FINAL REPORT:
Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Infrastructure

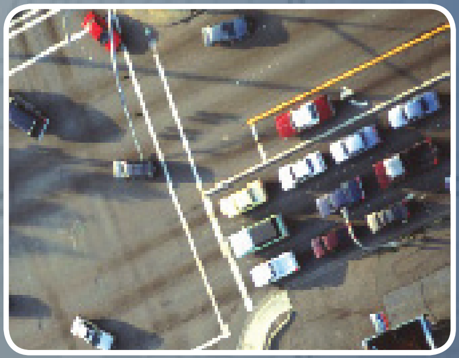

VOLUME 2B

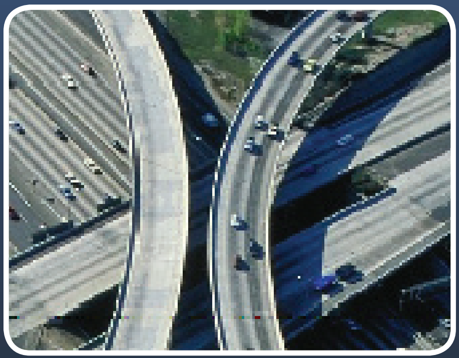

U.S. Department of Transportation
Research and Innovative Technology Administration

Notice

The U.S. Department of Transportation provides high-quality information to serve Government, industry, and the public in a manner that promotes public understanding. Standards and policies are used to ensure and maximize the quality, objectivity, utility, and integrity of its information. US DOT periodically reviews quality issues and adjusts its programs and processes to ensure continuous quality improvements.

Technical Report Documentation Page

1. Report No.

2. Government Accession No.

3. Recipient's Catalog No.

4. Title and Subtitle

Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Infrastructure

5. Report Date

February 2009

6. Performing Organization Code

7. Author(s)

Ram Kandarpa, Mujib Chenzaie, Justin Anderson, Jim Marousek, Tim Weil, Frank Perry, Ian Schworer, Joe Beal, Chris Anderson

8. Performing Organization Report No.

9. Performing Organization Name and Address

Booz Allen Hamilton
8283 Greensboro Drive
McLean, VA 22102

10. Work Unit No. (TRAIS)

11. Contract or Grant No.

12. Sponsoring Agency Name and Address

Research and Innovative Technology Administration (RITA)
U.S. Department of Transportation
1200 New Jersey Avenue, SE
Washington, DC 20590

13. Type of Report and Period Covered

14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

In 2005, the US Department of Transportation (DOT) initiated a program to develop and test a 5.9GHz-based Vehicle Infrastructure Integration (VII) proof of concept (POC). The POC was implemented in the northwest suburbs of Detroit, Michigan. This report describes the overall approach undertaken to prove the VII concept through a structured testing program and describes the overall experimental design utilized in proving the VII concept, by providing an overview of the system architecture and the design of systems, subsystems, and components, as well as the public sector applications developed to prove some of the system concepts. It outlines the purpose and procedures for various tests, identifies the test articles, and documents the results of that testing. It also discusses the implications of those test results relative to the overall viability of the VII concept and identifies recommendations for future work, including if and how designs and standards may need to be modified.

17. Key Word

Vehicle Infrastructure Integration (VII), Connected Vehicle, Cooperative Vehicle Infrastructure, IntelliDrive^{SM}

18. Distribution Statement

19. Security Classif. (of this report)

20. Security Classif. (of this page)

21. No. of Pages

22. Price

TABLE OF CONTENTS

1.0 VEHICLE INFRASTRUCTURE INTEGRATION PROOF OF CONCEPT TECHNICAL OVERVIEW..............1

1.1 Concept of Operations 1

1.2 POC System Architecture Description 2

1.3 Dedicated Short-Range Communications 4

1.4 Security Subsystem 6

1.4.1 Security Subsystem Objectives 7

1.4.2 Security Architecture 7

2.0 INFRASTRUCTURE NETWORK DESCRIPTION 9

2.1 Service Delivery Node 11

2.1.1 SDN Architecture 11

2.1.2 SDN Interfaces 14

2.1.3 SDN NAP 15

2.1.4 Physical SDN Setup 16

2.2 ENOC 18

2.2.1 Managed Entity 19

2.3 Roadside Equipment (RSE) 22

2.3.1 Electrical/Mechanical Architecture 24

2.3.2 RSE Processing Unit 27

2.3.3 RSE Software Architecture 28

2.4 Infrastructure Services 38

2.4.1 Communications Service 39

2.4.2 ILS 40

2.4.3 AMDS 42

2.4.4 PDS 46

2.4.5 MEDS 49

2.4.6 IdAM/Security Services 52

2.4.7 NMS 52

2.4.8 POS 52

2.5 Security Subsystem 54

2.5.1 IdAM 54

2.5.2 Security Services 55

2.5.3 Certificate Authority 56

2.6 DTE 58

2.6.1 Physical Layout/Map 58

2.6.2 RSE Installations 61

2.6.3 Backhaul Distribution 63

2.6.4 Backbone Network 64

3.0 PUBLIC APPLICATIONS 65

3.1 Signal Timing Optimization 66

3.1.1 Signal Timing Optimization Application Objectives 66

3.1.2 Signal Timing Optimization Overall Architecture 66
3.1.3 Signal Timing Optimization Flow of Events 67
3.2 Ramp Metering 67
3.2.1 Ramp-Metering Application Objectives 67
3.2.2 Ramp-Metering Overall Architecture 68
3.2.3 Ramp-Metering Flow of Events 69
3.3 Traveler Information 69
3.3.1 Traveler Information Application Objectives and Use Cases 70
3.3.2 Traveler Information Overall Architecture 71
3.4 Weather Information 71
3.4.1 Weather Information Application Objectives 72
3.4.2 Weather Information Application Overall Architecture 72
3.5 Corridor Management Planning Assistance (CMPA) 73
3.5.1 CMPA Application Objectives 73
3.5.2 CMPA Application Overall Architecture 73
3.6 Corridor Management Load Balancing (CMLB) 74
3.6.1 CMLB Application Objectives 74
3.6.2 CMLB Application Overall Architecture 75
3.7 Public Applications Developed and Tested in the POC 75
4.0 SUMMARY 77
APPENDIX A: ACRONYMS A-1

Table of Figures

Figure 1-1: POC VII System Architecture 2
Figure 1-2: Overall System Structure 3
Figure 1-3: DSRC Channel Management Concept 6
Figure 1-4: Identity and Access Management Conceptual Architecture 8
Figure 2-1: Network Subsystem 9
Figure 2-2: Network Access Point 10
Figure 2-3: Overall VII Network 11
Figure 2-4 : SDN Architecture 12
Figure 2-5: SDN Software Architecture 14
Figure 2-6: SDN Interfaces 14
Figure 2-7: NAP 15
Figure 2-8: Herndon SDN 17
Figure 2-9: Herndon SDN Close-up 17
Figure 2-10: ATT OptiMan Node 18
Figure 2-11: Detroit SDN 18
Figure 2-12: ENOC Workstation 19
Figure 2-13: Fault Management Screen 20
Figure 2-14: Network Traffic Monitoring Screen 21
Figure 2-15: RSE Status Monitoring Screen 21
Figure 2-16: RSE Interfaces 23
Figure 2-17: RSE Assembly Mounted on Pole for Test 24

Figure 2-18: RSE Electrical Architecture........25
Figure 2-19: RSE Assembly Components........26
Figure 2-20: MCNU and Power Supply........26
Figure 2-21: RSE Mounted to Pole in DTE........27
Figure 2-22: RSE Software Architecture........28
Figure 2-23: Proxy Manager Architecture........29
Figure 2-24: Radio Handler Architecture........31
Figure 2-25: DSRC/WAVE Radio Architecture........32
Figure 2-26: DSRC Radio Architecture........32
Figure 2-27: DSRC Radio Mini-PCI Card........33
Figure 2-28: WAVE Upper Layer Software Architecture........34
Figure 2-29: HMM Service........37
Figure 2-30: Publish and Subscribe Architecture........39
Figure 2-31: ILS Architecture........41
Figure 2-32: Large-Scale AMDS Architecture........43
Figure 2-33: Local AMDS Architecture........44
Figure 2-34: AMDS Proxy Components........46
Figure 2-35: PDS Architecture........47
Figure 2-36: Probe Data Proxy Architecture........48
Figure 2-37: Probe Data Message Structure........49
Figure 2-38: Probe Data Representation of a Freeway Interchange........50
Figure 2-39: MEDS Architecture........51
Figure 2-40: POS Architecture........53
Figure 2-41: VII Network Security Architecture........54
Figure 2-42: Arterial RSE Locations........58
Figure 2-43: Freeway RSE Locations........59
Figure 2-44: RSE Regions Used for PDC Test........60
Figure 2-45: RSE Regions Used for Advisory Message Delivery and Presentation Test........60
Figure 2-46: RSE Regions Used for Tolling Test........61
Figure 2-47: WiMAX RSE Installation........62
Figure 2-48: 3G RSE Installation........62
Figure 2-49: T1 RSE Installation........62
Figure 2-50: Typical RSE Field Installation........63
Figure 2-51: POC Backhaul Networks........64
Figure 3-1: Ramp Metering Application Architecture........68
Figure 3-2: Traveler Information Application Architecture........71
Figure 3-3: CMPA Application Architecture........74

1.0 VEHICLE INFRASTRUCTURE INTEGRATION PROOF OF CONCEPT TECHNICAL OVERVIEW

1.1 Concept of Operations

As part of the overall Vehicle Infrastructure Integration (VII) program[1], various stakeholder groups developed a set of approximately 100 use cases. These use cases addressed the key issues that the stakeholders felt were relevant for the VII system. The use cases included safety, mobility, management, and commercial applications. In general, the use case descriptions did not fully articulate the use cases in the context of the system, but they did provide insight into the needs and priorities of the various stakeholders. From this set, 20 use cases expected to be available at the initial deployment of the system were identified and articulated in more detail. This initial set was selected because they exercised the breadth of the system.

To support these use cases, the following core system requirements were identified. These requirements address all of the basic needs of the expected use cases (if the system performs these core functions, then it is expected to be able to also perform the use cases, given a suitable application implementation). These core functions formed the basis for the proof-of-concept (POC) system design.

- The system delivers broadcast messages from network providers to on-board equipment (OBE) at specified geographic locations.
- The system delivers broadcast messages from local systems such as traffic signals or toll stations to OBEs at specified geographic locations.
- The system delivers broadcast messages between OBEs.
- The system collects data from OBEs and distributes topical information extracted from the data to network-connected subscribers.
- The system provides OBEs with access to remote private service providers, and this access can be carried over from one roadside equipment (RSE) to the next without disrupting the service.
- The system provides security functions to protect against attacks and to protect the privacy of the individual users.

Because developing and testing all 20 use cases would have been impracticable, the POC program identified a subset of use cases that exercise the core functions described above. These were then implemented and tested in ways to assess both the functionality of the system and the baseline performance, under the assumption that the system would provide these core functions in the same way regardless of the specific details of the application.

In some cases (e.g., several of the safety applications), the use cases were scaled back to allow assessment of key architectural and system aspects without requiring development of a full-blown application.

[1] The program name, Vehicle Infrastructure Integration (VII), was in official use at the inception of and during the execution of the work described in this report. The US Department of Transportation has initiated a new program titled "IntelliDrive^{sm}," which now encompasses all activities that were previously part of VII.

1.2 POC System Architecture Description

The system includes mobile terminals that, in the POC, were typically installed in vehicles. In the POC, these units are known as OBE (the reader is encouraged to read VIIC document "Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle" for more information). OBEs send and receive messages between each other (for vehicle-to-vehicle applications) and exchange messages with stationary roadside terminals known as RSE. The link between OBEs and between OBEs and RSEs is the dedicated short-range communications (DSRC) radio system. The RSEs are connected to and remotely managed from a Service Delivery Node (SDN) and an Enterprise Network Operations Center (ENOC). The SDN provides a variety of services that are described in more detail in subsequent sections.

Figure 1-1 presents an overview of the POC system architecture.

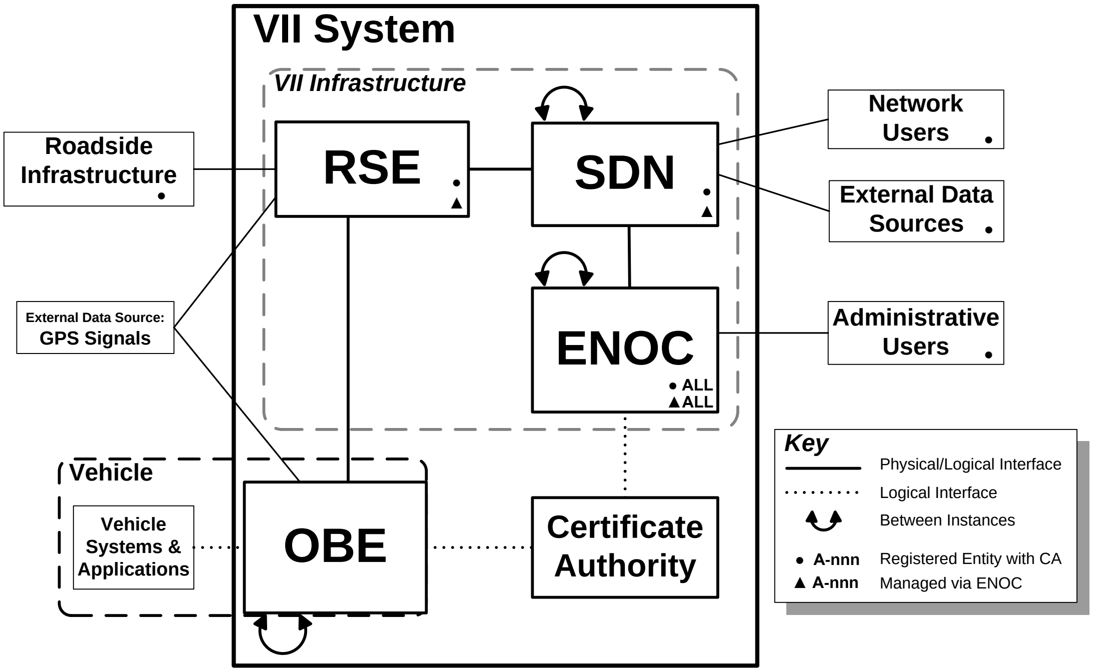

Figure 1-1: POC VII System Architecture

A critical aspect of the VII architecture is the management of scale. The system needs to be designed to support 100-percent vehicle deployment, which translates to just over 200 million vehicles. In operation, this means that applications such as probe data collection may be handling tens of millions of messages per second, across the entire network. Similarly, management of, for example, security credentials may require rapidly sending thousands of pieces of data to hundreds of millions of cars. At the other performance extreme, the system must allow a single non-vehicular user to post a warning sign in the vicinity of a particular road hazard.

While the POC was not expected to operate at these scales, the eventual requirement for scalability meant that the basic architecture needed to be structured in a manner that could scale without significant redesign. As a result, the POC implementation is a small-scale version of a fully scalable system.

To manage these large-scale extremes, the system uses a tiered architecture, as shown in Figure 1-2. As can be seen in the figure, any given RSE needs to be capable of interacting with up to 250 vehicles at any given time. This is determined primarily by the number of vehicles that can fit inside a typical radio frequency (RF) footprint, which provides a range of approximately 250 meters.

Each RSE is connected to a regional SDN via a backhaul link, and each SDN is connected to all other SDNs via a wideband backbone network. Using this architecture, any RSE is accessible from any SDN, and this is a key feature of the scalability of the system, since any user connecting to the local SDN can interact with any RSE.

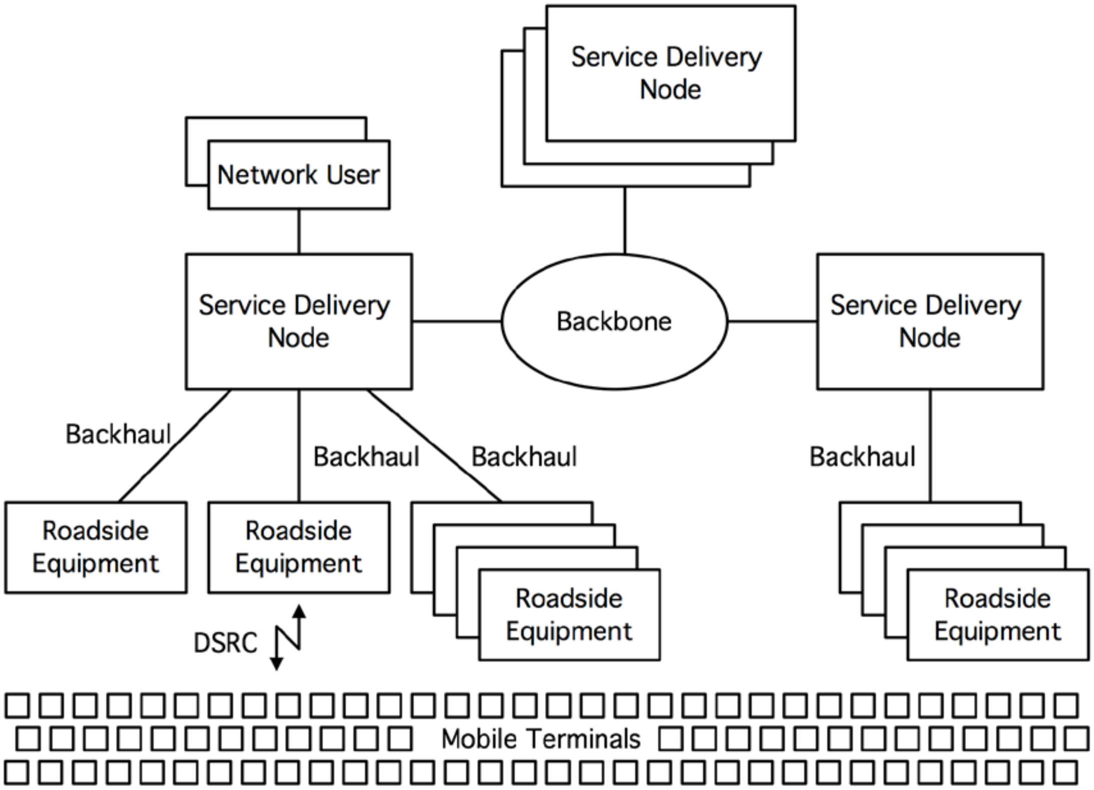

Figure 1-2: Overall System Structure

While the POC included only 55 RSEs, a fully deployed SDN would be expected to support about 1,000 to 2,000 RSEs, so there would be between 100 and 200 SDNs for a nationwide deployment. In the POC implementation, two SDNs were used, although the systems to support scaling were implemented to assure a valid assessment of performance.

The SDN provides a variety of services, but key to the discussion of scaling are the Advisory Message Distribution Service (AMDS), the Probe Data Collection (PDC) Service, the Probe Data Service (PDS), and the Network User Gateway.

The RSE is the terminal node (actually one of many terminal nodes) of the network. In the VII architecture, the RSE is intended to be located at roadway intersections and other strategic points in the road network in order to deliver important safety- and mobility-related messages and to provide a connection point for OBEs to execute various types of data transactions.

Because the services provided at any given RSE are generally not unique to that specific RSE (some safety services being an important exception), the RSE is arranged as a "pass through" terminal that extends SDN services out to the edges of the network. This architecture greatly simplifies the operation of the RSE, since one can store, for example, an advisory message at an SDN, and the system will replicate that message out to the appropriate RSEs according to the delivery instructions. The alternative would be to have every network user interact directly with every RSE, and this approach would be unmanageable at large-scale levels.

To illustrate this management of scale, the AMDS serves as a link between network users who have advisory messages for distribution and the entire installed base of RSEs (and therefore the entire OBE population). It is expected that under full-scale deployment, the number of signage providers would be less than about 10,000, so for signage the system allows 10,000 providers to efficiently interact with 200,000 RSEs and deliver signage messages to 200 million OBEs. In the reverse direction, the PDC would collect probe data from all of the RSEs attached to the SDN. This data is parsed into data "topics," and then data for any given topic is distributed to network users who have subscribed to that particular topic. A typical topic might be instantaneous road speed at a particular location on a particular road. This scheme would allow the full-scale system to collect vast amounts of data from vehicles on the road, and sub-divide the data passing only those parts of interest to any given subscriber. It is expected that under such a deployment, there may be about 10,000 to 100,000 probe data subscribers, so as with AMDS, the system effectively scales from over 200 million vehicles (generating roughly 50 GBps) to about 50,000 users of this data. The SDN also provides a straightforward routed communications system that links vehicle users with private service providers in either one-to-one or one-to-many relationships. In this role, the SDN is effectively like a mobile Internet system linking users to Web service centers such as navigation providers.

The POC implementation of the system included 55 RSEs placed at various locations in the northwestern Detroit suburbs. These RSEs were linked to the SDN using a variety of different backhaul technologies. One SDN was located in Novi, Michigan, and the other was located in Herndon, Virginia. The Herndon facility also included an ENOC to support system management function, and the certificate authority required to support security functions. The POC implementation was thus a minimalist version of the national system architecture, and this allowed the program to test all aspects of operation with a small-scale implementation.

1.3 Dedicated Short-Range Communications

The 75 MHz band in the 5.9 GHz frequency range allocated by the Federal Communications Commission (FCC) offers significant data transfer capacity. However, to make use of this spectrum in a mobile environment required development of new communications protocols. The core radio protocol used is based on the well-known 802.11a/b/g wireless Ethernet standard, often referred to as Wi-Fi. Because of the unique mobile environment, the 802.11a standard was modified to allow what is known as an "association-less" protocol, identified as 802.11p. This means that the system does not establish a conventional network with all of the mobile terminals as nodes, maintaining network-level "state" information about each other. The reason this is not done is that the mobile terminals (OBEs in the POC) are entering and leaving the hot spot rapidly, and there is insufficient time available to set up a new network identity for each new arrival and inform all other nodes in the network before the network changes again because a terminal has left or a new one has arrived. On the surface, this approach may seem to limit the functionality of the system since it means that any given mobile terminal cannot interact uniquely with another terminal (the way computers on an office network might), but this is not the case. Because the system is radio based, all terminals can hear all messages sent. So, under most circumstances, one can simply broadcast a message in the local area, and all terminals (OBEs and RSEs) can receive it, thus there is no need to establish a low-level network identity.

The higher levels of the protocol are defined in a suite of standards known as IEEE 1609 Wireless Access in Vehicular Environments (WAVE). This suite addresses security (1609.2), networking and messaging (1609.3), and channel management (1609.4). In particular, 1609.3 defines a WAVE Short Message Protocol (WSMP) that allows a simple way for a terminal to send messages in the local vicinity. WSMP allows for direct message addressing based on the Medium Access Control (MAC) address of the intended recipient; however, in practice, most WAVE Short Messages (WSMs) are broadcast and, therefore, are not addressed to any specific recipient.

The current DSRC standards divide the 75 MHz spectrum into 10 MHz channels. This allows RSEs in local proximity of each other to provide services without causing interference. It also allows for use of existing commercial 802.11 radio components. Since it is critical for safety reasons to ensure that all terminals can hear each other, and the standards developers did not want to assume the use of multiple radio receiver systems (or very wide band receiver systems), a method for channel management was developed and described in 1609.3 and 1609.4. The approach separates terminal operations into two modes—“provider mode” and “user mode”—and splits the use of channels into two time intervals called the Control Channel (CCH) interval and the Service Channel (SCH) interval, as shown in Figure 1-3. In the current standard, these intervals are set to 50 milliseconds (msec) each, although the specific durations could be modified. All terminals are required to monitor the CCH during the CCH interval. In provider mode, the terminal transmits a WAVE Service Advertisement (WSA) on the CCH during the CCH interval, and since all terminals are monitoring this channel at that time, they all receive the WSA. The WSA contains a list of the services that the provider (typically an RSE) will provide during the SCH interval, along with the SCH channel number that they will be using. The services are identified by a code number known as a Provider Service Identifier (PSID). If a terminal in user mode receives a WSA that contains a PSID of interest (e.g., a message associated with an application that is active on that terminal), the terminal will switch to the appropriate SCH during the SCH interval and will make use of that service.

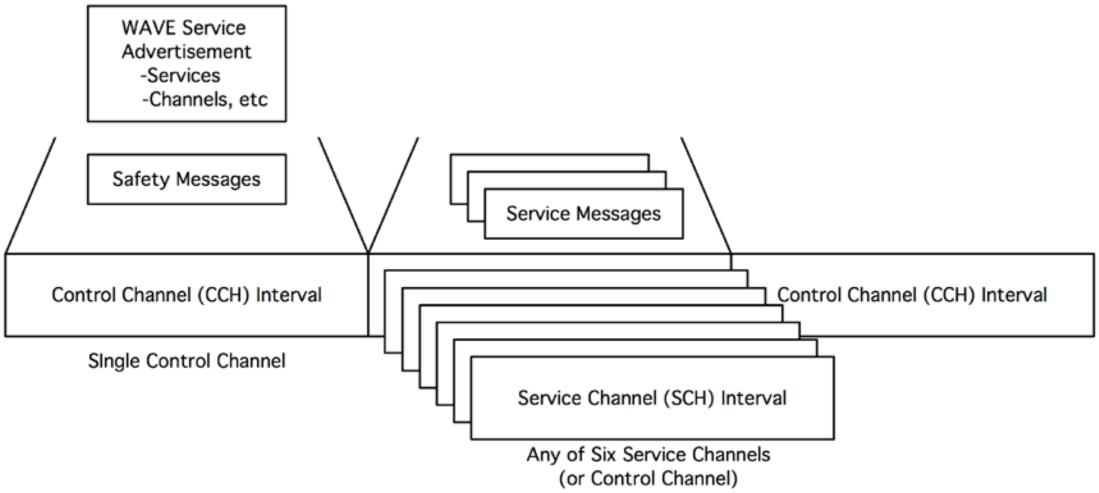
WAVE Service Advertisement
-Services
-Channels, etc
Safety Messages
Service Messages
Control Channel (CCH) Interval
SIngle Control Channel
Control Channel (CCH) Interval
Service Channel (SCH) Interval
Any of Six Service Channels
(or Control Channel)

Figure 1-3: DSRC Channel Management Concept

Because all terminals are required to monitor the CCH during the CCH interval, all high-priority safety messages are sent on the CCH during the CCH interval.

All low-priority services and all services using Internet Protocol (IP) are restricted to only use the SCH during the SCH interval. The result of this method is that all terminals have a high probability of receiving important messages, and less-important message traffic is distributed across the other channels, thereby reducing congestion.

IP transactions typically require some form of network setup, and, as described above, the DSRC protocol does not establish this. To support this type of traffic, the WSA also contains the IP address of the provider. In general, the standard does not describe the use of IP between OBEs because, in general, OBE-to-OBE messaging is safety related and will always use WSMP on the CCH. This avoids any issues with OBEs needing to route packets (although this sort of usage is not prohibited; it is just not defined in the standards). Once a user terminal has acquired the RSE IP address, it can then create its own IP address (using IPv6) rules and send IP packets to remote service providers. The RSE routes these packets through the backhaul network to the SDN, through the network gateway to a Virtual Private Network (VPN) via the Internet or a dedicated circuit, and then to the service provider.

While somewhat more complex than typical protocols, DSRC achieves the unusual feat of administering communications resources in real time in a way that assures that critical safety messages will have top priority, but that also allows lower-priority messages, both local messages and messages bound for distant servers, to simultaneously use the system.

1.4 Security Subsystem

The VII security subsystem is a complex set of functions and services that operate in parallel with the other elements of the system to ensure safe and verifiable system behavior and to prevent misuse of and attacks on the system.

1.4.1 Security Subsystem Objectives

The VII security subsystem is aimed at ensuring three basic objectives—privacy, authenticity, and robustness. The basic structure of the security subsystem is designed to provide assurance relative to the confidentiality of private message traffic, the authenticity of public message traffic, and the anonymity of private generators of public messages.

1.4.1.1 Privacy

Privacy is addressed through two fundamental aspects in the VII security subsystem, namely assurance of anonymity and confidentiality. While service providers outside the system may need to know the identity of a specific OBE, the VII system itself does not need to know this information. The system has been specifically designed to avoid requiring any form of traceable or persistent identification of any POC OBE. In addition, when identifying information is passed through the system to trusted service providers, the system provides mechanisms to encrypt this information so that none of the system elements or operators can access it. Finally, these encryption schemes are also used to suppress the opportunity for observers to correlate operational information (e.g., vehicle speed information) with physical observation, so the system also protects against misuse by external attackers.

1.4.1.2 Authenticity

In any system, it is desirable to require users to prove their authorization to access and/or use the system's resources. The VII system is unique since, for the POC OBE, this authentication must be accomplished without violating the mobile user's privacy. The VII security subsystem provides a sophisticated means for validating an OBE's legitimacy without identifying that OBE. This approach assures users that information provided by the system is legitimate and truthful and allows the system to prevent access to it by users with no authorization or OBEs that appear to have been tampered with.

1.4.1.3 Robustness

It is inevitable that the system will be attacked. These attacks may be full scale sophisticated attempts to disrupt the system, or they may be small-scale pranks. In any case, the system must make it very difficult to mount an attack; it must be capable of identifying and terminating a severe attack in progress, and it must provide a means for rapid recovery of full capability following any actions to terminate the attack. This is often referred to as "defense in depth."

1.4.2 Security Architecture

The VII security subsystem is composed of two independent subsystems. The OBE (private mobile user) subsystem is responsible for securing messages sent to and from mobile terminals (in the POC, these were OBEs). This system is described in detail in the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.

The security functions for the infrastructure elements of the POC system are built around the Identity and Access Management (IdAM) subsystem. The IdAM subsystem provides VII system security by enforcing entity authentication, authorization, and account management. The IdAM subcomponents are distributed among the SDNs and the ENOC as an integrated solution of access management, security credential provisioning, and entitlement management technologies. The ENOC element contains subcomponents that provide entity lifecycle management, data storage, authorization, and authentication services, while the SDN element contains authentication, authorization, and read-only data storage services. The IdAM uses a certificate-based scheme for access and authentication and uses the concept of a registered entity, which is any element of the system that has and uses certificates for these purposes. Figure 1-4 conceptually illustrates the IdAM.

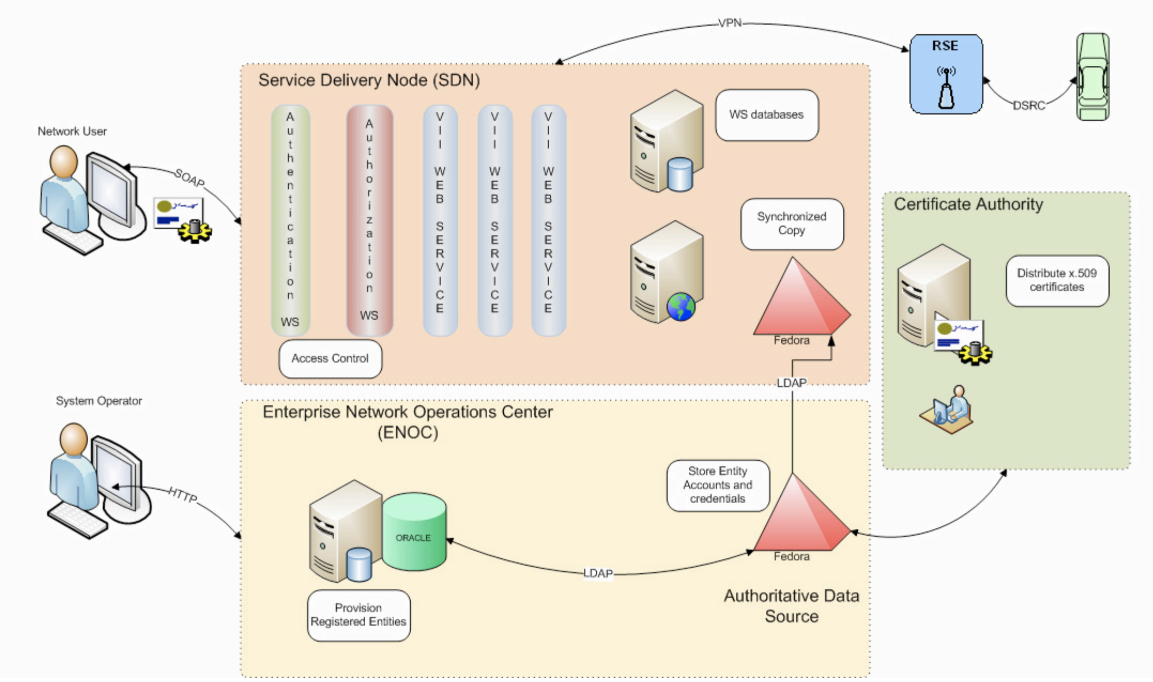

Figure 1-4: Identity and Access Management Conceptual Architecture

The IdAM subsystem is used to authenticate and authorize internal and external elements (all registered entities). Thus, each network subsystem established its VPN connections with other elements via the IdAM subsystem, and each external user (service provider, data subscriber, etc.) also uses the IdAM to access the system.

For the national deployment, there may be multiple regional ENOCs maintaining authoritative registered entity identity data. The SDN subsystems would contain local read-only copies of registered entity data that have been synchronized with the ENOC directories. For the POC, a single ENOC was deployed along with two SDN Subsystems. The ENOC synchronizes data with the local SDNs to replicate the data and allow for timely and accurate security transactions. This distributed architecture supports the geographically dispersed nature of the VII network and the large number of identities being managed by the system.

2.0 INFRASTRUCTURE NETWORK DESCRIPTION

Efficient and consistent delivery of VII services requires a robust, reliable infrastructure that interconnects the various components of the VII system. Vehicles must communicate with the roadside to reach services and applications; network users must be able to access data collected by the services; and the entire system must be managed and operated by an operating entity. The network subsystem is the communications infrastructure that enables data to be transported between RSE and users of the VII system. The Network Subsystem is expected to be a geographically distributed collection of interconnected nodes that provide localized and distributed network services, managed by a centralized ENOC, as seen in Figure 2-1.

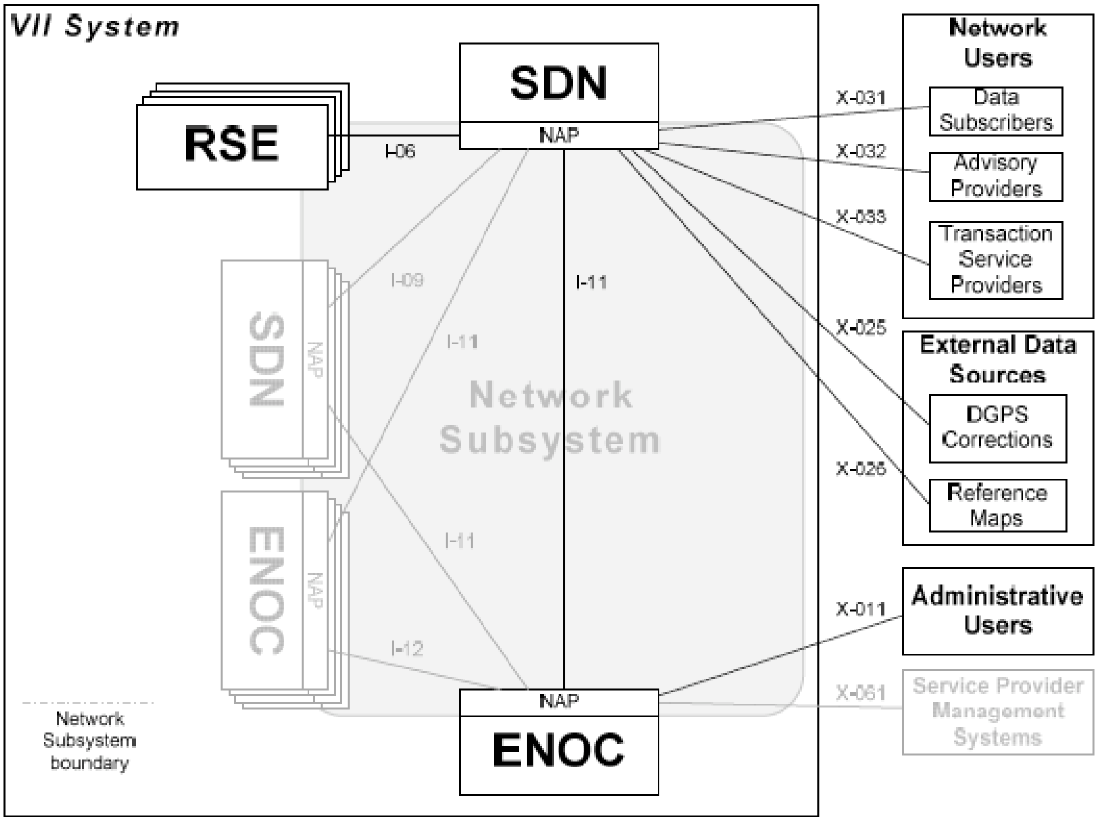

Figure 2-1: Network Subsystem

The SDN and ENOC nodes contain various servers that host applications and services for the VII system. The SDN servers host applications and services used by network users and external data sources, and the ENOC servers host applications and services used by administrative users. The term infrastructure servers refers to these two distinct sets of servers together.

The SDN also contains a Network Access Point (NAP) that supports routing of data packets to SDN services and to external services via an access gateway.

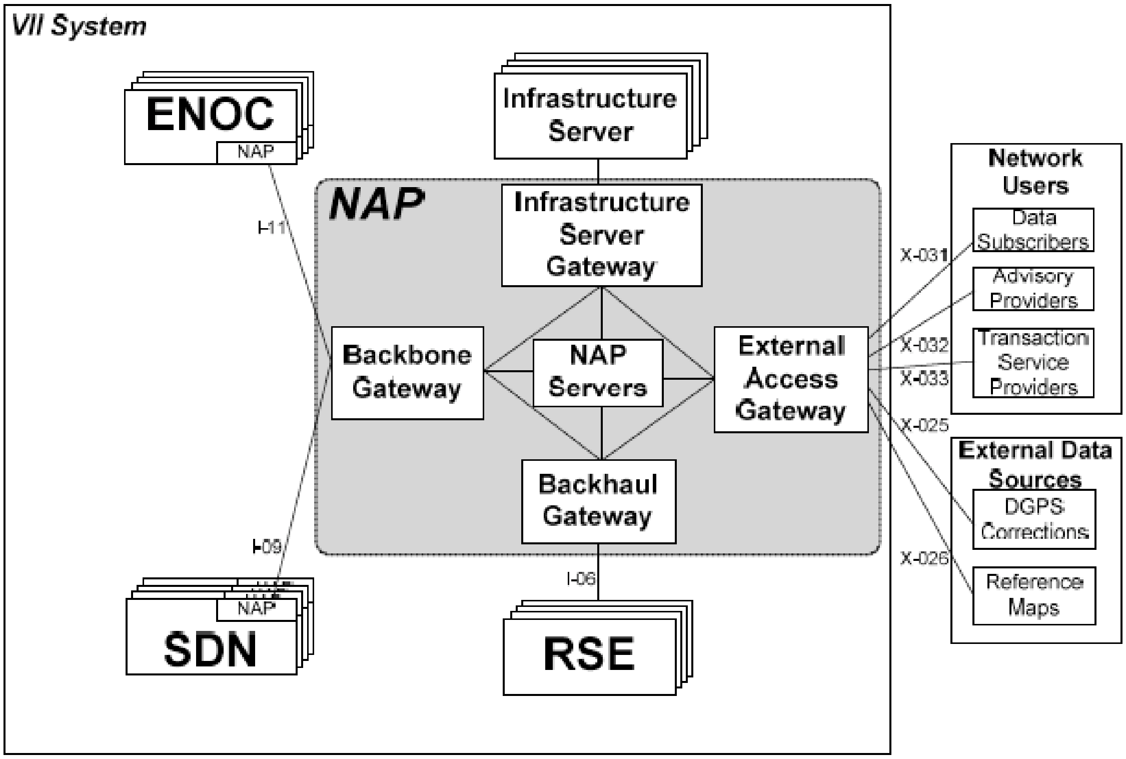

Figure 2-2: Network Access Point

As illustrated in Figure 2-2, the NAP is composed of interfaces to the backbone (to other SDNs), the backhaul (to RSEs), and the access gateway (to network users), routing functions to properly direct message traffic and a set of core services.

Figure 2-3 schematically illustrates the overall infrastructure network. While the POC system architecture supports an arbitrary number of SDN subsystems, the POC implementation is composed of two SDNs—one in Detroit, MI, and the other in Herndon, VA. These SDNs are connected to each other via a high-bandwidth data link known as the backbone network. The Herndon SDN was set up initially to test and validate the SDN design. It was used as the primary SDN to support the system. The Detroit SDN was set up to facilitate local operation of the RSEs and to validate the multi-SDN element of the architecture.

The SDNs are also constructed to allow a network user to access the system via any SDN, even if they are interacting with services and RSEs associated with another SDN. All of these interactions pass through the backbone network. As a result, a network user might access the system via the access gateway at a local SDN, and receive data from RSEs associated with a remote SDN through that same access gateway. Or the POC network users connect to the access gateway using the Internet.

The POC RSEs are connected to the Detroit SDN via three different backhaul links implemented using WiMAX, 3G cellular, and T1 lines, depending on the RSE location.

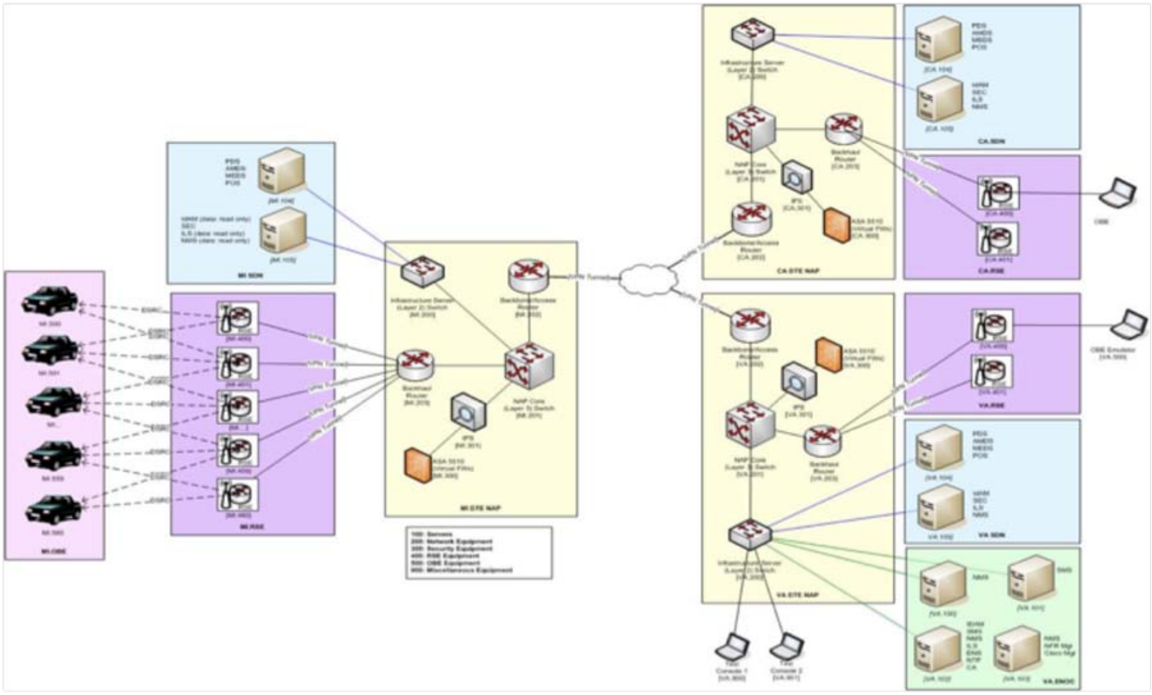

Figure 2-3: Overall VII Network

Both the Detroit and Herndon facilities also supported ENOC functions, which allowed operators and system administrators to control and configure the deployed RSEs, infrastructure components, and communications links and to manage the overall network and the deployed RSE systems.

The certificate authority issues security credentials to elements of the system that require them. It also manages the overall security state of the system.

The following sections describe the elements of the infrastructure network in detail.

2.1 Service Delivery Node

2.1.1 SDN Architecture

The SDN is a focal point of services and connectivity that supports a set of RSEs. The SDN includes management and security functions as well as a suite of VII-specific services.

Each SDN includes an NAP, which is a communications hub that facilitates and controls connectivity between VII system components and provides controlled access for network users. This NAP comprises multiple gateways consisting of routers, firewalls, and intrusion detection systems. Each gateway is dedicated to support the respective connectivity to the RSE subsystems associated with that SDN, and to remote RSEs via the backbone link to the SDN associated with the remote RSEs.

The routing functions are provided via normal IP routing mechanisms that pass packets received from RSEs (and in most cases originating at OBEs) to either the local SDN services or through an NAP gateway to network users. Most network users are expected to connect to the NAP via a conventional Internet connection.

Figure 2-4 illustrates the basic SDN architecture.

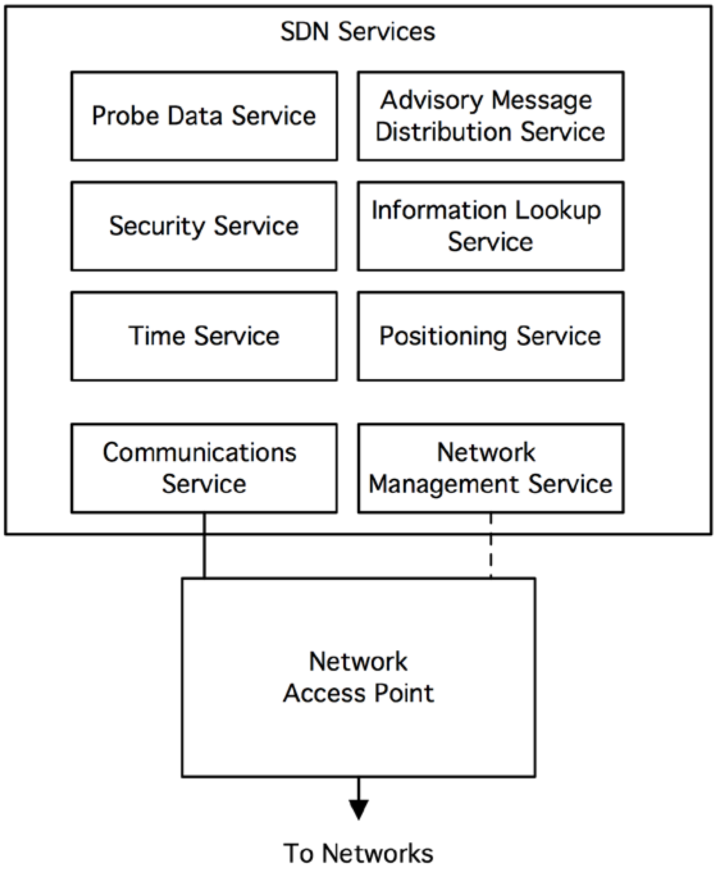

Figure 2-4 : SDN Architecture

To support the wide variety of services required by VII end user applications, while maintaining interoperability and flexibility, the SDN subsystem is configured as a service-oriented architecture (SOA). The SOA defines the use of loosely coupled services to support the requirements of the system processes and users. In an SOA environment, resources on a network are made available as services that can be accessed without knowledge of their underlying platform implementation. An SOA is not tied to a specific technology and may be implemented using a wide range of interoperability standards.

Figure 2-5 shows the SDN software architecture. This software system provides a variety of services that are extended out through the backhaul networks to the RSEs assigned to that SDN. These services are described in the larger system context in Section 2.4, and are briefly described below.

2.1.1.1 Communications Service

Communications services refer to services that provide data transport between network users and vehicles. Communications services include service classes, packet prioritization, quality of service, DNS look-up capability, address assignment, and a consistent time source throughout the network.

2.1.1.2 AMDS

The AMDS accepts formatted advisory messages submitted by network users via the access gateway. These messages include delivery instructions such as RSE ID(s), repeat timing, and message lifespan. The AMDS then passes these messages to the AMDS proxies resident in the appropriate RSEs for local broadcast to OBEs in the vicinity.

2.1.1.3 Probe Data Service (PDS)

The PDS collects anonymous probe data messages from all equipped OBEs and sends probe data content to network user probe data subscribers according to their subscriptions. The OBE's probe data function generates standards-based (SAE J2735) probe data structures and formats, effectively sanitizing OEM or vehicle specific information. Additionally, probe sequencing information is intermittently randomized to assist in maintaining vehicle anonymity.

2.1.1.4 Information Lookup Service (ILS)

The ILS is a support service used by network users to determine information about the system. It is most often used to identify RSEs according to location, so that a subscriber or provider can then properly reference the RSE.

2.1.1.5 Map Element Distribution Service (MEDS)

The MEDS distributes localized mapping data containing accurate roadway and lane geometries to vehicles and public service vehicles in support of lane-based safety applications. The MEDS can also be used to generate or validate existing maps using highly accurate vehicle positioning data extracted from local probe data messages.

2.1.1.6 Network Management Service (NMS)

The NMS collects, aggregates, and forwards network management data from managed network elements (MNEs) to the ENOC. The NMS also transmits network management commands from the ENOC to the MNEs. The MNEs managed by the NMS at the SDN include the RSEs connected to the SDN, the backhaul links, the backbone links, the access links, and the SDN gateways, as well as the infrastructure servers within the SDN.

2.1.1.7 Security Service

The VII Security Service provides “defense-in-depth” protection of the network hardware, software, and information components against hackers, unauthorized users, and physical or electronic sabotage. The service prevents, detects, isolates, mitigates, and corrects violations of VII system security. The VII Security Service also maintains network administration control of authorized user privileges and serves to protect the confidentiality of user information.

2.1.1.8 Positioning Service (POS)

The POS collects Global Positioning System (GPS) correction data from external sources and distributes it to the OBEs in the system via the RSEs. Using this approach, the OBE GPS receivers can obtain high-quality, location-specific correction data over the DSRC radio link, thereby eliminating the need for an additional GPS corrections receiver in the vehicle.

2.1.1.9 Time Service

Time Service distributes accurate time information to all system elements to facilitate management and synchronization.

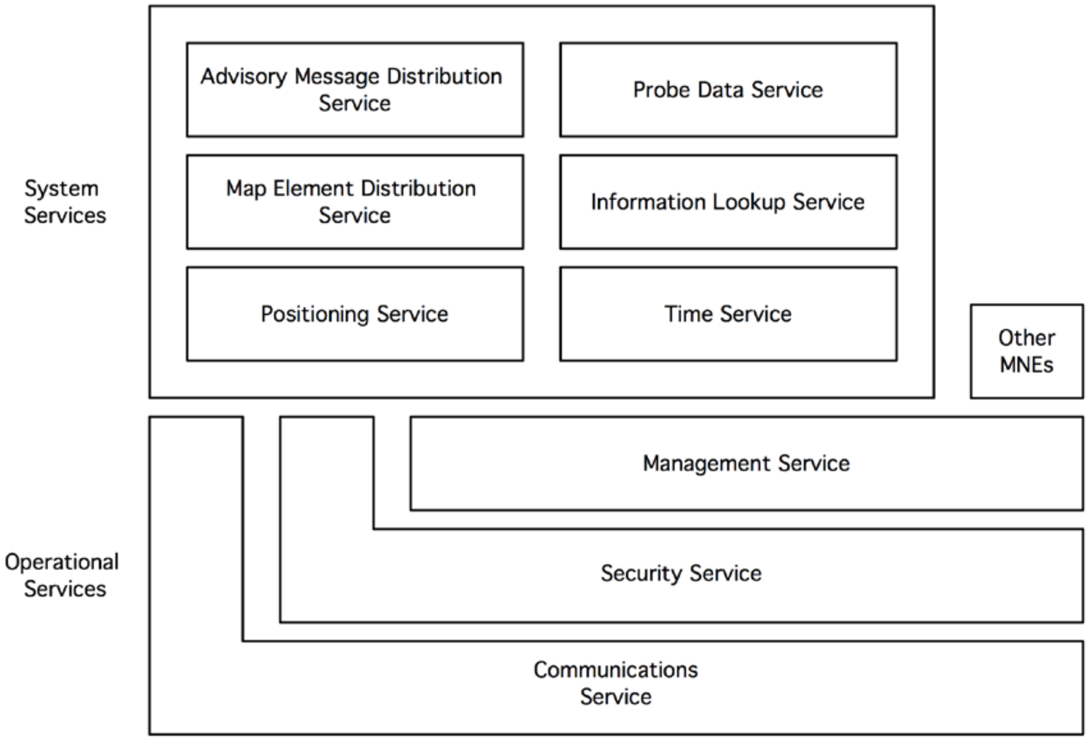

Figure 2-5: SDN Software Architecture

2.1.2 SDN Interfaces

Figure 2-6 illustrates the SDN interfaces. The SDN supports interfaces to network users via the access gateway, to the RSE functions via the backhaul interface, to the other SDNs via the backbone interface, and to the ENOC via a (typically locally implemented) management interface. Because of the backbone network, any SDN or RSE in the system may be managed from the ENOC through the SDN where the ENOC is located and then to the SDN or RSE via the backbone network.

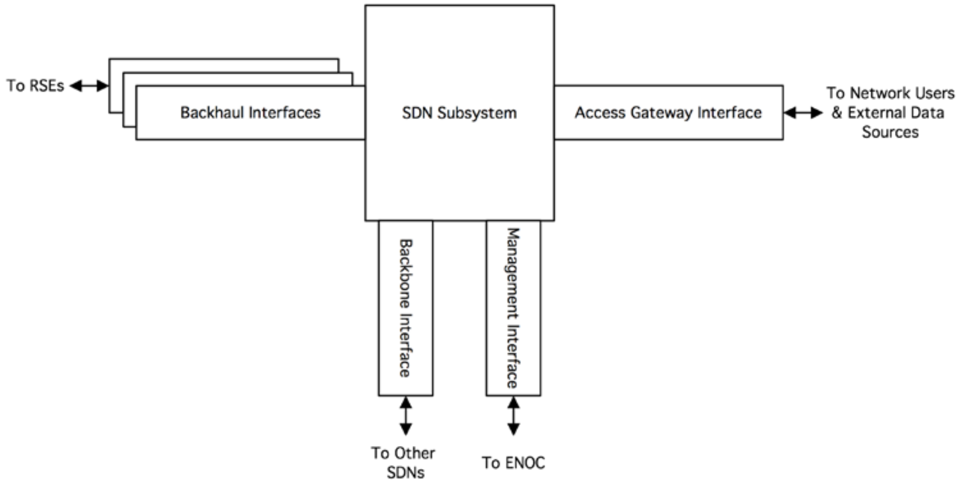

Figure 2-6: SDN Interfaces

2.1.3 SDN NAP

Key elements of the communications services are implemented by way of the NAP, which serves to link network users (via the access gateway) to the infrastructure servers (where the SDN system services are implemented), to the RSEs (via the backhaul gateway), and to other SDNs (via the backbone gateway). Figure 2-7 illustrates these NAP functions, which are described briefly below.

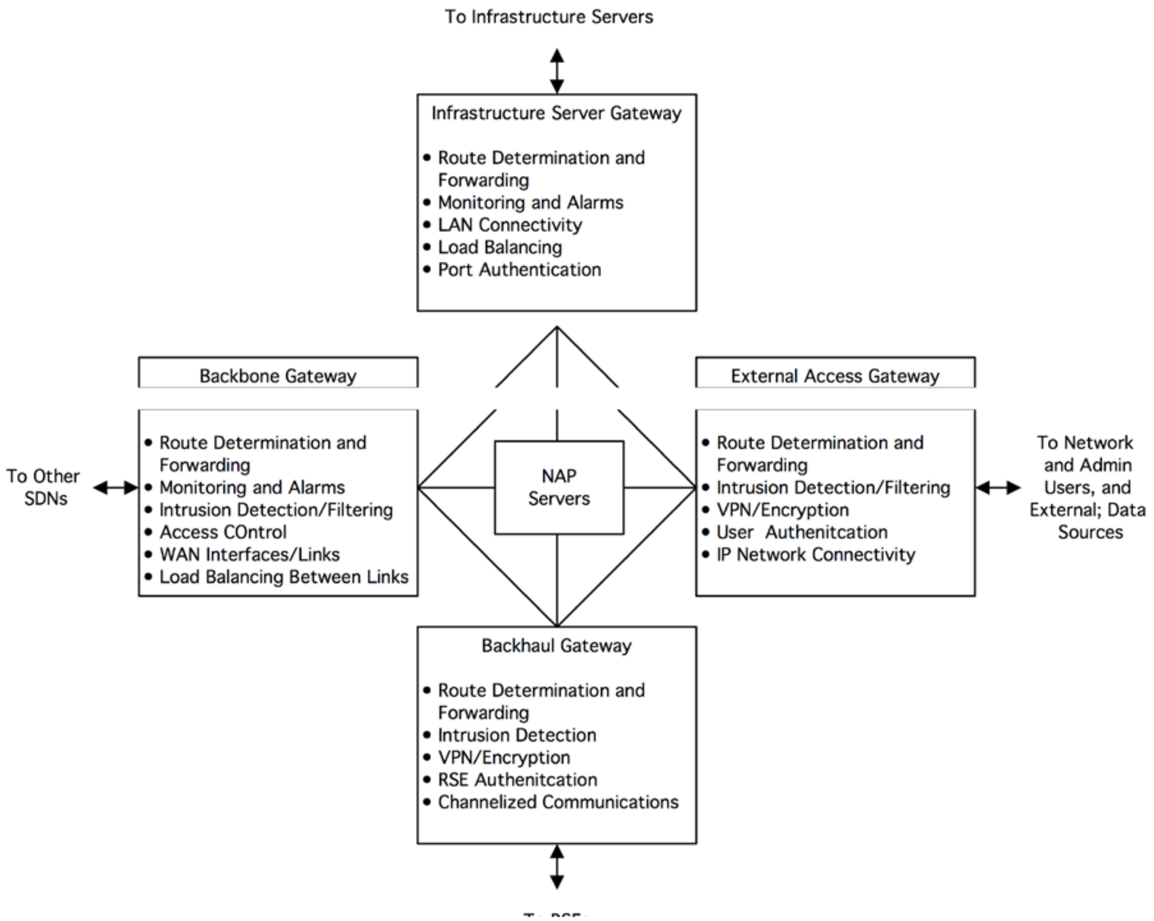
To Infrastructure Servers
Infrastructure Server Gateway
• Route Determination and Forwarding
• Monitoring and Alarms
• LAN Connectivity
• Load Balancing
• Port Authentication
Backbone Gateway
• Route Determination and Forwarding
• Monitoring and Alarms
• Intrusion Detection/Filtering
• Access COntrol
• WAN Interfaces/Links
• Load Balancing Between Links
To Other SDNs
NAP Servers
External Access Gateway
• Route Determination and Forwarding
• Intrusion Detection/Filtering
• VPN/Encryption
• User Authenitcation
• IP Network Connectivity
To Network and Admin Users, and External; Data Sources
Backhaul Gateway
• Route Determination and Forwarding
• Intrusion Detection
• VPN/Encryption
• RSE Authenitcation
• Channelized Communications

Figure 2-7: NAP

2.1.3.1 Infrastructure Server Gateway

Infrastructure server transport provides bi-directional transport of IP traffic between the NAP and associated infrastructure servers. Whenever the NAP receives a data packet that needs to be processed by an infrastructure server in the same SDN subsystem or ENOC subsystem instance, the server transport forwards that packet to the appropriate infrastructure server. Conversely, whenever an infrastructure server sends out a packet to a destination outside of the given SDN subsystem instance or ENOC subsystem instance, an infrastructure server gateway forwards that packet to one (or more) of the three other gateways.

2.1.3.2 Backbone Gateway

The backbone provides connectivity between geographically dispersed instances of the SDN and ENOC subsystems. Each backbone gateway consists of interconnected network devices and security devices operating as a part of the NAP within an SDN or ENOC subsystem instance. A backbone transport interface refers to two specific backbone gateways and the communications link between them. As shown in Figure 2-3, each SDN or ENOC subsystem instance is logically associated with a minimum of two other subsystem instances through physically diverse routes. For the POC, the two SDNs were associated only with each other.

2.1.3.3 External Access Gateway

The external access gateway is a combination of network devices and security devices that provide connectivity to parties external to the VII network subsystem. Figure 2-7 illustrates the role an external access gateway performs.

In the POC, external network users connect to the external access gateway via the Internet. From this connection point, data is routed either locally to the SDN or to other elements of the POC system via the backbone and backhaul networks.

2.1.3.4 Backhaul Gateway

The RSE backhaul provides bi-directional transport of IP traffic between the SDN subsystem and the RSE subsystem in support of all services provided by the VII system. Backhaul traffic enters and leaves each SDN subsystem instance through an RSE backhaul gateway. As shown in Figure 2-7, an RSE backhaul gateway is a combination of network devices and security devices operating at the NAP within each SDN subsystem instance. Technologies used in the RSE backhaul vary from location to location and depend on availability. In the POC implementation, the backhaul is implemented using 3G cellular, T1 wire line, and WiMAX links.

2.1.4 Physical SDN Setup

The SDN is typically housed in a standard computing rack resident in a typical computing facility. For the POC program, two SDNs were established—one at the Booz Allen Hamilton computing facility in Herndon, VA, and the other at the offices of the Road Commission for Oakland County (RCOC) in the Detroit, MI, area. The following sections describe these SDN facilities.

2.1.4.1 Herndon

The Herndon SDN is located at Booz Allen’s computing facility in Herndon, VA. This SDN initially was used to validate the design and to test the basic operations and services during development. The SDN functions are implemented using a set of servers and storage units housed in a rack at the Herndon facility, as shown in Figures 2-8 and 2-9. The certificate authority (both the X.509 and 1609.2) and all IdAM functions, including VPNs and user account management, are also implemented at the Herndon SDN location.

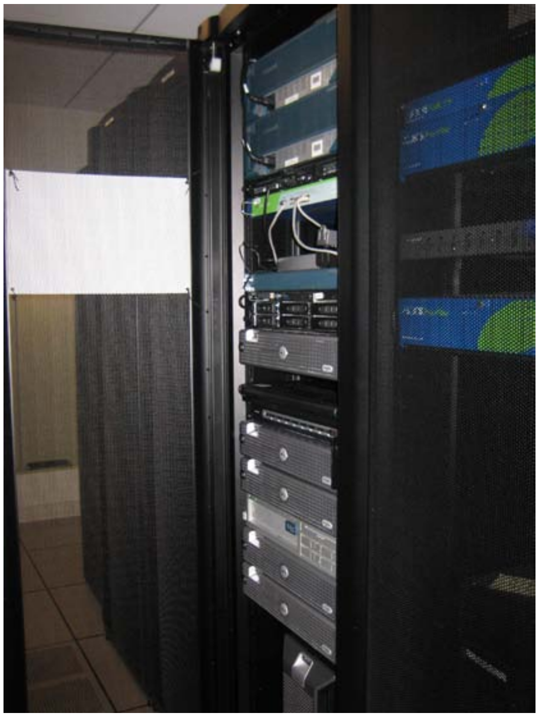

Booz Allen staff photo

Figure 2-8: Herndon SDN

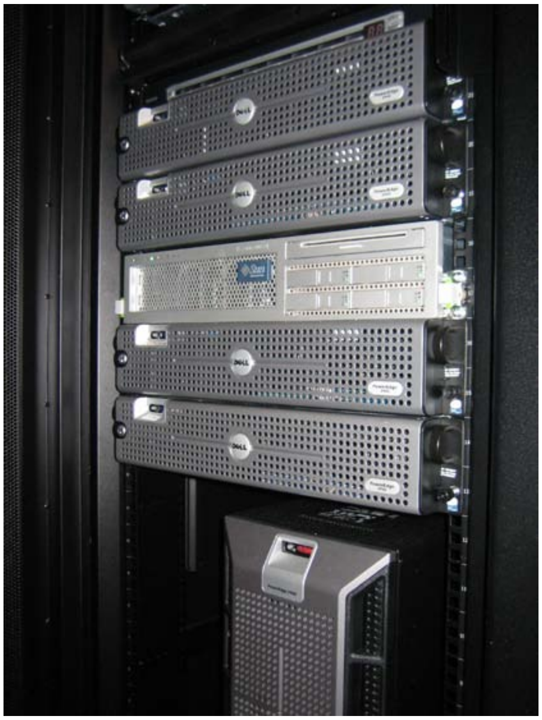

Booz Allen staff photo

Figure 2-9: Herndon SDN Close-up

The Herndon facility is connected to the Detroit development test environment (DTE) RSEs via the backbone network, making the Detroit and Herndon SDNs mutually redundant—a feature that is important for any future full-scale deployment.

2.1.4.2 Detroit

The Detroit SDN was located in a facility provided by the RCOC in Franklin, MI. This SDN is smaller, although it contains most of the same functions as the Herndon SDN. The Detroit SDN included a small-scale ENOC that allowed control and configuration of the RSEs locally while POC tests were underway. This approach proved to be highly efficient, as it allowed rapid modification and reconfiguration on site as tests were being performed.

The Detroit SDN hosted numerous RSEs, and these were connected using several different backhaul networks. These were linked to the SDN via a high-bandwidth switching/routing unit known as the Opt-E-Man. Three different Opt-E-Man switches were used—one for 3G links, one for wire line links, and one for WiMAX Links. Figure 2-10 shows the WiMAX Opt-E-Man switch. The SDN is composed of server storage units and power conditioning systems. Figure 2-11 shows the Detroit SDN.

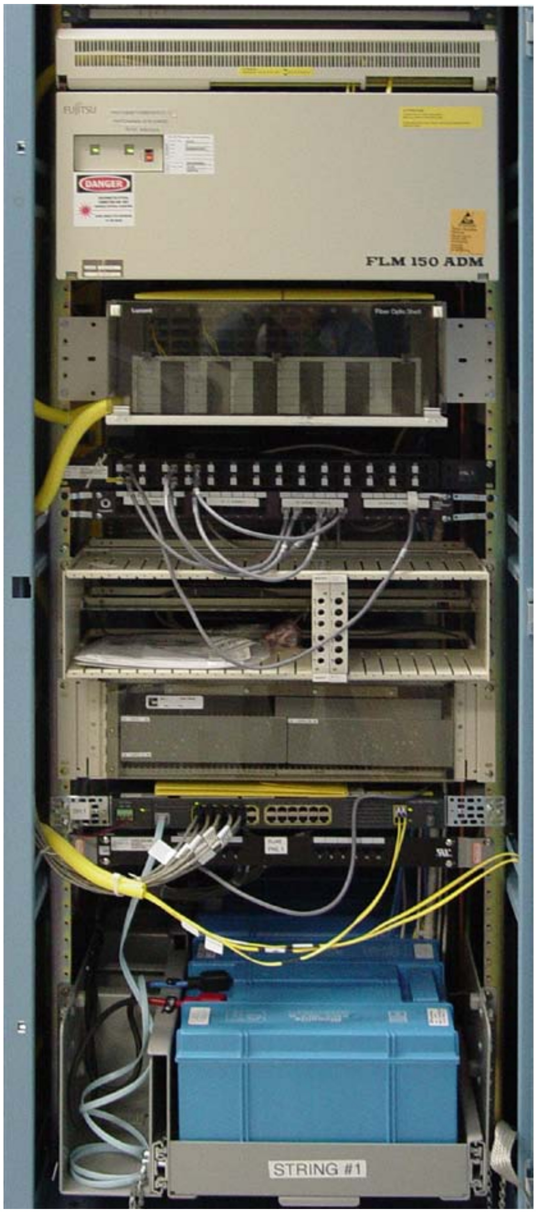

Booz Allen staff photo

Figure 2-10: ATT OptiMan Node

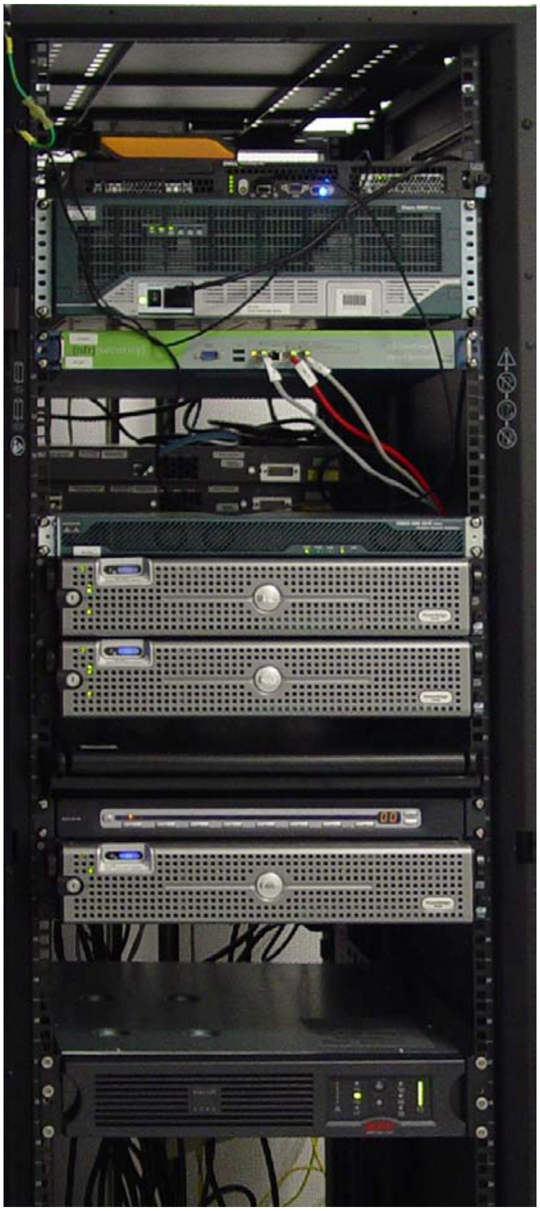

Booz Allen staff photo

Figure 2-11: Detroit SDN

2.2 ENOC

The VII infrastructure system is composed of many different hardware and software components that are managed through an ENOC. The ENOC subsystem provides a centralized monitoring and control function for the VII network and all (non-mobile) VII system components. The ENOC subsystem also provides for analysis and report generation in support of long-term trending and planning activities. The ENOC is staffed by operations personnel who use the ENOC systems to monitor and manage the VII network, operations, assets, and communications infrastructure.

The ENOC is responsible for the effective operation, availability, reliability, performance, and security of the VII system. Specific responsibilities of the ENOC subsystem include monitoring, provisioning, configuring, repairing, maintaining, and securing the RSE subsystem and the SDN subsystem, as well as monitoring the communications infrastructure that interconnects RSEs and SDNs. The ENOC supports the establishment and management of all internal VPNs that are used to secure data traffic between elements of the infrastructure network.

The ENOC systems are the lynchpin of a hierarchical management architecture that includes managed entities and MNEs, each requiring different support and levels of control.

Figure 2-12 shows a typical ENOC workstation.

Booz Allen staff photo

Figure 2-12: ENOC Workstation

2.2.1 Managed Entity

The ENOC subsystem provides a centralized authentication and authorization capability to the users (human and machine) that interact with the SDN subsystem. Each of these users is referred to as an "entity" to indicate that this capability is provided to more than just human users. Specifically, a managed entity is a person, organization, hardware device, or software process that is known by the VII system.

Managed entities include:

- Systems operators
- SDNs and servers
- SDN software services
- RSEs
- Software services resident on the RSEs.

2.2.1.1 Managed Network Element (MNE)

A special sub-class of managed entity is the MNE. The MNEs are the physical infrastructure—the hardware and associated software deployed throughout the VII system. MNEs include the RSE, backhaul, backbone, access, and SDN gateways, as well as the SDN itself.

MNEs include:

- Infrastructure (SDN and ENOC servers)
- Routers, gateways, and other networking equipment in the NAP
- RSE
- Software services resident on the RSE.

2.2.1.2 Management Service (MGMT)

The primary function of the ENOC is to run the network and provide a means for the ongoing configuration and maintenance of the various network subsystems (the managed entities). The VII Management Service provides facilities to maintain and monitor the performance, security, and configuration of the non-mobile components of the VII system. This includes tracking and management of the VII system configuration; provisioning and configuration of non-mobile VII system components; detection, isolation, and correction of VII network infrastructure and service problems; and monitoring of VII system and subsystem performance.

2.2.1.3 Fault Management Function

The VII Fault Management Function detects, isolates, and corrects faults in MNEs. The Fault Management Function collects alarm event data from the MNE, correlates and de-duplicates the alarm data, identifies the exact location of the faulty MNE, and repairs the faulty element through remote configuration or technician dispatch. Figure 2-13 presents an example of the POC fault management screen.

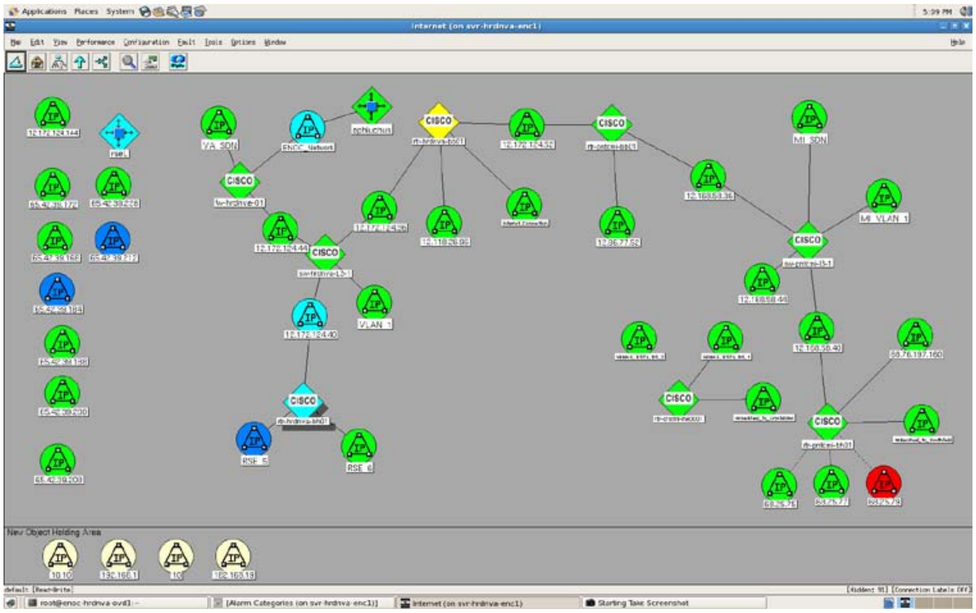

Figure 2-13: Fault Management Screen

2.2.1.4 Performance Management Function

The VII Performance Management Function collects performance data from the MNEs and generates performance reports. This includes performance monitoring of MNEs, automated analysis of system performance data, and generation of alerts to the fault management system when user-defined thresholds have been breached.

Figure 2-14 and Figure 2-15 show typical VII POC system performance monitoring screens.

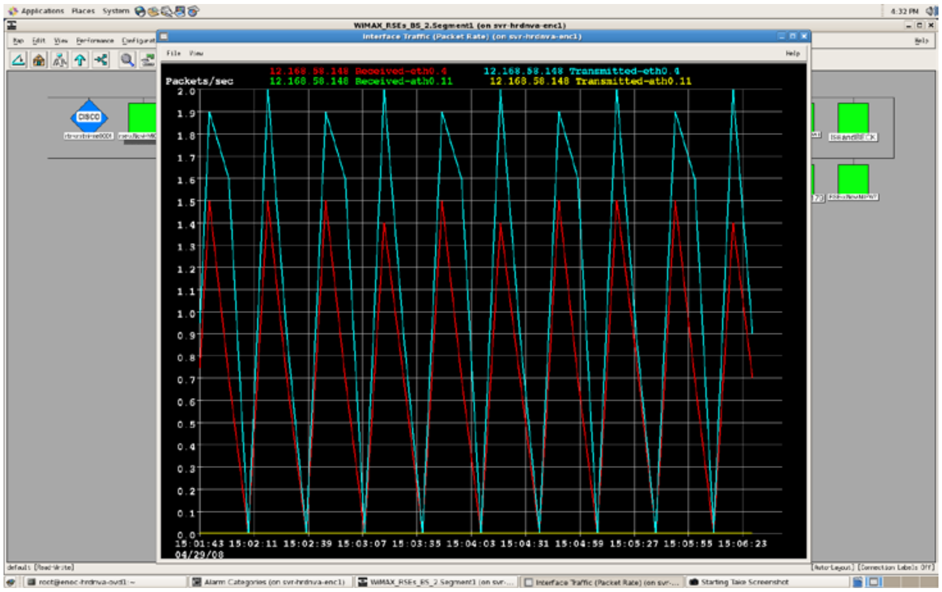
WIMAX_RSEs_BS_2.Segment1 (on svr-hrdnva-encl)
Interface Traffic (Packet Rate) (on svr-hrdnva-encl)
File View
Help
Packets/sec
12.168.58.148 Received-eth0.4
12.168.58.148 Received-ath0.11
12.168.58.148 Transmitted-eth0.4
12.168.58.148 Transmitted-ath0.11
2.0
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
15:01:43 15:02:11 15:02:39 15:03:07 15:03:35 15:04:03 15:04:31 15:04:59 15:05:27 15:05:55 15:06:23
04/29/08

Figure 2-14: Network Traffic Monitoring Screen

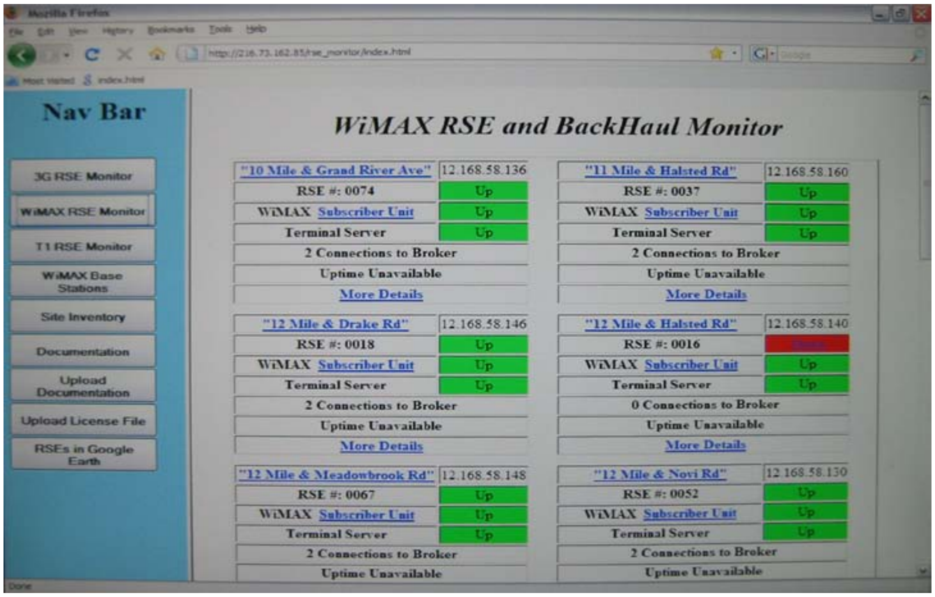

WiMAX RSE and BackHaul Monitor

"10 Mile & Grand River Ave"	12.168.58.136
RSE #: 0074	Up
WiMAX Subscriber Unit	Up
Terminal Server	Up
2 Connections to Broker	
Uptime Unavailable	
More Details	

"11 Mile & Halsted Rd"	12.168.58.160
RSE #: 0037	Up
WiMAX Subscriber Unit	Up
Terminal Server	Up
2 Connections to Broker	
Uptime Unavailable	
More Details	

"12 Mile & Drake Rd"	12.168.58.146
RSE #: 0018	Up
WiMAX Subscriber Unit	Up
Terminal Server	Up
2 Connections to Broker	
Uptime Unavailable	
More Details	

"12 Mile & Halsted Rd"	12.168.58.140
RSE #: 0016	[illegible]
WiMAX Subscriber Unit	Up
Terminal Server	Up
0 Connections to Broker	
Uptime Unavailable	
More Details	

"12 Mile & Meadowbrook Rd"	12.168.58.148
RSE #: 0067	Up
WiMAX Subscriber Unit	Up
Terminal Server	Up
2 Connections to Broker	
Uptime Unavailable	

"12 Mile & Novi Rd"	12.168.58.130
RSE #: 0052	Up
WiMAX Subscriber Unit	Up
Terminal Server	Up
2 Connections to Broker	
Uptime Unavailable	

Figure 2-15: RSE Status Monitoring Screen

2.2.1.5 Configuration Management Function

The VII Configuration Management Function remotely manages the configuration of MNEs. This includes establishing new MNE functionality in the VII system, as well as updating the existing configuration of an MNE. It also includes setting performance management thresholds on MNEs, inventory discovery, and asset management of MNEs. The Configuration Management Service supports the ILS by creating, updating, and deleting the identifying information for RSEs (capabilities, locations, connectivity, and political boundaries).

2.2.1.6 Identity and Access Management/Security

The ENOC also hosts core elements of the IdAM subsystem. These IdAM functions are distributed among the Certificate Authority, the ENOC, and the SDN. Please see Section 2.5.1 for a complete description.

2.3 Roadside Equipment (RSE)

The RSE is intended to serve as a network end point that provides connectivity between the SDN and other remote services and the OBE. The RSE is implemented as a self-contained unit that can be mounted on a signal controller, lamppost, gantry, or other suitable roadside structure. Because of its remote location, the RSE includes a variety of internal management and control functions that facilitate both autonomous operation and remote configuration and provisioning.

The RSE includes several data interfaces, as illustrated in Figure 2-16. The DSRC interface is used to exchange wireless data with OBEs in the local vicinity of the RSE. The backhaul interface is used to connect the RSE back to the SDN, and the Ethernet “Craft Port” is used to connect the RSE to a local safety system or service provider. This last interface is important for low-latency applications where it would be impractical or inefficient to route information from the provider back to the SDN only to then route it back to the same general location. Examples of these local services are traffic signal status and toll processing—services that are both time critical and highly location specific. The backhaul network provides physical transport to the SDN, but is actually comprised of several logical interfaces, as also shown in Figure 2-16. Specifically, the backhaul supports a Simple Network Management Protocol (SNMP) interface used to audit and manage the operation of the various RSE elements, an XML/RPC (XML-based Remote Process Control) interface used to deliver content to be broadcast by the RSE, and a general purpose TCP/UDP/IP interface used to transfer data packets between the RSE and the SDN services.

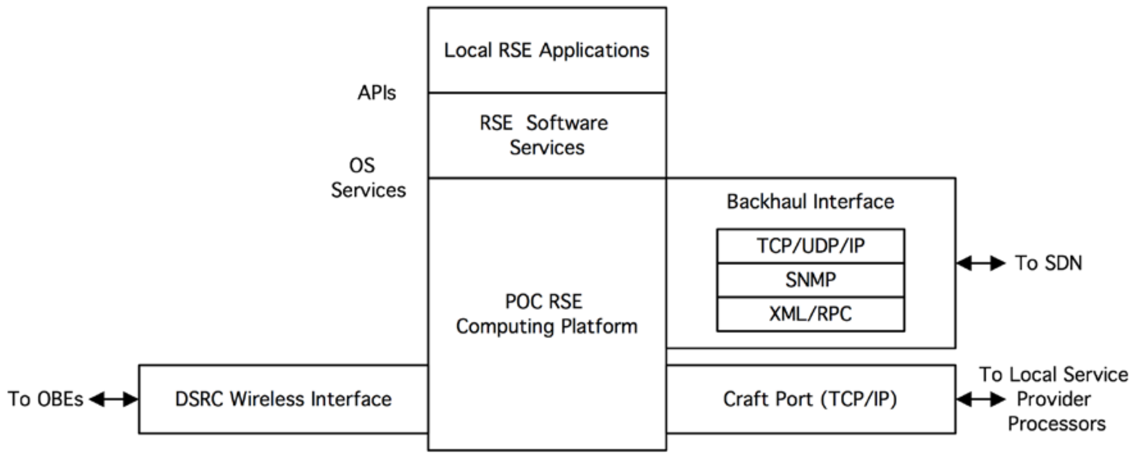

Figure 2-16: RSE Interfaces

Figure 2-17 shows the RSE subsystem attached to a lamppost for testing. This subsystem is based on an Intel-based processing unit, known as the TechnoCom Multiband Configurable Networking Unit (MCNU), that includes internal support for the DSRC radio and a GPS receiver. All software services are hosted on the MCNU platform. In addition, the RSE subsystem includes a separate power-conditioning module that provides conditioned local power (typically from the local roadway infrastructure systems).

The RSE is composed of the DSRC radio subsystem, a routing function, and a set of proxy applications that extend the services residing at the SDN (described above) out to each RSE associated with that SDN. The proxies essentially pass messages to and from their counterpart SDN services and interface to the RSE radio subsystem. The radio subsystem includes a DSRC radio and a radio handler that accepts or sends messages to and from the various proxies. The radio handler also constructs or updates a playlist that contains all broadcast messages to be transmitted.

Depending on the situation, an RSE may be connected to a local transaction processor (LTP) or a local safety system (LSS). This may be, for example, a local tolling system or a traffic signal controller. In operation, the LTP/LSS sends and receives messages to OBEs and to network users through the RSE functions. These messages usually have local relevance (as in tolling or signals) and thus need to originate local to the RSE.

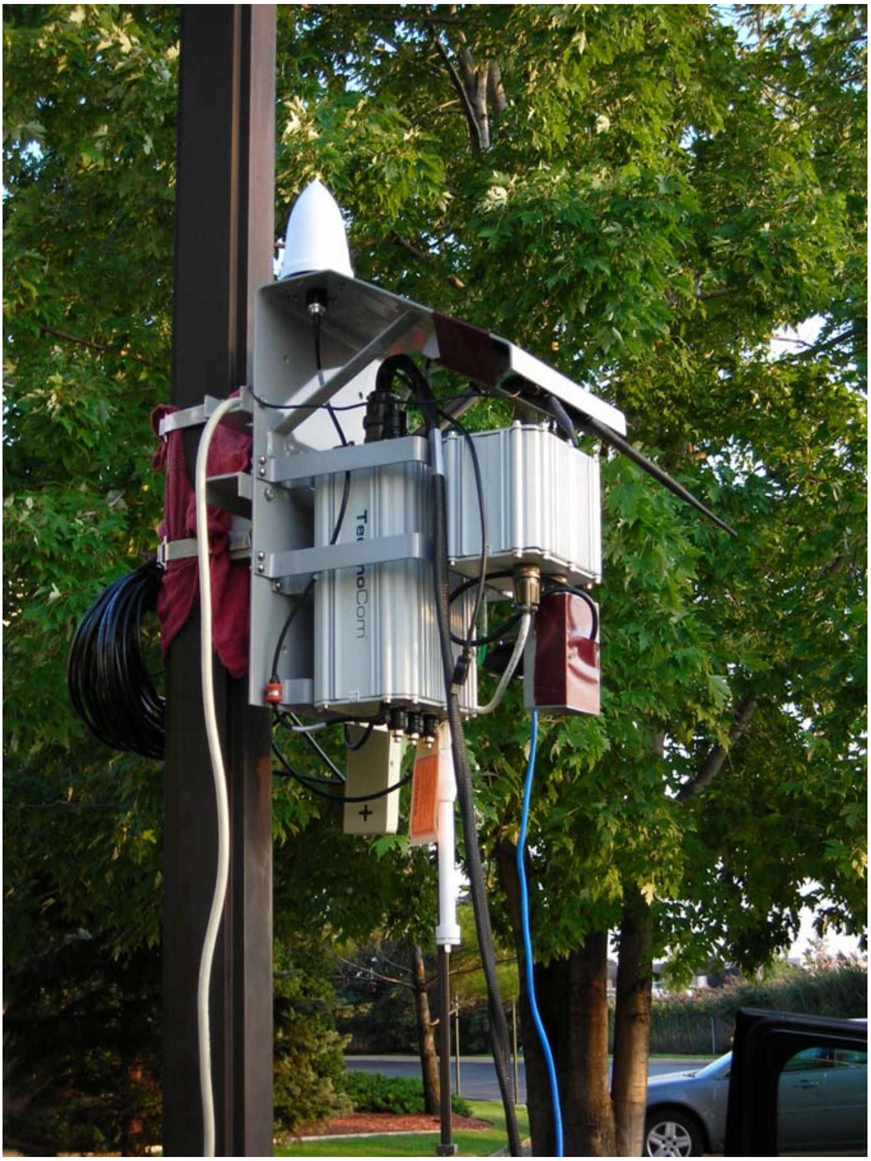

Booz Allen staff photo

Figure 2-17: RSE Assembly Mounted on Pole for Test

2.3.1 Electrical/Mechanical Architecture

2.3.1.1 Electrical Architecture

Figure 2-18 shows the RSE electrical architecture. This figure shows the separate power supply that conditions available 110 VAC power for the other RSE components. The MCNU unit provides ports for connecting the DSRC and GPS antennas and for connecting the various data input and output (I/O) lines.

The MCNU unit software is loaded via the USB thumb drive. This drive is not attached when the unit is operational, but this approach allows a simple and convenient way to install major software updates.

The backhaul transceiver is used to connect one of the MCNU Ethernet ports to the backhaul. Backhauls used in the POC included wire line (T1), WiMAX, and 3G.

The craft port provides a conventional Ethernet connection to local devices.

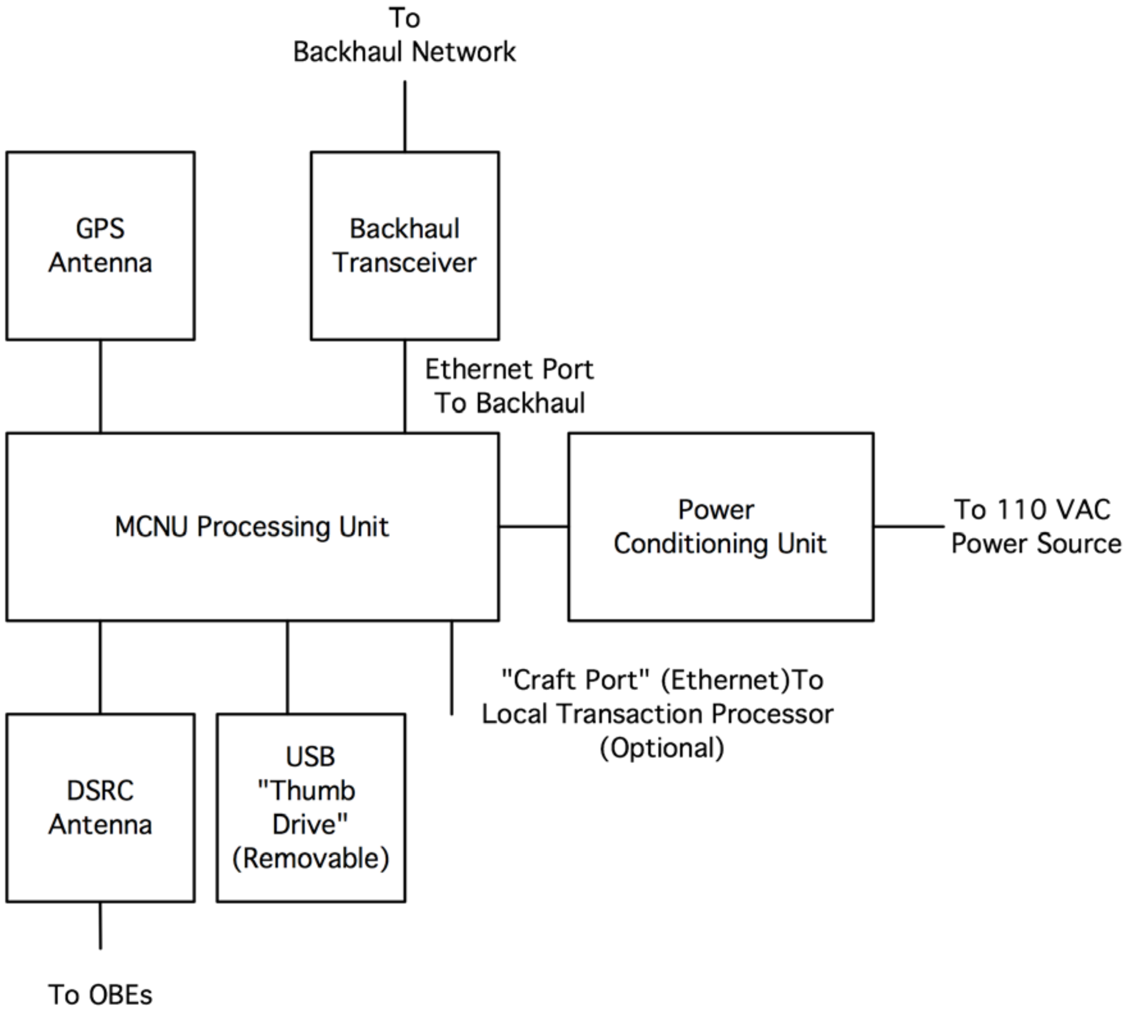

Figure 2-18: RSE Electrical Architecture

2.3.1.2 Mechanical Layout and Mounting

The RSE assembly is intended to be mounted about 5 meters above the ground, typically on a pole or a beam attached to a pole.

Figure 2-19 shows the various mechanical and cabling elements laid out. The MCNU unit is mounted to the RSE bracket (sled) and then the power supply is attached over it. The sled includes a cover to shield the RSE from sun and rain. In practice, the sled did not provide sufficient protection and a canvas cover was added to more fully enclose the unit during rainy weather.

The GPS antenna is mounted to the top of the sled where it has a clear view of the sky, and the DSRC antenna is mounted directly to the bottom of the MCNU, where it has a clear view of the roadway.

For clarity, the MCNU and the power supply are shown separately in the figures below.

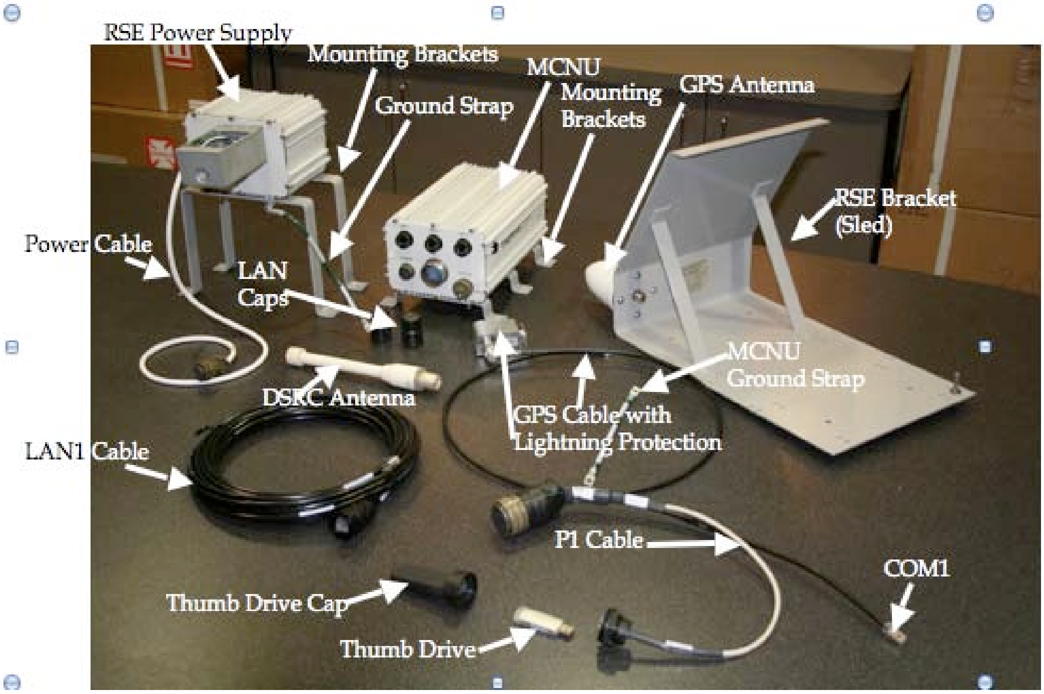

Booz Allen staff photo

Figure 2-19: RSE Assembly Components

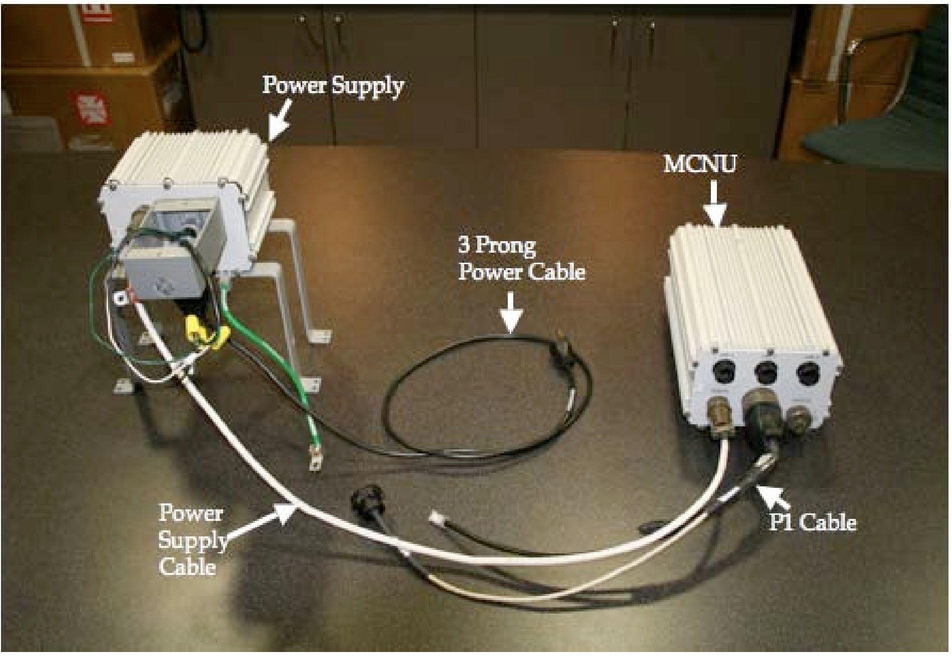

Booz Allen staff photo

Figure 2-20: MCNU and Power Supply

Figure 2-21 shows a typical RSE mounted to a pole in the DTE.

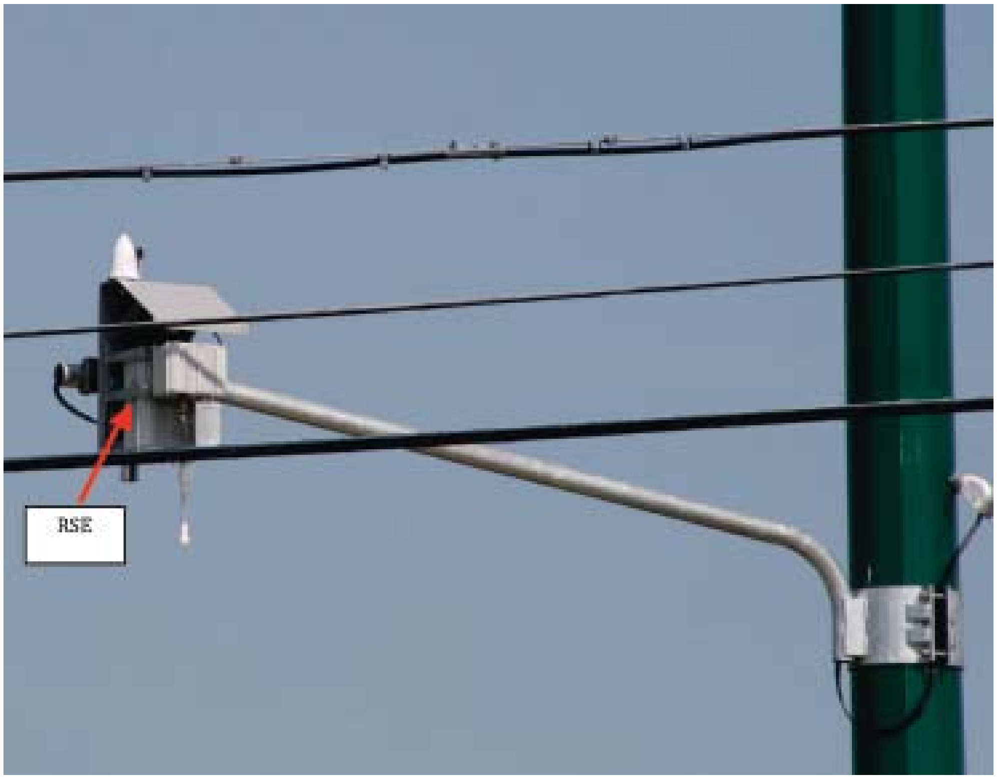

Booz Allen staff photo

Figure 2-21: RSE Mounted to Pole in DTE

2.3.2 RSE Processing Unit

The TechnoCom MCNU performs the primary processes in the RSE.

The MCNU is a Pentium 4-class computer in a rugged enclosure. It includes an internal GPS receiver and a DSRC radio card, as well as various data I/O. The motherboard uses a VIA C7 processor with a clock speed of 1 GHz. The front side bus speed is 533 MHz.

The MCNU is designed with an integrated conductive cooling system. This passive cooling system allows operation at the high end of the specified temperature range without the need for a fan. Fewer moving parts mean higher reliability.

The MCNU provides the following interfaces:

- Three Ethernet ports with RJ–45 interfaces
- Two RS–232 COM ports
- USB 2.0 port, accessed through a 24-pin connector
- Three status LEDs, visible through an external weatherproof window
- 5-pin external power input and +5V and +12V output
- Two sets of antenna connectors for wireless LAN (DSRC and WiFi), with main and auxiliary ports, and a GPS SMA connector.

2.3.3 RSE Software Architecture

The RSE processing is divided into management (MGMT) and proxy services.

The MGMT service performs all of the background tasks associated with the RSE. This includes bootstrap, security, routing, hardware interface, health monitoring and management (HMM), and life cycle support services.

The proxy service, composed of proxies, performs tasks supporting VII system services, such as AMDS, MEDS, POS, and PDS. The proxy service also provides the capability for managing individual proxies, such as starting and stopping.

Status and fault messages are forwarded to the HMM service for processing and distribution to the ENOC. The HMM service monitors the health of RSE components and provides for fault diagnosis, as well as providing an SNMP agent to update various management information blocks (MIBs) associated with performance and fault isolation monitoring. The HMM provides an interface to the ENOC to report on the health and status of RSE components.

At startup, bootstrap initializes a VPN for the backhaul link to the SDN. Once the VPN is established, all inbound backhaul message traffic is considered to be authorized and authenticated. The security firewall authorizes ENOC messages for maintenance functions and permits inbound messages only from selected IP addresses on selected IP ports. Other non-VPN IP ports also authorize and authenticate the source IP address and requested function. Once each function or proxy receives an authorized and authenticated message, the function or proxy performs the associated function.

Figure 2-22 shows the RSE software architecture. The primary elements are the DSRC radio stack, the radio handler, the proxy manager, the HMM service, and the bootstrap service. These services are described in detail in the following sections.

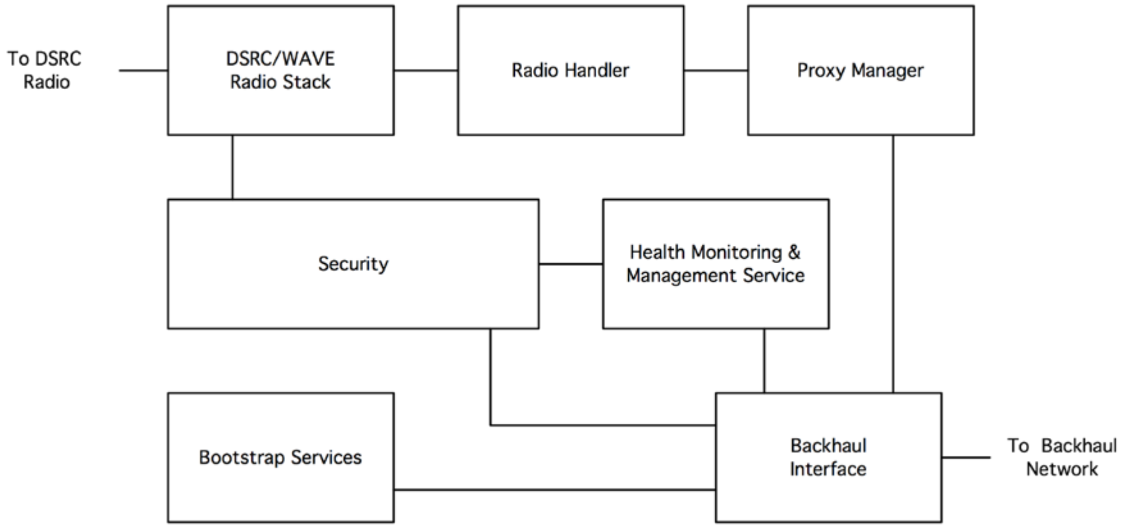

Figure 2-22: RSE Software Architecture

The overall service behavior of any given RSE is defined by its active proxy applications. Proxy applications are generally lightweight extensions of applications running at the SDN. The SDN applications send messages to the proxy applications for broadcast to the OBEs in the general vicinity of the RSE, and the proxy applications forward messages received from OBEs back to their respective SDN applications. This approach significantly minimizes message traffic for some applications since those messages that may be broadcast many times (e.g., an advisory message) do not need to be resent (end to end) through the entire system. Instead, they can be cached at each RSE by the coordinated activities of the SDN application and the RSE-based proxy. In other cases, the proxy applications serve to control the flow of data to mitigate congestion problems caused when, for example, a large number of OBEs all converge on a single RSE.

The RSE also contains various housekeeping functions that allow it to be easily managed remotely from the ENOC.

2.3.3.1 Proxy Manager

The proxy manager provides a simple management service that controls the various proxy applications resident in the RSE. It allows each proxy application to be remotely started and stopped from the ENOC without rebooting or otherwise affecting the RSE operations or the other proxies.

The proxy manager consists of a command listener that waits for commands from the ENOC. When it receives a command to start a proxy application, it activates the initializer that starts and sets up the application based on a stored configuration file. When it receives a command to stop a proxy application, it notifies the application and the initializer so that the application configuration can be saved and the application can shut down gracefully.

Figure 2-23 illustrates the proxy manager architecture.

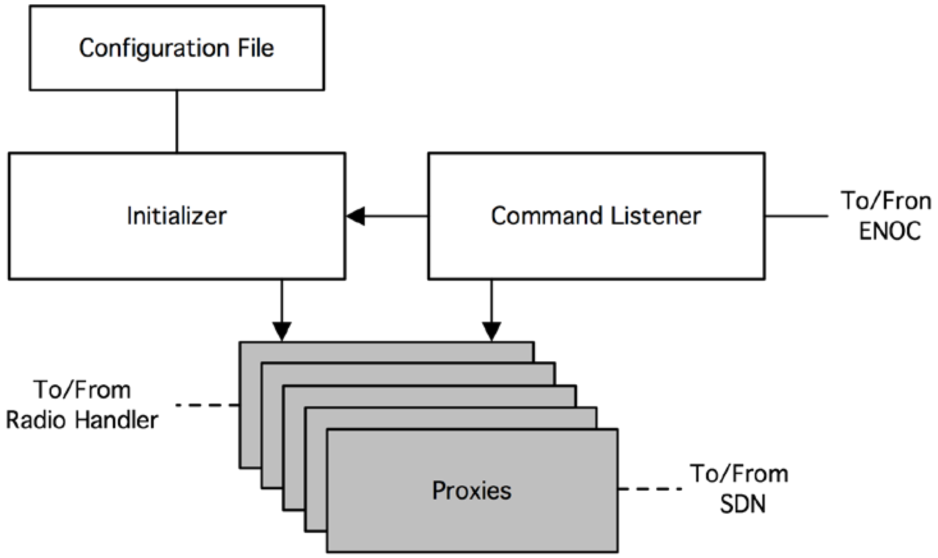

Figure 2-23: Proxy Manager Architecture

The RSE is capable of supporting an arbitrary number of proxy applications running simultaneously (subject to the processing limitations of the RSE itself). In the POC, typically all RSEs supported all proxy applications described below. The configuration file is used to provide data to inform the initialization process. For example, the lightweight proxy that is used to pass IP traffic through the RSE must be set up with the proper PSIDs to allow it to properly register with the radio handler, and thereby cause the proper application-oriented PSIDs to be included in the WSA.

2.3.3.2 Radio Handler

The radio handler provides comprehensive communications management for the various proxy applications resident in the RSE.

The DSRC/WAVE radio stack in the RSE operates in what is known as “provider mode.” This means that it sends out WSA and provides access to services on a service channel. To populate the WSA, the radio stack must know what services are available. In addition, to receive messages, it must know what message types it should accept and deliver to the top of the radio stack. The mechanism for this is the PSID.

As described above, the RSE only supports a small number of local proxy applications, but the VII system is designed to support a nearly endless array of services from private providers. In addition, any local applications attached to the RSE via the craft port need a way to register their services and also establish PSIDs for received WSM routing. These services are implemented as a “pass through” at the RSE. This means that they do not use any form of dedicated proxy application. Instead, they use a shared application known as the lightweight proxy that registers PSIDs for them and passes message traffic from the OBE to either the local craft port (local transaction processor) or the access gateway at the SDN. In addition, the RSE routing functions establish the correct IP routing information so that packets received for that remote service are properly routed out the backhaul network and on to the remote service.

The radio handler provides a simple way for local proxies to register their services with the radio and to establish the proper configurations for message handling. This approach also significantly simplifies the DSRC/WAVE radio interface, since it effectively only needs to deal with one application (i.e., the radio handler).

As illustrated in Figure 2-24, the radio handler is composed of a connection manager that registers PSIDs for messages and services with the radio. The message bridge translates WSMs into IP packets and directs them to the routing function where they can be routed to the appropriate service. IP packets received by the radio do not pass through the radio handler since they pass directly from the radio into the IP stack in the RSE, and are routed according to the packet IP address. The message bridge also maintains a playlist of messages to send to the radio for broadcast at regular intervals. The playlist allows a proxy application to submit a message that is intended to be repeatedly broadcast (e.g., a speed limit advisory) one time, and the radio handler then submits it at configurable intervals to the radio for transmission.

The Provider User Control function is a Web service used by the local RSE proxies and also by remote services to register PSIDs to be advertised in the WSA.

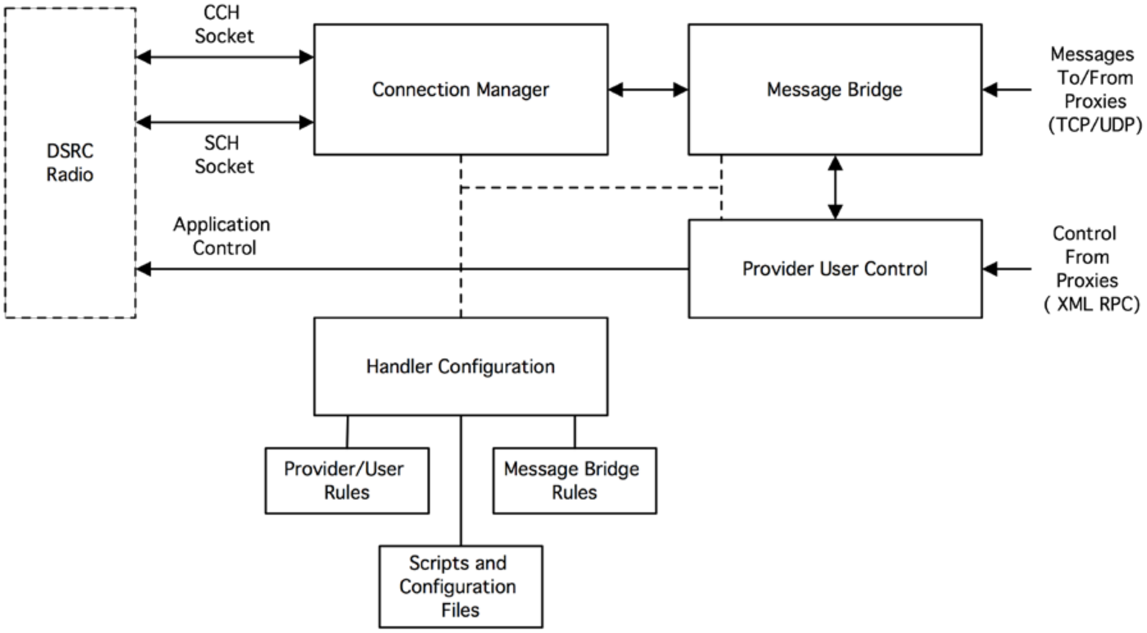

Figure 2-24: Radio Handler Architecture

2.3.3.3 DSRC Radio Stack

As described above in Section 1.3, the DSRC/WAVE system is based on 802.11p and IEEE 1609 standards. The OBE DSRC radio is implemented as a hybrid hardware and software system, as illustrated in Figure 2-25.

The physical layer and the supporting 802.11p protocols are implemented using a commercial WiFi radio packaged on a Mini PCI card. This card contains firmware that was modified to conform to the 802.11p standard. The upper layers of the DSRC protocol defined in IEEE 1609, known as WAVE protocols, are implemented in software running within the RSE software system.

The upper and lower layers of the radio subsystem are managed by a software element known as the WAVE management entity (WME). This forms what is known as the “management plane” of the radio, while the layers that operate on the messages themselves are called the “data plane.”

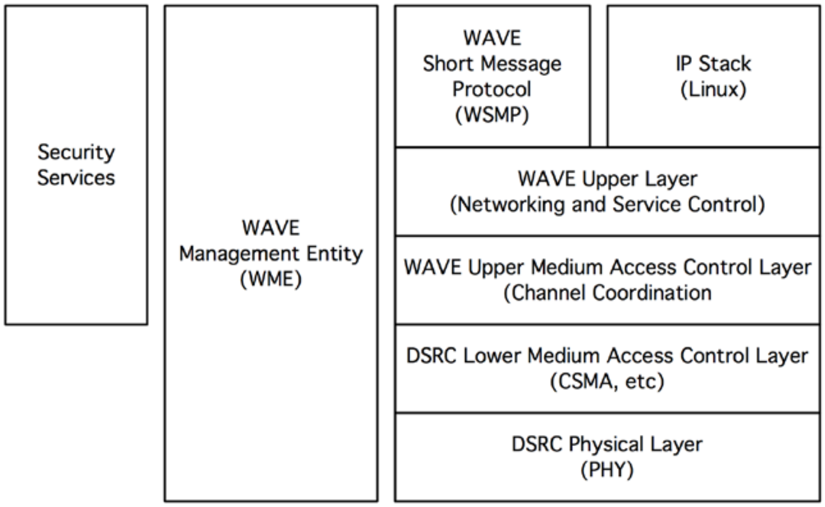
Security Services
WAVE Management Entity (WME)
WAVE Short Message Protocol (WSMP)
IP Stack (Linux)
WAVE Upper Layer (Networking and Service Control)
WAVE Upper Medium Access Control Layer (Channel Coordination
DSRC Lower Medium Access Control Layer (CSMA, etc)
DSRC Physical Layer (PHY)

Figure 2-25: DSRC/WAVE Radio Architecture

2.3.3.3.1 *DSRC Layers*

The DSRC radio physical layer and lower MAC layer are responsible for physically generating and receiving the RF signals and for controlling the basic operations associated with sending and receiving these signals. The requirements for this operation are specified in the IEEE 802.11p standard that defines DSRC radio.

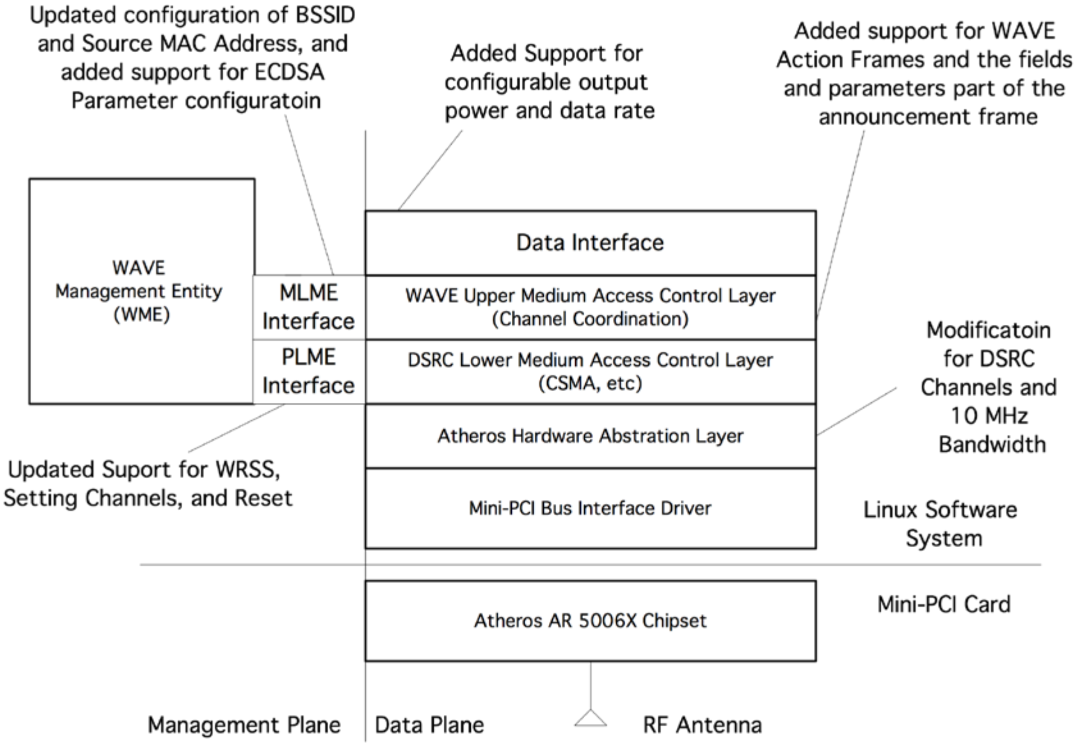
Updated configuration of BSSID and Source MAC Address, and added support for ECDSA Parameter configuratoin
Added Support for configurable output power and data rate
Added support for WAVE Action Frames and the fields and parameters part of the announcement frame
WAVE Management Entity (WME)
MLME Interface
PLME Interface
Data Interface
WAVE Upper Medium Access Control Layer (Channel Coordination)
DSRC Lower Medium Access Control Layer (CSMA, etc)
Atheros Hardware Abstration Layer
Mini-PCI Bus Interface Driver
Modificatoin for DSRC Channels and 10 MHz Bandwidth
Updated Suport for WRSS, Setting Channels, and Reset
Linux Software System
Atheros AR 5006X Chipset
Mini-PCI Card
Management Plane
Data Plane
RF Antenna

Figure 2-26: DSRC Radio Architecture

The physical and lower MAC layers are implemented using an Atheros radio subsystem implemented on a Mini-PCI card. The base radio card is designed to support the 802.11 a Wi-Fi standard that operates at a slightly lower frequency band, and operates using slightly different protocols. The basic 802.11a operation is not changed, but key elements have been added to allow the system to operate effectively in the high-speed vehicle environment where it is not possible (or necessary) to set up a full-blown network prior to communicating. Figure 2-26 summarizes the changes to the protocol stack. Figure 2-27 shows the Atheros Mini-PCI radio card.

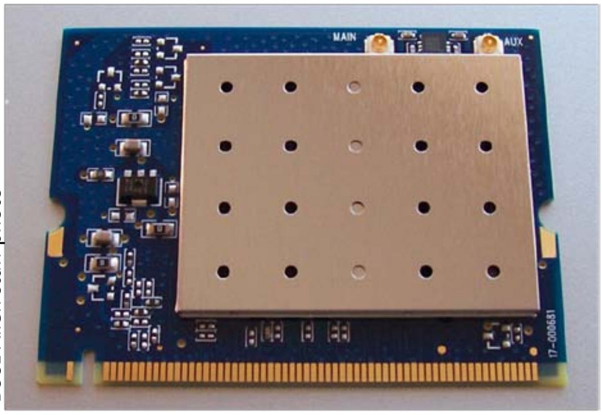

Booz Allen staff photo

Figure 2-27: DSRC Radio Mini-PCI Card

2.3.3.3.2 WAVE Layers

The upper layers of the DSRC/WAVE radio implement the WAVE part of the overall protocol as described in the IEEE 1609 standards. These include the overall service management logic that determines how a WAVE radio decides what services from which providers to use, the WSMP, and logic to manage the seven different DSRC channels defined for use in the United States by the FCC.

The WAVE layers support two different types of message elements—conventional IP packets and WSMs. As can be seen in Figure 2-28, this somewhat complicates the upper layers since normal Wi-Fi radios simply pass incoming packets to the IP stack provided by the operating system. In the case of IP communications, the VII implementation is not particularly different from this, but in the case of WSMs, there is no native function to route packets to the intended applications. As a result, the upper layer WAVE implementation also provides an application programming interface (API) that allows the user applications to both register (as user or provider) to support service and channel decisions (see below), and to send and receive WSMs. This is also illustrated in Figure 2-28.

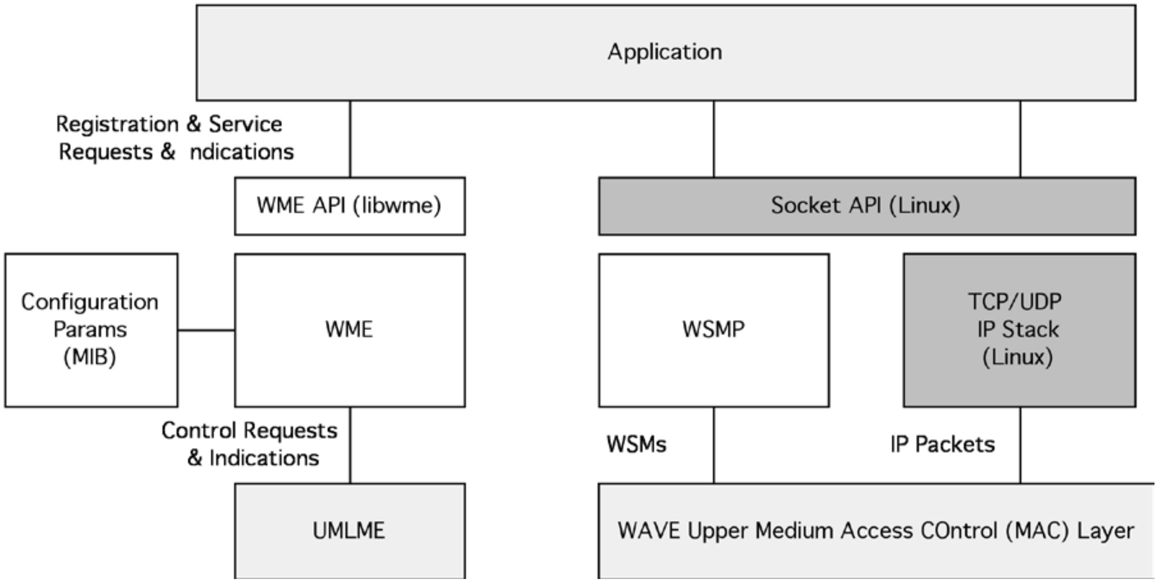

Figure 2-28: WAVE Upper Layer Software Architecture

As described in Section 1.3 above, the DSRC/WAVE protocol uses a CCH and SCH interval concept. By requiring that any radio monitor the CCH during the CCH interval, the system assures that any given radio will be tuned to the right channel (the CCH) at the right time to hear important messages such as safety messages and service announcements. A DSRC/WAVE radio may operate as either a user or a provider. While both a user and a provider operate the same from a message communications perspective, a provider is also able to issue a WSA. The WSA is broadcast on the CCH during the CCH interval to announce or advertise the services that the provider is offering, and to indicate on which of the other DSRC channels these services may be found. In general, OBEs operate in the user mode, although this is not a requirement.

The WME is responsible for receiving any WSAs and for deciding which channel (if any) to use during the SCH interval. This is done on the basis of what services the OBE applications have registered for, what services have been advertised by RSEs, and the relative priority of those services. The WME also interacts with the security services to verify the digital signatures of any received WSAs and, in provider mode, to digitally sign any outgoing WSAs.

The channel coordination layer is responsible for controlling the channel used by the radio, and for routing messages into queues for the channels on which they are intended to be sent. This is a key function since the radio must be synchronized to all other radios so that the CCH and SCH intervals line up, and messages intended for a specific service must be held until the radio is tuned to that channel. The channel switching operation is synchronized to the pulse per second provided by the GPS receiver.

2.3.3.4 Security

The RSE supports three basic security elements. For communications with the SDNs, it includes a VPN client managed by the IdAM subsystem. This both authorizes the RSE on the VII infrastructure network and secures data communications between the RSE and the network. The RSE proxy applications are not exposed to these security functions.

The RSE also hosts implementation of the 1609.2 security libraries and the certificate manager, and a local IP-based encryption/decryption capability used to protect probe data over the air.

2.3.3.4.1 1609.2 Security Libraries

The security libraries in the RSE implement the IEEE 1609.2 Security Protocol. These libraries are described in detail in document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*. The RSE implementation is based on the same software as used in the OBE. The mechanisms used in IEEE 1609.2 are based on general public key infrastructure (PKI) security principles. The mechanics of this will not be described here.

Of particular importance to RSE security operations are the scope elements included in the certificate. These elements restrict the time, geography, and function associated with any given signature so that a receiving OBE can be assured that a message signed by the RSE is allowed from a functional perspective (this RSE is authorized to send this type of message), is not being replayed (it is being sent within a minimum time interval from when it was signed), and is being sent from an authorized geographic location. To accomplish this, the IEEE 1609.2 headers include transmission time, transmission location, and the PSID of the originating application. If the message is subsequently received at a different time or location, then it will be considered invalid. In addition, if the message is sent from an originator who is not authorized to send messages of that type, then the PSID of the message will not match any PSID in the certificate, and the message will be considered invalid.

Outbound Security Operations

The primary RSE security operations are to sign WSAs and outbound WSMs (e.g., advisory and local safety system messages).

Both WSAs and WSMs are signed using a 1609 certificate that is linked to the identity of the RSE. This certificate is obtained by the RSE certificate manager from the 1609 CA.

Most IP operations rely on end-to-end (application layer) security, so the RSE does not, in general, support Internet Protocol Security (IPSec) transactions. There is an important exception to this statement, however, that is discussed in below.

Inbound Security Operations

There are currently no applications that involve the RSE receiving inbound WSMs that must be verified locally (that is, WSMs that are verified by the RSE itself). However, the security libraries are capable of performing this operation, so the system can support such a requirement if such a need arises in the future.

Because probe data snapshots contain information about the behavior of the vehicle at locations other than where the data is uploaded to the system, it is important to protect the data from local eavesdropping. This approach prevents, for example, a police officer from intercepting probe data indicating that the vehicle was speeding at some earlier location, and thereby issuing a citation. While this practice might prove to be illegal, encryption of the local radio link was seen as a sure way to avoid such issues from the start, so the concept of a locally encrypted link was introduced.

This system, known as VII Datagram Transport Layer Security (VDTLS), builds on the well-known Datagram Transport Layer Security (DTLS) system used in the Internet. Upon entering the radio coverage zone of an RSE, the OBE based probe data application receives a unique identifier from the RSE, for example, the IP address of the RSE. This is included in the WSA advertising PDC services. The probe data application then uses this identity to encrypt the probe data message, and signs it using the OBE application's 1609.2 anonymous certificate. Since the IP address of the RSE is always sent in the WSA, this approach requires no additional steps in the setup of the secure link. The key to this approach is that the RSE has developed a private key that reverses the encryption done by the OBE using the RSE identity as a key. This method is not as secure as full-blown asymmetric keys, but it is highly efficient, avoids the time required for a complete secure key (e.g., AES) exchange, and requires no compromise of the OBE's anonymity.

2.3.3.4.2 Certificate Management

The RSE certificate management subsystem has two elements. The first element interacts with the 1609 certificate authority to request and process certificates. The second element controls access to and use of certificates in the vehicle system.

As part of the installation process, the RSE application registers with the security services. The security services then generate an application and RSE specific key pair, which can be used to request certificates for that specific instance of that application. This key pair, and all other cryptographic and security material for the application, is stored in an application security file on the RSE known as a WAVE Security Context (WSC).

Based on the WSC information, the certificate manager then requests the appropriate security credentials from the certificate authority for each requesting application. The operation of the RSE certificate manager is effectively the same as the OBE certificate manager, except that the RSE does not include the provision for requesting or using anonymous certificates.

The certificate manager continually checks the state of the certificates for each active application, and routinely requests replacement certificates when older certificates expire or are found to be revoked.

The detailed structure of the 1609.2 certificate authority is described in the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.

2.3.3.5 Health Monitoring and Management

The HMM service is part of the management services suite. The HMM monitors the health, provides for fault diagnosis, and provides an SNMP agent to update various MIBs associated with performance and fault isolation monitoring. System operators at the ENOC may register with HMM. Registration allows both operator requests, in which an operator queries a specific RSE relative to some operational parameter. This allows an operator to proactively examine the state of the network. In addition, registered operators automatically receive alerts when HMM detects a fault condition. Figure 2-29 shows the architecture of the HMM.

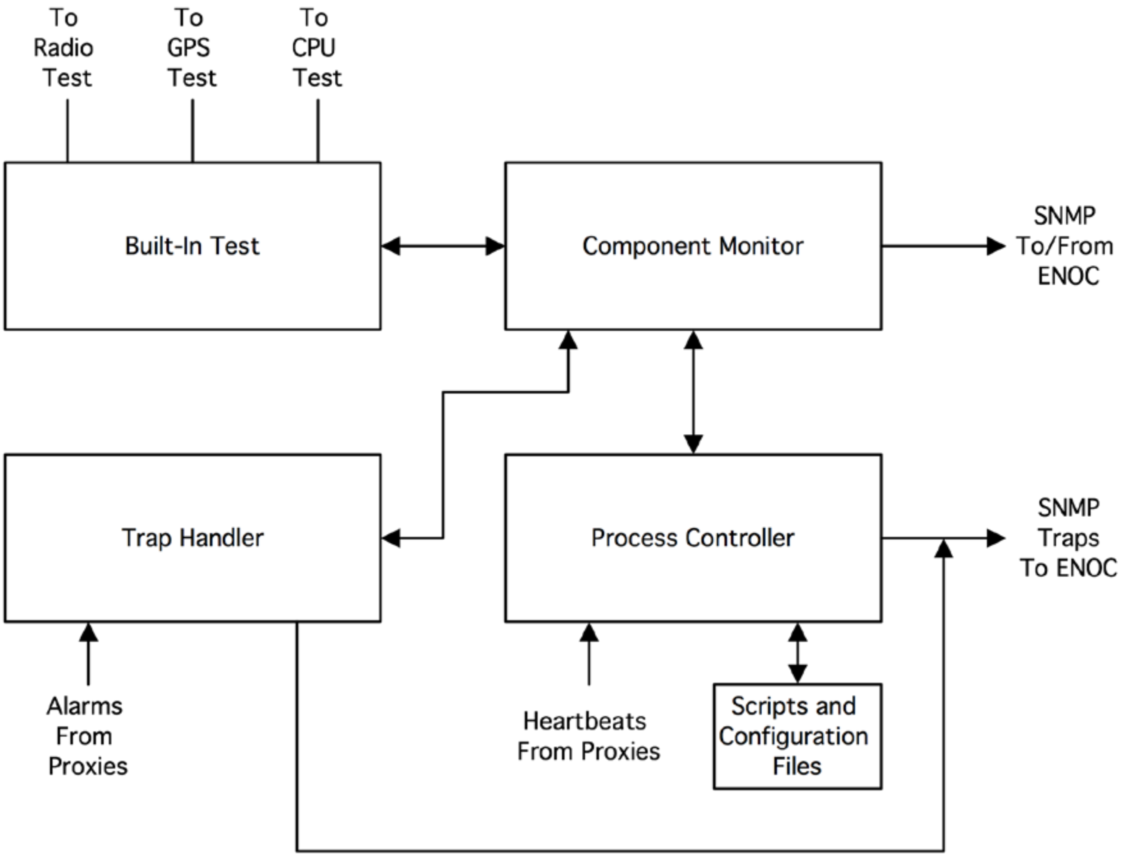

Figure 2-29: HMM Service

2.3.3.5.1 Built-In Test

The built-in test element runs one or more tests on a given hardware component within the RSE. Tests may or may not require the setting of test parameters by the test requester. Test results and expected values are returned along with an indication of which parameter(s) failed. Tests are timed for completion and execute sequentially because of limited processing resources. When the RSE runs diagnostic tests, it is not available to normal operations.

2.3.3.5.2 Component Monitoring

The component monitoring element monitors hardware components and determines their operational status and configuration for the purpose of evaluating and maintaining the RSE. The hardware components are examined through hardware ports and system logs to determine status. Key configuration data is reported along with status. Monitored components include DSRC radio, watchdog timer, memory, CPU, temperature sensor, and IP ports. While the requests originate from (or responses are sent to) the ENOC, the requests and responses flow through the HMM status request manager to the HMM component monitoring capability.

2.3.3.5.3 Process Controller

The process controller element monitors software processes and controls their runtime status and configuration for the purpose of evaluating and maintaining the RSE. The runtime status of software processes are monitored for heartbeat and controlled through system calls to the operating system. System logs are also examined for runtime warnings or errors. Key configuration data is reported along with process status. The RSE application processes are monitored along with critical operating system processes. The process controller can stop, start, restart, or reinitialize any RSE application process and change any configuration parameter within its control.

2.3.3.5.4 Trap Handler

The trap handler element monitors error and warning events on RSE proxies, services, and hardware handlers. When an error event is detected and is persistent, an unsolicited error trap is generated and sent out. When a critical system parameter, such as memory usage, reaches a pre-established threshold, a warning trap is automatically sent out. Traps contain enough information on the type of error/warning to enable external programs to begin resolution.

2.3.3.6 Bootstrap

The bootstrap service is part of the management services suite. This service orchestrates the boot process of the various proxies and services in the RSE. This service also provides the CPU status to the health monitoring service. It is composed of the two low-level processes.

2.3.3.6.1 Startup Loader

The startup loader element loads, initializes, and tests the RSE computing environment. Upon completion, the RSE is ready to start the service/proxy applications. Power-on integrity tests are run on the RSE hardware and operating system components. If any of the power-on tests fails, then the startup loader process records the failure and stops the bootstrap process. After successful completion of the power-on tests and operating system checks, the startup loader signals the start of applications.

2.3.3.6.2 Startup Element

The startup element launches the RSE services, handlers, and proxy applications. Before the applications are launched, the GPS configuration is checked and the RSE position data is stored. Startup then establishes a VPN connection to the SDN and obtains the initial RSE to SDN communication status. Once the RSE to SDN VPN is in place, the RSE can communicate to the ENOC (assuming the SDN to ENOC VPN connection is in place). Each RSE application is started in priority order, and each application is checked for successful start up. The SDN is sent a list of services (i.e., which applications are running). The RSE MIB is updated with startup results and a startup notification (trap) is sent to the ENOC. The startup notification trap signals the ENOC to fetch the completed startup results.

2.4 Infrastructure Services

While the SDN hosts most of the system services, these services are somewhat distributed across the network. Primary components are in fact hosted at the SDNs, but these components typically interact with network users and support services located at the RSEs. Because of this distributed nature of many of the services, they are described here in the larger distributed system context.

Many of the infrastructure services are based on what is known as a “publish-and-subscribe” architecture. This approach allows the decoupling of users and providers and significantly simplifies the process of managing a large number of both users and providers without creating an enormous amount of custom software. Using commercial-off-the-shelf (COTS) software elements in this publish-and-subscribe architecture, a wide variety of services can be created by developing only the provider elements. All message handling and transactional activities are provided by COTS messaging systems (often known as message queues, message switches, or brokers). Using this scheme, it is also possible to have brokers to act as both users and providers, so large distributed systems can be configured and/or extended without significant custom transactional software. This is shown in Figure 2-30.

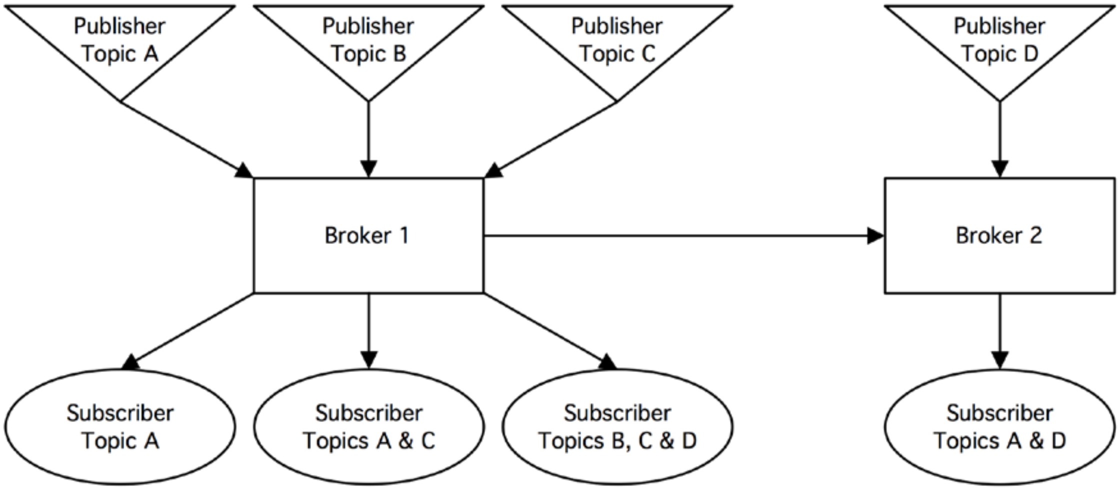

Figure 2-30: Publish and Subscribe Architecture

2.4.1 Communications Service

Communications services refer to network subsystem services that provide data transport between network users and public and private vehicles. Communications services include support for classes of service, domain name service (DNS), address assignment, and a consistent time source throughout the network.

The system must support different classes of service to enable the prioritization of network traffic to assure that safety and critical operational message traffic have priority over less-critical messages. The network subsystem uses conventional quality-of-service (QoS) mechanisms to support differential treatment of packets that flow thought the network. During periods of network message congestion, packets marked with a higher class of service, such as a flash-flood or hurricane warning, may receive priority over a lower class of service, such as a vehicle service message from an auto manufacturer. Likewise, messages intended to clear the way for an ambulance may be transmitted before packets executing a commercial electronic transaction. Packets carrying system (internal) control messages will receive the highest priority.

For the network subsystem to provide a differentiated treatment of packets in the VII network, the packets need to be classified and the headers appropriately marked so network devices and security devices can process the packets according to the specified class of service.

The POC network supports two classes of service (although this could be expanded as necessary).

Traffic entering and exiting the VII network subsystem may be subject to one or more traffic conditioning actions to support prioritization of specific classes of service. This is to ensure that the incoming traffic profile conforms to any rules specified in the network management policies. Traffic conditioning can involve delaying packets to bring the data stream in compliance with a traffic profile as well as marking of packet headers with values that translate into specific forwarding behaviors.

2.4.1.1 Traffic Forwarding

From a practical perspective, the service classes are implemented by a guaranteed amount of bandwidth to each class of service. In times of congestion, the network devices drop packets exceeding the bandwidth allocated for each class of service, so those classes with a higher bandwidth allocation are less likely to suffer packet losses or delays. Unused bandwidth assigned to a class of service may also be used to send traffic marked with other classes of service during periods of non-congestion—thus, optimizing the use of resources in the system. This QoS header-based packet forwarding behavior is used at QoS aware nodes throughout the network.

2.4.1.2 DNS

Because of the number of devices and the size of IPv6 addresses, an efficient DNS is essential for the proper end-to-end operation of the VII environment. Not only is it important that devices be given easy to remember names, but OBEs need to quickly determine how to access various services. Given that the physical addresses may vary from RSE to RSE, it is essential that names be clearly delimited. To facilitate this, the network uses a DNS system. The VII network DNS supports both IPv4 and IPv6 in the POC network. Requests from managed entities, including users and network devices, are authorized and filtered to maintain any required separations (such as between network users). Administrative users and system operators have the ability to create, read, update, and delete authorized DNS entries.

2.4.2 ILS

The ILS provides a way for external users and system administrators to obtain information about the system configuration. Figure 2-31 shows the ILS architecture.

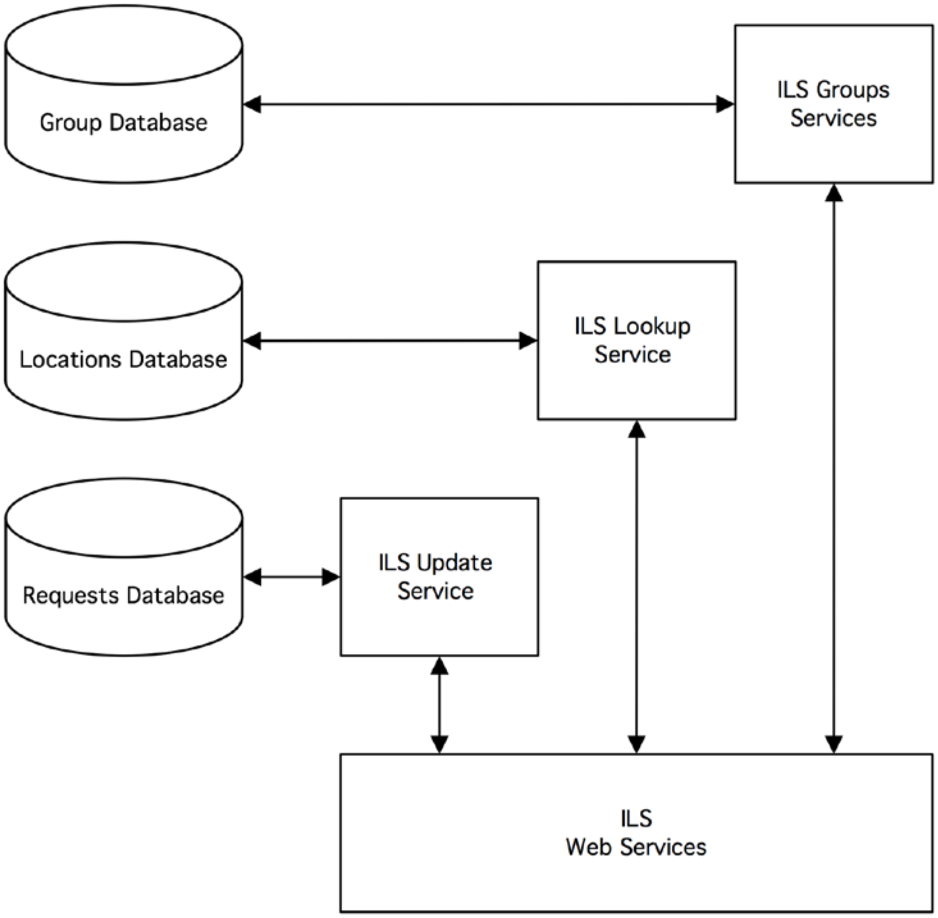

Figure 2-31: ILS Architecture

2.4.2.1 ILS Services

As shown in Figure 2-31, the ILS is composed of three basic services tied together using a conventional Web services broker.

2.4.2.1.1 ILS Lookup Service

The ILS Lookup Service provides functions to allow users and services to retrieve information about RSEs so they can effectively use the RSE network to distribute messages or to set up services at RSEs. The ILS Lookup Service allows users to search for information about the RSE network on the basis of geographic or political boundaries, by pre-defined groups, or by list requests. The information is provided in XML form validated by a shared data schema.

2.4.2.1.2 ILS Groups Service

The ILS Groups Service allows users to define groups of RSEs on the basis of various RSE parameters such as political boundary affiliation, geographic location, road affiliation, etc. By defining these groups, the user can then obtain information about the group or obtain updates about changes within that group automatically, and without needing to re-define the group again.

By allowing users to define groups of RSEs according to their needs, the ILS flexibly accommodates a wide range of user situations. For example, one user may be interested in RSEs within a given county or town boundary (political group), while another might be interested in a group of RSEs along a defined set of roads (geographic group). This system supports a wide range of choices defined and stored by the users.

2.4.2.1.3 ILS Updates Service

The ILS Updates Service automatically pushes information to internal VII services regarding important updates and changes in the RSE network. These updates are provided at defined and configurable intervals. Typical updates include:

- Changes in RSE privileges
- Changes in geographic boundaries
- Changes in political boundaries
- Addition of new RSEs
- Removal of existing RSEs.

2.4.2.2 ILS Databases

The ILS uses three database structures—the Groups Database, the Locations Database, and the Requests Database.

The Groups Database stores information about RSE groups custom created by a user. It uses a relational data structure that allows for groups of RSEs, groups of groups, etc. Ownership of the groups is maintained and synchronized with IdAM subsystem (see Section 2.5.1). Groups are given unique identifiers as well as human readable names to ensure data integrity.

The Locations Database maintains the current geographic location of RSEs within the VII system. This enables geographic query via a COTS plug-in and allows the ILS to provide a Geographic Information System (GIS) interface to administrative users. Requests can be made using a standard geographic schema, such as Geographic Markup Language (GML).

The Requests Database stores information about any request made to services that requires RSE information. Since the requests are stored, it can then interact with the Updates Service to inform VII services when the data about one or more RSEs in a request has been changed.

2.4.3 AMDS

The AMDS provides network users with the ability to send advisory messages to vehicles and public service vehicles in targeted geographic locations for a specified duration. Message dispatch is prioritized by type. The AMDS also supports cancellation of advisory messages by authorized network users.

2.4.3.1 Large-Scale AMDS Architecture

The AMDS uses the Java Messaging Service to implement its application transport protocol. At the heart of this architecture lies a message broker responsible for managing and routing TCP/IP messages. The message broker at each SDN subsystem allows for clustering to handle heavy message traffic loads. Each message broker advertises global subscriptions, allowing geographically disparate AMDS subsystems to intercommunicate with each other transparently. Meaning, the AMDS publishes to topics without requiring any geographic information for the message destination, effectively making a single logical publish-and-subscribe architecture across the VII system. This is illustrated in Figure 2-32.

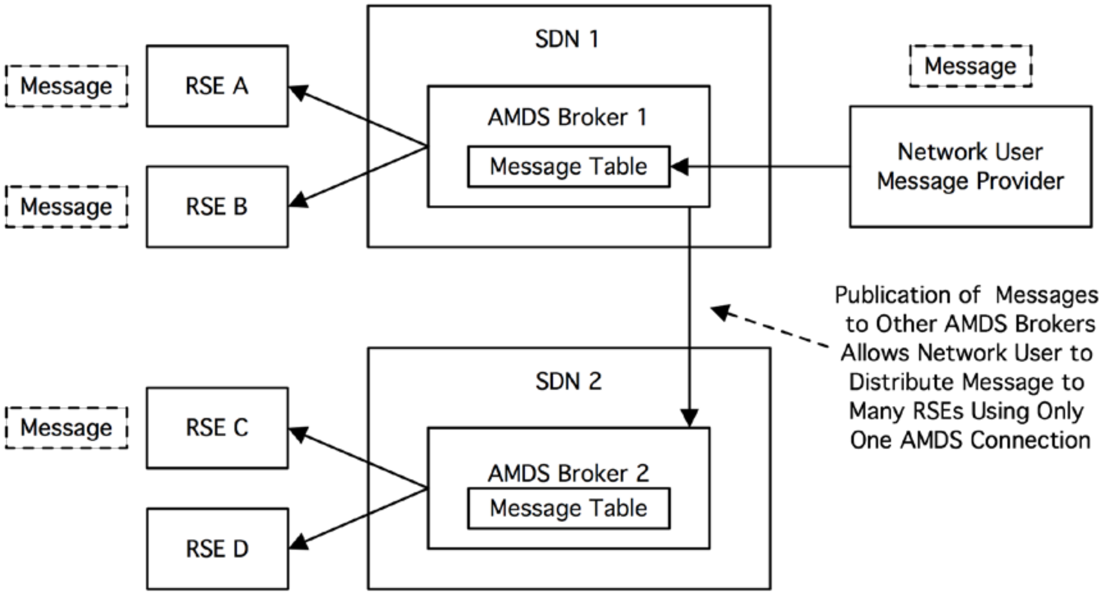

Figure 2-32: Large-Scale AMDS Architecture

2.4.3.2 Local AMDS Architecture

Figure 2-33 shows the AMDS architecture local to an SDN. This figure shows the message broker used to implement publication and subscription processes both locally to the RSEs and globally to AMDS instances at other SDNs.

Localized queuing is used at each SDN to buffer requests. Queuing provides an easy means of prioritizing messages at the application layer. During heavy messaging periods, as messages accumulate in the queue, higher-priority messages will preventively move to the front of the queue and be processed for delivery to RSEs before lower-priority messages. Additionally, a temporary response queue is created for each request destined to an RSE subsystem. This gives AMDS a simplified method of uniquely associating each response with a specific request and allows for the reduced complexity of managing messages and threads.

Since the message broker manages the actual transport and management of the messages, the actual interaction with the broker is transparent to the AMDS subsystem. The AMDS must maintain a connection to the message broker and monitor its queues to operate successfully. Thus, throughout the remainder of the document, the actual message broker discussion assumes that a message is either received on a temporary queue or published to a global topic.

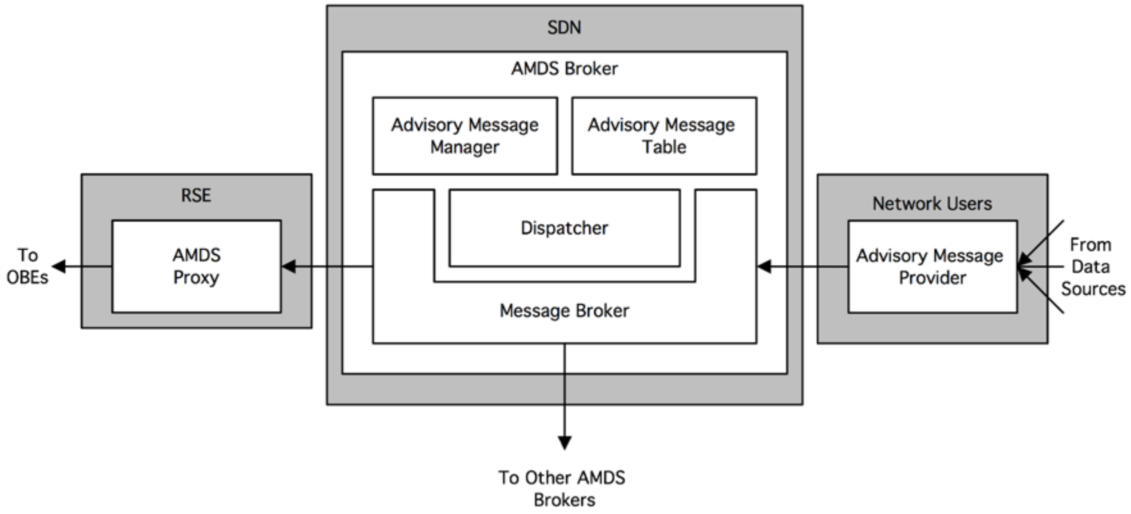

Figure 2-33: Local AMDS Architecture

The AMDS subsystem is designed to guarantee the delivery of valid messages received. After validation, the AMDS subsystem stores the message and its state or status in the Advisory Message Table (AMT). The AMT is the set of database tables related to all of the AMDS processing. The dispatcher and manager update their respective tables with parameter values pertinent to the data available and processing occurring for their respective roles. The manager maintains versions of the advisory messages, along with an indicator for the last operation attempted with the RSE and response details from the RSE, as well as the time a message was sent and when a response was received. This level of detail facilitates the guaranteed delivery of messages, since the advisory message processing state of the RSE is maintained within the AMT. If the delivery of the message is not possible, the AMDS subsystem retrieves the message from the AMT and delivers the message at a future date. This design aspect is particularly important for safety-related advisory messages, and every effort is made to deliver the advisory messages in a timely manner. Additionally, the AMDS subsystem emphasizes the importance of diligently tracking the state or status of every message request.

While the AMDS is a “guaranteed message delivery system,” a given RSE could be overwhelmed with messages from some large number of message providers. In this circumstance, the AMDS would deliver messages to the AMDS proxy at the subject RSE, and the combination of delivery timing and number of messages could be greater than the RSE would have time to actually transmit. In this case, lower-priority messages might never be sent. This situation is currently not addressed in the system. To do so will require revision of the AMDS and possibly creation of some sort of RSE load management subsystem.

2.4.3.3 Message Broker

The AMDS message broker is a message-oriented publish-and-subscribe system that interacts with network user advisory providers and the various RSEs it is associated with to receive and messages for publication, and to direct those to the AMDS proxies at the appropriate RSEs. The message broker works in conjunction with the other elements of the AMDS subsystem.

2.4.3.4 AMT

The AMT stores advisory messages published by providers together with the current state of the message. This information includes:

- Advisory message ID
- Targeted RSEs
- Advisory message contents
- Message duration
- Message activation time
- IP address of the advisory provider
- Port of the advisory provider.

2.4.3.5 Advisory Message Dispatcher

The advisory message dispatcher receives messages and delivery instructions from the advisory provider via the message broker. It then validates the messages and instructions and submits the message to the message broker with the proper topical information so that the message broker can publish it to the AMDS proxy subscribers.

The delivery instructions dictate when the SDN subsystem sends the advisory message to the corresponding AMDS proxies resident at the RSEs for distribution to passing OBEs. The delivery instructions also dictate the duration of distribution for the advisory message.

The advisory message dispatcher also supports updates, cancellations, and other administrative functions requested by the advisory provider.

2.4.3.6 AMDS Proxy

An instance of the AMDS proxy resides in each POC RSE. The AMDS proxy is an AMDS client that subscribes to the topics for its host RSE at the AMDS broker in the SDN that that RSE is associated with. When the AMDS broker receives an advisory message that contains delivery information indicating that it should be published to a particular RSE, the broker sends (publishes) the message and its associated delivery instructions to the AMDS proxy for that RSE.

Figure 2-34 illustrates the detailed structure of the AMDS proxy.

The RSE proxy manager controls the operation and configuration of the AMDS proxy (see Section 2.3.3.1).

2.4.3.6.1 Java Message Service (JMS) Connection Manager

The AMDS JMS connection manager subscribes to topics at the AMDS broker. When the AMDS broker has a message matching the AMDS proxy subscription, the message is sent to the AMDS proxy JMS connection manager. The JMS connection manager then passes the message and its associated instruction to the message handler.

2.4.3.6.2 Message Handler

The message handler receives the message elements from the JMS connection manager. The message content is separated from the delivery instructions and instructions are used to then either create a new playlist entry or update an existing one. The message content and the playlist information are then passed to the playlist manager.

2.4.3.6.3 Playlist Manager

The playlist manager maintains the active advisory message playlist based on instructions from the message handler. When the playlist changes, it sends the updated message content and broadcast information to the radio handler. The radio handler then advertises the AMDS service and broadcasts the message content on the CCH or SCH at the intervals defined in the playlist.

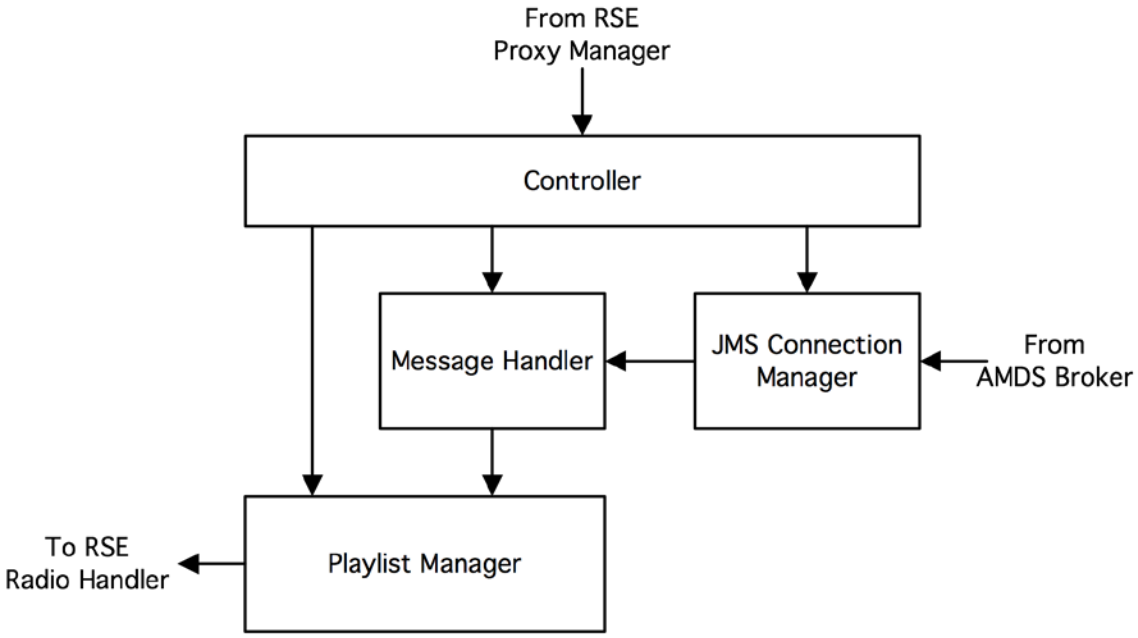

Figure 2-34: AMDS Proxy Components

2.4.3.7 Advisory Message Manager

The advisory message manager updates and controls the content of the AMT in order to delete messages that have expired and to manage the overall state of the AMDS broker.

2.4.4 PDS

The PDS consists of a geographically distributed group of logically connected servers that implement a publish-and-subscribe architecture for the collection of probe data from OBEs and the distribution of elements of that data to network users. These servers reside within the SDN subsystem and have the ability to communicate and share data with one another and to serve data to network users. The PDS accepts a stream of probe data messages collected from OBEs at the RSEs. It then parses these messages and places the various content elements (different probe data parameters such as speed, vehicle status, events, etc.) into queues that are structured along these topical categories. The data in these topical queues is then sent to network users that have established subscriptions on the basis of these topics. So, a network user that subscribes to Topic A at locations X, Y, and Z will receive any data associated with Topic A that is collected at any of the specified locations.

Figure 2-35 illustrates the general architecture of the PDS.

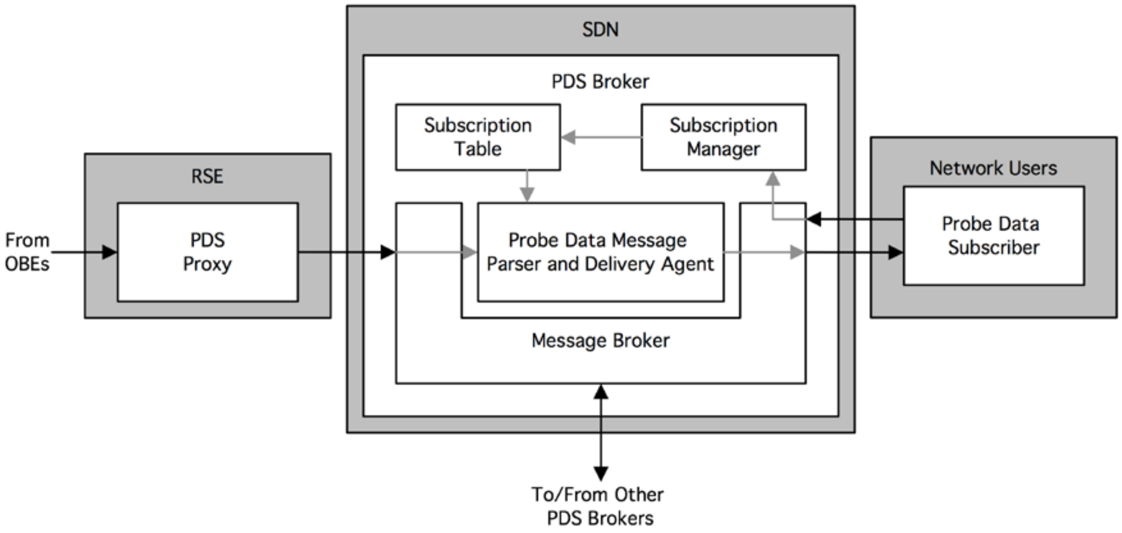

Figure 2-35: PDS Architecture

Other than temporary buffering to support parsing and subscription fulfillment, none of the probe data is stored in the system. It simply passes into the system, is transformed according to its content, and leaves the system. This avoids issues with data management, privacy, and scalability.

The PDS is comprised of five components—the probe data subscription manager, the probe data subscription table, the probe data message parser and delivery agent, the PDS proxy, and a message broker. The following sections discuss each individually.

2.4.4.1 Probe Data Proxy

The probe data proxy resides in the RSE (and will be active at any RSE currently configured to collect probe data). The probe data proxy receives probe data messages from passing OBEs. As it receives these, it publishes the message using the publish-and-subscribe architecture described in Section 2.4. In this case, the PDS system at the SDN is the subscriber, and the "topic" is the specific RSE that hosts the probe data proxy. Figure 2-36 shows the architecture for the probe data proxy.

When the probe data message receiver receives OBE-originated probe data messages from the RSE radio handler, it places them in a message queue. This message queue allows the probe data proxy to balance the different rates at which the probe data messages may be received against the rate that the RSE backhaul and PDS system can support. In situations where many OBEs are delivering probe data at a single RSE, the message rate may exceed the ability of the backhaul and PDS to accept them. In this case, the messages simply collect in the queue and are then passed out at whatever rate the back end of the system can support. No messages are lost. When it has sent one message, the probe data message publisher picks the next message off the top of the queue (thus, moving all messages in the queue up one spot), and publishes the message to the PDS broker at the SDN.

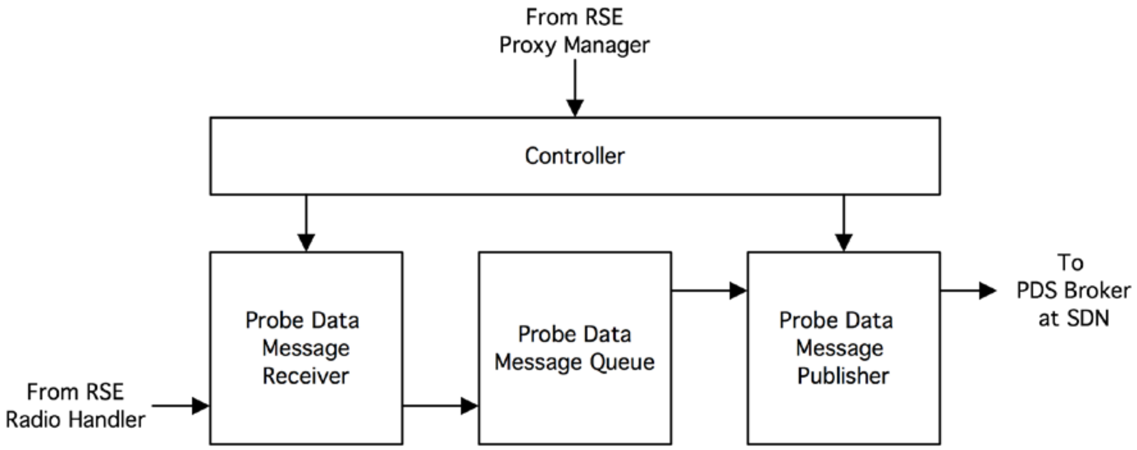

Figure 2-36: Probe Data Proxy Architecture

2.4.4.2 Probe Data Subscription Manager

The probe data subscription manager provides an interface for network users to establish, update (edit), query, or delete probe data subscriptions. In addition to the usual user (subscriber) registration information, subscriptions consist of definitions of topics that the probe data system can provide. Topics typically include the various data parameters contained in the probe data message (such as speed or wiper status, or event indications such as hard braking or traction control system activation). All content returned for a topic typically includes location and time so that they are meaningful in the roadway context, but subscriptions can also include location as a parameter. For example, a network user might subscribe to speed at a particular location, or might subscribe to all traction control events within a given region (e.g., their region of responsibility).

Once established by the network user, the probe data subscription manager stores and maintains the subscription information in the probe data subscription table.

2.4.4.3 Probe Data Subscription Table

The probe data subscription table provides persistent storage of network users’ subscription requests. This allows users to maintain subscriptions over long periods of time.

2.4.4.4 Probe Data Message Parser and Delivery Agent

The probe data message parser and delivery agent is the core of the PDS. As probe data messages are received from the probe data proxies hosted in the RSEs, the probe data message parser separates the components of the probe data message payload. As shown in Figure 2-37, each probe data message consists of a series of probe data snapshots, and each snapshot includes the various parameters reported by the source vehicle (not all vehicles support all data parameters).

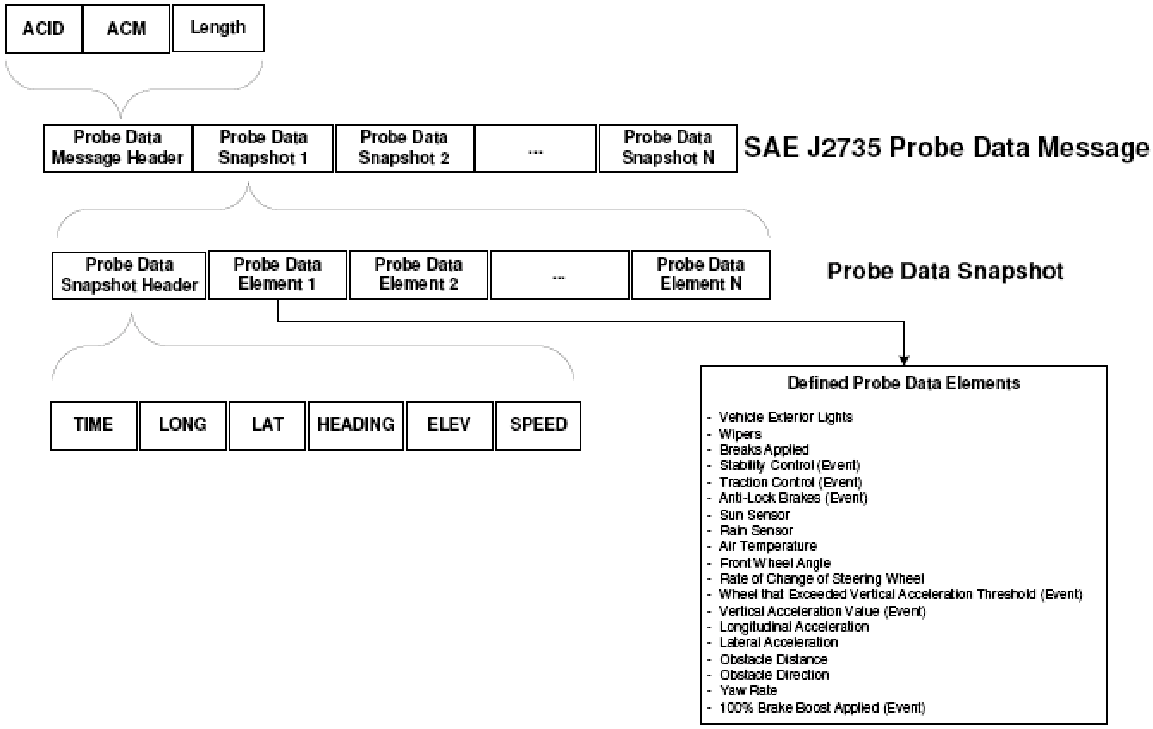
ACID
ACM
Length
Probe Data Message Header
Probe Data Snapshot 1
Probe Data Snapshot 2
...
Probe Data Snapshot N
SAE J2735 Probe Data Message
Probe Data Snapshot Header
Probe Data Element 1
Probe Data Element 2
...
Probe Data Element N
Probe Data Snapshot
TIME
LONG
LAT
HEADING
ELEV
SPEED
Defined Probe Data Elements
- Vehicle Exterior Lights
- Wipers
- Breaks Applied
- Stability Control (Event)
- Traction Control (Event)
- Anti-Lock Brakes (Event)
- Sun Sensor
- Rain Sensor
- Air Temperature
- Front Wheel Angle
- Rate of Change of Steering Wheel
- Wheel that Exceeded Vertical Acceleration Threshold (Event)
- Vertical Acceleration Value (Event)
- Longitudinal Acceleration
- Lateral Acceleration
- Obstacle Distance
- Obstacle Direction
- Yaw Rate
- 100% Brake Boost Applied (Event)

Figure 2-37: Probe Data Message Structure

The separated, or parsed, data elements are then published to the message broker where they are distributed to the network users on the basis of their subscriptions.

2.4.4.5 Message Broker

As described in Section 2.4.4, the PDS message broker is implemented using a publish-and-subscribe system. In the PDS, the probe data message parser and delivery agent subscribes to a topic that corresponds to each RSE associated with the SDN that hosts the PDS. In addition, probe data message parser and delivery agents in other SDNs may also subscribe to those same RSE topics (assuming they have a network user subscription that requests them). When the message broker receives a message from a probe data proxy at a given RSE, it then publishes that message to any subscribers, which includes the local probe data message parser and delivery agent, and may include other SDNs.

The message broker also acts as a publisher. When it receives topical content from the probe data message parser and delivery agent, it sends that content to each of the network users that have subscribed to that particular content.

2.4.4.6 Probe Data Management

Probe data management is a secondary function of the PDS. This subsystem uses the AMDS to distribute a probe data management directive message (a special type of advisory message) that is used by the OBE to change the parameter content and collection rates for probe data messages in a given area.

2.4.5 MEDS

The MEDS is a meta-service that relies on the PDS and the AMDS. The role of the MEDS is to use collected probe data in a given region (e.g., an intersection or a road curve) to dynamically generate maps indicating the actual average geometric travel lines for vehicles on that road section. Using this data, the MEDS generates a Geometric Intersection Description (GID) that includes lane positions and limit lines (the places where vehicles stop prior to entering the intersection). These GIDs are expected to be used for intersection collision avoidance functions as well as various types of dynamic traffic load based intersection signal control, and other similar applications.

A MEDS instantiation is located at each SDN subsystem. The core MEDS functionality is the geostatistical analysis of probe data elements to derive new road geometry descriptions. Probe data elements received from the PDS are analyzed against a baseline transportation dataset (currently NAVTEQ) in order to determine potential geometry changes. As changes are detected and approved by authorized administrative users, the AMDS compiles the map updates into the GID format and then distributes them to RSEs at the appropriate roadway locations. A MEDS proxy at the RSE handles broadcasting the appropriate GID to the OBE for use in safety applications.

Figure 2-38 illustrates this concept of using vehicle probe data to generate road geometry information by showing raw data used in the MEDS representing a freeway interchange and surrounding roads.

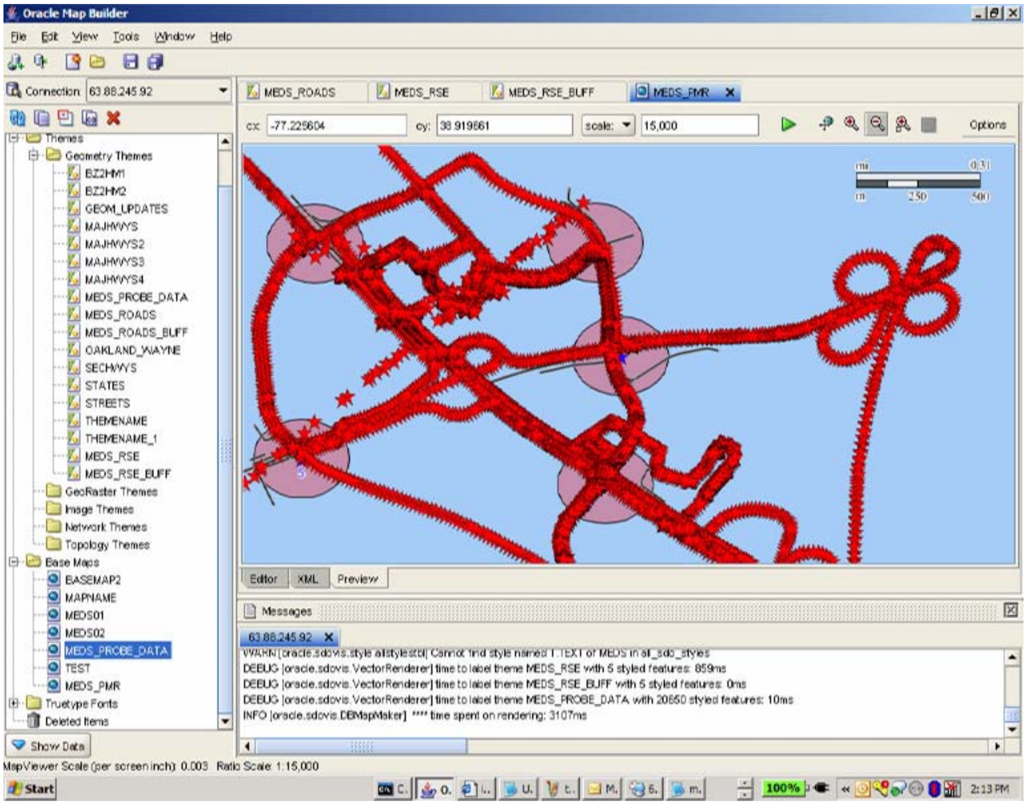

```
WARN [oracle.sdovis.style allstylestbl] Cannot find style named T.TEXT of MEDS in all_sdo_styles
DEBUG [oracle.sdovis.VectorRenderer] time to label theme MEDS_RSE with 5 styled features: 859ms
DEBUG [oracle.sdovis.VectorRenderer] time to label theme MEDS_RSE_BUFF with 5 styled features: 0ms
DEBUG [oracle.sdovis.VectorRenderer] time to label theme MEDS_PROBE_DATA with 20650 styled features: 10ms
INFO [oracle.sdovis.DBMapMaker] **** time spent on rendering: 3107ms
```

MapViewer Scale (per screen inch): 0.003 Ratio Scale: 1:15,000

Figure 2-38: Probe Data Representation of a Freeway Interchange

2.4.5.1 MEDS Architecture

Figure 2-39 illustrates the MEDS architecture.

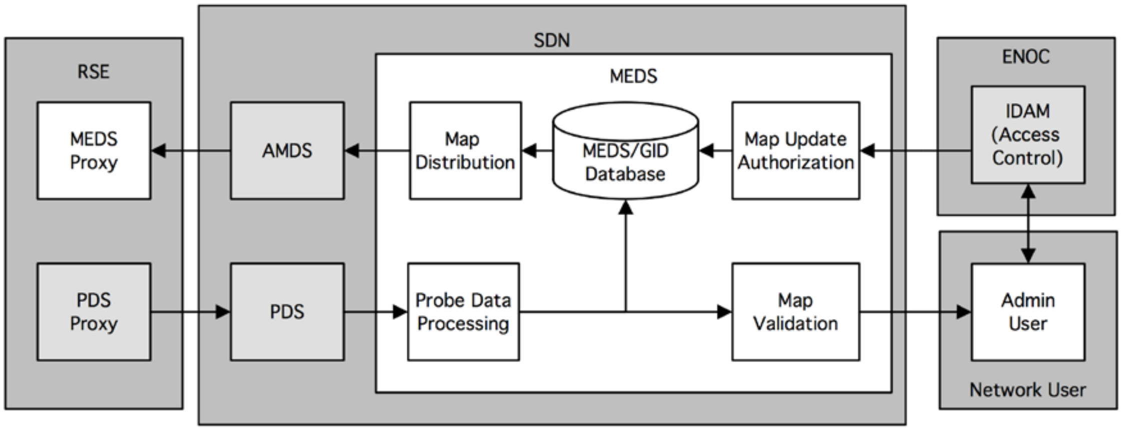

Figure 2-39: MEDS Architecture

The MEDS consists of the following elements.

2.4.5.1.1 Probe Data Processing

The probe data processing element subscribes to the PDS to receive probe data associated with the intersection or road segment in question. The data parameters are probe sequence number, latitude, longitude, elevation, speed, heading, and time.

The processing element temporarily stores this data until a statistical threshold is reached, at which time it can assure a reasonable level of accuracy and validity. Using a geostatistical algorithm, it identifies the road segments associated with consistent statistical trends for vehicle paths, and then converts this into a GID form. It deletes the probe data elements upon completion of geospatial analytics. The resulting GID is delivered to the map validation and update element.

2.4.5.1.2 Map Validation and Update

The map validation and update element stores the GID provided by the processing element as a pending map update in the administrative approval queue. An administrative user is required to inspect the pending map update before it can be delivered to the system for active use. This step assures that maps used for safety purposes are accurate and dependable.

2.4.5.1.3 Map Update Authorization

The map update authorization element provides the user interface and notifications for the administrative user to inspect and approve pending maps updates. It interfaces with the ENOC to request administrative user approval and delivers a notification message and “pending” updated map in GML format. Upon administrative user approval, it writes pending GIDs to the MEDS database.

2.4.5.1.4 MEDS Database

The MEDS database stores all updated GIDs for distribution.

2.4.5.1.5 Map Distribution

The map distribution element interacts with the AMDS to arrange for delivery of updated maps and GIDs to the appropriate RSEs. Using the approach described in Section 2.4.3, the AMDS delivers the GIDs to the MEDS proxy resident in each of the RSEs identified by the map distribution element.

2.4.5.1.6 MEDS Proxy

The MEDS proxy receives the GID and delivery instructions from the AMDS. It then interacts with the radio handler in the RSE as described in Section 2.4.3 to arrange for the GID/map to be broadcast to local OBEs.

2.4.6 IdAM/Security Services

The SDN also hosts certain elements of the IdAM subsystem. These functions are distributed between the certificate authority, the ENOC, and the SDN. Please see Section 2.5 for a complete description.

2.4.7 NMS

The NMS collects, aggregates, and forwards network management data from MNEs (see Section 2.5) to the ENOC. The NMS also transmits network management commands from the ENOC to the MNEs. The MNEs managed by the NMS at the SDN include the RSEs connected to the SDN, the backhaul links, backbone links, access links, and SDN gateways, as well as the Event Notification System (ENS), AMDS, MEDS, POS, and ILS components within the SDN.

The NMS facilitates local on-site management and troubleshooting at the SDN as well as remote management and troubleshooting from the ENOC.

The management services supported by the SDN are as follows:

- ***Fault Management*** – The detection, isolation, and correction of faults in VII MNEs
- ***Performance Management*** – The collection of performance data from the VII MNEs and generation of performance reports
- ***Configuration Management*** – The remote management of the configuration of VII MNEs for provisioning, fault correction, and asset management purposes
- ***Accounting Management*** – Measurement of the utilization of information technology resources by VII system users
- ***Security Management*** – The prevention, detection, and correction of security-related events in VII MNEs.

2.4.8 POS

The POS is located at each SDN and is intended to provide a network-based source of differential global positioning system (DGPS) data.

The POS receives data from external sources (in the POC, this was a small-scale High-Accuracy Differential GPS, or HA-NDGPS, station) and delivers it via the RSEs to OBEs that can then use the data to improve the accuracy of their GPS position estimates. This eliminates the need for the OBE to support additional radio interfaces in order to receive differential corrections.

Figure 2-40 illustrates the POS architecture.

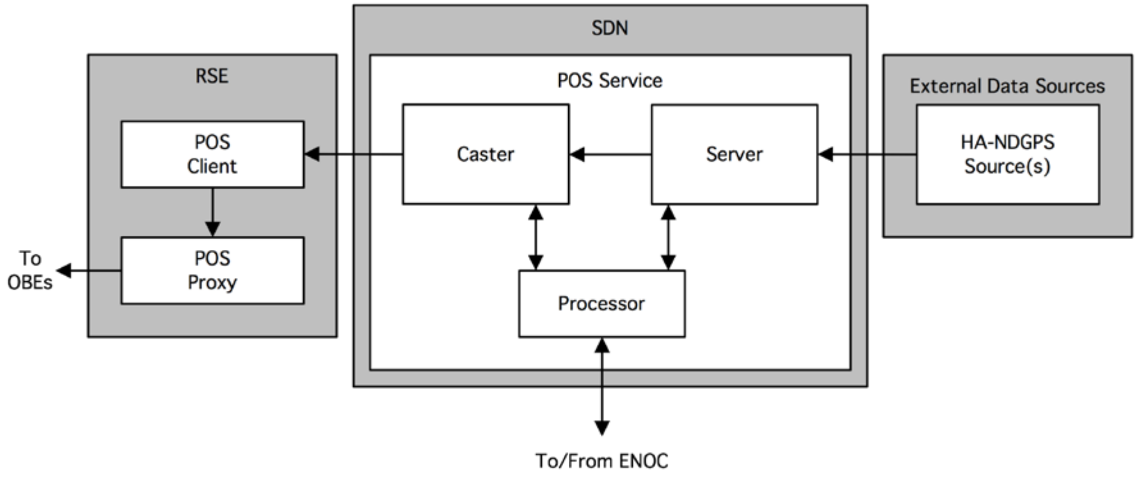

Figure 2-40: POS Architecture

2.4.8.1 Caster

The caster maintains a list of available external data source correction mount points and monitors the availability of the various correction sources. When it receives a request for a particular source from an RSE client, it directs the server to connect to that source. When a data stream is available, it streams the corrections from the source using HTTP, and sends that data to the requesting POS client in the RSE.

2.4.8.2 Server

The server establishes and manages the communications with the external data sources under direction from the caster. It also manages internal transfer of the Radio Technical Commission for Maritime Services formatted data stream.

2.4.8.3 Processor

The processor interacts with the ENOC and logs fault conditions.

2.4.8.4 Client

The client resides on the RSE. Based on the RSE position and the currently available external data sources provided by the caster, the client requests a specific data source. It then passes correction data received from the caster to the POS proxy for broadcast to the OBEs.

2.4.8.5 Proxy

The POS proxy resides on the RSE. It establishes the link to the RSE radio handler and RSE proxy manager. It also registers the position of the RSE and forwards this information to the ENOC via the processor in the POS service. This is used to support the geographic scope element of the security functions. If the location of the RSE changes, its security certificates will have an incorrect geographic scope, and the messages will be considered invalid.

2.5 Security Subsystem

The POC system uses a distributed security subsystem to provide access control, user and message verification and authentication, security monitoring, and attack mitigation and response. The security subsystem is composed of four major components.

- IdAM
- Security management services
- Security services
- Certificate authority(ies).

Figure 2-41 shows the high-level architecture of the security subsystem. The following sections describe the interrelated elements.

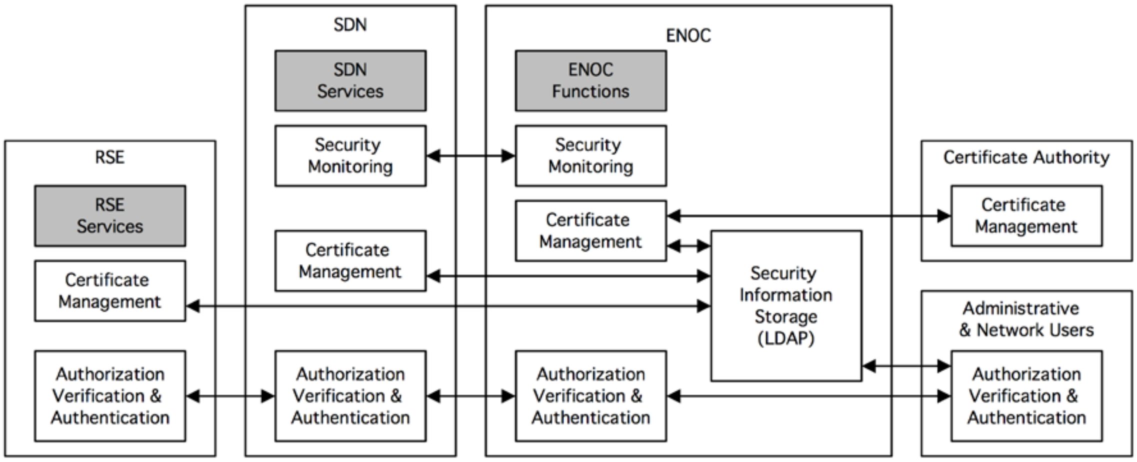
SDN
SDN Services
Security Monitoring
Certificate Management
Authorization Verification & Authentication
ENOC
ENOC Functions
Security Monitoring
Certificate Management
Security Information Storage (LDAP)
Authorization Verification & Authentication
RSE
RSE Services
Certificate Management
Authorization Verification & Authentication
Certificate Authority
Certificate Management
Administrative & Network Users
Authorization Verification & Authentication

Figure 2-41: VII Network Security Architecture

2.5.1 IdAM

The IdAM subsystem overlays security functions on the SDN, the RSEs, the ENOC, and external users (administrative users and network users). The IdAM provides security management functions as well as certificate-based access control for users.

The IdAM works with registered entities. A registered entity is a person, organization, hardware device, or software process that is known by the VII system. It includes network users, administrative users, RSE subsystems, the ENOC subsystem, SDN subsystems, certificate authority subsystems, reference map providers, and service provider management systems. The primary role of the IdAM is to provide security services for registered entities within the VII infrastructure.

From an administrative perspective, the IdAM provides a repository for registered entity data while ensuring that the privacy and security of identity data is maintained. It provides secure credentials to registered entities through an x.509 certificate authority, and it provides a framework for managing entities throughout their lifecycle (from initial registration to de-provisioning).

Operationally, the IdAM provides access management including authentication and authorization services to network users, administrative users, subscription providers, and RSEs. In addition, it provides a framework for enforcing access control policies on users, RSE groups, political boundaries, and geographic boundaries.

The IdAM service does not include the 1609.2 certificate authority that is used to secure messages that travel over the DSRC link. The 1609.2 security functions are described in detail in the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.

2.5.2 Security Services

The VII security service manages the security of VII assets and data transmitted across the network. It consists of an integrated suite of tools, applications, and systems that remotely monitor, provision, configure, repair, and maintain the managed security elements of the VII system. It detects, isolates, and takes action to mitigate security breaches. The VII security service provides for "defense-in-depth" protection of its hardware, software, and information components against hackers, unauthorized users, and physical or electronic sabotage.

The VII security service also maintains network administration control of authorized user privileges and protects the confidentiality of user information. It supports centralized and decentralized IdAM services for the network. Synchronization of storage repositories to support access management services provide for failover and scalability across the network.

Access management services provide authentication and authorization services to requesting managed entities that are delivering real-time authentication and policy- and rule-based authorization.

Provisioning services support managed entity life-cycle processes, including certificate life-cycle management for X.509 PKI certificates.

2.5.2.1 Security Event Monitoring Function

The security event monitoring function allows the ENOC to receive, aggregate, correlate, and report indications of malicious activity from MNEs. This function communicates with MNEs via the ENOC-to-MNE interface to obtain security event notifications.

2.5.2.2 Security Authentication Function

The security authentication function provides the authentication to internal and external VII entities. The ENOC subsystem accepts a request for authentication in the form of a credential. In most cases, the authentication token used throughout this network is either an IEEE 1609.2 or an X.509 certificate. To authenticate the certificate holding entity, entities are required to authenticate to the ENOC subsystem that will ensure the token's validity. The ENOC subsystem will return a message indicating the entity's authentication status.

2.5.2.3 Security Authorization Function

The security authorization function provides authorization to entities across the VII network. This function can be called by various components to determine access rights and authorization of entities requesting information or services throughout the network. The authorization decisions are based on defined policy stored within the ENOC subsystem.

2.5.2.4 Security Entity Lifecycle Management Function

The security entity lifecycle management function provides security management capabilities at the SDN. The management process provides the service for SDN registered entities to have security credentials provisioned, de-provisioned, updated, and revoked. The SDN subsystem also controls the registration or de-registration of constituent network users relative to the access of SDN-based services. Authorization policies surrounding access control are stored locally for the SDN and managed within the SDN subsystem. Through the management process, administrative users may request changes in service offerings within their jurisdictions.

2.5.2.5 Security Synchronization Function

The security synchronization function passes data between the ENOC subsystem and the SDN subsystem. A local instance of the data resides at the SDN subsystem and is maintained on a regularly scheduled basis to remain consistent with the ENOC subsystem. The synchronized data includes identity data to be used for authentication, authorization, and entity life-cycle management purposes. In general, the ENOC-based security information storage system provide a central mechanism for the storage and update of security credentials as part of the managed entity functions that form the core of the ENOC operations.

2.5.3 Certificate Authority

The certificate authority is a central point of trust in the system. The certificate authority provides certificates to OBEs that attest to the authenticity and legitimacy of an OBE for use in signing both identified and anonymous messages, and provide certificates to other users to allow them to exchange signed and encrypted messages with OBE applications.

In the POC, the certificate authorities are separated into those that provide conventional Internet-style certificates defined by the X.509 standard and those that provide certificates used for over-the-air messages. The reason for using these two different categories of certificates is that X.509 certificates are standard and in wide use across the Internet; therefore, most equipment and COTS software are compatible. For over-the-air credentials, the X.509 certificate is too large (about 1,400 bytes). Since many VII DSRC messages are only a few hundred bytes (e.g., the Heartbeat message is typically less than 200 bytes), the smaller certificates defined by the 1609.2 standard were utilized. This results in approximately 100-percent security overhead for these smallest message types (in contrast, using a 1,400-byte X.509 certificate would result in an 800-percent security overhead). The 1609.2 certificate authority is discussed briefly in Section 2.5.3.1.4 below, and in detail in the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.

2.5.3.1 X.509 Certificate Authority

The certificate authority subsystem functions closely with the ENOC subsystem in the management of X.509 and IEEE 1609.2 certificates. The certificate authority supports the enablement of secure VII communication by providing X.509 certificates to managed entities in the VII system. The lifecycle management of the certificates is initiated by the ENOC subsystem and directly interacts through the certificate authority subsystem. The ENOC is responsible for X.509 certificate lifecycle management, certificate storage, and maintaining updated copies of the certificate revocation lists (CRLs). Updated CRLs ensure that the VII system security can restrict revoked entities from accessing VII services and resources. The ENOC subsystem interacts closely with the certificate authority subsystem to manage the lifecycle of the PKI certificates. The certificate authority is responsible for issuing and registering certificates to requesting entities. The certificate authority also maintains an updated CRL, whereas the

ENOC subsystem maintains a copy of this CRL for access control purposes. Incoming certificate lifecycle requests are intercepted by the ENOC subsystem and ultimately reach the certificate authority subsystem, where the decision is made by a registration authority (RA) on whether to issue, revoke, or update the entity's certificate.

2.5.3.1.1 X.509 Certificate Lifecycle Management

Certificate lifecycle management issues, revokes, and updates certificates and key pairs for VII devices. Similar to key management requirements for VII, certificate management includes integrated lifecycle management (provisioning and de-provisioning) for anticipated large numbers of certificates.

The X.509 certificates are used by network users, administrative users, RSEs, and subsystems to securely send messages through the non-wireless elements of the VII environment. Lifecycle management for X.509 certificates includes initial registration and certification of certificate(s), key pair recovery, key pair updates, certificate updates, and revocation requests.

2.5.3.1.2 X.509 Certificate Standards

The X.509 certificate authority manages the certificates in correspondence with the standards set forth by the Internet Engineering Task Force in the *Request for Comments (RFC) 4210 Internet X.509 Public Key Infrastructure Certificate Management Protocols*. More information surrounding the certificate lifecycle process is also detailed in the *VII System Infrastructure System Security Plan (SSP)* and the *RFC 5280 Internet X.509 Public Key Infrastructure (PKI) Proxy Certificate Profile*.

2.5.3.1.3 X.509 Certificate Lifecycle

Within the VII network, X.509 certificate lifecycle management requests include requests for updates, creation, registration, and revocation of certificates used throughout the VII system. Requests may be submitted to the ENOC subsystem by network and administrative users, as well as VII components or subsystems. The IdAM component within the ENOC subsystem validates the request. Upon authorization, the message is forwarded to the certificate authority. The certificate authority reviews and processes the request. In the case of request for entity registration, an RA or local registration authority assures that the identity is in fact who they are claiming to be (identity proofing) and is authorized to perform the requested certificate action. The end result of this identity proofing process is the issuance of a certificate for the end entity's public key. The certificate is posted in the IdAM repository.

Entity registration is only one of the certificate lifecycle activities that occur within the certificate authority subsystem. Requests for certificate registration and certification, key pair recovery, key pair updates, certificate updates, and revocation requests are all supported by the certificate authority subsystem with reliance on the ENOC subsystem for lifecycle management activities.

2.5.3.1.4 IEEE 1609.2 Certificate Lifecycle Management

IEEE 1609.2 certificates are used to support vehicle-to-infrastructure, as well as vehicle-to-vehicle communications, for a variety of vehicle safety applications and transportation operations. These certificates are supplied by a special 1609.2 certificate authority set up to support the POC. The functional elements of the Native 1609.2 Security Services Secure Messages include certificate request, receive certificate response, prepare outgoing secured message, process incoming secured message, prepare outgoing secured WSM, process incoming secured WSM, process root certificate update message, and process CRL. Lifecycle management for certificates includes certificate issuance, renewal, and revocation. The 1609.2 certificates are used by RSE and OBE systems.

For a detailed description of the 1609.2 certificate authority as well as the DSRC security operations, please see the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.

2.6 DTE

For the POC project, a small-scale version of the VII system was implemented in the greater Detroit area. This implementation included an SDN and 55 RSEs placed at arterial intersections and at various freeway locations. The RSE locations cover 45 square miles and represent 75 centerline miles of roadway.

The DTE also included a mini ENOC located in Novi, MI (adjacent to the Vehicle Infrastructure Integration Consortium, or VIIC, Headquarters). This allowed efficient interchange between the VIIC development and test teams and the Booz Allen infrastructure team development and test teams, and allowed rapid configuration and control of the RSEs during system integration.

2.6.1 Physical Layout/Map

Figure 2-42 shows the location of the DTE arterial RSEs, and Figure 2-43 shows the Freeway RSEs.

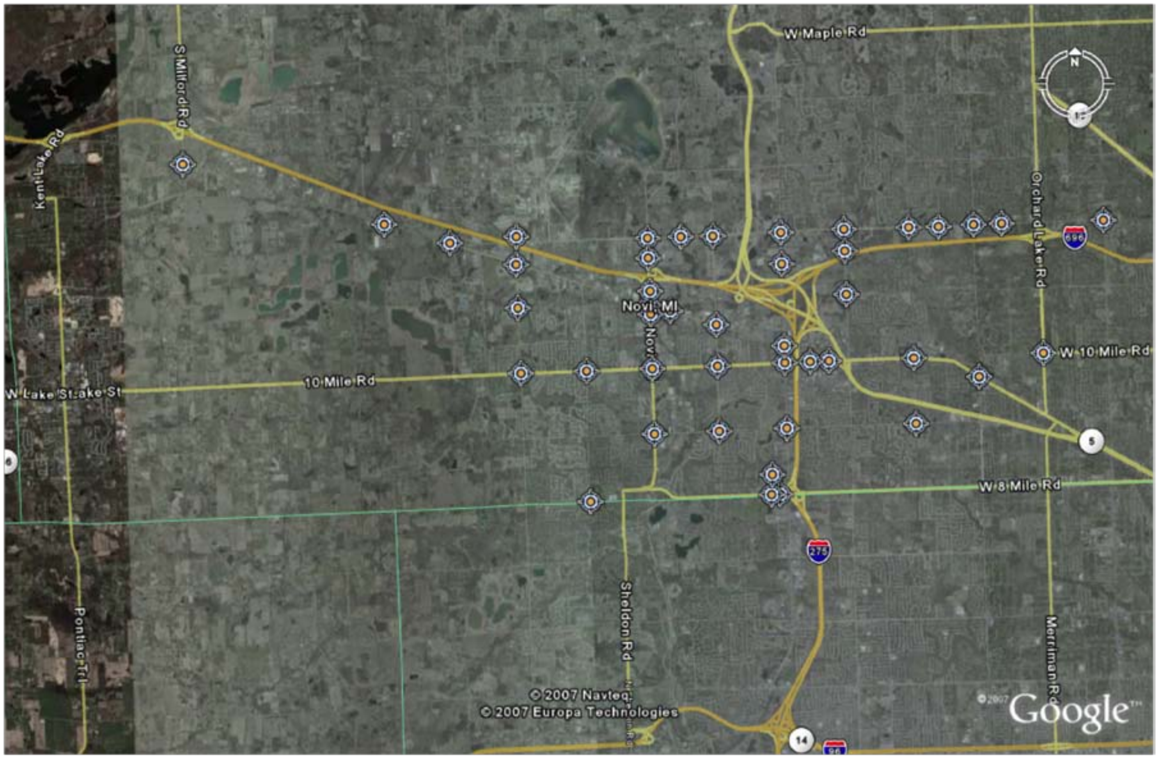

Booz Allen generated plot points on Google map

Figure 2-42: Arterial RSE Locations

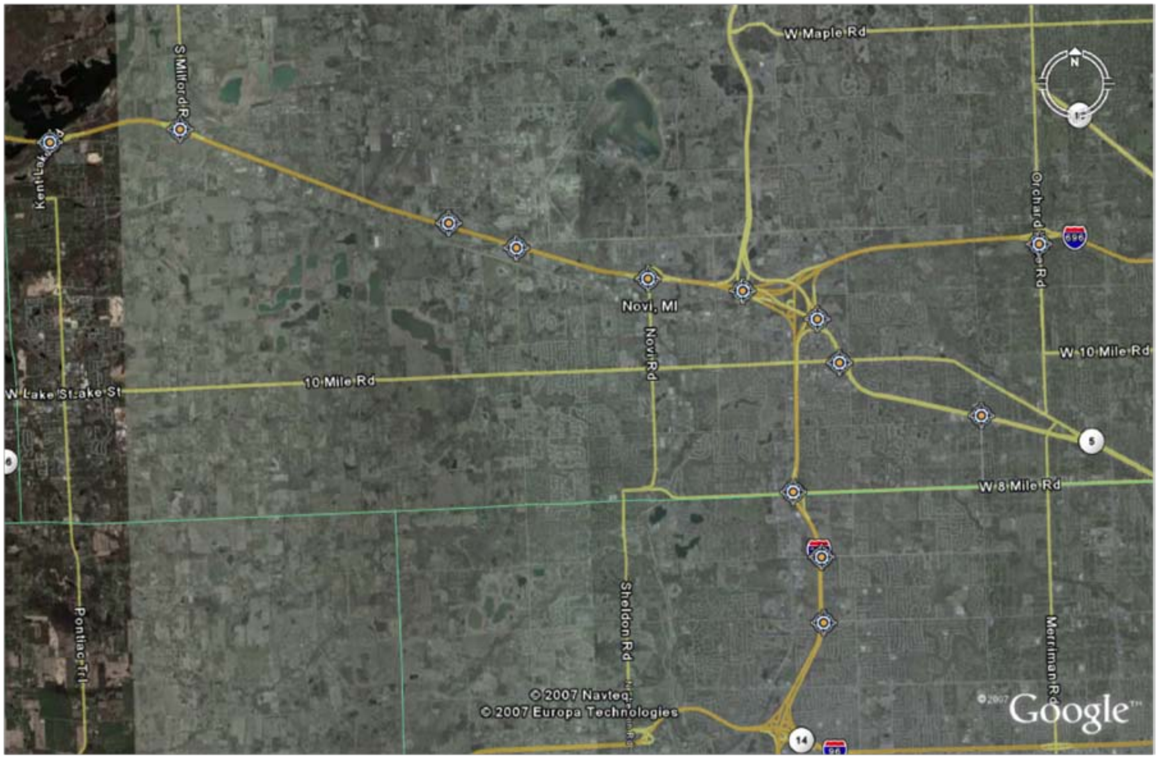

Booz Allen generated plot points on Google map

Figure 2-43: Freeway RSE Locations

The RSE locations were chosen to facilitate a variety of application and system tests. The DTE included two relatively long freeway routes that allowed for interactions with OBEs passing at high speeds. Many RSEs were located at arterial intersections. This provided an opportunity to collect arterial probe data (together with freeway probe data) and allowed testing of a variety of other applications.

Figure 2-44 shows the RSE regions used to test PDC. Figure 2-45 shows the RSEs used to test advisory message delivery and presentation, and Figure 2-46 shows the freeway on-ramp region used for tolling testing.

Other tests were performed using a variety of arterial intersections in the DTE.

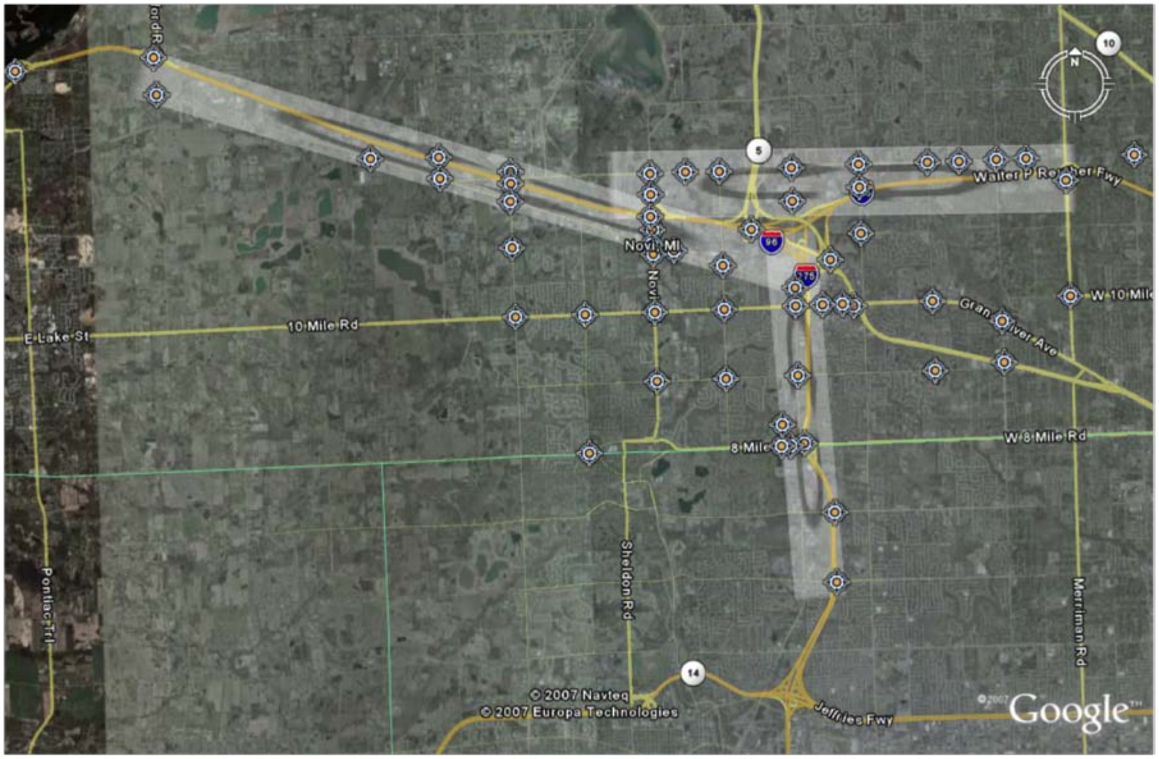

Booz Allen generated plot points on Google map

Figure 2-44: RSE Regions Used for PDC Test

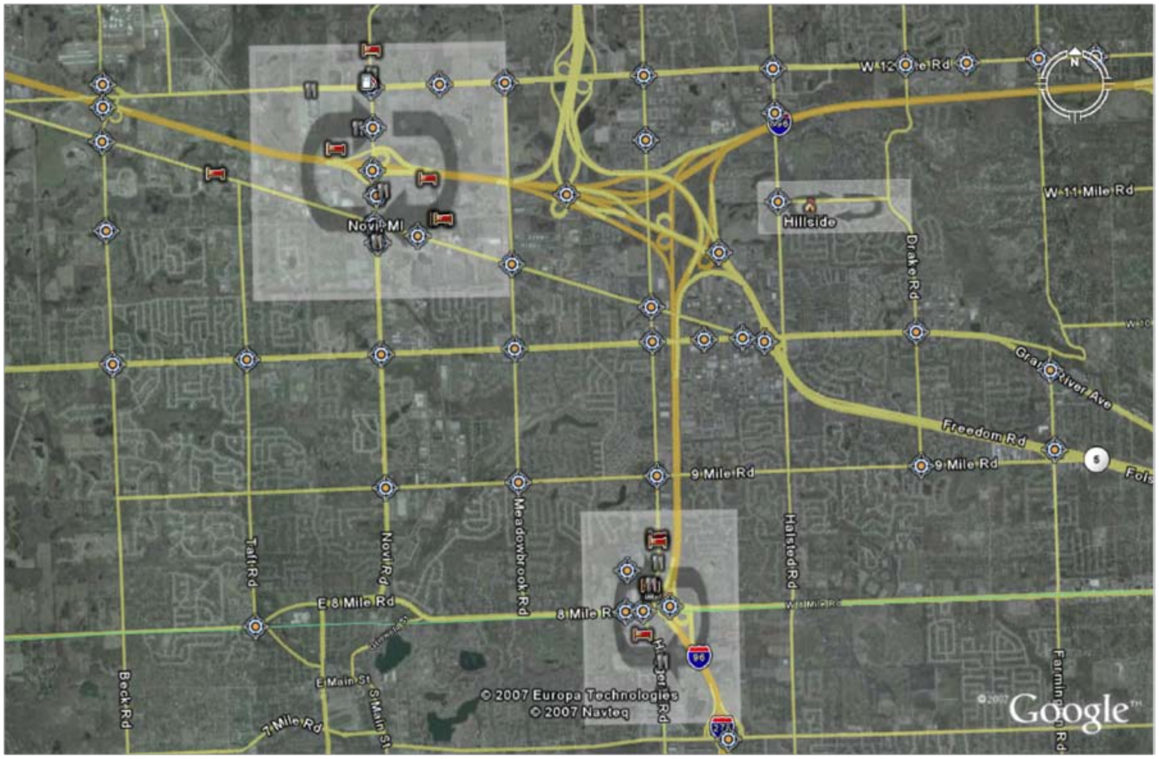

Booz Allen generated plot points on Google map

Figure 2-45: RSE Regions Used for Advisory Message Delivery and Presentation Test

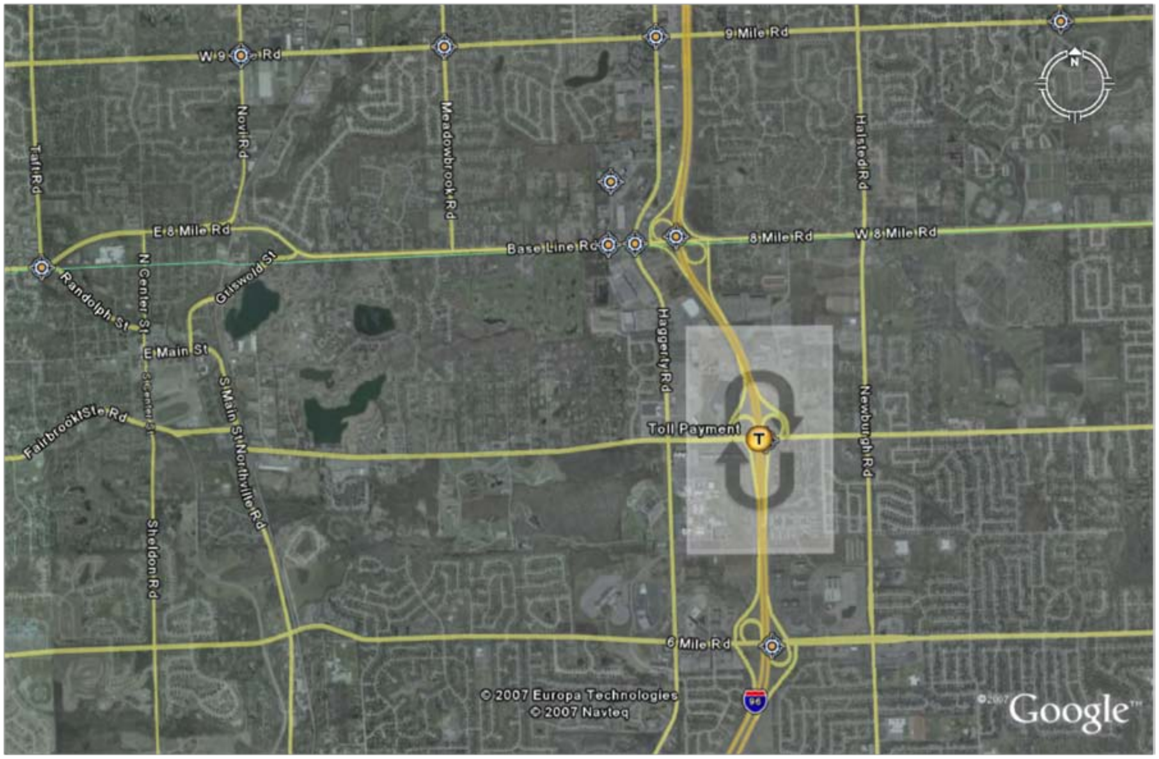

VIIC generated plot points on Google map

Figure 2-46: RSE Regions Used for Tolling Test

2.6.2 RSE Installations

RSEs are typically mounted on existing poles used for lighting or signals along the roadway. The details of the installations are site specific and vary slightly depending on the method used for backhaul. Figure 2-47, Figure 2-48, and Figure 2-49 provide examples of RSE installations used in the DTE.

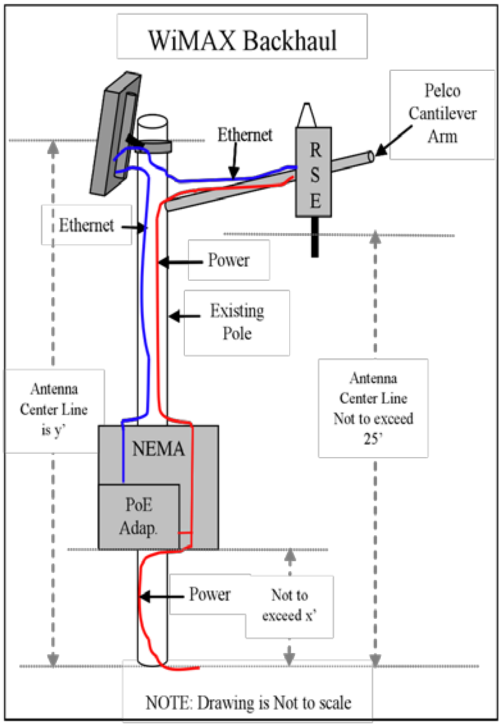

Figure 2-47: WiMAX RSE Installation

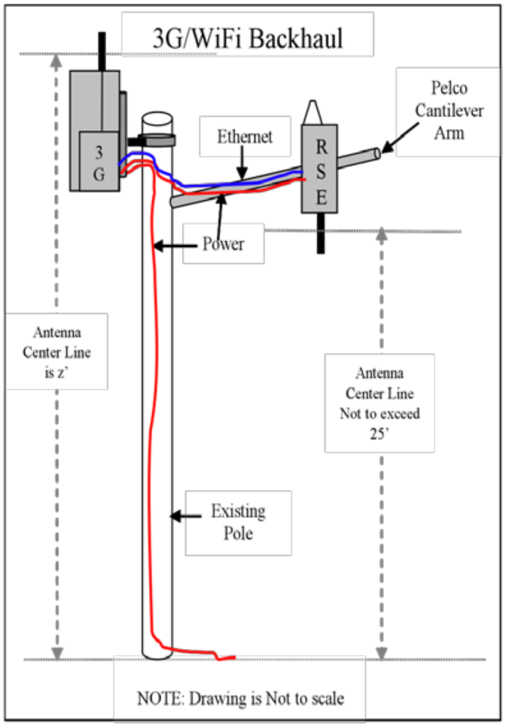
3G/WiFi Backhaul
Pelco Cantilever Arm
3 G
Ethernet
R S E
Power
Antenna Center Line is z'
Antenna Center Line Not to exceed 25'
Existing Pole
NOTE: Drawing is Not to scale

Figure 2-48: 3G RSE Installation

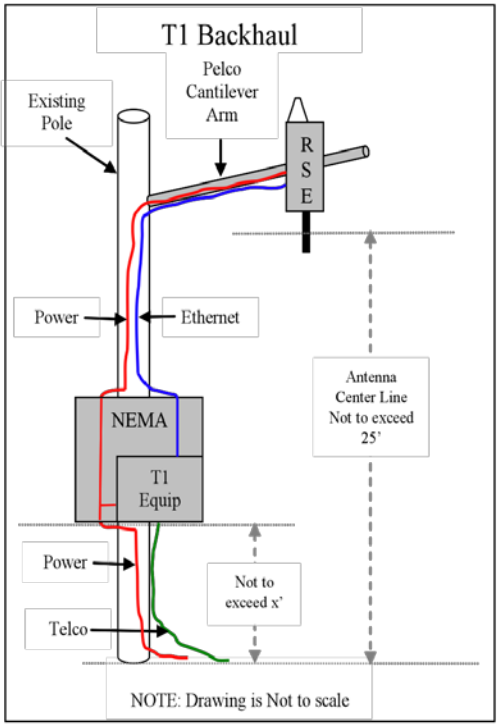

Figure 2-49: T1 RSE Installation

Figure 2-50 shows a typical RSE field installation.

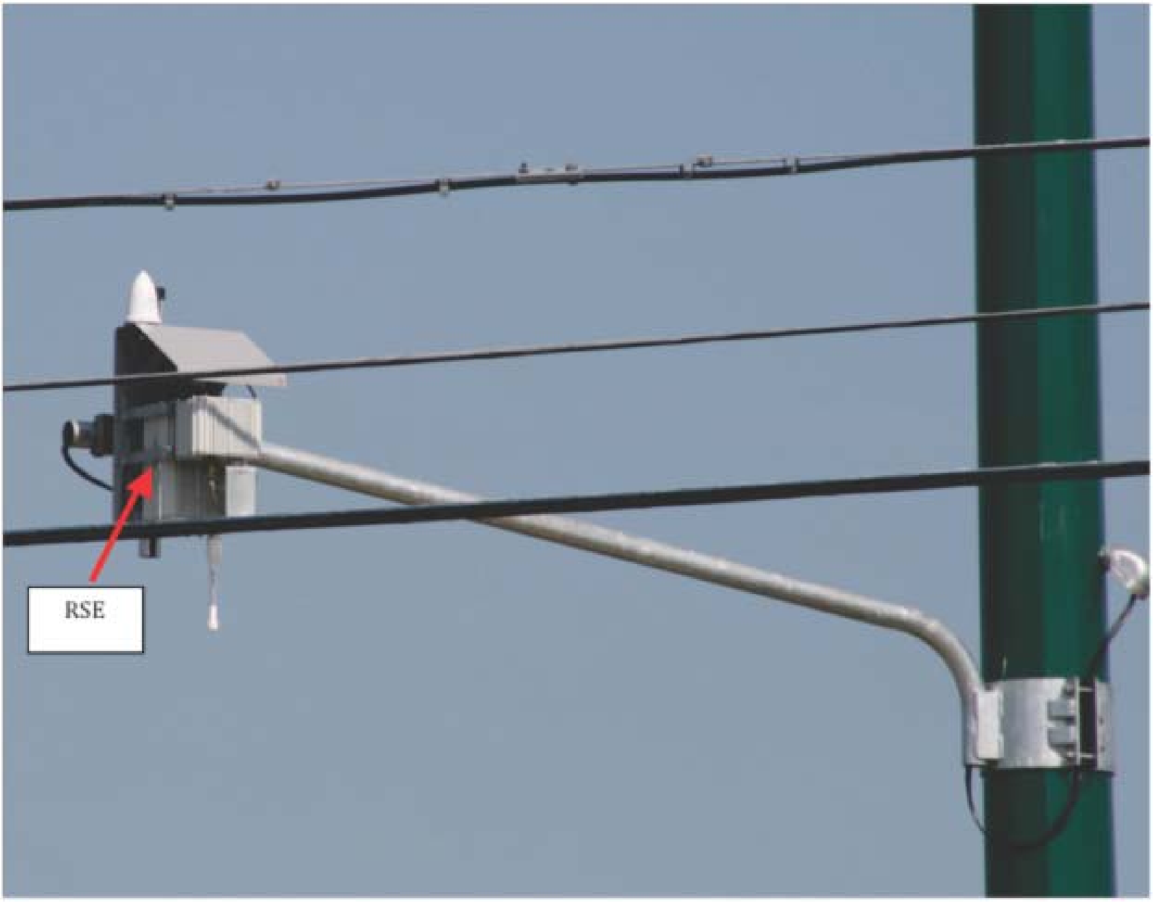

Booz Allen staff photo

Figure 2-50: Typical RSE Field Installation

2.6.3 Backhaul Distribution

The backhaul network is used to connect an RSE to its home SDN. The POC used three different types of backhaul network—T1 wire line, 3G wireless, and WiMAX wireless. Figure 2-51 illustrates the POC backhaul setup.

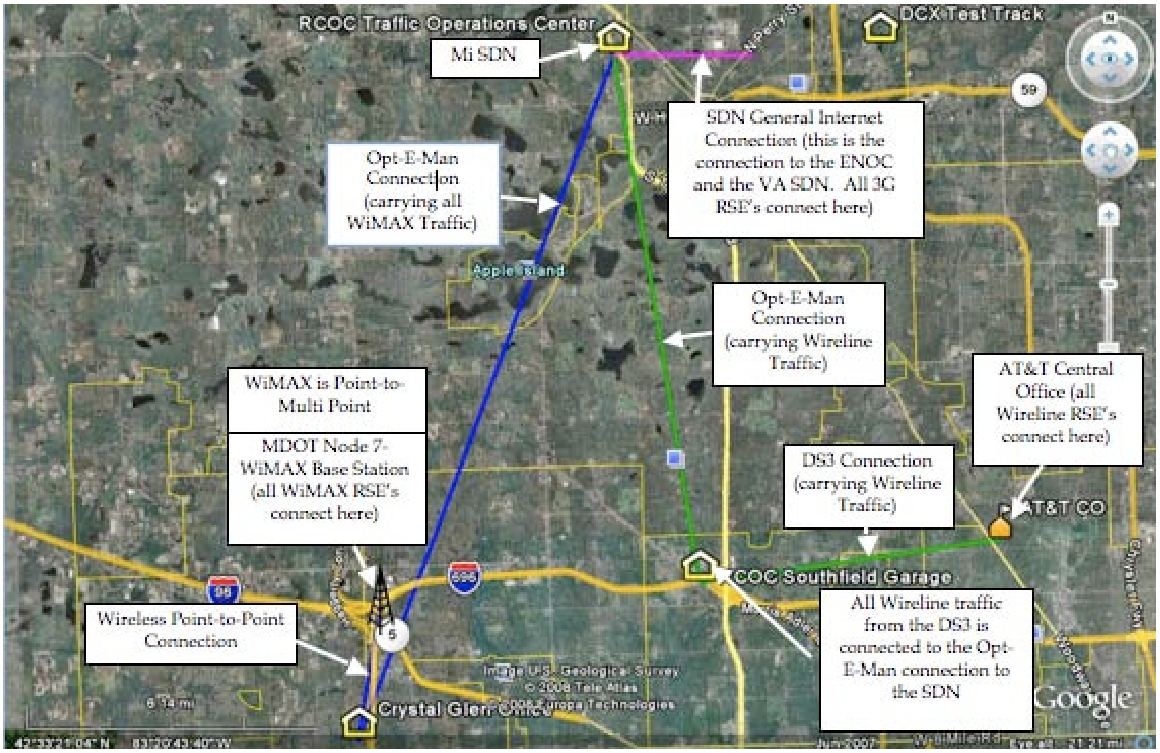
RCOC Traffic Operations Center
Mi SDN
DCX Test Track
SDN General Internet Connection (this is the connection to the ENOC and the VA SDN. All 3G RSE's connect here)
Opt-E-Man Connection (carrying all WiMAX Traffic)
Apple Island
Opt-E-Man Connection (carrying Wireline Traffic)
AT&T Central Office (all Wireline RSE's connect here)
WiMAX is Point-to-Multi Point
MDOT Node 7- WiMAX Base Station (all WiMAX RSE's connect here)
DS3 Connection (carrying Wireline Traffic)
AT&T CO
RCOC Southfield Garage
All Wireline traffic from the DS3 is connected to the Opt-E-Man connection to the SDN
Wireless Point-to-Point Connection
Crystal Glen Office
96
696
5
59
Image U.S. Geological Survey
© 2008 Tele Atlas
© 2008 Europa Technologies
Google

Booz Allen generated plot points on Google map

Figure 2-51: POC Backhaul Networks

WiMAX-connected RSEs all link to a WiMAX hub located at the POC headquarters at Crystal Glen in Novi. An ATT Opt-E-Man link (provided as part of the ATT backhaul service) is used to connect the WiMAX hub to the SDN located at the RCOC facility in Farmington Hills.

The T1-line-connected RSEs all link back to the ATT central office in Franklin. A DS3 connection carries this traffic to the RCOC Southfield Garage, where it is then carried by a second Opt-E-Man link to the SDN.

All 3G RSEs link to the nearest 3G base station where the message traffic is delivered to the Internet. An Internet connection at the SDN completes this link.

2.6.4 Backbone Network

The POC backbone network links the SDN in Herndon and the SDN in Detroit. This network is implemented by a leased partial T3 circuit.

3.0 PUBLIC APPLICATIONS

The initial POC plan included extensive testing of the following public sector applications:

- ***Signal Timing Optimization –*** The goal of this application is to determine whether the VII system can collect data useful for optimizing traffic signals. The application uses collected probe data to analyze traffic flows in a particular area and then, based on that analysis, adjusts the signal timing. This process can be carried out on an ongoing basis to maintain optimal flow over time.

- ***Ramp Metering –*** The goal of this application is to determine whether the VII system can collect data useful for controlling freeway on-ramp metering lights to maximize the freeway throughput. This application is similar to signal timing optimization except the traffic situation is significantly different. One element to be measured is whether the PDS can provide data to accurately model the merging behavior of vehicles at the ramp region.

- ***Traveler Information*** – The goal of this application is to determine whether the VII system can deliver data that can then be used to generate useful advisory messages. This application uses the PDS to collect a variety of information and then, based on that information, generate advisory messages that are then distributed back out to the roadway environment using the AMDS subsystem.

- ***Weather Information –*** The goal of this application is to determine whether the VII system can deliver data that can then be used to generate a useful assessment of weather conditions in a given region. The results of this assessment can then be converted into weather advisory messages as appropriate, which can then be distributed to the appropriate roadway regions using the AMDS.

- ***Corridor Management Planning Assistance –*** The goal of this application is to determine whether the VII system can deliver travel pattern data that can then be used to support and improve road-planning activities. By understanding the general origin and destination (O-D) information of travelers in a region, it is expected that road planners can better determine how to manage the roads, as well as where and how to invest in road improvements.

- ***Corridor Management Load Balancing –*** The goal of this application is to determine whether the VII system can deliver road data that can then be used to support and improve transportation management activities. Corridor management load balancing is principally concerned with making the best use of the roadway infrastructure through the redistribution of traffic through advisory messages distributed using the AMDS.

The original objectives of the POC public sector applications were to: a) validate the technical viability of the VII system and b) provide limited demonstration of the potential utility of VII to determine various transportation management measures. All six public applications described above were developed based on detailed requirements specifications. However, the full-scale testing of these applications was de-scoped from the VII program. As a result, the analytical software (i.e., algorithms) necessary to assess the data and produce the transportation management measures was developed but not tested. As an alternative, extensive POC data was collected under controlled conditions so that general assessments could be made relative to the viability of these applications and the ability of the VII system to support them. The data sets acquired are expected to be useful in the future to fully test these applications.

For completeness, the following sections describe in detail each application.

3.1 Signal Timing Optimization

Traffic engineers modify signal timings to provide the highest possible level of service on arterial and surface streets. Modifications rely on extensive collection of traffic data to characterize signal performance. VII offers the possibility of expanding the breadth of data collected and reducing the difficulty and cost in obtaining information.

The signal timing optimization application is geared toward proving the viability of using VII to support enhanced signal timing applications in the future. The POC efforts concentrate on determining the traffic measures with potential applicability to signal timing that can reliably be determined from VII probe data.

The process for development of this application concept was to take the use case definition of signal timing optimization and determine the aspects of the use case that warranted a POC and were practical to develop in that POC timeframe. This concept of operations presents this subset of the use case.

3.1.1 Signal Timing Optimization Application Objectives

The objective of the signal timing optimization application is to determine whether the VII system can provide the measures used for determining the effectiveness of signal systems and help efficiently generate new signal timing plans, including:

- Queue length
- Stops and stop locations
- Stop delay
- Cycle failures
- Time and position trajectories
- Queue overflow
- Arterial travel time
- Volumes (by lane)
- Turning movements.

3.1.2 Signal Timing Optimization Overall Architecture

The signal timing optimization application is a network side application that uses data collected by the probe data element of the VII system. The signal timing optimization application user sets up the application by using the VII ILS to determine the RSE associated with the signal of interest. Using this information, the user sets up a probe data subscription for that particular RSE. In operation, the VII system delivers probe data from the designated RSE. The probe data archive stores this data, and the signal traffic characteristics generator analyzes the data to derive the parameters described in Section 3.1.1. These parameters are then compared with ground truth data to determine whether the VII system can provide useful information to be used in signal timing operations.

3.1.3 Signal Timing Optimization Flow of Events

Preconditions include:

1. The application has an up-to-date digital centerlines map by lane of the area 1,000 feet around the intersections under study.
2. The signal timing optimization knows what data to subscribe to in order to fulfill the needs of the application.
3. Vehicles are transmitting probe vehicle data to the PDS.

The flow of events includes:

1. The signal timing optimization application sets the probe generation rate for the area under test.
2. The signal timing optimization application subscribes to VII probe data for the area under test, including data elements speed, location, and heading.
3. The signal timing optimization application receives VII probe data for which it is subscribed from the PDS.
4. The signal timing optimization application archives VII probe data.
5. The signal system operator measures signal system performance metrics directly in the field, including queue length, stops, stop locations, stop delays, time and position trajectories, queue overflow, cycle failures, speeds, and arterial travel times through the study area.
6. The signal system operator requests metrics from the signal timing optimization software for the same time periods as data was recorded in the field.
7. The signal timing optimization application calculates signal system performance metrics from archived VII probe data for the requested time period.
8. The signal system operator compares the metrics based on VII probe data with the same metrics measured in the field to determine the performance of calculations based on VII data.

3.2 Ramp Metering

Freeway operators use ramp-metering systems to moderate the flow of traffic onto freeways. This can help to keep traffic flowing on the mainline while still providing access to the freeway. Dynamic modification of ramp-metering rates is difficult, as it requires knowledge of freeway performance, queue length on the ramp, and arterial performance. The VII system offers the potential to gather a more robust and richer data set describing the freeway, ramp, and arterial conditions, leading to more timely and accurate evaluation of ramp-metering rates.

The ramp-metering application is geared toward proving the viability of using VII to support enhanced ramp-metering applications in the future. The POC efforts concentrate on determining the traffic measures with potential applicability to ramp metering that can reliably be determined from VII probe data.

3.2.1 Ramp-Metering Application Objectives

The objective of the ramp-metering application is to determine whether the VII system can produce the data elements needed for ramp metering. Specifically:

- Link travel times on adjoining freeways
- Mainline speed by lane
- Ramp queue length
- Queue spillback
- Stops and stop locations on both ramps and mainline
- Stop delay on both ramps and mainline
- Time in queue.

3.2.2 Ramp-Metering Overall Architecture

The ramp-metering application is a network side application that uses data collected by the probe data element of the VII system. Figure 3-1 shows the basic architecture. In this figure, the details of the PDC system have been omitted for clarity.

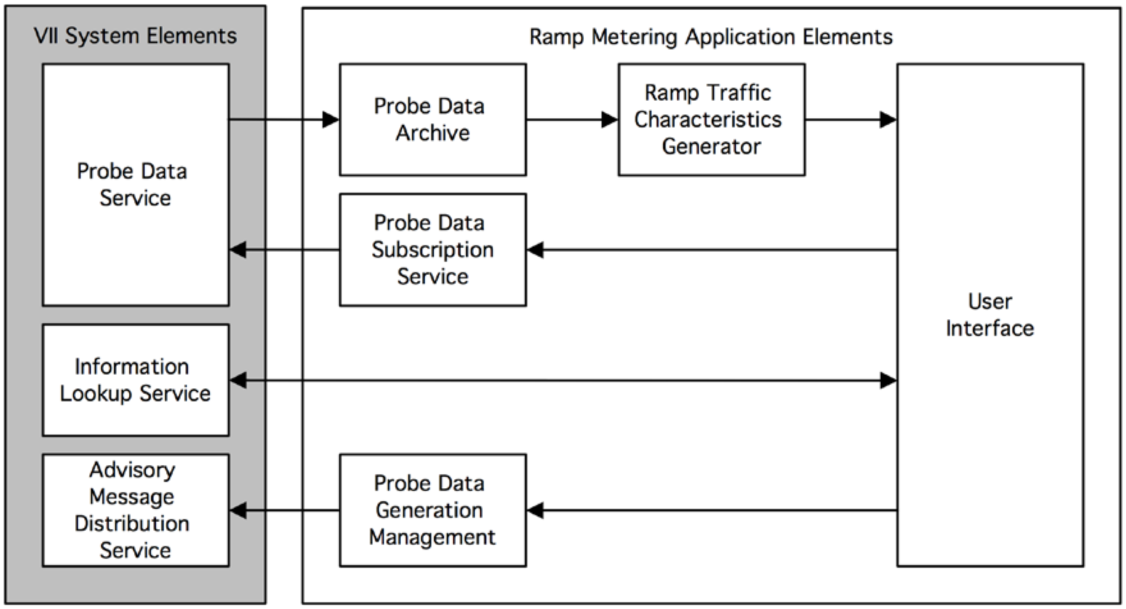
VII System Elements
Probe Data Service
Information Lookup Service
Advisory Message Distribution Service
Ramp Metering Application Elements
Probe Data Archive
Ramp Traffic Characteristics Generator
Probe Data Subscription Service
Probe Data Generation Management
User Interface

Figure 3-1: Ramp Metering Application Architecture

The ramp-metering application user sets up the application by using the VII ILS to determine the RSE associated with the ramp of interest. Using this information, the user sets up a probe data subscription for that particular RSE. In operation, the VII system delivers probe data from the designated RSE. The probe data archive stores this data and the ramp traffic characteristics generator analyzes the data to derive the parameters described in Section 3.1.1. These parameters are then compared with ground truth data to determine whether the VII system can provide useful information to be used in ramp-metering operations.

3.2.3 Ramp-Metering Flow of Events

Preconditions for ramp metering include:

1. The application has an up-to-date digital centerlines map by lane of the area 1,000 feet around the ramp under study, including ramp upstream and downstream, related arterial, and mainline.
2. The ramp-metering application knows what data to subscribe to in order to fulfill the needs of the application.
3. Vehicles are transmitting probe vehicle data to the PDS.

The flow of events includes:

1. The ramp-metering application requests VII probe data for the area under test, including data elements—speed, location, and heading
2. The ramp-metering application receives VII probe data for which it is subscribed from the PDS.
3. The ramp- metering application calculates ramp performance metrics from the probe data stream, including link travel times, mainline speeds, ramp queue length, delay on both ramps and mainline, any instances of queue spillback, stops, stop locations and stop delay on ramps and mainline, and time in queue.
4. The freeway system operator measures ramp performance metrics directly in the field, including link travel times, mainline speeds, ramp queue length, delay on both ramps and mainline, any instances of queue spillback, stops, stop locations and stop delay on ramps and mainline, and time in queue. All applicable measurements are taken by lane.
5. The freeway system operator compares the metrics based on VII probe data with the same metrics measured in the field to determine the real-time performance of calculations based on VII data.

3.3 Traveler Information

Publicly provided traveler information systems provide information that may be relevant to a traveler’s current direction of travel, including, for example, travel times, incident alerts, road closures, and work zones. The VII system offers the potential to provide this information in a way that is targeted to a driver’s current direction of travel and to update the information in real time as the vehicle progresses along a route.

Within the scope of traveler information, public entities (both state and local) provide information derived from public data and provide geographically relevant information. Private entities can combine the public information with information from other public and private data sources to provide a more individually targeted service, providing enhanced “personalized” traveler information, depending on the trip details provided by a motorist subscribing to such information.

Public traveler information provided in the POC includes travel times, non-recurring congestion information, and incident information generated from VII probe data or from existing traveler information available from the RCOC’s FAST-TRAC system. This information is provided through the VII system to POC test vehicles that are within the geographic range for which the information may be relevant. Interpretation, prioritization, and presentation of public traveler information within the vehicle are included as a part of the in-vehicle signage POC application.

Incident information notifications use the same mechanisms as other road status notifications, except that these messages are generated and disseminated as soon as they are detected by the traveler information application.

3.3.1 Traveler Information Application Objectives and Use Cases

The objective of the traveler information application is to demonstrate that VII probe data can be used to determine appropriate measures for traveler and incident information. Specifically:

- Freeway link travel times
- Freeway link speed profiles
- Arterial link travel times
- Arterial link speed profiles
- Link delays
- Lane occupancy and vacancy
- Incident detection time
- Incident location.

The traveler information application is based on an array of use cases that interact with multiple VII system services. These use cases are:

- ***Collect General Vehicle Probe Data:*** This use case is responsible for communicating with the vehicle and collecting vehicle probe data.
- ***Acquire Published Probe Data:*** This use case is responsible for acquiring probe data from the appropriate locations to be used for analysis.
- ***Analyze Real-Time Data and Determine Travel Times and Roadway Conditions:*** This use case is responsible for analyzing the raw probe data and determining road hazards and conditions that will become the subject of advisory messages.
- ***Acquire Published Traffic Management Center (TMC) Traffic/Incident/Surveillance Data:*** This use case acquires traffic, incident, and surveillance (e.g., closed-circuit television) data from TMC and WWW sources for analysis to determine road hazards and conditions that will become the subject of advisory messages.
- ***Provide Civil Emergency Notifications:*** This use case is responsible for providing simulated civil alerts that will be distributed to all vehicles in a geographic region.
- ***Provide Traveler Information to Network End Users:*** This use case distributes the results of the probe data and external data analysis to the off-board navigation POC application.
- ***Disseminate Geographically Focused Traveler Information:*** This use case distributes advisory messages, derived from the analysis use case above, to vehicles at locations relevant to the situation/incident.
- ***Provide Traveler Information Messages:*** This use case provides the generated advisory messages to the AMDS for broadcast from the appropriate RSEs.

- ***Present Traveler Information:*** This use case is responsible for presenting traveler information to the driver. It is implemented by the in-vehicle signage application described in the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*.
- ***Record Performance Measures:*** This use case is responsible for validating the generated messages against ground truth information and for recording the potential utility of the traveler information POC application. The use case is also responsible for monitoring the performance of the application via a graphical user interface.

3.3.2 Traveler Information Overall Architecture

Figure 3-2 shows the traveler information application architecture. This figure illustrates how the various elements of the application interact with, and tie together the VII system services.

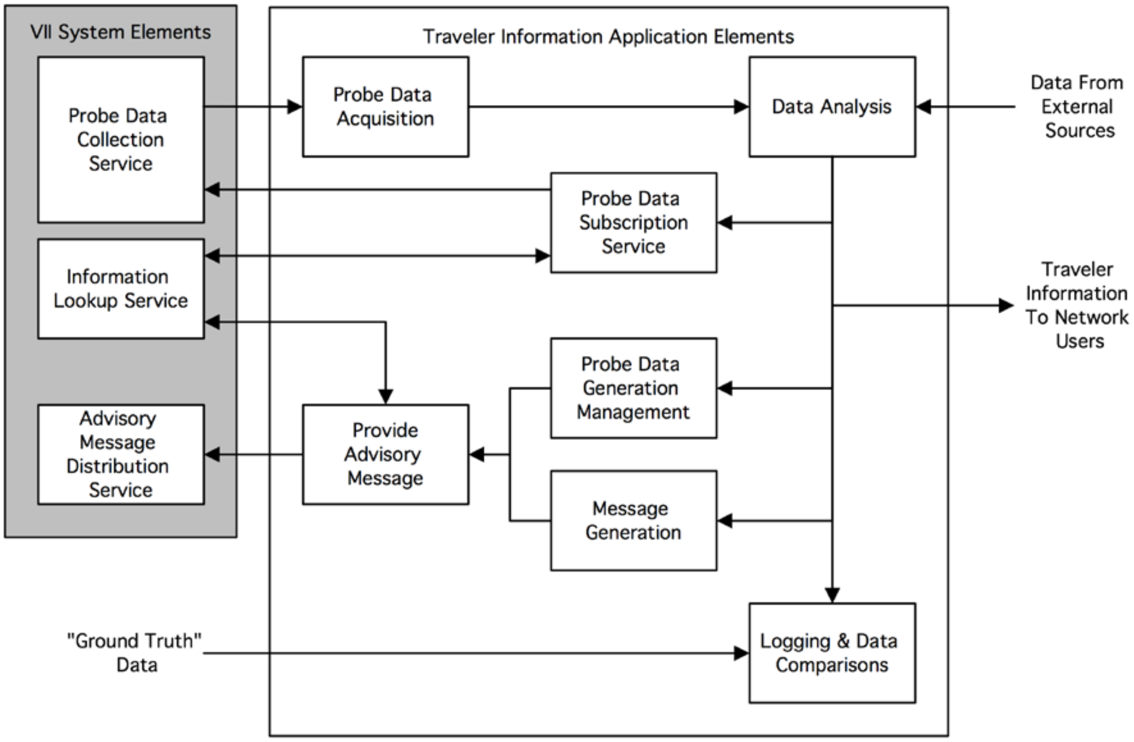
VII System Elements
Probe Data Collection Service
Information Lookup Service
Advisory Message Distribution Service
Traveler Information Application Elements
Probe Data Acquisition
Data Analysis
Data From External Sources
Probe Data Subscription Service
Traveler Information To Network Users
Probe Data Generation Management
Provide Advisory Message
Message Generation
"Ground Truth" Data
Logging & Data Comparisons

Figure 3-2: Traveler Information Application Architecture

3.4 Weather Information

For the VII POC, the three weather-related use cases (the Weather Information Traveler Notification, the Weather Information Improved Weather Observing and Forecasting, and the Winter Maintenance) were to be combined into the weather information application. The Weather Information Traveler Notification use case focuses on providing road weather advisories to travelers using *Clarus* weather data. The Weather Information Improved Weather Observing and Forecasting use case looks at the back-end processing of probe vehicle environmental data in the research arena developing new algorithms and thresholds for generating weather-related alerts. The Winter Maintenance use case involves the collection of probe data, the conversion of the probe data to weather observations, and the subsequent management of winter maintenance vehicles with regard to their routing and treatment recommendations. The three use cases were primarily differentiated by their backend processing, namely, the inclusion of road weather technologies such as *Clarus* and the Maintenance Decision Support System; service provider processing; and subsequent generation of messages for winter maintenance vehicles, travelers, and research institution(s). It was decided that the primary goal at the VII POC was to capture and archive the vehicle probe data including the important weather-related data elements for subsequent analysis by research institution(s).

The captured vehicle data along with supplemental in situ observations (for ground truth comparison with the VII provided data) were to be archived and made available to approved research institution(s) to study the functionality of the weather-related use cases within the VII deployment timeframe (i.e., 2007 to 2011). A second part of the weather information POC application testing was to include sending road weather information advisories to the POC test vehicles from the research institution(s) if the data from the initial tests proved to be useful in developing such advisories. Road weather advisories can range from general regional "nowcasts" and forecasts to event-driven weather watches, warnings, and advisories to route-specific weather-related data.

3.4.1 Weather Information Application Objectives

The objectives of the weather information application are to test the effectiveness of using probe data collected by the VII POC system in determining weather-related road conditions and specifically to:

- Demonstrate that VII can provide vehicle probe data that is useful for determining:
  - Ground-level atmospheric conditions
  - Pavement surface conditions
- Automate the collection of vehicle-based atmospheric and pavement data
- Enable the weather community to understand how vehicle-based atmospheric and pavement data compare to in situ observations and climatological outliers
- Provide the weather community with some knowledge of the latency issues involved with vehicle probe data collection and processing
- Improve the weather community's ability to quantify the meteorological state variables based on vehicle probe data
- Enable the weather service provider to generate enhanced traveler weather advisories including:
  - General watches, warnings, advisories, and forecasts
  - Regional roadway atmospheric and pavement condition information
  - Value-added route-specific atmospheric and pavement condition information
- Allow for latency issues to be explored by analyzing the dissemination of weather advisories to vehicles through use of the VII network.

3.4.2 Weather Information Application Overall Architecture

The general architecture for the weather information application is the same as for traveler information (Figure 3-2). The key difference is that the analysis and ground truth information are weather related.

3.5 Corridor Management Planning Assistance (CMPA)

Planning for new roads and road improvements requires the collection of travel data, so that planners can characterize trip volumes from point to point. Traditionally, this is done through surveys and fixed-point volume counts that can require specialized installations.

The VII concept provides the possibility for directly gathering this information from vehicles, eliminating the need to establish data collection locations and increasing both the speed and flexibility with which planning data can be acquired. The CMPA application is geared toward proving the viability of using VII to support capital planning activities. The POC efforts concentrate on determining the traffic measures with potential applicability to planning that can reliably be determined from VII probe data and demonstrating the use of an opt-in application to collect additional data.

3.5.1 CMPA Application Objectives

The goal of this application is to determine whether the VII system can deliver travel pattern data that can then be used to support and improve road-planning activities. By understanding the general O-D Information of travelers in a region, it is expected that road planners can better determine how to manage the roads, as well as where and how to invest in road improvements.

3.5.2 CMPA Application Overall Architecture

Figure 3-3 shows the overall CMPA application architecture. The trip path application in the vehicle subsystem accumulates location and time data during each operational cycle for the vehicle. When the vehicle passes an RSE that supports download of the data, and it has a trip record file to transfer, the trip path vehicle application sends the complete trip record to a network side data accumulator. This is accomplished by using the routing capabilities of the RSE and the SDN.

The trip path data accumulator stores the records containing complete origin-to-destination logs for each trip made by a trip-path-equipped vehicle.

In parallel, the CMPA subscribes to the PDS for the road sections about which it needs to gather overall roadway information. As vehicles drive on the subject road, they deliver probe data that is forwarded by the PDS, according to a set of subscription parameters, to the probe data archive element of the CMPA. The probe data archive stores the data for further analysis by the road network metrics generator. The road network metrics generator determines aggregate parameters for the road section under consideration, for example, road segment speeds versus time of day and day of the week.

The CMPA application periodically acquires O-D data sets from the trip path data accumulator and aggregated roadway information from the road network metrics generator, and analyzes these to identify travel patterns useful for future road planning or roadway management.

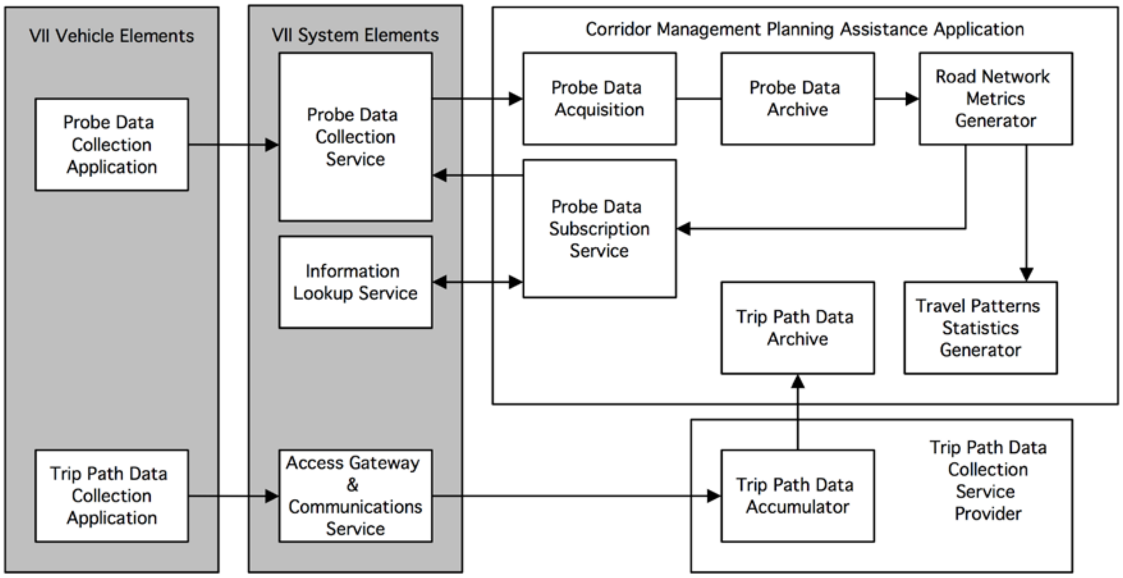
VII Vehicle Elements
Probe Data Collection Application
Trip Path Data Collection Application
VII System Elements
Probe Data Collection Service
Information Lookup Service
Access Gateway & Communications Service
Corridor Management Planning Assistance Application
Probe Data Acquisition
Probe Data Archive
Road Network Metrics Generator
Probe Data Subscription Service
Trip Path Data Archive
Travel Patterns Statistics Generator
Trip Path Data Accumulator
Trip Path Data Collection Service Provider

Figure 3-3: CMPA Application Architecture

3.6 Corridor Management Load Balancing (CMLB)

Traffic managers work with freeway operators, signal system operators, and transit system operators to help move travelers and goods in the most economical fashion. This requires extensive coordination between agencies with traditionally disparate goals, operating procedures, and data collection methods. VII presents the possibility for wide-area data collection that these agencies typically do not have access to and further provides the potential for information dissemination that bridges the gaps between operational entities. CMLB is principally concerned with making the best use of the roadway infrastructure through the redistribution of traffic.

For the POC, the CMLB application was intended to concentrate on determining the practicality of achieving performance measurement of roadway resources using VII data. The process for development of this application concept was to take the use-case definition of CMLB and determine what aspects of the Day 1 use case warranted a POC and were practical to develop in that POC timeframe.

3.6.1 CMLB Application Objectives

The goal of this application is to determine whether the VII system can deliver road data that can then be used to support and improve transportation management activities. CMLB is principally concerned with making the best use of the roadway infrastructure through the redistribution of traffic through advisory messages distributed using the AMDS subsystem. The overall objectives for the POC were to:

- Test the VII system’s capability to deliver anonymous probe data to an external data subscriber, based on the subscriber’s subscription profile
- Test the VII system’s capability to distribute broadcast advisory messages to VII-equipped vehicles.
- Determine whether VII can provide the measures that can be used to support strategies for corridor load balancing, including:

- Travel times by link on freeways and arterials
- Speed profiles by lane on freeways and arterials
- Volumes on freeways and arterials
- Stop locations on freeways and arterials
- Stop delay on freeways and arterials
- Lane closures on freeways and arterials.

3.6.2 CMLB Application Overall Architecture

The CMLB architecture is effectively the same as the traveler information architecture. The key differences being the types of probe parameters subscribed to and the use of the results from analysis of those parameters. For example, the results of traffic analysis might be to adjust the timing of traffic signals on arterials or to change routing of certain types of traffic either through variable lane and turn restrictions or through route suggestions and directives.

3.7 Public Applications Developed and Tested in the POC

As mentioned earlier, the public applications work in the POC was de-scoped. As a result, only a subset of the initially planned developments and subsequent tests were performed. These tests emphasized safety and generally used existing developed services and applications. The application and system service functions examined in the POC public applications testing are briefly described below. Volume 3b includes a detailed discussion of the tests and the results.

- ***Probe Data Management:*** Examine the potential for enhancing intersection safety by directing vehicles to provide specific data at specific intervals for a given intersection.

- ***Advisory Broadcast Strategies:*** Provide in-vehicle warnings and advisories under a variety of traffic and other conditions to see how quickly the system can provide localized information directly to vehicles.

- ***Probe Data Reception:*** Determine whether probe data can be collected in near real time under a variety of traffic and other conditions, which could potentially help monitor and detect traffic and weather conditions that affect driver safety.

- ***Probe Data Vehicle Trajectories – Turning Movements:*** Determine whether probe data can be collected in near real time at intersections, which could potentially help monitor and detect traffic conditions (such as queue overflows and suboptimal signal phase and timing) at the intersection that adversely affect safety at that intersection.

- ***Probe Data Vehicle Stop and Start:*** Determine whether probe data can be collected in near real time in stop-and-go traffic, which could potentially help monitor and detect roadway congestion and, in turn, help motorists avoid the congested roadway segments and potential accident locations.

- ***Probe Data Vehicle Trajectories: Lane Change:*** Determine whether probe data can be collected in near real time from vehicles making unusual lane changes in the roadway, which could potentially help monitor and detect traffic conditions such as roadside hazard that adversely affect safety along the roadways.

- ***Probe Data Event-Generated Snapshots/Icy Road:*** Determine whether probe data, based on vehicle kinematics data (such as anti-lock braking system, traction control, etc.) can be collected in near real time, which could potentially help monitor and detect adverse weather conditions that affect driver safety.

- ***Trip Path:*** Collect trip path (O-D) data for future analysis.

4.0 SUMMARY

The POC project included the implementation of a small-scale version of the conceptual VII infrastructure system. While limited to only 55 RSEs and two SDNs, the POC network was implemented using a unique suite of architectural concepts that can allow the system to be easily scaled to any size. This approach resulted in slightly more complexity. However, since it is unknown how large a deployed system might be, this “scale free” approach promises to assure that the concepts proven during the POC remain applicable as the system is expanded.

All key service elements and interfaces of the system were implemented, and these were coordinated with OBE application and service developments (see the document titled *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Technical Description – Vehicle*) to allow end-to-end test and assessment of each of the key operational concepts, as described in Section 1.1.

All elements and interfaces of the system were tested through direct functional and performance testing and through indirect applications test and assessment. The *Final Report: Vehicle Infrastructure Integration Proof-of-Concept Results and Findings – Infrastructure* (Volume 3) describes the system test and assessment.

APPENDIX A: ACRONYMS

3G	Third-Generation Cellular (Wireless Data System)
AASHTO	American Association of State and Highways Transportation Officials
AMDS	Advisory Message Delivery Service
AMT	Advisory Message Table
BMW	Bavarian Motor Works
CCH	Control Channel
CEP	Circular Error Probable
CPU	Central Processing Unit
CRL	Certificate Revocation List
DGPS	Differential Global Positioning System
DSRC	Dedicated Short Range Communications
DTE	Development Test Environment
ENOC	Enterprise Network Operations Center
ENS	Event Notification System
ESB	Enterprise Service Bus
FCC	Federal Communications Commission
Gbps	Gigabit Per Second
GHz	Giga-Hertz
GID	Geographic Intersection Description
GIS	Geographic Information System
GM	General Motors
GML	Geographic Markup Language
GPS	Global Positioning System
HA-NDGPS	High-Accuracy National Differential GPS
HB	Heartbeat
HIP	Host Identity Protocol
HMI	Human Machine Interface
ILS	Information Lookup Service
I/O	Input/Output
IP	Internet Protocol

IPv6	Internet Protocol Version 6
JMS	Java Message Service
LSS	Local Safety System
LTP	Local Transaction Processor
MAC	Medium Access Control
MCNU	Multiband Configurable Networking Unit
MDOT	Michigan Department of Transportation
MEDS	Map Element Distribution System
MNE	Managed Network Element
MHz	Mega-Hertz
NMS	Network Management Service
OBE	On-Board Equipment
O-D	Origin and Destination
OEM	Original Equipment Manufacturer
PDC	Probe Data Collection
PDS	Probe Data Service
PDM	Probe Data Management
PER	Packer Error Rate
PSID	Provider Service Identifier
QoS	Quality of Service
RA	Registration Authority
RCOC	Road Commission for Oakland County
RF	Radio Frequency
RSE	Roadside Equipment
SCH	Service Channel
SDN	Service Delivery Node
SIT (Tunnel)	Simple Internet Transition (Tunnel)
SPAT	Signal Phase And Timing
TMT	Technical Management team
UDP	Universal Datagram Protocol
USDOT	United States Department of Transportation
V2I	Vehicle to Infrastructure
V2V	Vehicle to Vehicle

VDTLS	VII Datagram Transport Layer Security
VII	Vehicle Infrastructure Integration
VIIC	Vehicle Infrastructure Integration Consortium
VPN	Virtual Private Network
VW	Volkswagen
WAAS	Wide Area Augmentation System
WAVE	Wireless Access in Vehicular Environments
WO	Work Order
WSA	WAVE Service Advertisement
WSMP	WAVE Short Message

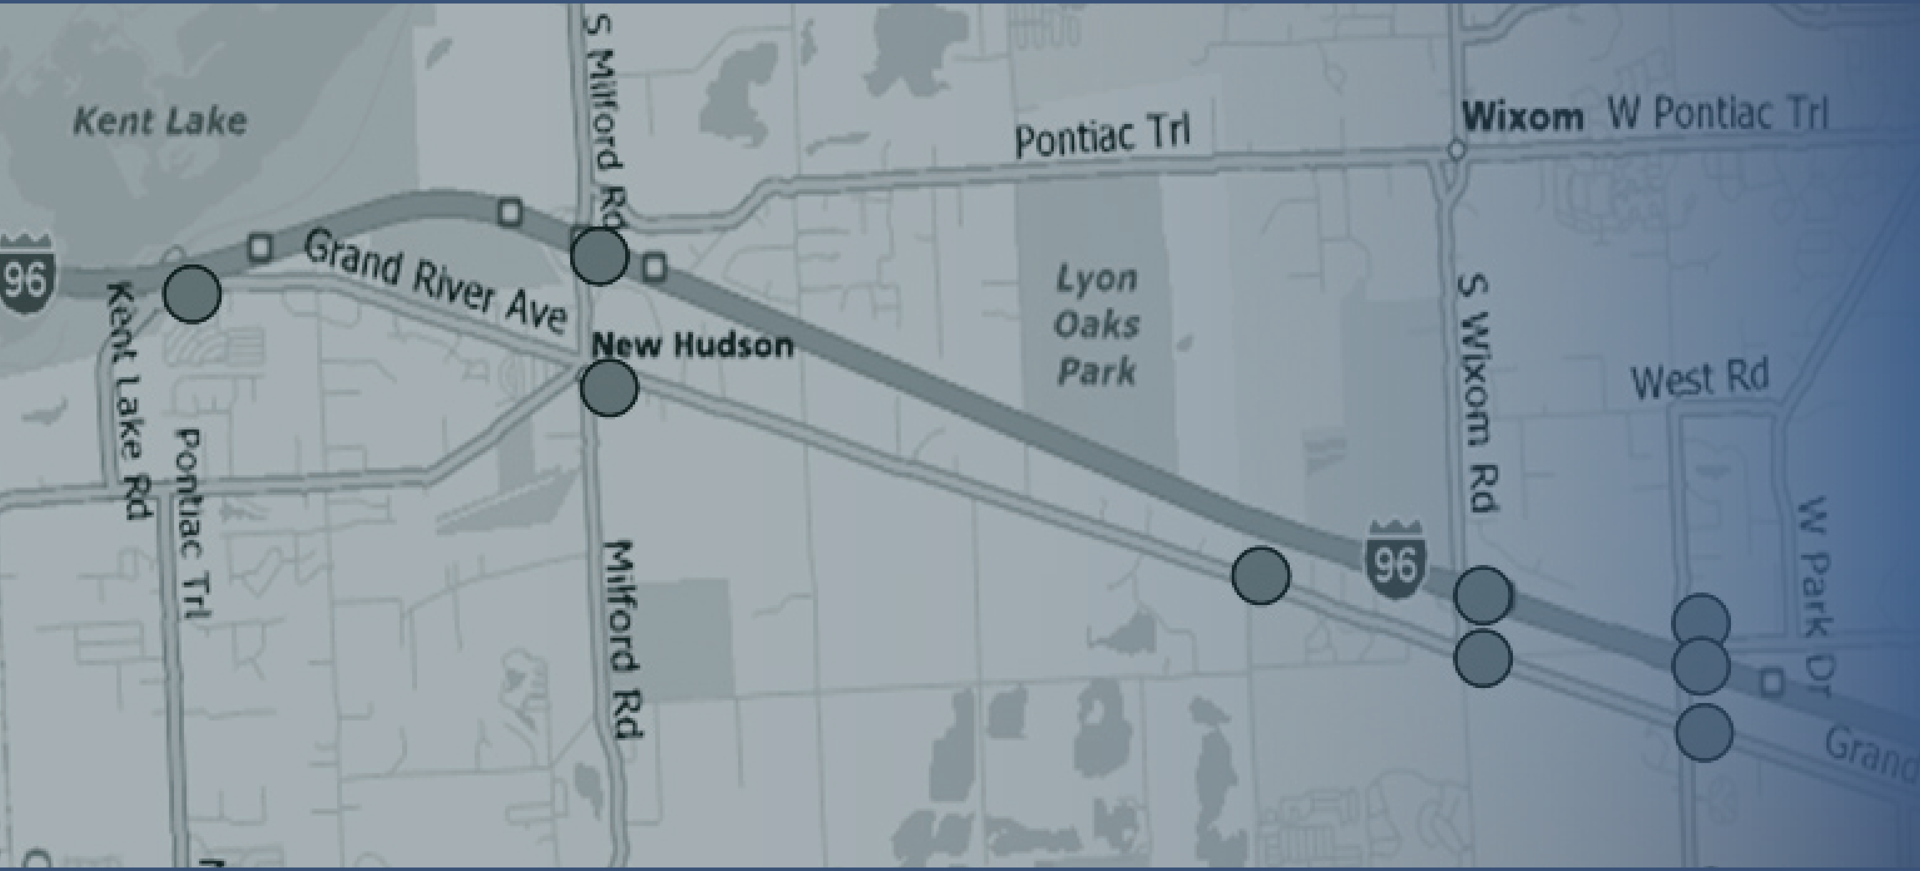

U.S. Department of Transportation
Research and Innovative Technology Administration

U.S. Department of Transportation | 1200 New Jersey Avenue, SE | Washington, DC 20590