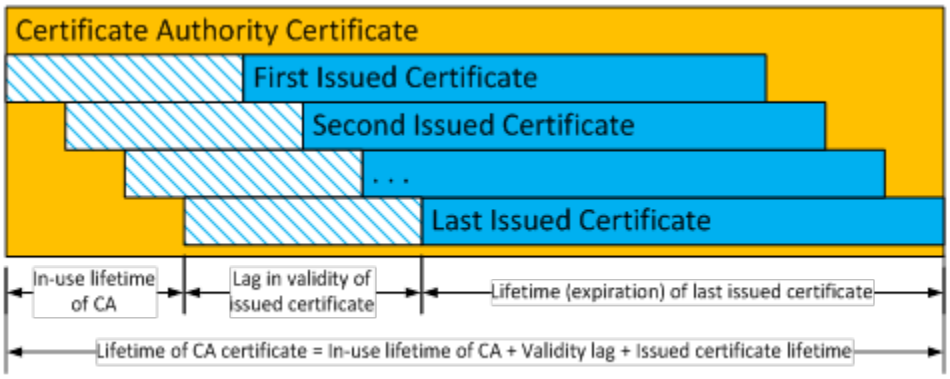

Figure 6: Impact of Lag in Validity of Issued Certificates
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

As additional layers are added to the certificate hierarchy, this process is repeated up to the root CA. When operational factors and the requirement to be able to issue new certificates at any time are considered, the required lifetime of each CA certificate is further increased going up the trust chain.
For an estimated vehicle lifetime of 30 years, it will be necessary to renew the enrollment certificate multiple times. An enrollment certificate lifetime of 6 years greatly reduces security concerns due to certificate longevity and allows an automatic renewal mechanism that can accommodate EEs with infrequent network connectivity. As better and more frequent network connectivity becomes available to EEs, it may be possible to further reduce these lifetimes. Figure 7 illustrates the impact of issued certificate lifetime, certificate validity lag, and operational factors on the PKI hierarchy.

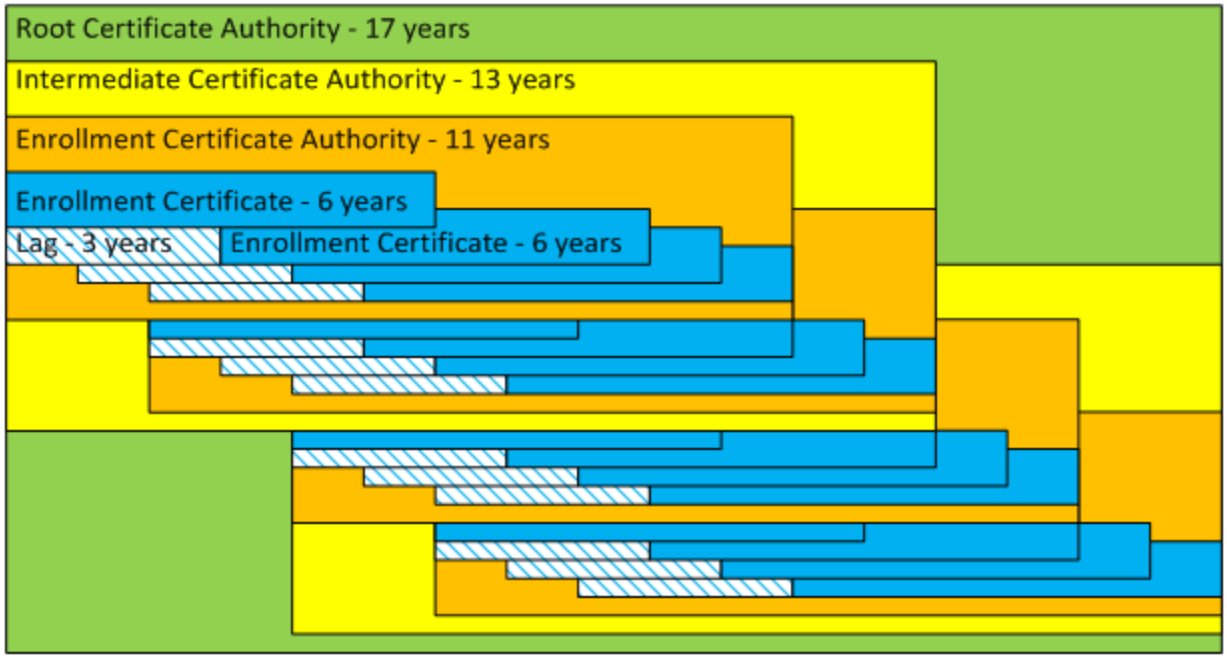

Figure 7: Relationship Between Enrollment and CA Certificate Lifetimes
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

Establishing a fixed schedule for the expiration of elector certificates, root CA certificate(s), ICA certificates, and enrollment CA certificates is recommended to reduce operational complexities. For offline CAs, this procedure increases security by minimizing the frequency that they are required to be accessed. Certificates issued in the middle of this fixed schedule (due to revocation or new instances) will expire according to the defined schedule and will have a reduced overall lifetime.

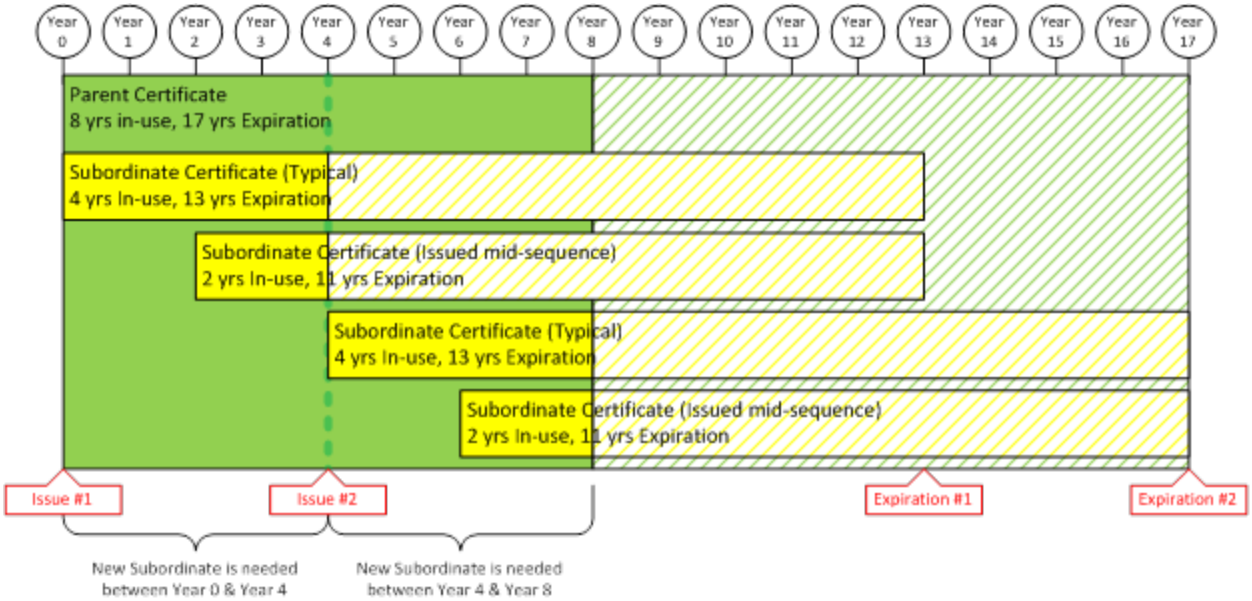

Figure 8: Example of Mid-Sequence CA Certificate
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

All roots have a specific lifespan and will retire at some point. Additionally, new requirements (e.g., increased key size) may necessitate standing up a new root. There must be an effective method to retire and replace roots without negatively impacting SCMS system operations. There are many ways to manage root retirement to limit operational interference. For example, an old root can "sign" a "roll over" certificate (contains the new root's public key) which allows the new root to be trusted without a lot of effort. Depending on the reason for the new root, the two roots can exist simultaneously (e.g., for algorithm updates) or, after the new root is established, it signs new issuing CA certs for the existing CA. The old root is revoked and the new issuing CA certs are distributed (they use a Global Certificate Chain File [GCCF] to distribute the new issuing CA certs) – the issuing CA now belongs to the new root. There is also the elector method further described in the next section.

Root retirement further emphasizes the need for a standard approach to trust anchor management, which is discussed in Section 3.2.4. Replacing a root is not inherently any harder (or easier) under other Trust Anchor Management regimes. The SCMS Manager and stakeholders will need to define the actual method to manage root retirement within the PKI policies discussed in Chapter 4. To ensure the overall integrity of the SCMS, the minimum and maximum lifetime of each certificate type will be defined and enforced by SCMS manager policy. Operators will have some degree of flexibility in defining the actual certificate lifetimes.

The addition of a new root (via whatever means) typically has no impact on previously issued end entity enrollment and pseudonym certificates unless there is some security issue being addressed (e.g., compromise of an algorithm). Previously issued certificates may continue to be relied on until they expire or are revoked and replaced as normal.

3.2.4 Elector Establishment and Management

The trust anchor management method (i.e., a trusted mechanism to manage roots) within the current SCMS design is the elector concept. However, there are other trust anchor management options (e.g., a European Commission concept has a trust list manager as the trust anchor) that should be considered. The resulting overall PKI infrastructure will not differ significantly from other methods. There will be some mechanism for adding/deleting trust anchors. Each method has tradeoffs such as cost, policy needs, and operational constraints that should be fully evaluated before committing to a method. The SCMS Manager will need to eventually oversee the selected trust anchor management process. Refer to Section 4.2.2.5 for further information on the need for specifying the trust anchor management method, processes, and procedures within the PKI policies.

The IP.com "Elector-Based Root Management System to Manage a Public Key Infrastructure" paper and CAMP SCMS End Entity Requirements document describe the elector concept in detail. In summary, electors operate at a higher level than the root CA by signing Trust Management Messages to be used by other PKI components. Essentially, electors authorize themselves and root CAs to operate within the PKI. Trust management messages are signed by one or more electors and can add a root CA certificate, add an Elector Certificate, revoke a root CA certificate, and revoke an Elector Certificate. End entities and other PKI components know the necessary number of such signed trust management messages (e.g., 2 of 3) from non-revoked Electors that will authorize the action contained in the messages (e.g., revoke root CA "A"). These messages contain a time frame for the operation to occur.[22]

[22] Reference(s): Elector-Based Root Management System to Manage a Public Key Infrastructure

The signature on the elector certificate does not have any cryptographic value as the signature is by the elector itself, and, therefore, the initial trust in an elector certificate is established outside of the PKI. For this reason, the integrity of the initial set of elector certificates must be ensured by means other than the cryptography used in generating the certificate itself. The initial provisioning of elector certificates in end entity systems is completed offline in a secure environment during enrollment. Subsequent updating of elector certificates can be completed within the PKI through revocation and adding by using the elector model described in the previous paragraph.[23]

The advantage of the elector model is that there is no single point of failure. However, each elector is essentially a specialized, stand-alone PKI that has a single certificate. Deploying an elector could cost as much as standing up a root. The SCMS Manager, as the industry governance organization could bear the cost of deploying and operating the electors or any other trust anchor management method. The question, which has yet to be answered, is how many electors are necessary to make it very unlikely that a quorum cannot be achieved without being cost prohibitive (e.g., Is three electors enough? Is six cost prohibitive?). This will need to be a discussion of risk to determine the initial deployment of electors. However, as the team understands, additional electors could be added as the system matures to ensure resiliency.

Electors are only one answer to establish trust across a framework with multiple roots. Table 7 contains a brief list and description of other methods employed within PKI systems.

Table 7: Additional Trust Anchor Management Methods (Non-Exhaustive)

Public Interest Objective	Description
Trust List Manager	The European Cooperative Intelligent Transport Systems (C-ITS) is being deployed with a multiple root architecture. It uses a mechanism called a Trust List to inform end entities of new and revoked roots. In the C-ITS model, there will be a single Trust List Manager – a specialized entity – with a certificate trusted by end entities. An updated trust list will be published periodically. The Certificate Authority/Browser Forum (policy body for publicly trusted PKIs as implemented in the world wide web) also uses trust lists, but the management and distribution of the lists is decentralized. Each browser/operating system vendors decides what roots to trust and for what purposes. They each have a proprietary method to conduct trust list updates. The decision on what roots to include is established by the trust list manager.
Cross Certificates between Peers	A cross certificate is a specialized certificate issued by a CA to another PKI. Typically, they are in pairs, although that isn't a technical requirement. The end entity consuming a certificate issued by the foreign PKI traces a path from that certificate through the cross certificates to a trusted root. The U.S. Government Federal Bridge is operated by the General Services Administration and is intended to facilitate interoperability among member organizations. There is a central policy authority which has established procedures for evaluating and determining the level of trust to be placed upon an applicant PKI against a standard set of criteria in the Federal

[23] Reference(s): SCMS PoC Implementation EE Requirements and Specifications Supporting SCMS Software Release 1.2.1

Public Interest Objective	Description
	Bridge Certificate Authority (FBCA) Certificate Policy (CP). Once the level of comparability between PKIs is determined, the FBCA and the entity PKI issue certificates to each other which is how trust between PKIs is managed.

3.3 Compliance: Auditing and Certification

The SCMS Manager will play a role in CME auditing and end entity certification no matter the technical build design, ownership model, or governance model. However, the extent of that role and oversight of auditing and certification will greatly depend on governance and ownership models along with the level of authority bestowed upon the SCMS Manager. This section discusses some of the needs for auditing and certification within the ecosystem and how these important activities could be accomplished.

3.3.1 Audits

There must be a method to ensure CME, as well as external entities such as DCM, compliance with the CP and other SCMS policies – PKI auditing programs. The SCMS Manager would create an auditing program that would meet the requirements of the PKI Certificate Policy (CP) and any additional standards. Commercial PKIs use the following standards to guide audits. These standards audit to the appropriate policies and practice statements, and are not technology specific.

1) WebTrust for Certification Authorities v2.0
2) ETSI TS 102 042/ ETSI EN 319 411-1.

Auditing validates that the security measures in the CP are in practice at the organizational level by CMEs. Auditing is a common practice and effectively used to encourage compliance with standards. The SCMS Manager may outsource the auditing function to a third-party provider or providers that have specific expertise in ITS and PKI, and that could provide training and assistance to CMEs that do not meet the security standards set in the CP. Usually the auditee pays for the audit and must account for these costs when developing their funding streams, such as fee structures. Some PKIs base audit frequency on the last audit performance to reduce costs and prioritize audit activities. Depending on the level of overall government involvement within the SCMS ecosystem, the government may also be able to require intermediate inspections of CMEs between full audits. Enforcement of penalties for noncompliance may go beyond the authority of the SCMS Manager, especially if criminal activity is involved. General oversight by the SCMS Manager will ensure that CMEs are sharing information in accordance with the CP.[24]

3.3.2 Device Certification (and Potential Re-Enrollment) Criteria for End Entities and Back Office Systems

End entities will need to meet certain PKI requirements, and functional and performance requirements, for initial enrollment and to maintain enrollment status with the SCMS regardless of the ownership and governance model. However, the level of control of the SCMS Manager over the certification process will depend on the established policies.

[24] Reference: “Organizational and Operational Models for the SCMS”

- The SCMS Manager will need to establish the certification requirements to ensure the end entity can adequately protect keys, etc. Device certification ensures that EEs operate as per mandated (V2V) specifications, and possible local safety inspection regulations.
- The SCMS Manager will set and publish its certification lab accreditation policy and process.
- A private company or other organization may then set up the required test lab facilities and request accreditation from the SCMS Manager. The test lab completes the published accreditation process and, if it meets the stated criteria, receives accreditation. This grants the test lab the ability to certify devices and to refer to itself as accredited.[25]
- The particular governance model and associated policies will influence how the certification requirements and processes are funded, enforced, and audited.

There are additional polices with respect to the wider SCMS ecosystem that need consideration, for example:

- The SCMS will likely need to consider whether re-enrollment certification criteria should require electronic proof that the end entity has met state and/or local safety inspection requirements.
- The SCMS Manager should consider policies regarding implications of end entities that are enrolled but subsequently fail state or local inspection requirements.

3.4 Communications Options for Providing SCMS Services

The SCMS will need to support communication with end-entities and with supporting services. The SCMS Manager will need to establish, administer, manage, and fund communications necessary to support SCMS Services, and will need to establish policies and guidance for communications with the SCMS.

3.4.1 Impact of Commercial Communication Services

The SCMS will most likely support a hybrid communications approach with both unicast and broadcast data distribution. Unicast sessions will be used with end entities for Enrollment, Provisioning, Authorization, Misbehavior Detection, and for other supporting services. While unicast sessions could potentially be used for Revocation services, this is an inefficient (retransmission of identical data to multiple users) and bandwidth intensive service, and is far better suited to broadcast services. Unicast sessions will be implemented as TCP/IP connections from the SCMS to the end entity with the last communication hop to/from each end entity over any viable communications media supported by the end entity including IEEE WAVE/DSRC, 3GPP 3G, 4G, and Wi-Fi. End-to-end payload encryption, and digitally signing of data messages will be employed as specified in IEEE 1609.2. Unicast sessions could be established over commercially provided services such the 3G and 4G services currently available from the telecommunications service providers, or via public network access as, for example, backhauled from a public RSE or through a public or private Wi-Fi wireless access point.

In this hybrid communications model, the communications services used by each end entity will be dictated by the communications media installed and communications services activated within each end entity device. Any costs (e.g., hardware, software, service provider fees) necessary to implement the end entity's options for communicating with the SCMS will be borne by the end entity.

[25] Reference: CAMP Technical Design of the PoC SCMS for V2X Communications

From the perspective of the SCMS, the different options will have no direct impact, as all TCP/IP connections will funnel though the ISP connection utilized by the SCMS. The SCMS Manager would need to determine how to administer, manage, and fund ISP services. Furthermore, the SCMS Manager will need policies and guidance for all communications services potentially used by end entities and supporting services to communicate with the SCMS.

Broadcast data distribution, which is ideally suited for Revocation services, can be established with end entities over commercially available services such as satellite radio (e.g., SiriusXM), and digital FM radio services. This offers significantly increased operation efficiencies and reduced costs, as the SCMS need only provide a single signed copy of each CRL to be distributed to a broadcast service provider. Each CRL can be broadcast by the broadcast service provider to the geographic area specified by the SCMS for the interval specified by the SCMS. As the SCMS is a content provider to the broadcast service provider, there may be additional costs for the broadcast services provided, in addition to costs for the ISP services needed to connect to the broadcast service provider. The SCMS Manager would need to determine how to administer, manage, and fund these additional broadcast communications services. Likewise, the SCMS Manager will need policies and guidance for all broadcast communications service providers which may be used by end entities to receive information from the SCMS.

3.4.2 SCMS-Provided Certificate Usage

As described in Section 3.4.1, there will be a variety of communications methods and services used to support SCMS operations. Additionally, going forward, new communications services will emerge (e.g. 5G) which will augment, and potentially supplant existing communications services. This does not necessarily complicate SCMS functionality, as for the most part, there are no strong dependencies between the communication media used between end-entities and the SCMS. For all SCMS Communications described above, IEEE 1609.2 specified security services and constructs can be used to sign and encrypt (when necessary) data payloads using SCMS provided certificates, and is decoupled from any security services employed by the various communications service providers.

The exception to the lack of dependency, is the inherent reliance on the IEEE 1609.2 as the basis for the definition of SCMS certificates and services. Significant changes to the IEEE 1609.2 specification, such as the addition of different subscriber types would impact, and potentially complicate SCMS functionality.

For the majority, if not all other V2I communications, IEEE 1609.2 specified security services and constructs can be used to sign and encrypt (when necessary) data payloads using SCMS provided certificates, and is decoupled from any security services employed by the various communications service providers.

Table 8: SCMS Areas of Interest and Ownership, Governance, or Operational Model Impact Summary

Areas of Interest	Associated Risks	Mitigation of Risks by the Ownership/Governance/ Operational Model or Policies
Single vs. Multiple Root Certificate Authorities	• Single point of failure with a single root • Additional cost and attack vectors with multiple roots	Based on current research, multiple roots are likely to be required to ensure there is no single point of failure, increased redundancy, and increased interoperability. The SCMS Manager will need to include requirements and policies for the establishment and management of roots within the Certificate Policy. The number of roots does not necessarily change the main principles of the governance model. However, if there are multiple roots, there will likely be multiple owners within the ecosystem providing certificate services.
Single vs. Multiple Pseudonym Certificate Authorities	• Scaling issues with a single PCA • Potential privacy issues with multiple PCAs if an RA routes request to a single PCA	For scalability reasons, it will be necessary to have multiple PCAs at some point during, or after the initial National SCMS deployment. The number of PCAs should not affect the governance model. However, similar with multiple roots, there will likely be multiple owners within the ecosystem providing certificate services if the option is available and providers can receive value from providing these services. The SCMS Manager will need to include requirements within the certificate policy for each PCA to be used by multiple RAs, and for RAs to distribute individual certificate requests from a single vehicle between the different PCAs at random.
Certificate Authority Retirement	• Compromise of a long-lived, in-use public/private key pair • Cryptanalysis as computing power increases and technologies improve over time	Certificate authority retirement emphasizes the need for a standard approach to trust anchor management. The SCMS Manager and stakeholders will need to define the actual method to manage root retirement within the PKI policies. To ensure the overall integrity of the SCMS, the minimum and maximum lifetime of each certificate type will be defined and enforced by SCMS manager policy. Operators will have some degree of flexibility in defining the actual certificate lifetimes.
Elector Establishment and Administration	• Compromise of a root without the ability to remove the compromised root and add a new root	The SCMS Manager will need to include the trust anchor management method (i.e., the elector model based on current research) and polices within the Certificate Policy. Electors could potentially be managed by the SCMS Manager. The trust anchor management method should not affect the ownership and governance models, other than there needs to be policies and controls to manage trust anchors.

Areas of Interest	Associated Risks	Mitigation of Risks by the Ownership/Governance/ Operational Model or Policies
Audits	• SCMS components not following the certificate policy or protecting information per requirements	The SCMS Manager will need to establish auditing policies in accordance with established standards (e.g., WebTrust for Certification Authorities v2.0 and ETSI TS 102 042/ ETSI EN 319 411-1). The governance model will need to ensure the SCMS Manager has the authority to penalize noncompliance.
Device Certification	• Non-certified devices enrolling in the SCMS and introducing inconsistent performance and potential vulnerabilities into the system which could compromise other devices and SCMS components	The SCMS Manager will need to establish the certification requirements to ensure the end entity can adequately protect keys, etc. Additionally, the certification requirements will need to ensure that the end entity conforms to all required specifications, standards, regulations, etc. The SCMS Manager will set and publish its certification lab accreditation policy and process. The particular governance model and associated policies will influence how the certification requirements and processes are funded, enforced, and audited.
Commercial Communication Services	• Inability to utilize certain communications services within the V2X communications system	The SCMS Manager would need to determine how to administer, manage, and fund ISP services and additional broadcast communications services required for National SCMS component internal communications. The SCMS Manager will also need policies and guidance for all communications services potentially used by end entities and supporting services to communicate with the SCMS, as well as all broadcast communications service providers which may be used by end entities to receive information from the SCMS.
SCMS-provided Certificate Usage	• SCMS services and certificates are incompatible with certain communications services	Within its policies, the SCMS Manager must include directives that require all SCMS communications use IEEE 1609.2 specified security services and constructs.

Chapter 4: SCMS PKI Policy

A comprehensive PKI policy is necessary no matter the structure of the final technical SCMS build, the SCMS Manager owner, CME owner/operator, or governance model. The PKI policy should be structured similar across the board according to industry best practices. However, the content and guidance within the policies will differ based on technical build, ownership model, and governance model. Also, vice versa, these models may be developed to support desired policies identified by stakeholders during the model development process. This chapter describes the policy needs and various examples of how to develop and implement the appropriate policies to ensure a functional and secure PKI. Setting a shared understanding of the required PKI policies for a National SCMS early in the development process ensures all stakeholders are aware of the impact particular models could have on policy. For example, a completely public model with a single root may simplify certain aspects of the PKI policy but also result in an inflexible system that cannot rapidly respond to threats and compromises.

4.1 Importance of Policy

A PKI Certificate Policy (CP) describes the operational and security requirements that will be implemented within the PKI. A CP does not say how the requirements are met. That is described by the implementer in a Certification Practice Statement (CPS). The CP is typically made publicly available so that any interested party can examine the requirements under which a specific implementation of the policy is operated.

Having an overarching CP is especially important in a distributed environment as is anticipated for the SCMS PKI. As currently envisioned, portions of the PKI must be operated by separate organizations while still meeting the overall functional, security, and privacy requirements. Without an overarching CP to which all elements of the PKI align, it will not be possible for the SCMS Manager to provide appropriate guidance and oversight to the implementing organizations.

Regardless of the model chosen for the SCMS Manager, the development and oversight of the implementation of policy will be one of its most crucial functions. The remainder of this section describes approaches to policy implementation that the SCMS Manager could take and lays out a specific foundation for the formulation of policy that will apply regardless of the SCMS model chosen or the approach the SCMS Manager decides to take.

4.2 Policy Development and Implementation

4.2.1 Overall Approach

4.2.1.1. CP Development

One role of the SCMS Manager is development and approval of the overarching CP. There are several methods that can be chosen for the development of the CP. How the CP content is developed will impact community acceptance of the CP and the role of the SCMS Manager in oversight and governance. Two fundamentally different approaches are:

- **Imposed from the top.** The SCMS Manager has sole authority over the policy. While changes may be proposed by anyone, the decision to adopt a change and, if adopted, when that change goes into effect is the decision of the SCMS Manager.

- **Developed through community consensus.** The SCMS Manager orchestrates recommendations and various stakeholders vote on the changes and when an accepted change goes into effect. The voting membership and percentage of votes needed to approve a change (e.g., simple majority, super majority – 60 to 75 percent, unanimous) would be established as part of the SCMS Manager's charter or another organizing document.

Solicitation of input can be done in a closed group (e.g., only distributed to voting members) or can be done publicly (using systems like GitHub).

4.2.1.2. CP Implementation

There are two basic models of CP implementation:

- **Unified.** All elements operate under the single CP
- **Distributed.** Various elements have different CPs

Where all operations are defined in a single CP, the SCMS Manager will be responsible for ensuring that each element has a CPS that has been reviewed and approved as complying with the single CP.

Where there are many CPs, the SCMS Manager is responsible for ensuring that each CP is comparable to the overarching CP maintained by the SCMS Manager. Each CP will have a designated Policy Management Authority that is responsible for ensuring that elements aligned under that CP operate under a CPS that is approved as complying with the local CP.

4.2.2 Format of SCMS Policies

Adopting and requiring participating entities to use a standard format will make compliance and comparability analysis significantly easier for the SCMS Manager. The PKI industry standard is RFC 3647. While this RFC was developed for PKIs that issue a X.509 certificate, the format is not X.509 specific and provides a well-recognized and accepted structure to address the broad range of topics necessary to ensure that a PKI is operated in a secure manner while meeting the PKI's operational objectives. This is also the format that has been adopted by the European Commission for European C-ITS.

While all elements of the PKI operate under the CP, their CPSs need to be tailored to their specific function. The use of a specific format requires each entity to specifically address each requirement to ensure that nothing is inadvertently overlooked. Some of the CP requirements will apply to all while others can be marked as "not applicable" in the practice statement. For example, many requirements related to operating a Pseudo-CA would not be applicable to the Misbehavior Authority. The entity statement that a requirement is not applicable in its CPS demonstrates that the requirement was considered and not that it was overlooked. The following subsections describe the type of content necessary within each policy, but do not identify the actual policies for each element of the SCMS. These policies will eventually be developed and structured based on the final technical structure of the SCMS ecosystem and the selected ownership and governance model(s).

4.2.2.1. Overview of the PKI

This section of the policy addresses the PKI in general. It identifies the policy owner, the scope and applicability of the PKI, identifies the kinds of entities that the CP addresses, appropriate uses of certificates issued by the PKI, and the administration of the policy and practice statements.

This section would be the primary place where the SCMS model and implementation choices (e.g., unified versus distributed) would have the most impact.

The overview section describes how the SCMS Manager oversees and manages the CP, what parts of the overall infrastructure are subject to the policy, and the mechanisms for approving entities for their role in operating portions of the SCMS. It could also place limits on the appropriate use of the certificates issued by the SCMS (e.g., where the use of enrollment certificates is appropriate and where their use would be prohibited).

4.2.2.2. Operation of Repositories

Repositories are systems operated by the PKI that provide information to entities that consume the PKI services. Items such as CA certificates, CRLs, and the CP itself are posted to the PKI repositories. This section describes what kind of information is to be posted, access control (public or private), availability requirements, and integrity of the information posted to the repositories.

At a minimum, the SCMS would need repositories for CRLs. It may also require that various documents (e.g., CPs, CPSs, audit report) be posted. There may be a single location where all of this is posted the SCMS or individual entities may generate and post information in a distributed manner.

4.2.2.3. Identity and Authentication

One of the core functions of the PKI is to link the public key in a certificate with the holder of that key. This section addresses how the PKI verifies the proper identity is being bound to the public key. This section specifies the requirements for identity proofing that must be satisfied before the PKI issues any certificate. There will be specific requirements related to the certificates issued to the PKI entities themselves (e.g., CAs, RAs), as well as the enrollment and pseudonym certificates to be issued to the V2X consumers.

The SCMS will have policies how the DCM identifies the equipment to receive enrollment certificates and the requirements related to that operation. It will also specify the requirement for the RA to use the enrollment certificate to authenticate the request for pseudonym certificates and authentication of other required interactions within the SCMS.

4.2.2.4. Certificate Life Cycle

Every Certificate issued by the PKI has a life cycle. It is generated and issued to the entity that holds the private key and needs to be managed until it expires or is revoked. This section provides the framework for the issuance process for different types of certificates, how they are requested, generated, and provided to the subscriber. It also addresses how expiring certificates are replaced, how revocation is done, and how often revocation information needs to be published.

The SCMS will specify requirements for the interaction between LAs, RAs, and CAs in the generation and distribution of certificates. It will also specify the requirements related to the Misbehavior Authority, when certificates need to be revoked, how often CRLs are published, etc.

4.2.2.5. Facility, Management, and Operational Controls

This section provides specific requirements that ensure that the physical and operational environment is properly implemented. It is particularly important in the SCMS PKI because of the need for control and separation of functions and data to ensure the goal of consumer privacy is achieved.

Physical Controls

This section addresses the requirements related to the physical plant. It specifies requirements for the physical security of the PKI entity equipment.

The SCMS will specify any requirements related to the facilities that house SCMS components. Typically, this would provide direction on how hardware is secured (e.g., guards, locks, intrusion detection) as well as other physical aspects (e.g., fire, waste, disaster recovery).

Procedural Controls

Procedural controls specify the roles that will be necessary to operate and maintain the various PKI entities, how those individuals authenticate to the PKI, what requirements are in place to ensure separation (both within an entity and between entities).

The SCMS will define, at a high level, the roles of the people who operate the SCMS components. More importantly, it will specify the requirements for separation of roles within specific components (e.g., system administrator cannot be a security officer on the same component) and for separation between components (e.g., no individual may hold any role on more than one RA, LA, or CA).

Personnel Controls

This section addresses how people are selected, vetted, trained, and approved to perform their function within the PKI. The SCMS would specify any requirements related to background checks, training, job rotation, etc.

Audit Log Management

Audit data consists of electronic and physical records collected during PKI operations to demonstrate that the PKI is operating as required.

The SCMS would specify the events that are required to be captured in manual or automated audit logs, the security applied to those logs, how often audit data is reviewed, and how long audit data needs to be retained.

Archiving of Data

Archives are the long-term records that need to be captured and retained by the PKI for an extended period of time. The SCMS would specify what those records are, how long they need to be retained, and how they are protected during the required retention period.

Trust Anchor Management

This section will specify the requirements related to the management of trust anchors in the V2X equipment. While there are several different options available, the SCMS Manager will need to settle on a single method to ensure that deployed devices are able to implement it properly. Current designs are based on the elector method as described in Chapter 3. *It will be difficult, if not impossible to change the chosen method once large-scale deployment begins.*

Compromise Recovery

Related to trust anchor management, the question of how the SCMS will recover from the compromise of one of the SCMS entities needs to be worked through early.

SCMS Element Termination

This section will specify requirements for SCMS elements that end their service in the SCMS PKI. Regardless of the reason, entity termination will have requirements such as turnover of audit/archive records and timely notification of intent to terminate.

4.2.2.6. Technical Security Controls

Technical security controls include requirements for SCMS elements and subscribers related to key generation, key size, key and certificate lifetimes, software and hardware development controls, certifications, and network security controls.

4.2.2.7. Certificate and CRL Profile

This section provides specifications for minimum standards for certificates and CRLs created by the SCMS. It describes mandatory, optional, and prohibited certificate fields for each kind of certificate to be issued. Conformance to these requirements is necessary to ensure interoperability when certificates are issued by different organizations.

4.2.2.8. Compliance Audits

A Compliance audit is an independent review of SCMS element performance against the CP and its approved CPS. The compliance auditor reviews audit and archive records and reviews/observes a representative sample of actions performed by the element to ensure it is operating properly. This independent review precludes the need for the SCMS Manager to directly inspect the operations of the SCMS element while retaining an understanding of how well the SCMS is meeting the requirements of the policy. In addition, the compliance audits can provide input into the CP change process, letting the SCMS Manager know what requirements are not working well or may not be achieving their intended result.

4.2.2.9. Legal and Other Matters

This section is intended to cover matters that are not directly related to SCMS operation but critical to the overall implementation of the SCMS. The topics covered range from liability, fees, applicable law, to protection of privacy information obtained by the SCMS.

4.2.3. Other Policy-Related Topics

4.2.3.1. Tailoring (How the Policy Will be Tailored and Applied to Each Element)

The exact determination of how each element of the SCMS will apply the CP will be a function of the structure. As an example, the following presents a partial section by section mapping for each element.

Table 9: Notional CP Mapping by SCMS Function

Section	Root	ICA	PCA	ECA	RA	LA	MA	DCM	LOP
1 PKI Overview	X	X	X	X	X	X	X	X	X
2 Operation of Repositories							X		
3 Identity and Authentication	X	X	X	X	X			X	
4 Certificate Lifecycle	X	X	X	X	X		X		
5 Fac., Mgmt., Ops Controls	X	X	X	X	X	X	X	X	X
6 Technical Security Controls	X	X	X	X	X	X	X	X	X
7 Certificate and CRL Profile	X	X	X	X	X	X	X	X	
8 Compliance Audits	X	X	X	X	X	X	X	X	X
9 Legal and Other	X	X	X	X	X	X	X	X	X

Although a specific section may be required, the subsections and information required will also be based on the nature of the SCMS element. E.g., the MA has a role and needs to address the revocation functions which would be described in the Certificate Lifecycle Section. Other portions of the Certificate Lifecycle which deal with certificate issuance and maintenance would not be applicable to the MA.

4.2.4. Existing PKI Policy Models

There are several different models of PKI governance and oversight that exist in the X.509 PKI space. Generally, these models are differentiated by the number of roots, the number of CPs, and how CP compliance is ensured.

4.2.4.1. Single CP, Single Root

The simplest PKI structure has all elements of the PKI operating under a single root CA operating under a single certificate policy. This structure lends itself well to environments which support a single organization, where the organization controls both the policy and the implementation of the policy. This is also a basic building block for all other models. Examples of this would be the original Common Policy Framework and Committee on National Security Systems (CNSS) PKIs operated by the U.S. Government. In this model, CP compliance is ensured by having each element of the PKI operating under a CPS that is reviewed, approved, and audited by a central policy management authority. That policy management authority has complete control over the content of the CP.

4.2.4.2. Single CP, Multiple Roots

Often, a single CP/single root become a single CP/multiple root implementation. The organization's requirements evolve over time at it determines it needs two or more roots, established for different purposes, operating under a single CP. This could be for things like different key sizes and algorithms as requirements for implementing strong cryptography evolve. For example, a legacy root stays in existence while a new one is introduced which meets more stringent cryptographic requirements. The older one is removed after all its issued certificates expire or are revoked. The separate roots could also be for different purposes. For example,

the SCMS could establish a separate root for the Enrollment and Pseudonym CA structures. Both the Common Policy framework and the CNSS PKIs have already evolved from their original single root structure to multiple roots under the same policy. As with the single CP/single root model, CP compliance is ensured by having each element of the PKI operating under a CPS that is reviewed, approved, and audited by a central policy management authority (which in this case would be the SCMS Manager).

4.2.4.2. Multiple CPs, Multiple Roots

In this architecture, there are a number of single CP/single (or multiple) root PKIs all joined together in an extended trust framework. While each of these separate structures could operate completely independent of each other, the community needs are better served by having them all operate under an overarching framework. While the trust relationship could be bi-lateral (i.e., each PKI independently evaluating every other PKI to decide on comparability), as the number of PKIs grow, that becomes unmanageable. Examples of this are the PKIs operating under the U.S. Federal Bridge CP, the Public Trust PKI implemented under the CA/Browser (CA/B) Forum Baseline Requirements and the EC's V2X architecture[26]. There are also several commercial bridges that operate similarly to the Federal Bridge, and many of those commercial bridges are members of the Federal Bridge.

The Federal Bridge CP operates under a U.S. Government Policy Authority (PA) which analyses the CP of each potential member of the bridge to ensure that the potential member's CP provides a comparable level of security to the Federal Bridge CP. The PA operates a specialized root that only exists to issue cross certificates to the entity roots. These cross certificates allow user systems to dynamically create a trust path to a root the user system already trusts. Changes to the Federal Bridge CP are controlled by the PA after coordination with the membership.

The CA/B Forum establishes the policy baseline for what are referred to as the Public Trust PKIs that issue most of the web server certificates seen by people on the internet. The Forum leaves decisions on which PKIs meet the requirements to individual browser vendors. These vendors each establish a process for review and approval of candidate PKIs into the vendor controlled trust store and provide a vendor specific process for maintaining the trust list in the user's computer systems. The changes to baseline requirements are vetted publicly and then formally voted on by the membership under rules established in the CA/B Forum charter.

The EC has established its PKI governance model as multiple policy, multiple roots. Like the Federal Bridge, the EC determines policy comparability and publishes changes to the policy after consultation with member PKIs. Unlike the Federal Bridge, the EC uses a Trust List mechanism, like the CA/B forum vendors, as the means to distribute the trusted roots to user systems. The Trust List Manager has a certificate that is trusted by all parties that need to consume updates to the trust list.

[26] Standard (EN) 302 665, "Intelligent Transport Systems (ITS); Communications Architecture."

Chapter 5: SCMS PoC Description

Currently, the USDOT is leading the SCMS PoC to support CV pilots and other federally-funded V2X related efforts. While the National SCMS will look substantially different from the SCMS PoC, the government and industry can make use of SCMS PoC practices, policies, lessons learned, and potentially even SCMS PoC infrastructure when deploying the National SCMS ecosystem. As the SCMS PoC effort continues to advance, the National SCMS Deployment Support team will stay engaged to implement new information and lessons learned during the development of potential ownership and governance models as appropriate. This chapter provides a more detailed overview of the current SCMS PoC proposed functionality and status. Later iterations of this report will include a description of what and how can be reapplied for a National SCMS.

5.1 SCMS PoC Summary[27]

As a part of the USDOT's commitment to ensuring that CV technologies operate securely, an SCMS Proof of Concept (PoC) has been created in partnership with the CAMP. It uses a PKI-based approach that permits authorized system participants to use digital certificates issued by the SCMS PoC to authenticate and validate V2V and V2I messages. The SCMS PoC is being established as one of the technologies being integrated into the CV pilot sites to demonstrate how the system can operate in realistic environments. It supports a set of use cases (defined in Section 5.3) that comprises a subset of those needed for a National SCMS deployment. At this time, the expected lifespan of the SCMS PoC is through December 2020. Deployment sites funded by the USDOT and beyond the CV pilots are also candidates to request enrollment in the SCMS PoC.

Two SCMS systems will operate as part of the PoC: A Quality Assurance (QA) SCMS and a Production SCMS. The QA SCMS is intended to permit EE device vendors and CV deployers to help develop and test their devices. The interfaces to the QA SCMS will be identical to the Production SCMS; however, there will be a few key differences in policies, procedures, and architecture. The first key difference is that the QA SCMS has less stringent security requirements to enable easier use for device developers and deployers. This includes relaxing the secure environment requirements for bootstrapping devices and removing the need for devices to be certified before connecting to the QA SCMS, although it is still recommended that bootstrapping occur in a secure environment. The second key difference is that the QA SCMS will have its own, separate root CA, which is driven by the less rigorous security requirements. This ensures that if there is a compromise of one of the QA SCMS CAs, it will not have an operation impact. Future versions of the QA SCMS may also add new capabilities and functionality that does not exist within the Production SCMS.

The Production SCMS is the system intended to connect to deployed and operational devices so they can receive their operational certificates. A key difference between the Production SCMS and QA SCMS is that the Production SCMS will not own and manage the root CA. The Production SCMS will have an intermediate CA (ICA) that the USDOT will control the policies and procedures for; however, the root CA will be owned and

[27] Refence(s): "SCMS PoC Root Access Management Policies and Protocols", the "Initial Set of SCMS PoC Governmental Management Policies and Organizational Documents", and the "SCMS PoC Governmental Management Concept of Operations (ConOps)"

operated by Integrity Security Services. There is a Subscriber Agreement and Certificate Policy that sets the policies and rules between the root CA and ICA. Having a separately owned and managed root CA will allow other CV deployment groups, such as automakers building vehicles with CV devices, to set up their own ICAs and allow secure authenticated communications between their vehicles and CV deployment devices.

The Connected Vehicle Deployment Support is another key part of the SCMS PoC system in which end users will interact. Connected Vehicle Deployment Support will be the initial technical support team that end users will interact with when they have issues with their CV deployments. This support will primarily be a Tier 1 level support service that provides some initial troubleshooting help before escalating issues to the SCMS Operations team. This team also operates and maintains a tool that provides automated workflows and trouble ticket tracking.

5.2 Roles and Responsibilities Within the SCMS PoC

Stakeholders in the SCMS PoC can be divided between users and the four teams providing governance, operations, and management of the PoC. Table 10 lists and defines these four teams.

Table 10: Overview of the Roles and Responsibilities of the PoC Teams

PoC Team	High-Level Roles and Responsibilities
SCMS PoC Governmental Management Team	Develops, approves and enforces policies and procedures in the context of the SCMS PoC, and is comprised of the USDOT stakeholders and its contracted partners
SCMS Operations Team	Responsible for the day-to-day operations of the QA and Production SCMS including ensuring operational uptime, troubleshooting user issues, managing bug fixes, developing new capabilities, and providing configuration management
CV Deployment Support Team	Provides first tier technical support to the CV deployment teams operating with the SCMS, including initial troubleshooting and operating and maintaining the workflow and trouble ticketing software system
Device Certifiers	External organizations that certify CV devices, which is a requirement for the device to be enrolled with the Production SCMS. Certification is done according to the Certification Operating Council (COC) guidance

5.3 Description of Use Cases and Capabilities of the SCMS PoC

The SCMS PoC has been designed with a set of use cases that represent a subset of the use cases needed in an eventual production deployment. The PoC use cases' focus on the operation functions where users outside the SCMS system interact with the system. Table 11 provides a list and brief description of these use cases, including use cases for implementation in version 1 of the PoC, as well as future releases.

Table 11: SCMS Proof of Concept Use Cases

PoC Use Case	Definition
Device bootstrapping: QA SCMS	For the PoC Version 1, bootstrapping for the QA SCMS will be a manual process requiring the organization that is bootstrapping the device to collect enrollment certificate requests, send those requests to USDOT for approval, and after approval have the SCMS Operations staff generate the required certificates and send them back to the organization.
Device bootstrapping: production SCMS	Device bootstrapping for the production SCMS is also a manual process requiring the user to collect enrollment certificate requests, send those requests to the USDOT for approval; and after approval generating the required certificates and sending them back to the EE
Certificate distribution: OBU pseudonyms	This defines the certificate distribution for OBU pseudonyms
Certificate distribution: OBU identifications	This defines the certificate distribution for OBU identification certificates
Certificate distribution: RSU applications	This defines the certificate distribution for RSU applications
Certificate distribution: back office system applications	This defines the certificate distribution for back office system applications
Certificate distribution: device CRL download	This defines the device CRL download. The process is identical for both the QA and Production SCMS and for all device types
Certificate revocation: OBU pseudonyms	Certificate revocation is not a use case that will be included in Version 1 of the PoC; however, in a future version OBE revocation will be integrated with the to-be-awarded “Misbehavior Authority Integration” subproject
Certificate revocation: RSU application certificate and OBU identification	Certificate revocation is not a use case that will be included in Version 1 of the PoC; however, in a future version RSE application and OBE identification certificate revocation will be integrated with the to-be-awarded “Misbehavior Authority Integration” subproject
Misbehavior reporting	Misbehavior reporting in the QA and Production SCMS systems will be a manual process for version 1 of both systems. Device misbehavior reporting will rely on the CV pilot deployment sites identifying suspect devices through their individual data collection efforts. It is anticipated that a more formal misbehavior detection process will be implemented in version 2.
SCMS technical support	This defines the three-tiered support mechanism to address users who request technical support. It is modeled after ITIL-based support schemes.

PoC Use Case	Definition
Local policy file change request	CV deployment sites have the option of defining different policies or configurations with respect to the SCMS. This use case defines how the SCMS Operations team will process the request, including updating the local policy file and uploading changes to the RA for dissemination, if the request is approved.
CV pilot reserved PSID request	Once the SCMS Production SCMS system is initialized, new PSIDs cannot be added without creating a new Intermediate CA. 16 PSIDs were procured for the PoC that will be part of the SCMS Production SCMS system and reserved for future use. This use case defines how requests to use these 16 PSIDs will be handled.
Device re-enrollment	Device re-enrollment will not be supported by version 1 of the QA and Production SCMS systems. Affected devices will need to be re-bootstrapped as per use cases QA SCMS Device Bootstrapping and Production SCMS Device Bootstrapping.

5.4 Description of SCMS PoC Policies and Procedures

The PoC includes a series of governance policies and procedures that are currently being finalized by the SCMS PoC Governmental Support project. These policies and procedures meet the needs of the set of use cases as presented in Section 5.3. The list of policies and procedures, and their mapping to the use case needs, is defined in Table 12.

Table 12: Policies and Procedures Established for the Proof of Concept

Policy or Procedure	Definition
Certificate policy	Identifies the roles and duties of each of the key actors in the SCMS's PKI. This policy relates to Chapter 4 (SCMS PKI Policy).
Certificate practices statement	Describes the practices and policies associated with issuing and managing public key certificates within the SCMS. This policy also relates to Chapter 4 (SCMS PKI Policy)
Certification of devices policy	This defines what is required to approve CV devices for enrollment with the QA and Production SCMS systems. For the Production SCMS system, this includes the tests from the Certification Operating Council test suite that are required for enrollment. The scope of this policy includes OBUs, RSUs and back office systems. This policy relates to the discussion in Section 3.3.2 (Device certification).
Addition of applications	This details how a developer would go about requesting one of the PSIDs that have been reserved for the CV pilots
Connected Vehicle Core System (CVCS)	This defines the request and reporting forms to implement the defined use cases, and includes both the information fields for each form as well as the workflow that is initiated when a form is submitted

Policy or Procedure	Definition
SCMS forms and workflows	
SCMS troubleshooting procedures	This defines steps the Connected Vehicle Deployment Support team will take to address reported issues, including issue escalation procedures

Finally, the PoC establishes how these policies and procedures map to the needs from the PoC's defined use cases. This mapping is shown below in Table 13.

Table 13: Mapping of Policies and Procedures to Use Case Needs

PoC Use Case	Policy and Procedure Need	Policy and Procedure Mapping to Need
Device bootstrapping: QA SCMS	QA SCMS enrollment request form	CVCS SCMS forms and workflows
	QA SCMS enrollment request workflow	CVCS SCMS forms and workflows
	QA SCMS enrollment approval criteria	Certification of devices policy
Device bootstrapping: production SCMS	Device certification request form	CVCS SCMS forms and workflows
	Device certification criteria (OBU/RSU)	Certification of devices policy
	Device certification criteria (back office systems)	Certification of devices policy
	Production SCMS enrollment request form	CVCS SCMS forms and workflows
	Production SCMS enrollment request workflow	CVCS SCMS forms and workflow
	Production SCMS enrollment approval criteria	Certification of devices policy
Misbehavior reporting	Misbehavior report form	CVCS SCMS forms and workflow
	Misbehavior report workflow	CVCS SCMS forms and workflow
SCMS technical support	SCMS technical support form	CVCS SCMS forms and workflow
	SCMS technical support workflow	CVCS SCMS forms and workflow

PoC Use Case	Policy and Procedure Need	Policy and Procedure Mapping to Need
	CV deployment support troubleshooting procedure	SCMS troubleshooting procedures
Local policy file change request	Local policy file change request form	CVCS SCMS forms and workflow
	Local policy file change request workflow	CVCS SCMS forms and workflow
CV pilot reserved PSID request	CV pilot reserved PSID request form	CVCS SCMS forms and workflow
	CV pilot reserved PSID request workflow	CVCS SCMS forms and workflow

Appendix A. Connected Vehicle Overview

This Appendix provides a high-level overview of the connected vehicle concept and Public Key Infrastructure (PKI) implementation including an overview of the SCMS concept. The content in this Appendix is intended as a primer or refresher for readers who are unfamiliar with these concepts, to help provide context and understanding of the material in this report.

A.1 The Connected Vehicle Concept

The CV concept aims to improve roadway safety and efficiency by providing a mechanism for situational information to be shared between vehicles, and between vehicles and other entities in the roadway environment such as traffic signals, warning signs, and even non-vehicular users such as pedestrians and cyclists.

Conceptually the system is quite simple. A radio link is provided that allows messages bearing a variety of content to be sent from any equipped roadway entity to other nearby roadway entities in order to provide additional information about the current situation in that area. Examples include:

- One vehicle broadcasting privacy-protected information about its current operational state (e.g., position, speed, heading, and other operational data) to other vehicles, pedestrians, and cyclists in the immediate vicinity;
- A roadside infrastructure system such as a traffic signal, broadcasting its current signal state and near-term timing information to vehicles in the vicinity of the traffic signal;
- Pedestrians or cyclists broadcasting their privacy-protected position to vehicles nearby to allow drivers to better “see” them.

Ideally, this information can then be used by the recipients to improve situational awareness. For example, if a vehicle is braking hard, this information may be useful to vehicles several cars behind so that the drivers can be warned that this event, which may not be visible to the driver, is occurring. By providing this advanced warning, or in some cases simply by announcing their presence, drivers can be more aware of what is happening around them, and can thus be more prepared and informed so they can choose their driving actions appropriately.

While the dominant focus of the CV concept has been on information exchanges between vehicles, it is also understood that other applications will emerge as the population of equipped vehicles grows. For example, vehicles can periodically report their speed and location to remote servers to provide enhanced traffic management information; traffic signal systems can provide information that allows vehicles to smoothly traverse a city center, hitting all green lights; and emergency vehicles can preempt traffic signals and inform everyone on the road that they are approaching, including what direction they are approaching from and where they may turn. Once vehicles are able to communicate wirelessly, the potential range of applications is substantial.

Though conceptually simple, the CV concept involves a wide array of technologies, and will require a considerable level of coordination to emerge efficiently and rapidly. Much of this coordination has been

accomplished through the development of standards that ensure that the messages broadcast from one entity can be received and understood by other entities. This interoperability is essential, since the devices (vehicular, roadside, or personal) will likely be manufactured many different companies, and there is no practical way to test every new device against every other device.

It is important to note that while the basic CV concept does not depend on any particular communications technology, the concept does impose some communication requirements.

First, since in most cases information sent from any device is intended to be made available to any other user in the vicinity, it is essential for all users to be equipped with compatible communications devices. If not, users equipped with one type of device will not be able to receive messages from users equipped with a different type of device. For example, a user equipped with a cellular phone would be unable to receive messages broadcast by a user using some non-cellular radio system. The specific choice of communications technology can vary considerably (limited as described below), but to enable intercommunication between *all* road users, they must *all* be using the same communications technology.

Second, since vehicles move relatively quickly and hazardous situations can develop rapidly, the system must provide communications with minimal delay. The tolerable delay depends on the application. For example, a traffic congestion report does not change rapidly, so that sort of information could be obtained, for example from a remote server with a delay of several seconds (or even minutes). On the other hand, automated vehicles operating in very close proximity under closed loop control may require data exchanges with less than a few milliseconds of delay. And any applications that involve information that may change spontaneously (for example a change in the timing plan for a traffic signal), needs to be communicated quickly enough (e.g., a few hundred milliseconds) that the users can respond appropriately and safely.

Third, most messages are intended to be useful to all users in the vicinity (assuming they are suitably equipped to receive them, as noted above). Since it is impractical to determine addresses or identifiers of all users (which may number in the hundreds in some cases) in the immediate vicinity given requirements on delay times, the system must provide some mechanism to allow a user in any given location to determine the operational state of those users in the local vicinity without the necessity of one-to-one communications exchanges[28]. There are various ways to address this issue. For example, if the system was fast enough, each vehicle could post its status on a server, and then query the server for the status of all vehicles in the local area. Alternatively, and much more practically, messages can simply be broadcast in nature, that is, not addressed to any particular recipient, but receivable by all users in some vicinity. From a practical perspective, the current system is envisioned as broadcast based.

Lastly, since the system is primarily intended for safety applications, it must exhibit high availability. This means that it must be independent of extraneous interference from other radio systems, and it must be capable of supporting message traffic volumes compatible with a typical congested roadway. This last point can be rather complex, since there is a tendency to assume that longer radio range is “better,” since it assures the widest distribution of information. However, as the system’s range increases, the number of vehicles in range grows (typically proportional to the square of the range), and this increases the volume of messages the system must support (or, said differently, increases the potential communications link congestion). The operational state of two vehicles that are hundreds of meters apart is clearly not as safety critical as that of vehicles within a few

[28] Note, even if it were possible to learn the addresses of all local users, if each user were to directly query every other user on a one-to-one basis, the number of data exchanges would be significant (for N vehicles, the number of one-to-one exchanges would need to be N(N-1).

hundred feet, or a few tens of feet. So, the range of the system must be designed to be sufficient to support the envisioned applications, and not so long as to create message congestion problems.

A.1.1 Information Exchanges

As described above, the current CV system uses broadcast communications to distribute information to all other users in the local area (the area being defined by the range of the system). This does not mean that the system cannot also support one-to-one data exchanges. As currently envisioned, the system supports two types of communication: "broadcast" and "unicast."

The current system design also includes the provision for communicating using single messages, or using the internet protocol (IP). Single message communication uses what is known as the Wireless Access in Vehicular Environment (WAVE) Short Message (WSM). This is a single packet message that may or may not include an address for the recipient. IP communication follows the Internet Protocol, and may involve multiple packets, each including the sender and recipient IP addresses.

The WAVE protocol, which defines the higher-level layers of the communications process, is defined by the IEEE 1609.x suite of standards. The lower layers, including the physical radio channel, is defined within the 802.11 suite of standards for the United States. This is also known as dedicated short range communications, or DSRC.

Broadcast messages are always sent using the WSM protocol. These include no addressing (since they are intended for all recipients). Instead they include information that allows the recipient to determine how to interpret and use the message. This supports a wide variety of messages intended to support a wide variety of applications, and provides substantial flexibility in message and application design. As will be described below, it also provides a mechanism for providing permissions, since not every user can be allowed to broadcast messages of any given type[29].

Unicast messages do include an address. These messages are also broadcast over the radio link, so all users in range can technically "hear" them, but based on the communications protocols, only the user to whom the message is sent will process the message. Other recipients will receive the message and immediately discard it because it does not bear their address. In the case of a WSM, the message includes the Medium Access Control (MAC) address of the recipient. It is important to note that to determine the MAC address of a user the sending device must either know this in advance (and know, somehow that the recipient device is in range), or it must start the communications process by sending a broadcast message.

The CV system supports three basic types of data exchanges. These are implemented using one of the messaging mechanisms described above.:

- **V2V** – These are WSMs that typically carry operational state information. In most applications, these are broadcast (unaddressed) from one vehicle and received by any other vehicles in radio range.
- **V2I** – These may be either WSMs or IP messages. When used in a localized area, they are typically WSMs sent from a roadside system such as a traffic signal, and are received by vehicles in range of the roadside transmitter. In some cases, a vehicle may exchange data with a remote server using some sort of roadside system as an intermediary. In this situation, the roadside system acts as an

[29] For example, a private vehicle should not be allowed to broadcast a traffic signal preemption messages used by emergency vehicles.

internet access point, and the vehicle would use IP communication to access and exchange data with a remote server.

- **V2X** – V2X messages are an outgrowth of "the Internet of things". This area and its applications are only recently emerging, but conceptually, V2X would, for example allow for a user's smart phone to receive messages from cars or from the roadside, and would also be capable of sending messages to cars or the roadside. In the roadway environment, these are presumably WSMs much as described to V2V, but the content and nature of the messages would be somewhat different.

These data exchanges are not necessarily mutually exclusive, but are generally described in accordance with their primary function. For example, a roadside system may receive V2V messages broadcast by passing vehicles, and take some other action based on those messages, such as adjusting signal timing to smooth traffic flow. Similarly, a user's personal device may receive V2V messages and provide a caution to the user, for example warning them of an oncoming vehicle.

While there are a wide variety of CV applications, and many different messages associated with those application, the core vehicle status message used is known as the BSM. This message is defined in the SAE standards J2735, and J2945 which specify the various vehicle operating parameters that comprise the message, and how these parameters are encoded. At this time, there has been no development of a personal safety message (PSM). It is likely that messages of this type may emerge soon to provide for the evolution of V2V into V2X. In this model, different device types (e.g., vehicles, pedestrians, cyclists, children, pets) might generate messages associated with their particular activity, and yet all other device types would be configured to receive and properly interpret these different message types to provide for a broad and sophisticated level of roadway safety.

A.1.2 Terminal Types

The CV concept identifies a variety of different terminal types. These include:

- Private vehicles, which will typically include a variety of safety applications for broadcasting vehicle status information to other vehicles and roadside systems and may include other applications for providing safety and mobility information to the driver.
- Public vehicles, which may include most of the safety applications found in private vehicles, but may also include specialized applications. These could include warning beacons based on the type of vehicle (e.g., emergency vehicle, snow plow, utility) and potentially specialized control applications, for example traffic signal preemption.
- Roadside systems, which may provide public services (e.g., a traffic signal or a warning beacon), or private services (e.g., internet access, payment transactions).

A.2 Overview of Public Key Cryptography

The CV system uses a security mechanism based on public key cryptography. This will be more fully described below, but a basic understanding of public key concepts is helpful to understand the overall system.

Conventional cryptography typically relies on a key that is shared by the two parties seeking to communicate privately. In this model, blocks of data are encrypted using the key, and then sent. The recipient has the same key and uses that to decrypt each block of data. This is known as symmetrical encryption, because the key used to encrypt the message is the same as that used to decrypt it.

A problem with symmetrical cryptography is that it assumes that the communicating parties have some means for sharing the symmetric key. This may be the case if the parties can, for example, physically meet and agree on a key, and then later use the key to communicate. However, in the CV application, any given pair of communicating vehicles are unlikely to have ever encountered one another, and thus there is no practical way to have established a common symmetric key. In addition, a common key used by all vehicles would simply obviate the use of encryption. Symmetric encryption is in fact used in some CV applications, but it is only possible by exchanging the symmetric key using a public key system, as described below.

Public key cryptography is based on *asymmetrical* key encryption. Here the key used to encrypt data is different from the key that is used to decrypt it. There are two mathematically related keys, each one able to encrypt or decrypt data. Data that is encrypted by one key can only be decrypted by the other, and vice versa. These key pairs are derived using mathematics that presents a simple problem in one direction, and a very difficult problem in the other. For example, it is easy to multiply two large numbers together, but it is much more time consuming to factor the resulting large number to arrive at the two original numbers. In practice, today's public key systems use very sophisticated algorithms based on number theory to generate asymmetric key pairs that are exceedingly difficult to reverse (i.e., derive the matching key from only one key), yet do not require excessively long keys. The public key cryptography system used in the CV system is based on what is known as "elliptic curve" cryptography.

Because the keys are different, and are very difficult to derive from one another, one key can be published and the other can be kept secret. This gives rise to the name *public* key cryptography.
A critical element of this concept is the management and distribution of keys. For example, it could be possible that a bad actor might distribute a key claiming it was someone else's key. In this situation, it is important for users to be able to verify the ownership of a public key. This is accomplished by including a certificate that includes a signature over the key (actually over a hash of the key) that is signed by a trusted third party. In some cases, there may be several layers of such third parties creating what is known as a "chain of trust". This chain of trust depends on a set of policies that define how the various parties in this chain mange keys and certificates. The overall set of entities that make up this chain of trust are known as a "public key infrastructure", or PKI. The last link in this chain of trust is an organization that is known and trusted by both the sender and the receiver of the message, or more commonly by the organizations that form the chain of trust, is the "trust anchor", or "root of trust" (commonly just called "the root").[30] A public key infrastructure (PKI) supports the distribution and identification of public encryption keys, enabling users and computers to both securely exchange data over networks such as the Internet and verify the identity of the other party.

Without PKI, sensitive information can still be encrypted (ensuring confidentiality) and exchanged, but there would be no assurance of the identity (authentication) of the other party.

A typical PKI includes the following key elements:

- A trusted party, called a certificate authority (CA), acts as the root of trust and provides services that authenticate the identity of individuals, computers and other entities
- A registration authority, often called a subordinate CA, certified by a root CA to issue certificates for specific uses permitted by the root
- A certificate database, which stores certificate requests and issues and revokes certificates
- A certificate store, which resides on a local computer as a place to store issued certificates and private keys.

[30] http://searchsecurity.techtarget.com/definition/PKI?

The CV PKI is somewhat more complex than a typical PKI because it must provide certificates without disclosing the identity of the user, and this added anonymity requirement results in added functions that are discussed below.

A.3 Connected Vehicle Performance Considerations

While the CV system includes a wide array of technical requirements that are beyond the scope of this report, several high-level requirements warrant discussion. These are message accuracy, message validity, privacy, and system recovery.

A.3.1 Message Content Accuracy

A key function of the CV system is to provide operational information to road users to improve safety and mobility. Since the applications that provide these features depend directly in messages received from other vehicles and from roadside systems, the effectiveness of these applications depends directly on the accuracy of the content of the messages.

In general, the accuracy of message content depends directly on the accuracy of the sensors that provide this information. To assure system effectiveness, it is thus possible to simply specify minimum required levels of accuracy for the various data elements, and then rely on the manufacturers to meet these requirements. This approach assumes that devices (primarily vehicle devices) have not been altered or otherwise tampered with. This issue is discussed below.

A.3.2 Message Validity

Assuming the content of a message is accurate, there is a need to be able to determine that a message is valid. Specifically:

- The originator of the message has the authority and permissions to send it;
- The message is being sent in a location where the originator can operate;
- The message was not recorded at some other location and/or time, and then replayed at the current location;
- The message has not been altered in some way from when it was originated.

To provide the above features each message is digitally signed. The signature includes a selected portion of the original message (called a digest) which is then encrypted using a special type of cryptographic key (as will be explained below). The key is unique in that it can be used to encrypt the digest, but it cannot be used to decrypt it. This signature is then appended to the message together with a certificate. The certificate includes the public key required to decrypt the digest. The recipient of a signed message verifies that the message has not been falsified by generating a digest of the received message, and decrypting the digest in the signature using the key provided in the certificate. If the digests match, then the message has not been changed since it was signed. In addition, the signature usually also includes a time and location stamp, so that the recipient can verify that the message was sent at the current time, and location (if the signed message had been recorded somewhere else at another time, then the time and location in the signature would not match the current time and location). The certificate also includes the permissions for the sender of the message so that the recipient can determine, by examining the certificate, if that sender is authorized to send that particular type of message.

To assure that the certificate itself was not falsely created, it is also digitally signed, in this case by an organization that is well-known (specifically that is known by anyone receiving such a message). This entity is

known as a “certificate authority”. Each terminal thus must include a key for each certificate authority so that it can decrypt the certificate signature, verify the digest, and prove to itself that the certificate is valid.

A given message recipient may, at its discretion, choose to validate as much of any given message as they see fit (there is no requirement that a message be validated by the recipient). For example, since different messages may be sent with the same certificate, it may not be necessary to re-validate the certificate for every message signed using that certificate. Similarly, once a vehicle receives status information from another vehicle, it may not be necessary to validate the content of the message unless the content lies outside what the recipient would expect to see. For example, if a vehicle status message describes the vehicle as traveling at 60 mph (27 meters/sec) in a particular location, the recipient would expect the next message, sent 100 milliseconds later, to show the sending vehicle to have moved 2.7 meters. If the position provided in the next message was substantively different from this expected value, then the sender may be changing speed, or the message content could have been changed. A validation of the message will confirm which of these is the case.

A.3.3 Privacy

Since the CV system assumes that individual users (e.g., vehicles, users of personal devices) will transmit operational status messages, there is a significant concern about privacy. Specifically, since many of the messages are transmitted regularly and include position and speed data, they could be used to track an individual vehicle and determine the path it had taken, or to identify and cite vehicles that are violating traffic laws without physically observing the violation (for example by placing a receiver along the roadside and simply recording the messages any vehicle that is, for example exceeding the speed limit) *if they include any sort of identifying information*. In the first instance, the system would be violating the user’s privacy, and in the second, it would be using information obtained from the vehicle to accuse a driver. There have been a wide range of opinions on the seriousness of these potential issues, but the consensus has been to avoid them altogether by simply eliminating identifying information from broadcast messages.

Obviously for information exchanges that include identifying information, this data can be encrypted, so the privacy issue moves from the CV system itself to a matter between the two parties engaged in the encrypted exchange. Broadcast messages, however, are generally agreed to be anonymous to avoid infringing the privacy of the sender.

The combination of anonymity while also seeking to maintain message validity presents a challenge to the design of the CV security system. In most systems trust is based on identity. In the CV system, trust is based on certification of the terminal. This certification is only indirectly and cryptographically tied to the identity of the device owner, making it numerically difficult to determine the originator from any given message. In this way, it is possible to validate that a message was sent by an authorized terminal that had permissions to send that message at that location, but there is no practical way to determine who sent the message. This is a key feature of the CV system, and, as will be discussed below, it has significant impact on the overall security design.

As an initial matter, privacy can be assured by simply not providing any identifying information in the broadcast messages. Thus, for example the messages are designed to not contain the Vehicle Identification Number (VIN) or any other identifier that could be used to link the message to a particular vehicle and thence to the owner of the vehicle.

However, because the messages are also signed, it is also necessary to eliminate identifying information in the signature and certificate. This is more challenging, because the signature is formed and verified using a pair of keys that are unique to the vehicle that originated the message and each certificate has information that allows

it to be revoked if necessary (described further in the system recovery section). To avoid being able to track a vehicle based on its certificate, the CV system uses a scheme where each vehicle is furnished with a set of certificates. The vehicle terminal only uses each certificate for a short period of time. Because vehicles generally may not be in contact with one another for more than a few tens of seconds, it becomes very difficult to track a vehicle based on its certificates. For example, if the same certificate was always used, a malicious actor could place receivers strategically along the road way to possibly track a vehicle as it moved across town. In this same situation, if the vehicle periodically changed its certificates, noting the certificate used at one location would provide no ability to determine the movements of the vehicle by observing its certificates at other locations.

A.3.4 System Recovery

Because the vehicle population in the United States is large (>500M vehicles), the CV system represents an attractive target for mischief and attack. Attacks and misuse of the system may range from individual vehicle owners tampering with their vehicles so that they send out erroneous or fake messages, to more sophisticated attacks where traffic signal information is spoofed and sent out to create confusion among drivers at an intersection. Because it is impossible to accurately predict every possible attack that may emerge, the system is designed with internal mechanisms to recover from attacks by removing vehicles from the system. This is accomplished through a CRL. The CRL is periodically compiled by the PKI and distributed to all user devices in the system. When a message is received, one step in the verification process is to check the certificate identifier against the CRL. If that certificate is listed on the CRL, the message is ignored. Used in this way, a device with revoked certificates may still be able to transmit messages, but every message it sends will be ignored by all of the other recipients.

A.4 The Connected Vehicle Security Subsystem

To address the performance concerns described above, the CV system uses a fairly complex PKI. The elements of this are described below. Because of the scale of the auto industry, the CV PKI will be larger than any other PKI ever established. Part of this is because of the large number of vehicles on the road and part of this is due to the large number of certificates used. To assure privacy and non-trackability, each transmitting device cycles randomly through its cache of certificates. The current design uses about 1000 certificates which are replaced each year, meaning the PKI must manage about 5 trillion certificates per year. Added to this scale is the complex structure of the PKI, which is necessary to mitigate internal compromises, and to further prevent accidental identification or trackability of the vehicles. These elements are described briefly below, and in greater detail in the following chapters. Many of the descriptions and figures used in this section are extracted from the document “Technical Design of the Proof-of-Concept Security Credential Management System for V2X Communications” authored by the Collision Avoidance Metrics Partnership (CAMP).

A.4.1 Message Security Mechanisms

As described above, all WAVE Short Messages are signed using special certificates that do not contain any identifying information and that cannot be easily tied back to either the identity of the vehicle (or owner) or to the other certificates used by that vehicle. These certificates are known as pseudonym certificates. In use, the signing process is no different than was described for Bob and Alice above. The vehicle OBE generates a hash of the message, encrypts the hash with its private key, and then appends a certificate containing its public key, its permissions, and the CA signature over that information. A receiving vehicle generates a hash of the message using the same hash algorithm, decrypts the signature provided with the message using the public key in the certificate to obtain the hash created by the sending vehicle OBE, and compares the two

hashes. If they match, the message is verified. As a further check the receiving OBE may also verify the certificate by verifying each certificate in the chain of trust until it arrives at a CA that it trusts (i.e., the root CA).

The permissions associated with a certificate typically attest to the applications and regions of operation that the OBE is authorized to execute. This means that either the certificate must be updated if a new application or operating area is added, or each application must have its own set of certificates.

A.4.2 Message Privacy Mechanisms

Privacy and non-trackability on the road is accomplished as described above, by using pseudonym certificates that do not include any information that can be easily linked to the vehicle identity, and by cycling through certificates so that the vehicle cannot be easily tracked by tracking a single certificate identifier. Since the certificate changes, the vehicle essentially changes its identify randomly.

Assuring that the process of issuing certificates to the vehicle does not subvert the privacy gained by the use of pseudonym certificates is quite complicated, and results in a complex PKI.

The provisioning process starts when the device is manufactured. At that time, a Device Configuration Manager (DCM) provides keys and certificates that allow the OBE to trust SCMS components. These include certificates of some of the key SCMS elements, information about chains of trust, policies, and contact information for the elements of the SCMS in which the device will need to interact.

Once the OBE has this information, it is ready to be enrolled. During enrollment, the device receives an enrollment certificate that it can use in interactions with the SCMS, such as requesting application certificates.

Assuring the integrity of the DCM is critical to the overall integrity of the CV system. It is essentially the integrity of the OBE at this point that is being certified. It is thus essential that device firmware, and any other data to be injected to the device by the DCM, conform to the policies set forth by the SCMS Manager. This secure operational process involves some physical protection and assurance that only certified devices are provisioned with certified software.

This process is illustrated in Figure 9. The various other SCMS functions and components are further described in Chapter 2.

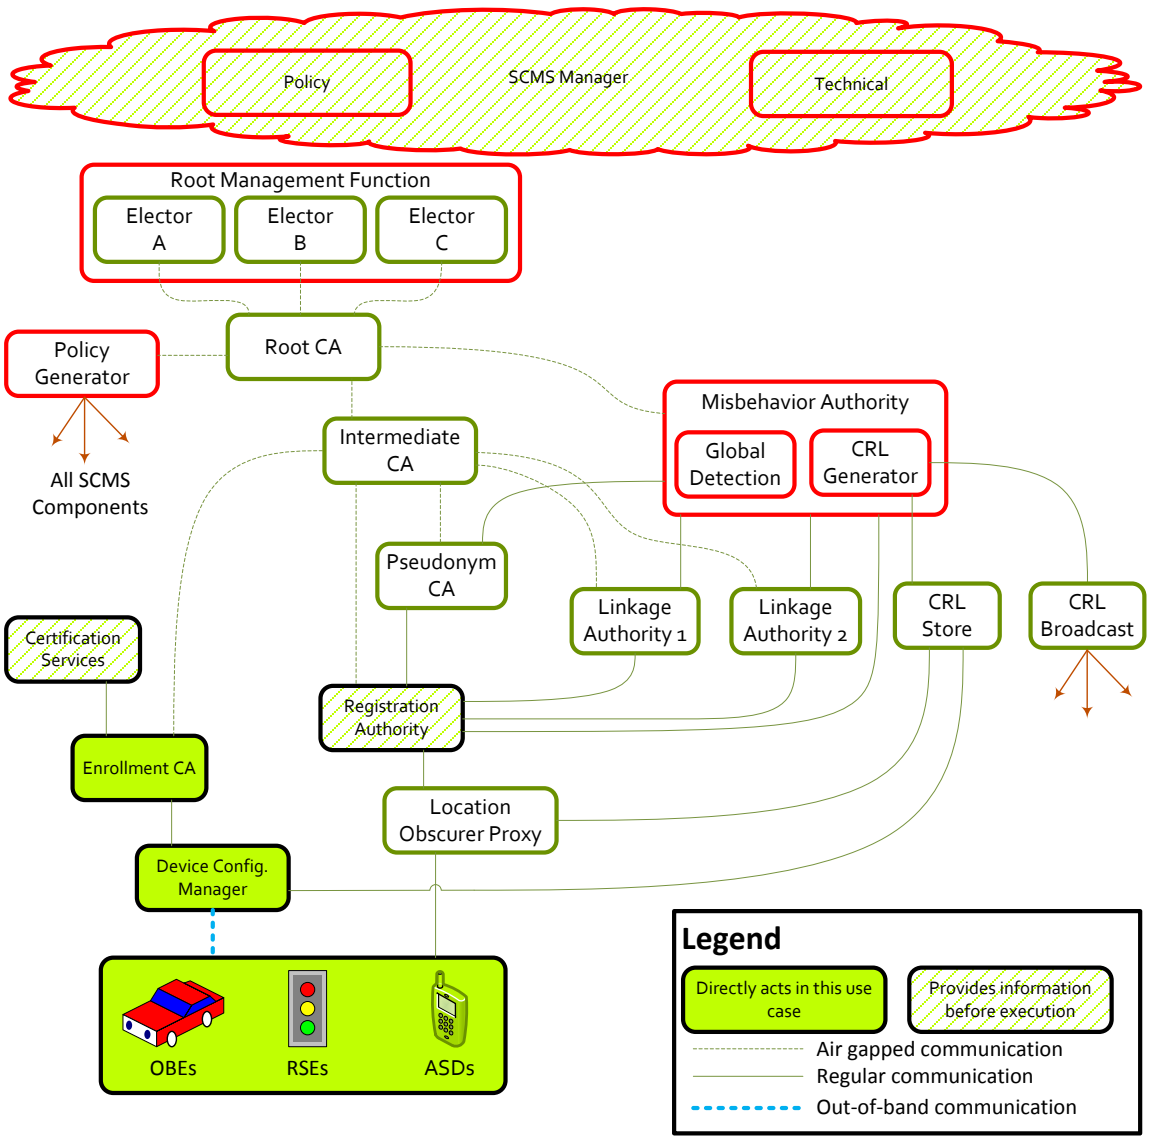

Figure 9: Device Provisioning and Enrollment within the SCMS
(Source: CAMP Technical Design of the PoC SCMS for V2X Communications)

The next step in the process is to provision the pseudonym certificates.
Using the information provided to it during the initialization and enrollment process, the OBE creates a certificate request, signs it with the enrollment certificate, encrypts the signed request for the RA and sends it to an entity known as the Location Obscurer Proxy (LOP). The LOP strips any information that could be used to determine the device's location, for example the identity of the communications access point the OBE is using to contact the SCMS via the internet, and forward the request to the Registration Authority (RA). This request uses a technique known as "butterfly keys" wherein the OBE generates its private keys, and creates a set of cryptographic values that the RA can use to generate the corresponding public keys. This approach avoids the need for the OBE to send separate requests for each key. Since the OBE will have roughly 1000 or so certificates for each application, and may have several applications that need certificates, it will have several thousand key pairs, and would otherwise need several thousand requests.

Upon receiving the certificate request, the RA checks to make sure that the certificate batch request is correct and authorized, and it then performs butterfly key expansion on the request to create a batch of public keys to be certified.

To assure that the certificates can be revoked at some future time if necessary, it is essential to create a mechanism for linking the certificates together. This is important because the messages are only identifiable by the certificate and, without this linkage, a malicious OBE might use one certificate to sign its malicious messages. When that certificate was revoked, it could move on to the next certificate, effectively drawing out the revocation period by 1000 times. By linking the certificates, it becomes possible to revoke all of them at once. This linkage must be cryptographically secure so that only the select components of the SCMS are able to determine what certificates are linked. This linkage is accomplished using a cryptographic technique to generate a linkage value for each certificate using the linkage seed value. Because the process is cryptographic, it is infeasible to determine that two linkage values are related unless one knows the linkage seed value. If one has the linkage seed value, it is a simple matter to generate the linkage values by repeatedly performing the cryptographic algorithm first on the linkage seed value and then on each subsequent linkage value to obtain the next linkage value. Thus, the linkage values for a batch of certificates are all related to one another, but the relationship can only be determined by generating the set starting with the linkage seed value. It is exceedingly difficult computationally to go the other way and derive one linkage value from another.

To assure that the entity generating the linkage value cannot undermine the integrity of the system by identifying how a set of certificates are related, the SCMS uses *two* Linkage Authorities. Each LA randomly creates a linkage seed, and from this it generates a pre-linkage value (PLV) using the cryptographic algorithm. Each LA then encrypts this PLV and passes it to the RA, who sends the two encrypted PLVs, one from each LA, together with a certificate request to the Pseudonym Certificate Authority (PCA). To avoid allowing the PCA to "see" which certificates are related to each other (because they are all requested together), the RA shuffles the requests for any given OBE with requests from other OBEs, so the PCA cannot determine which certificates belong to any single OBE. The PCA decrypts and combines the two PLVs to create a single linkage value (LV) for that certificate, signs the certificate, encrypts it using the requesting device's public key (which is in the certificate) and returns the signed and encrypted certificate to the RA, who then forwards them to the requesting device.

Using this scheme:

- Each LA knows the seed it used to generate its linkage value, but it does not know the seed value for the other LA, and this cannot determine the LV in the certificate. Because they are never sent to the LAs neither LA has any access to the certificates.
- The RA does not know the linkage seed values or the PLVs (because they were encrypted by the LA before being passed to the RA.
- The PCA knows both PLVs and the combined LV and can read the certificates, but it does not know which certificate goes with which vehicle, and does not know the linkage seed values.
- Since each certificate signed by the PCA is encrypted, the RA does not know anything about the signed certificate including the LV. Even though it knows which OBE to send the signed certificate, it cannot see that certificate, so it would have no way to link it when used by that particular OBE.

A.4.3 System Recovery Mechanisms

The primary recovery mechanism for misbehaving devices is to revoke the certificates of the device. The method to detect misbehaving devices is still in development. Based on research to date, it is presumed that

OBEs that are reporting incorrect data, or are otherwise sending erroneous or problematic messages will be reported to the Misbehavior Authority (MA), an entity charged with identifying misbehavior by OBEs, through some form of observation yet to be completely determined. This report would presumably include the certificate corresponding to the offending messages (i.e., the one that was sent with the message when the misbehavior was observed).

To revoke a batch of certificates generated by the PKI as described above, the MA sends the certificate associated with the offending message to the RA, the PCA, and the LAs. Using historical information at the PCA and RA, the LAs determine the linkage seed values corresponding to the LV in the offending certificate.

This linkage seed is published by the MA in a CRL. When an OBE receives a CRL, it simply uses the same cryptographic function used by the LAs and the PCA (described above) to regenerate the linkage values that correspond to the received linkage seed value (moving forward in the chain of values). It then compares the linkage values it has generated from the CRL to that contained in each received certificate. If there is a match, the certificate and the message that carried it are disregarded.

In parallel, the MA informs the RA to blacklist the OBE enrollment certificate, so that it cannot request new certificates.

The PKI must also provide policies and mechanisms to recover from more comprehensive threats, such as the disclosure of private keys at any level in the chain of trust.

Acronyms

Table 14: Acronyms

Acronym	Definition
ANPRM	Advance Notice of Proposed Rule Making
ARC-IT	Architecture Reference for Cooperative and Intelligent Transportation
ASD	Aftermarket Safety Device
ASN.1	Abstract Syntax Notation. One
BSM	Basic Safety Message
CA	Certificate Authority
CA/B	CA/Browser
CAMP	Collision Avoidance Metrics Partnership
CCMS	Cooperative ITS Credentials Management System
CIA	Confidentiality, Integrity, and Availability
C-ITS	Cooperative Intelligent Transport Systems
CME	Certificate Management Entity
CNSS	Committee on National Security Systems
COC	Certification Operating Council
CP	Certificate Policy
CPS	Certification Practice Statement
CRL	Certificate Revocation List
CRLG	CRL Generator
CTL	Certificate Trust List
CV	Connected Vehicle
CVCS	Connected Vehicle Core System
CVE	Connected Vehicle Environment
CVRIA	Connected Vehicle Reference Implementation Architecture
DCM	Device Configuration Manager
DOS	Denial of Service
DDOS	Distributed Denial of Service
DSRC	Dedicated Short Range Communications
EC	European Commission
ECA	Enrollment Certificate Authority
ECDSA	Elliptic Curve Digital Signature Algorithm
EE	End Entity
ETSI	European Telecommunications Standards Institute
FBCA	Federal Bridge Certificate Authority

Acronym	Definition
FHWA	Federal Highway Administration
FIPS	Federal Information Processing Standard
FM	Frequency Modulation
GCCF	Global Certificate Chain File
GD	Global Detection
HTG	Harmonization Task Group
HTTPS	Hypertext Transfer Protocol
ICA	Intermediate Certificate Authority
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
ISP	Internet Service Provider
IT	Information Technology
ITS	Intelligent Transport Systems
LA	Linkage Authority
LOP	Location Obscurer Proxy
LV	Linkage Value
MA	Misbehavior Authority
MAC	Medium Access Control
MOU	Memorandum of Understanding
NHTSA	National Highway Traffic Safety Administration
NPRM	Notice of Proposed Rule Making
OBE	On-board Equipment
OBU	On-board Unit
OCSP	Online Certificate Status Protocol
OEM	Original Equipment Manufacturer
PA	Policy Authority
PCA	Pseudonym Certificate Authority
PKI	Public Key Infrastructure
PLV	Pre-Linkage Value
POC	Proof of Concept
PSID	Provider Service Identifier
PSM	Personal Safety Message
QA	Quality Assurance
RA	Registration Authority
RFC	Request for Comments
RFP	Request for Proposal
RSE	Roadside Equipment
RSU	Roadside Unit
SCMS	Security Credential Management System
SPAT	Signal Phase and Timing

Acronym	Definition
TCP	Transmission Control Protocol
TIM	Traveler Information Message
TLS	Transportation Layer Security
USDOT	United States Department of Transportation
V2I	Vehicle-to-Infrastructure
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-Everything
VIN	Vehicle Identification Number
VPN	Virtual Private Network
WAVE	Wireless Access in Vehicular Environment
WSM	Wireless Access in Vehicular Environment Short Message

References

Crash Avoidance Metrics Partnership (CAMP). (November 15, 2016). Security Credential Management System Proof–of–Concept Implementation EE Requirements and Specifications Supporting SCMS Software Release 1.2.1. Federal Highway Administration (FHWA), USDOT.

Crash Avoidance Metrics Partnership (CAMP). (December 12, 2016). Technical Design of the Proof-of-Concept Security Credential Management System for V2X Communications. Federal Highway Administration (FHWA), USDOT.

European Telecommunications Standards Institute (ETSI). (September 17, 2010) Intelligent Transport Systems (ITS); Communications Architecture. European Standard (EN) 302 665.

IP.Com. (March 1, 2016.) Elector-Based Root Management System to Manage a Public Key Infrastructure. IPCOM000245336D. IP.Com.

USDOT. (December 13, 2016). Advanced Notice of Proposed Rulemaking (ANPRM) Federal Motor Vehicle Safety Standards; V2V Communications. NHTSA-2016-0126. National Highway Traffic Safety Administration (NHTSA), USDOT.

USDOT. Architecture Reference for Cooperative and Intelligent Transportation. (Last updated September 19, 2017). https://local.iteris.com/arc-it/index.html. Federal Highway Administration (FHWA), USDOT.

USDOT. (September 2017). Fundamental Principles and Research of the Security Credential Management System Proof–of–Concept. Presentation. Federal Highway Administration (FHWA), USDOT.

USDOT. (July 2016). Guidance Summary for Connected Vehicle Pilot Site Deployers Security Operational Concept. Final Report- FHWA-JPO-16-338. Federal Highway Administration (FHWA), USDOT.

USDOT. (August 31, 2017). Initial Set of Security Credential Management System Proof–of–Concept Governmental Management Policies and Organizational Documents. Final Report. Federal Highway Administration (FHWA), USDOT.

USDOT. (August 18, 2014). Notice of Proposed Rulemaking (NPRM) Federal Motor Vehicle Safety Standards; V2V Communications. NHTSA–2014–0022. National Highway Traffic Safety Administration (NHTSA), USDOT.

USDOT. (October 23, 2013). Organizational and Operational Models for the Security Credentials Management System (SCMS), Industry Governance Models, Privacy Analysis, and Cost Updates. Draft Revision Report. Federal Highway Administration (FHWA), USDOT.

USDOT. (June 1, 2017). Security Credential Management System Proof–of–Concept Governmental Management Concept of Operations (ConOps). Draft Report. Federal Highway Administration (FHWA), USDOT.

USDOT. (June 30, 2017). Security Credential Management System Proof–of–Concept Root Access Management Policies and Protocols. Draft Report. Federal Highway Administration (FHWA), USDOT.

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-18-688

U.S. Department of Transportation