VOICES Cooperative Driving Automation Proof-of-Concept Systems Integration Test 1

PUBLICATION NO. FHWA-HRT-24-072

APRIL 2024

U.S. Department of Transportation
Federal Highway Administration

Research, Development, and Technology
Turner-Fairbank Highway Research Center
6300 Georgetown Pike
McLean, VA 22101-2296

FOREWORD

The Turner-Fairbank Highway Research Center performs advanced research into several areas of transportation technology for the Federal Highway Administration. The Office of Safety and Operations Research and Development focuses on improving safety- and operations-related technology through research, development, and testing.

This report documents the results and technical performance analysis of the U.S. Department of Transportation's (USDOT) Systems Integration Test 1 (SIT-1)(1) of the prototype, Virtual Open Innovation Collaborative Environment for Safety (VOICES).(2) The USDOT VOICES program developed a prototype distributed virtual test platform to enable collaboration among participating entities (e.g., public sector including State and local governments, private sector, and academic institutions) in an intellectual property-protected virtual collaborative environment for research and interoperability testing of prototype cooperative transportation applications. The purpose of SIT-1 was to provide a proof of concept that demonstrated how four simulation nodes across three geographically distributed sites could communicate and coordinate using some of the prototype VOICES technologies. The SIT-1 test elements operated together to successfully perform two cooperative driving automation maneuvers during the test. The intended audiences of this report are researchers and developers interested in testing interoperability of a connected surface transportation system.

Carl Andersen
Acting Director, Office of Safety and Operations
Research and Development

Notice

This document is disseminated under the sponsorship of the U.S. Department of Transportation (USDOT) in the interest of information exchange. The U.S. Government assumes no liability for the use of the information contained in this document.

Non-Binding Contents

Except for the statutes and regulations cited, the contents of this document do not have the force and effect of law and are not meant to bind the States or the public in any way. This document is intended only to provide information regarding existing requirements under the law or agency policies.

Quality Assurance Statement

The Federal Highway Administration (FHWA) 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. FHWA periodically reviews quality issues and adjusts its programs and processes to ensure continuous quality improvement.

Disclaimer for Product Names and Manufacturers

The U.S. Government does not endorse products or manufacturers. Trademarks or manufacturers’ names appear in this document only because they are considered essential to the objective of the document. They are included for informational purposes only and are not intended to reflect a preference, approval, or endorsement of any one product or entity.

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. FHWA-HRT-24-072	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle VOICES Cooperative Driving Automation (CDA) Proof-of-Concept Systems Integration Test 1		5. Report Date April 2024	
		6. Performing Organization Code	
7. Author(s) Yingyan Lou (ORCID: 0000-0001-9913-572X), Andrew Loughran, Jonathan Smet, Ethan Slattery, Robin Laqui, and Danielle Chou (ORCID: 0009-0001-6134-1971)		8. Performing Organization Report No.	
9. Performing Organization Name and Address		10. Work Unit No. (TRAIS)	
Leidos Inc. 11251 Roger Bacon Drive Reston, VA 20190		11. Contract or Grant No. 693JJ321D000010	
12. Sponsoring Agency Name and Address Office of Safety and Operations Research and Development Federal Highway Administration Georgetown Pike McLean, VA 22101-2296		13. Type of Report and Period Covered Final for the SIT-1; not final for the program; August 2022–September 2022	
		14. Sponsoring Agency Code HRSO-40	
15. Supplementary Notes The contracting officer's representative was Danielle Chou (HRSO-40; ORCID: 0009-0001-6134-1971).			
16. Abstract This report documents the results and technical performance analysis of the proof-of-concept Virtual Open Innovation Collaborative Environment for Safety (VOICES)(2) Systems Integration Test 1 (SIT-1)(1) during the months of August and September 2022. VOICES SIT-1 provided a proof of concept that demonstrated how four simulation nodes across three geographically distributed sites could communicate and coordinate real-time testing using some of the planned VOICES technologies. VOICES SIT-1 featured two prototype cooperative driving automation (CDA) applications: work zone and platooning, both implemented using the Federal Highway Administration's CARMA℠ suite of CDA tools.(3,4) The results from the proof-of-concept VOICES SIT-1 verified that VOICES can be used for distributed testing of cooperative and connected transportation applications, including CDA prototypes and simulations.			
17. Key Words VOICES, distributed testing, CDA, CARMA, digital twin, simulation, co-simulation, collaborative testing, LVC	18. Distribution Statement No restrictions. This document is available to the public through the National Technical Information Service, Springfield, VA 22161. https://www.ntis.gov		
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 90	22. Price N/A

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

SI* (MODERN METRIC) CONVERSION FACTORS

APPROXIMATE CONVERSIONS TO SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
in	inches	25.4	millimeters	mm
ft	feet	0.305	meters	m
yd	yards	0.914	meters	m
mi	miles	1.61	kilometers	km
		AREA		
in^2	square inches	645.2	square millimeters	mm^2
ft^2	square feet	0.093	square meters	m^2
yd^2	square yard	0.836	square meters	m^2
ac	acres	0.405	hectares	ha
mi^2	square miles	2.59	square kilometers	km^2
		VOLUME		
fl oz	fluid ounces	29.57	milliliters	mL
gal	gallons	3.785	liters	L
ft^3	cubic feet	0.028	cubic meters	m^3
yd^3	cubic yards	0.765	cubic meters	m^3
		NOTE: volumes greater than 1,000 L shall be shown in m^3		
		MASS		
oz	ounces	28.35	grams	g
lb	pounds	0.454	kilograms	kg
T	short tons (2,000 lb)	0.907	megagrams (or “metric ton”)	Mg (or “t”)
		TEMPERATURE (exact degrees)		
°F	Fahrenheit	5 (F-32)/9 or (F-32)/1.8	Celsius	°C
		ILLUMINATION		
fc	foot-candles	10.76	lux	lx
fl	foot-Lamberts	3.426	candela/m^2	cd/m^2
		FORCE and PRESSURE or STRESS		
lbf	poundforce	4.45	newtons	N
lbf/in^2	poundforce per square inch	6.89	kilopascals	kPa

APPROXIMATE CONVERSIONS FROM SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
mm	millimeters	0.039	inches	in
m	meters	3.28	feet	ft
m	meters	1.09	yards	yd
km	kilometers	0.621	miles	mi
		AREA		
mm^2	square millimeters	0.0016	square inches	in^2
m^2	square meters	10.764	square feet	ft^2
m^2	square meters	1.195	square yards	yd^2
ha	hectares	2.47	acres	ac
km^2	square kilometers	0.386	square miles	mi^2
		VOLUME		
mL	milliliters	0.034	fluid ounces	fl oz
L	liters	0.264	gallons	gal
m^3	cubic meters	35.314	cubic feet	ft^3
m^3	cubic meters	1.307	cubic yards	yd^3
		MASS		
g	grams	0.035	ounces	oz
kg	kilograms	2.202	pounds	lb
Mg (or “t”)	megagrams (or “metric ton”)	1.103	short tons (2,000 lb)	T
		TEMPERATURE (exact degrees)		
°C	Celsius	1.8C+32	Fahrenheit	°F
		ILLUMINATION		
lx	lux	0.0929	foot-candles	fc
cd/m^2	candela/m2	0.2919	foot-Lamberts	fl
		FORCE and PRESSURE or STRESS		
N	newtons	2.225	poundforce	lbf
kPa	kilopascals	0.145	poundforce per square inch	lbf/in^2

*SI is the symbol for International System of Units. Appropriate rounding should be made to comply with Section 4 of ASTM E380.
(Revised March 2003)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

3D	three dimensional
BSM	basic safety message
CAV	connected and automated vehicles
CDA	cooperative driving automation
CoP	community of practice
CS	data collection system
CSV	comma-separated values
DCS	data collection system
DoD	Department of Defense
DMZ	demilitarized zone
DSRC	dedicated short-range communications
EM	Execution Manager
FHWA	Federal Highway Administration
GUI	graphical user interface
IPsec	Internet Protocol Security
IT	information technology
KB	kilobyte
LVC	live, virtual, and constructive
MOM	mobility operations message
OM	object model
OST-R	Office of the Secretary of Transportation-Research
PoC	proof of concept
PCAP	packet capture
ROS	Robot Operating System
RSU	roadside unit
SDO	stateful distributed objects
SPaT	signal phasing and timing
SRC	Scientific Research Corporation
STOL	Saxton Transportation Operations Laboratory
TCM	traffic control message
TCR	traffic control request
TDCS	TENA data collection system
TENA	test and training enabling architecture
TFHRC	Turner-Fairbank Highway Research Center
TRMC	Test Resource Management Center
USDOT	U.S. Department of Transportation
V2I	vehicle-to-infrastructure
V2V	vehicle-to-vehicle
V2X	vehicle-to-everything
VOICES	Virtual Open Innovation Collaborative Environment for Safety

CHAPTER 1. SIT-1 OVERVIEW

PURPOSE OF THIS DOCUMENT

The U.S. Department of Transportation (USDOT) Office of the Secretary of Transportation-Research (OST-R) is responding to cross-cutting departmental needs related to interoperability testing of connected and automated vehicles (CAVs) and connected infrastructure. As a result, OST-R collaborated with the Federal Highway Administration (FHWA) to establish the initial framework and instantiation of a Virtual Open Innovation Collaborative Environment for Safety (VOICES).(2)

This document serves as the VOICES Systems Integration Test 1 (SIT-1) report.(1) This chapter provides an overview of the SIT-1 integration test scope and test objectives and then highlights the innovations in SIT-1. Chapter 2 describes SIT-1 as executed. Chapter 3 provides a high-level summary of SIT-1 results and reflects on how SIT-1 satisfies the VOICES core system requirements defined by the community of practice (CoP) in the early stages of the project. Chapter 4 details SIT-1 network and cooperative driving automation (CDA) performance results, followed by discussion and conclusion in chapter 5.(4)

BACKGROUND

VOICES is a platform that enables collaborative testing among participating entities (e.g., public sector including State and local governments, private sector, and academic institutions) in a distributed virtual collaborative environment for research and interoperability testing of prototypical connected transportation applications, such as CDA.(4)

VOICES leverages the test and training enabling architecture (TENA) as a common language used by all participants.(5) TENA was originally developed by the U.S. Department of Defense (DoD) and is managed by the US DoD Test Resource Management Center (TRMC). For the USDOT VOICES prototype, SIT-1 test used TENA as an interfacing language between prototype CDA vehicles and simulations from entities that were geographically distributed. These entities interacted in real-time in a common testing environment that blended live, virtual, and constructive (LVC) simulations. Following are DoD terminology for the LVC simulations:(6)

- Live simulations refer to simulation instances with real roadway infrastructure, real physical vehicles, real other roadway entities, or both. Either human drivers or automated software systems can operate real physical systems.
- Virtual simulations refer to simulation instances with real human users, operators, or both in a simulated travel environment.
- Constructive simulations refer to simulation instances with simulated vehicles, other road users operating, or both in simulated environments following predefined driving logic.

TENA's capability to carry out distributed and blended LVC simulations is achieved through a common object model (TENA OM) that enables semantic interoperability and a high-performance communication infrastructure (TENA middleware) for real-time data exchange.(5) Instances of TENA OMs are called stateful distributed objects (SDOs), and they carry data that describe relevant attributes of objects in the LVC simulations.(6) The VOICES prototype used TENA to connect multiple geographically distributed research and development sites; thus, it enabled distributed testing of prototype ecosystems from various stakeholders to collaborate across their respective individual simulation environments. Each geographical site can have one or more simulation nodes that represent a specific LVC simulation instance. All nodes across all sites were integrated into the VOICES prototype through TENA adapters and TENA middleware.

The Saxton Transportation Operations Laboratory (STOL) developed the CDA prototype vehicles and CAV simulations as part of the FHWA CDA program and utilized them in SIT-1.(4) The main product of the FHWA CDA Program is the CARMA℠ ecosystem, which consists of a suite of open-source software (OSS) for CDA and includes both vehicle and infrastructure technologies.(3) CARMA vehicle technologies include CARMA Platform℠ (full-stack CAV software) and CARMA Messenger℠ (software for connected but nonautomated vehicles).(7,8) CARMA infrastructure technologies include CARMA Streets℠ and CARMA Cloud℠.(9,10) CARMA Streets represents the infrastructure piece of CDA at conflict areas (e.g., intersections). A part of CARMA Streets is the Vehicle-2-Everything (V2X) Hub, a set of software that facilitates data exchange needed in V2X communications by translating data elements into various standard protocols.(11) CARMA Cloud further supports regional transportation systems management and operations (TSMO) through cloud-based management of transportation systems, data exchange, and multiple simultaneous remote services. The VOICES project also leveraged the cosimulation capability developed as part of the CARMA everything-in-the-loop (XiL) project.(3) The CARMA XiL cosimulation tool uses the Eclipse MOSAIC cosimulation framework and incorporates the full-stack CARMA Platform in the loop.(7) It also integrates other simulation tools such as Cars Learning to Act (CARLA®), Simulation of Urban Mobility (SUMO™) and Network Simulator 3 (NS-3).(12,13,14) The VOICES project used CARMA-CARLA integration, which is a part of the CARMA XiL cosimulation tool.(15) Note that other CDA prototypes and CAV simulations can be integrated into VOICES as well if TENA adapters are developed.(5)

INTEGRATION TEST SCOPE

The purpose of SIT-1 was to demonstrate a proof of concept (PoC) in which CDA testing used a common platform to coordinate the interactions of multiple geographically distributed LVC simulation nodes in a synchronous synthetic testing environment.(4,6)

SIT-1 involved three sites running four LVC simulation nodes.(6) The three sites were connected through site-to-site Internet Protocol Security (IPsec) tunnels on the Internet 2.(16) Note that this network configuration is not the intended long-term design of the VOICES network. Instead, the SIT-1 temporary network was configured within existing networking capabilities and security constraints.

SIT-1 was composed of multiple distributed testing experiments involving two prototype CDA applications, work zone and platooning.(4) SIT-1 used CARMA-specific implementations of both CDA applications since some CARMA software had already been written and could be easily adapted by the technical team.(3)

The work zone application was selected to demonstrate distributed CDA testing with CDA-enabled infrastructure.(4) In SIT-1, an active work zone with reduced speed was set up in CARMA Cloud.(10) Vehicles equipped with CARMA Platform received notification about the work zone through vehicle-to-infrastructure (V2I) communications.(7,17) The vehicle would then slow down on entering the work zone.

In the platooning application, vehicles formed a platoon by leveraging vehicle-to-vehicle (V2V) communications to join the lead vehicle from the rear.(18) Platooning requires the constant exchange of multiple different messages at a rate of at least 10 Hz; therefore, platooning performance may suffer if there is any distributed test latency. The resultant sensitivities would depend on the platooning distance/time gap, platooning speed, vehicle response times, network performance, and multiple other technical performance indicators.

SIT-1 TEST OBJECTIVES

Primary Objectives

The primary test objectives of SIT-1 were to ensure the integrity of the test results and to verify the following:

- Each node was connected to the secure network.
- All relevant adapters functioned properly as intended. Digital twins of all live dynamic elements (traffic signal and vehicle) were created and updated within a CARLA simulator at all networked virtual and constructive simulation nodes, using data from the live simulation node.(12)
- Digital twins of the virtual vehicle at the virtual simulation node were created and updated within a CARLA simulator at all networked constructive simulation sites, using data from the virtual simulation node.(12)
- Digital twins of constructive CARMA vehicles hosted at each constructive simulation node were created and updated within a CARLA simulator at other networked virtual and constructive simulation nodes.(3,12)
- All live and constructive CARMA vehicles hosted at physically distributed locations communicated with each other through TENA middleware.(5) Since the CDA V2V application tested during SIT-1 was platooning, verification of V2V communications in SIT-1 was limited to SAE International J2735b basic safety messages (BSM) and customized J2735 test messages used for CARMA Platform platooning plugin (i.e., CARMA Mobility Request Message, CARMA Mobility Response Message, and CARMA Mobility Operations Message).(3,19,20)
- All vehicles (live and constructive) in the platoon communicated with the onsite CARMA Cloud instance at Turner-Fairbank Highway Research Center (TFHRC).(10) Since the CDA V2I application tested during SIT-1 was work zone slowdown, verification of V2I communications in SIT-1 was limited to customized J2735 test messages used for the

work zone application (i.e., CARMA Traffic Control Request (TCR) and CARMA Traffic Control Message (TCM) messages).(17,19,20)

- Data collection, data visualization, and test replay functionalities operated as intended.

- Output measurements and analysis statistics of network distribution network latencies, jitter, and packet loss across the distributed network were at baseline levels without running any CDA applications. (4)

Secondary Objectives

SIT-1 also had secondary test objectives to better understand the performance of specific CDA applications using the test platform. The following are the secondary objectives:(4)

- Measure and analyze statistical distribution of network latencies, jitter, and packet loss across the distributed network during performance of CDA applications.(4)

- Measure and characterize statistical distribution of platoon formation and steady-state performance by live and constructive CARMA vehicles hosted at physically distributed locations. Relevant performance metrics are the following:

    - Speed of each vehicle.
    - Vehicle response times to CARMA Mobility Request Messages and CARMA Mobility Operations Messages (MOMs).(20)
    - Vehicle response times to commanded vehicle control inputs.
    - Distance or time or both headway between two consecutive vehicles (live or constructive) in the platoon.

- Measure and characterize statistical distribution of quality of service from CARMA Cloud for the work zone slowdown use case. One relevant performance metric was the CARMA Cloud response times to CARMA TCR messages.(10)

- Measure and characterize statistical distribution of work zone responses by the platoon of CARMA vehicles (live and constructed) hosted at physically distributed locations. Relevant performance metrics were:

    - Vehicle response times to CARMA TCMs.
    - Vehicle response times to commanded vehicle control inputs.
    - Speed of each vehicle upon receiving TCM from CARMA Cloud and through the work zone.(10)

SIT-1 FIRSTS

Building upon VOICES Demo 0 (which integrated a live and a virtual node at the same physical site), SIT-1 included significant new features.

Whereas Demo 0 used a local area network located at TFHRC, SIT-1 established a temporary secure network across three sites: TFHRC in McLean, VA; MITRE Corporation in McLean, VA;

and Scientific Research Corporation (SRC) in Augusta, GA. This distributed testing network was created using site-to-site IPsec tunnels and was configured with existing networking capabilities and security constraints.(16)

Second, SIT-1 included LVC simulation nodes, whereas VOICES Demo 0 tested only live and virtual simulation nodes.(6) In each constructive simulation, the CARMA Platform software stack was fully integrated into CARLA simulation.(7,12) New VOICES adapters for these constructive nodes enabled the CARMA Platform instances at the constructive nodes to interact with live vehicles equipped with CARMA Platform. Simulation ground truth was replicated at each site in its local CARLA simulation; each simulation contained replicas (referred to as digital twins) of real and simulated roadway objects hosted at other sites.

Third, SIT-1 verified that distributed testing of cooperative and connected transportation applications, including CDA prototypes and simulations over a secure common network, is feasible.(4) SIT-1 showcased the distributed testing of two CDA use cases: vehicle platooning and work zone speed reduction. For the platooning use case, new VOICES adapters and OMs were developed to enable simulated communication of the required CARMA mobility messages among relevant nodes.(20) The work zone use case included a local instance of CARMA Cloud as part of SIT-1 to exchange TCR or TCM through relevant new VOICES adapters with all vehicles (live and constructive) in the platoon.(10)

CHAPTER 2. SIT-1 DESIGN

FUNCTIONAL ARCHITECTURE

SIT-1 included four LVC simulation nodes hosted at three geographically distributed sites.(6) The three sites were TFHRC, MITRE Corporation, and SRC. A live simulation node and a virtual simulation node were hosted at TFHRC. Two identical constructive simulation nodes were hosted at the other two sites, respectively. In addition, a local instance of CARMA Cloud and core TENA middleware applications were hosted at TFHRC.(10,5)

Figure 1 shows the functional architecture of SIT-1 as executed. Each simulation node had a dedicated computer running relevant CDA and CAV simulation software, as well as VOICES adapters developed using TENA technologies.(4,5) The SIT-1 plan contains more detailed block diagrams of VOICES adapters at each simulation node. Appendix A includes the detailed block diagrams for readers' convenience.

The functional architecture of the executed SIT-1 differs slightly from SIT-1 Plan's architecture. One difference is that the core TENA middleware (e.g., TENA Console, Entity Generator, Scenario Publisher) were installed and run from computer 4 (figure 1), the same computer that also ran the virtual simulation.(5) The SIT-1 plan had proposed to run the core TENA middleware from a separate computer at TFHRC. This change did not functionally affect SIT-1. The core TENA middleware can run from any computer connected to the VOICES platform. Another difference is that each simulation node ran an instance of the TENA Data Collection System (TDCS) during execution. In the SIT-1 Plan, only one instance of the TDCS was considered. Multiple instances of TDCS (one at each simulation node) were needed during execution to produce data necessary for latency analysis. In chapter 4, describes how latency analysis used data from multiple TDCS.

NETWORKING

A main component of the VOICES system will be the secure, efficient, and distributed network used to connect all participants. In the interim, the SIT-1 network architecture used secure site-to-site IPsec tunnels to connect all participants: a direct tunnel was created from every site to every other site using IPsec security protocols.(16) This network configuration was logistically and technically onerous for receiving information technology (IT) approvals and to physically establish.

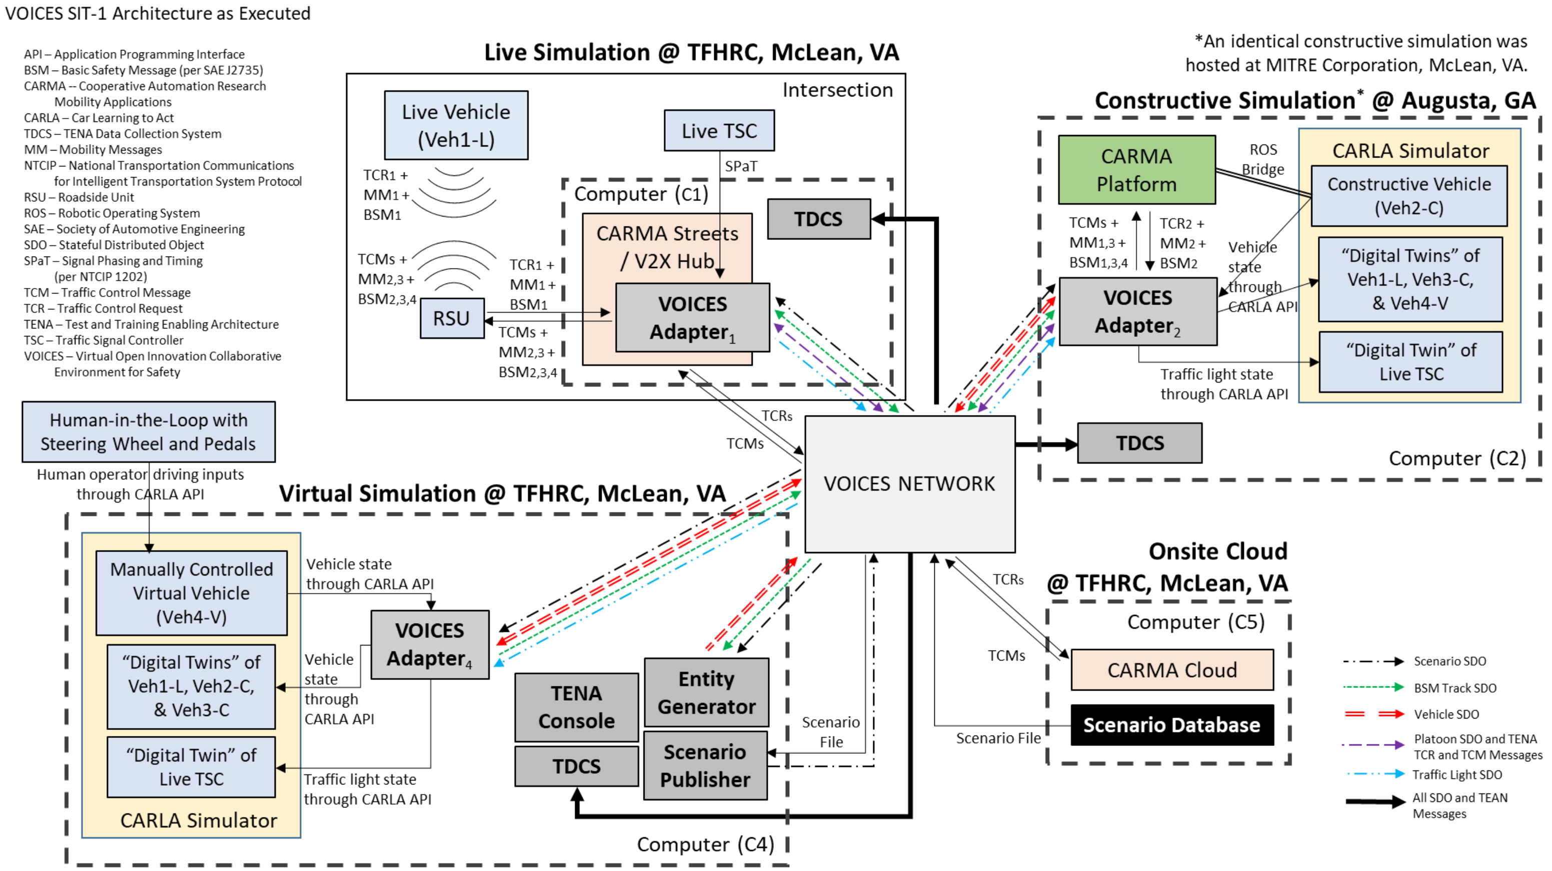

Source: FHWA.

API = application programming interface, TDCS =TENA Data Collection System, MM =mobility messages, NTCIP =National Transportation Communications for Intelligent Transportation System Protocol, RSU =Roadside Unit, ROS = Robotic Operating System, SAE = Society of Automotive Engineering, SDO = stateful distributed object, SPaT = Signal Phase and Timing, TCM = Traffic Control Message, TCR = Traffic Control Request, TENA = Test and Training Enabling Architecture, TSC =Traffic Signal Controller, VOICES = Virtual Open Innovation Collaborative Environment for Safety.

See references 5, 7, 9, 10, 11, 12.

Figure 1. Illustration. SIT-1 functional architecture.

A main component of the VOICES system is a secure, efficient, and distributed network that connects all participants. The SIT-1 network architecture used secure site-to-site IPsec tunnels to connect all participants.(16) A direct tunnel ran from every site to every other site using IPsec security protocols. This network configuration was logistically and technically onerous for receiving IT approvals and to physically establish.

Figure 2 presents an overview diagram of the SIT-1 network and shows that devices from a simulation site configured to be behind the same firewall and connected to the same gateway router. While SRC and MITRE simulation sites had only one computer connecting them to SIT-1 test event, TFHRC simulation site had multiple computers behind the same gateway router. In the middle of the diagram are the three site-to-site tunnels.

Each site went through comprehensive design, submittal, and approval processes within their internal IT security teams to allow and to open their respective tunnels. At TFHRC, engineers set up a special network demilitarized zone (DMZ) in the high speed I2 network. This DMZ was accessible only from TFHRC using specific ports and devices. Each site had a detailed network design document. Each document contains detailed hardware and software specifications for every device connected to the SIT-1 network, including IP addresses and contact information. For security reasons, these documents are not publicly available.

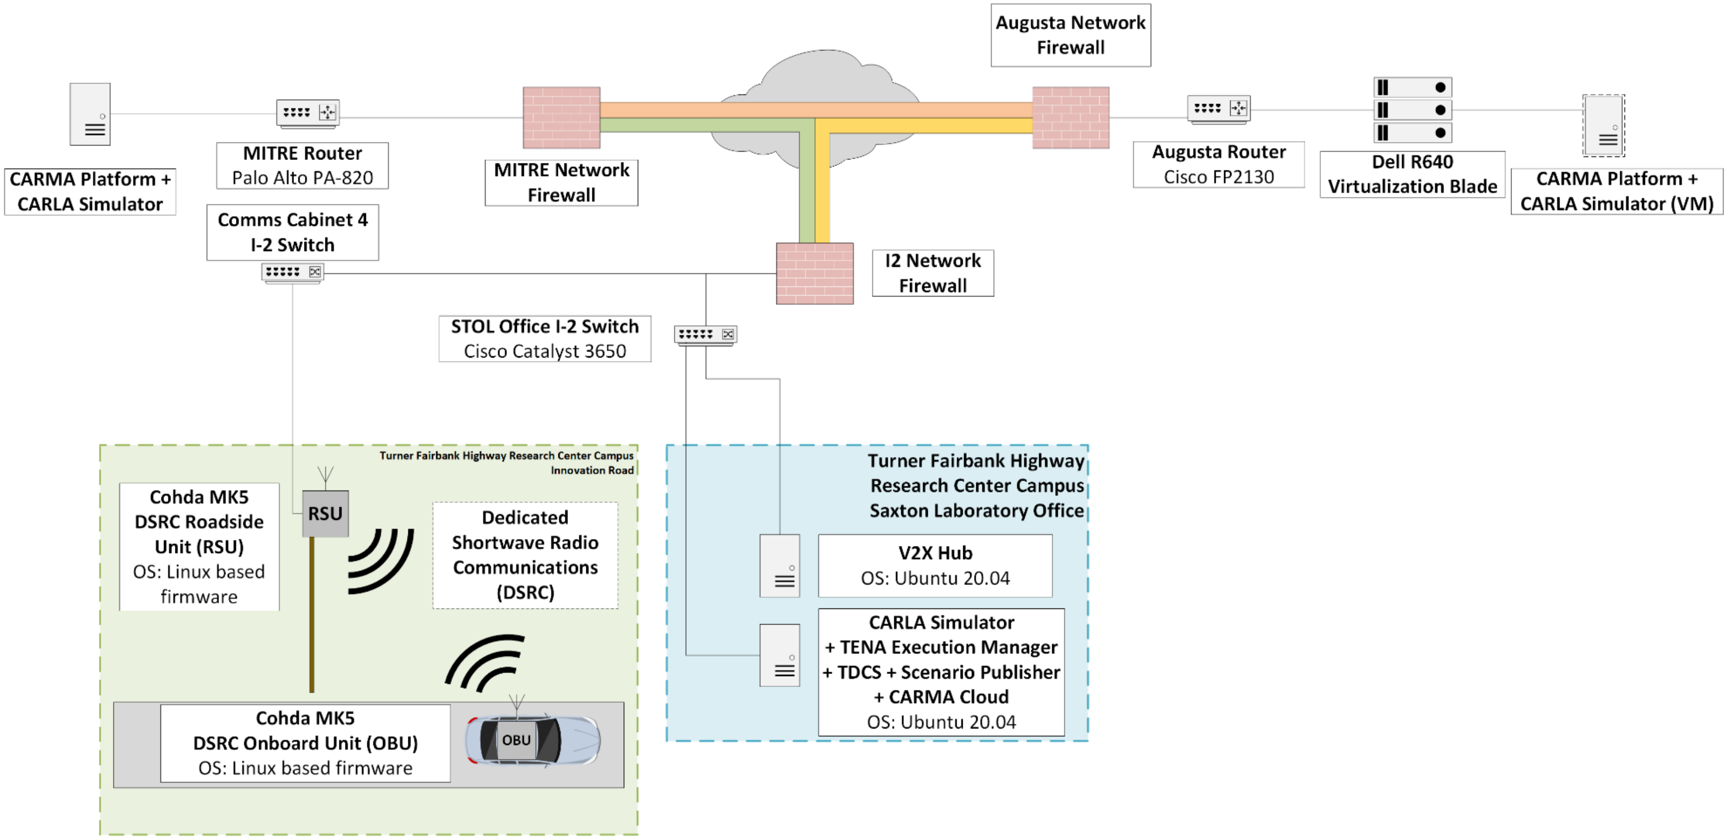

Source: FHWA.
See references 5, 7, 11, 12, 21.

Figure 2. Diagram. SIT-1 network.

SOFTWARE CONFIGURATION MANAGEMENT

This section addresses the software versions used for SIT-1.

CARMA Infrastructure Software

The following infrastructure software versions were deployed for SIT-1:

- V2X Hub: 399-voices-sit1-fix.(11)
- CARMA Cloud.(10)

CARMA Platform

To package CARMA Platform code and manage its dependencies, STOL uses Docker.(7) For all test runs, all live and constructive vehicles ran a customized version of CARMA Platform developed specifically for SIT-1, which is available on Docker.

The Docker image of CARMA Platform used during SIT-1 is primarily based on CARMA Platform 3.11.0 but contains the most recent update of the work zone feature released in CARMA Platform 4.0.3 and several software updates that addressed challenges and bugs discovered during SIT-1 test preparation.(7) At the time of this writing, the most up-to-date version of CARMA Platform is 4.2.0, released on July 29, 2022. SIT-1 used an older version of CARMA Platform because that was the most recent version the CARMA XiL cosimulation tool currently supports. The use of this older version of CARMA Platform caused some incompatibility issues, which were resolved through a customized solution. Chapter 5 presents details and approaches to future improvements.

CARMA-CARLA Integration

SIT-1 used the CARMA-CARLA Integration Tool, developed as part of the CARMA XiL cosimulation project.(15) The specific branch of the CARMA-CARLA Integration Tool used is carma-simulation-1.0.0.

TENA® Middleware

SIT-1 used TENA-MiddlewareSDK-v6.0.8.B, which is available on the TENA-SDA website.(5)

TENA Object Models for SIT-1

Nine TENA OMs were developed for and adopted in SIT-1.(5) The source code of these TENA OMs can be found on the TENA-SDA website.

VOICES Adapters

The source code of all VOICES adapters and plugins developed for SIT-1 are hosted on the private TRMC Bitbucket account using the tag SIT-1_RFR. Relevant adapter and plugin applications deployed for SIT-1 using the VOICES TENA adapter build script located on the VOICES-PoC GitHub®.(22) Instructions for this build script are in the build script readme file.

TEST SCENARIO

The TFHRC campus was selected as the main test facility for SIT-1. All virtual and constructive simulation nodes at the three geographically distributed sites used a three-dimensional (3D) CARLA map. Figure 3 shows the TFHRC campus, SIT-1 test route (highlighted), and relevant roadway features. The entire test route is within the communication range of the roadside unit (RSU) at the west intersection.

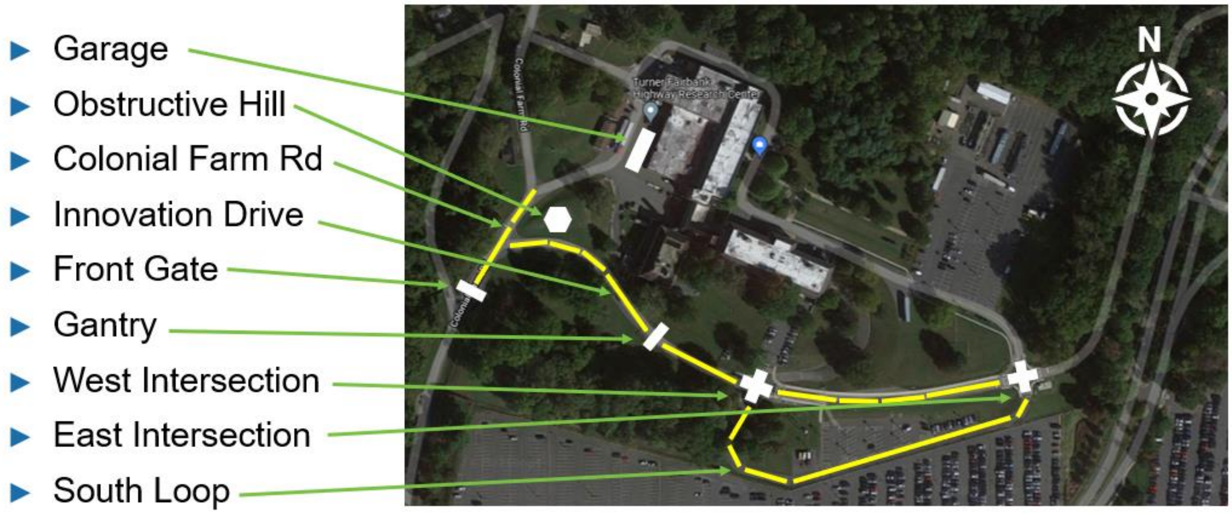

Original map: © 2022 Google Earth™. Modified by FHWA. (See Acknowledgements section.)

Figure 3. Map. SIT-1 Roadway environment and test route.(23)

SIT-1 included work zone and platooning applications implemented in CARMA.(3) An active work zone was configured in CARMA Cloud between the west and the east intersections with reduced speed limit.(10) For SIT-1, the work zone application ran simultaneously with the platooning application. At the beginning of the test scenario, the three vehicles (one live at TFHRC and two constructive at the other two sites) were staged near the T-intersection of Colonial Farm Road and Innovation Drive. The live vehicle was at the front of the group, followed by the constructive vehicle from SRC, and then the constructive vehicle from MITRE. The Scenario File specified the initial positions of the three vehicles. Test engineers at the three test sites coordinated to start engaging CARMA Platform one by one.(7) Once engaged, the three vehicles started traveling eastbound along Innovation Drive, driven by CARMA Platform. While engaged, CARMA Platform periodically broadcast TCRs to query roadway restrictions from the infrastructure. Upon receiving TCRs, CARMA Cloud sent back TCMs with relevant information such as geofence and associated roadway restrictions. The lead live CARMA vehicle registered the work zone geofence information and associated speed limit in its world model. As the three vehicles started to drive and came within a certain distance of each other, the platooning plugin of CARMA Platform automatically engaged. The three vehicles exchanged CARMA Mobility Request Messages, Mobility Response Messages, and MOMs to form a platoon.(20) Once the platoon was formed (around the gantry shown in figure 3), the platoon continued to travel eastbound on Innovation Drive. The west intersection light remained green in the east-west direction for the platoon to travel through without stopping. As the platoon of vehicles entered the active work zone, the lead live vehicle slowed down to comply with the reduced speed limit registered through received TCMs. The two constructive following vehicles also slowed down in response to the lead vehicle's deceleration, while maintaining the platoon. After the platoon left the work zone, the vehicles returned to command the normal speed set for the test route. Upon arrival at the east intersection, the safety driver in the live vehicle took manual control of the vehicle. The two constructive vehicles were configured to stop and disengage from the platoon at the east intersection. After the safety driver took control of the live vehicle, the safety driver maneuvered the live vehicle to make a right turn and drive westbound along the South Loop back to the west intersection.

EXPERIMENT DETAILS

Three sets of experiments were performed during SIT-1. Experiment 1 (August 16, 2022) included three runs with a Lexus® RX450H running CARMA Platform as the live lead vehicle.(7) Experiment 2 (September 13, 2022) included two tests with a Chrysler® Pacifica running CARMA Platform as the live lead vehicle. Experiment 3 (September 21, 2022) included an additional three tests with a Chrysler Pacifica running as the live lead vehicle. The two real vehicles were equipped with different low-level vehicle controllers: the Lexus RX450H was equipped with a PACMod controller, and the Chrysler Pacifica had a New Eagle controller.

In each set of experiments, the constructive vehicle at SRC was the second vehicle in the platoon, and the constructive vehicle at MITRE was the third vehicle in the platoon.

All eight runs across the three sets of experiments except one were successful, where all vehicles completed the route and carried out the two CDA applications as intended.(4) The unsuccessful run was experiment 3 run 2, in which the third platooning vehicle ran off the road halfway through the test route. Further analysis of data from this unsuccessful run (chapter 4 for details) indicated that the computer running the constructive simulation seemed to have been momentarily frozen, which could have prevented CARMA Platform from effectively controlling the vehicle.(7)

HARDWARE SPECIFICATION

This section provides detailed hardware specifications for every computer used in SIT-1 in table 1. All computer numbers mentioned in table 1 are in reference to the SIT-1 functional architecture diagram). SIT-1 post-test analysis highlighted the fact that high-performing computing hardware can help provide a smoother testing experience. SIT-1 did not conclusively determine minimum hardware requirements required by specific node(s) to successfully complete a test. Future researchers should design studies to test outcomes of hardware variations among one or more nodes.

Table 1. SIT-1 hardware.

Site: Node: Computer	CPU (Cores at Base Clock Speed)	RAM (Speed)	GPU (Video Memory)
TFHRC: Live simulation node: Live vehicle In-vehicle computer	Intel Xeon™ E3-1275 (4 cores at 3.4 GHz)	32 GB (2,133 MHz)	NVIDIA RTX 2080 (8 GB)
TFHRC: Live simulation node: computer 1–CARMA Streets/V2X Hub(9,11)	Intel Atom™ E3845 (4 cores at 1.91 GHz)	8 GB (1,333 MHz)	None
TFHRC: Virtual simulation node: computer 4	Intel Core™ i9-10900X (10 cores at 3.7 GHz)	64 GB (2,933 MHz)	NVIDIA RTX 2080Ti (11GB)
SRC: Constructive simulation node: Computer 2	Intel Core™ i9-11900K (8 cores at 3.5GHz)	32 GB (3,200 MHz)	NVIDIA RTX A4000 (16 GB)
MITRE: Constructive simulation node: Computer 3	Intel i7-6700HQ (4 cores 2.6 GHz)	16 GB (2,400 MHz)	NVIDIA GTX 1060 (6GB)

The TFHRC site used three computers supporting two simulation nodes. The in-vehicle computer used to run CARMA Platform for the live vehicle was about 3 yr old at the time of the test and performed well during SIT-1 testing.(7) The CARMA Streets/V2X Hub computer (computer 1) is a small form factor edge computer that had been in service for about 3 yr.(9,11) The V2X Hub computer is not meant to handle heavy computational load, as can be seen from its relatively low-end hardware specifications compared to other computers in table 1. Analysis results (chapter 4) suggest that the additional computational load of the adapters proved to overwhelm the CARMA Street/V2X Hub computer. The TFHRC virtual simulation computer (computer 4) was a high-performance computer at STOL that was about 3 yr old at the time of the test and performed well during SIT-1.

A single computer (computer 2) was used for the constructive simulation node at SRC in Augusta, GA. This computer was one of the two high-performance computers purchased for the project.

A single laptop computer (computer 3) was used for the constructive simulation node at MITRE in McLean, VA, and was 6 yr old at the time of the test. Like the V2X Hub computer, this laptop appeared to suffer from computational overwhelm, which likely was due to the computer having older hardware processor technologies as well as the cooling limitations from a laptop's form factor.(11)

Chapter 4 details results that revealed the importance of hardware. Chapter 5 discusses the hardware issue and next steps.

CHAPTER 3. SIT-1 HIGH-LEVEL ASSESSMENT

SIT-1 TEST RESULTS SUMMARY

SIT-1 verified that researchers could use the VOICES prototype for distributed testing of cooperative and connected transportation applications, including CDA prototypes and simulations over a secure common network.(4)

Primary Test Objectives

The primary test objectives of SIT-1 were achieved and verified in the following ways.

Objective A. Verify that each node was connected to the secure SIT-1 network.

Each node was connected using secure IPsec tunnels.(16) Connectivity was tested using the TENA ping test.(24) The TENA Ping Test also produced latency data for baseline network performance analysis (Objective I).

Objective B. Verify that all relevant adapters function as intended.

VOICES adapters deployed for SIT-1 were tested against the criteria specified in the Test Procedures section of the SIT-1 Test Plan. The test procedures are also included in appendix B for the readers' convenience. All adapters met their respective criteria and functioned as intended. The criteria specified in the SIT-1 Test Plan were those that test engineers could readily observe during an active run. Some examples of such criteria include increasing numbers of associated SDOs of relevant TENA plugins displayed in the TENA console and status of CARMA Platform instances shown through CARMA Platform UI.(5,7) Upon further investigation of all data collected, testers observed that some adapters did not generate certain expected outputs (certain message types). These missed messages, however, did not affect the CDA applications performed during SIT-1.(4) Chapter 4, Network Performance Results under Load, details results.

Objectives C–E. Verify that the digital twins of all elements were created and updated within a CARLA simulator at all networked virtual and constructive simulation nodes.

All digital twins were first verified visually. The visual verification was a quick sanity check during an active run to make sure that all expected digital twins were generated and that their status properly reflected that of the source entities to naked human eyes. All digital twins of vehicles and traffic signals passed the visual verification performed by STOL test engineers. Data logs were further processed and aligned as part of the latency analysis to achieve Objectives C-E. While testers observed some discrepancies in certain data fields (e.g., vehicle positions and speed), the discrepancies were very minor and likely arose due to computational precision (i.e., rounding and truncation). Chapter 4 and appendix C discuss this situation.

Objectives F–G. Verify that V2X and V2I communications required by the two CDA applications were successfully carried out through the SIT-1 network.

The specific V2X messages relevant to the two CDA applications are the following:(11,4)

- Customized J2735 test message used for platooning, namely CARMA Mobility Request Message, CARMA Mobility Response Message, CARMA Mobility Operations Message (Info), and CARMA Mobility Operations Message (Status).(19,20)
- Customized J2735 Test Message used for work zone, namely CARMA TCR message and CARMA TCM.(19,20)

For these messages to be communicated among different simulation nodes across the network, relevant VOICES adapters translated the message into TENA Stateful Data Objects (SDOs) and TENA messages, and then converted the messages back to their original formats by appropriate adapters at the receiving simulation node.(5) All relevant and available datasets (i.e., packet capture (PCAP) files, TENA SDO and TENA message datasets) were collected by the TDCS at each step of the communication and data conversion pipeline process and reviewed across all steps to ensure that each message was properly encoded and decoded.

Testers verified that all relevant adapters properly handled all messages related to platooning and were successfully passed among all three vehicles (one live and two constructive). Additionally, for SIT-1, successful demonstration of the three-vehicle CARMA platoon with mixed live and constructive vehicles verified test Objective F.(3)

For the work zone application, TCR and TCM exchanges between the live vehicle and the onsite CARMA Cloud were successful. For constructive vehicles, the last component in the communication and data conversion pipeline did not function.(10) The TENA TCM Messages meant for the two constructive vehicles were not translated back to J2735 CARMA TCMs for the associated CARMA Platform instances at the two constructive simulation nodes.(5,19,7) However, this missed last step did not affect the work zone application tested in SIT-1 because all vehicles were in a platoon when entering the work zone and the constructive vehicles slowed down following the lead live vehicle. The relevant adapter (carma-platform-tena-adapter in figure 30 in appendix A) should be updated to fix this issue.

Objective H. Verify that the VOICES data collection, data visualization, and test replay functionalities operated as intended.

VOICES simulation nodes communicated entirely using TENA SDO and TENA Message data across the network.(5) Testers captured, viewed, and replayed data using standard TENA utilities, including TDCS, TENA DataView, and TENA Playback. For SIT-1, testers used TDCS to collect the data, TENA DataView to view the data (including live graphical visualizations), and TENA Playback to replay the data. Figure 4 shows an example of a da TENA DataView graph, and figure 5 shows a TENA Playback example. Test engineers manually verified Objective H during and after SIT-1 experiments.

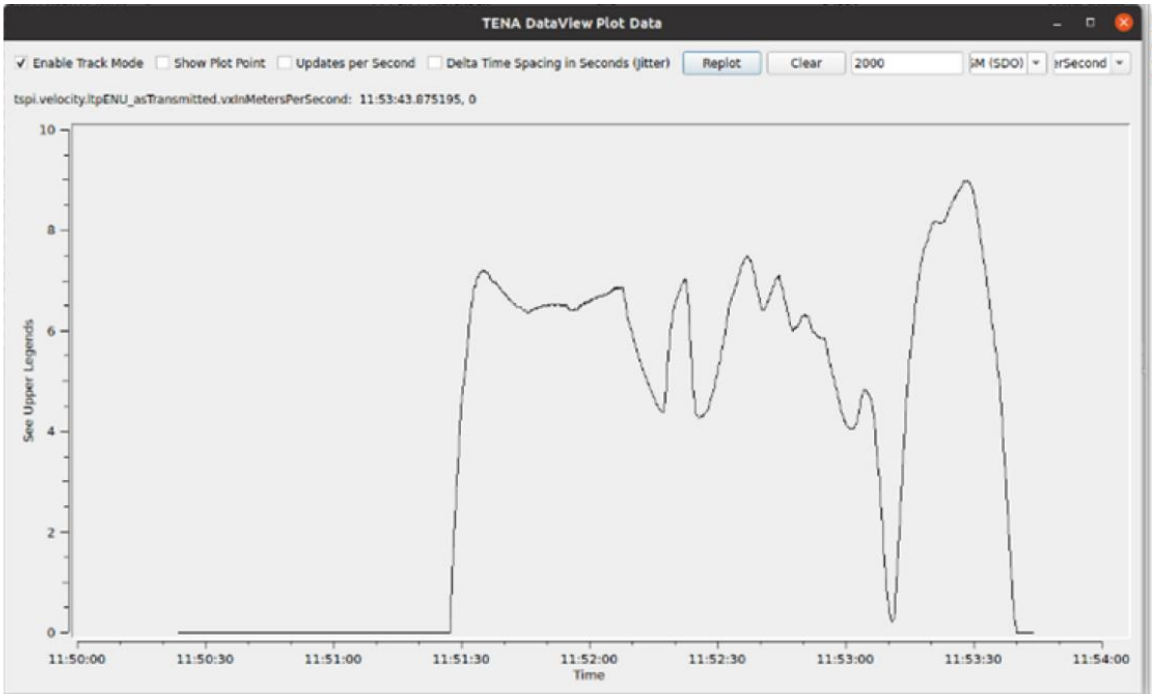

TENA DataView Plot Data

☑ Enable Track Mode ☐ Show Plot Point ☐ Updates per Second ☐ Delta Time Spacing in Seconds (jitter)

2000

```
tspi.velocity.ltpENU_asTransmitted.vxInMetersPerSecond:  11:53:43.875195, 0
```

Source: FHWA.

Figure 4. Graph. Example of a TENA DataView graph.(5)

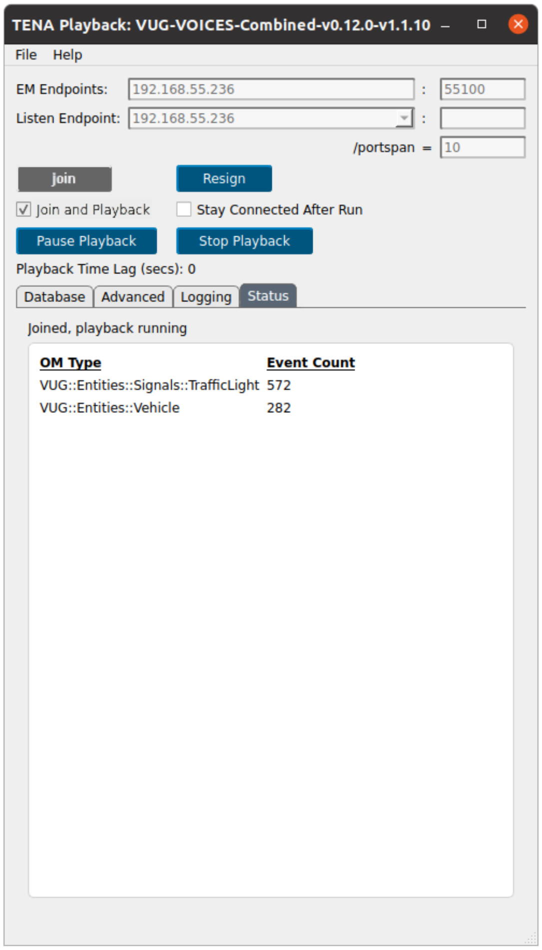

EM Endpoints: 192.168.55.236 : 55100
Listen Endpoint: 192.168.55.236 :
/portspan = 10
☑ Join and Playback
☐ Stay Connected After Run
Playback Time Lag (secs): 0

Joined, playback running

OM Type	Event Count
VUG::Entities::Signals::TrafficLight	572
VUG::Entities::Vehicle	282

Source: FHWA.

Figure 5. Screenshot. Example of a TENA playback tool example.(5)

Objective I. Measure and analyze statistical distribution of network latencies, jitter, and packet loss across the distributed network at baseline levels without running any CDA applications.(4)

Chapter 4, the Baseline Network Performance Results section, presents detailed network performance results at baseline levels without running any CDA applications.(4) Network latencies across the SIT-1 network are shown to be low under baseline conditions. The round-trip transmission latency between the two simulation sites in McLean, VA, was under 5 ms on average and under 20 ms on average between any of the two Virginia sites and the Georgia site. Jitter under baseline conditions was largely under 1 ms. During testing, there was no observed packet loss while transmitting data across the SIT-1 network.

Secondary Test Objectives

Objective J. Measure and analyze statistical distribution of network latencies, jitter, and packet loss across the distributed network during performance of CDA applications.(4)

Chapter 4 contains detailed network performance results while running the two CDA applications.(4) Transmission latencies across the SIT-1 network are like those under baseline conditions.

Objective K. Measure and characterize statistical distribution of platoon formation and steady-state performance by live and constructive CARMA vehicles hosted at physically distributed locations.(3)

Detailed analysis and results are addressed in chapter 4. Plots of individual vehicle speeds and space headways between vehicles show how well live and constructive vehicles follow commanded speeds issued by the CARMA Platform.(7) The results revealed the differences in controlling physical and constructive vehicles, due to the limited fidelity of physics simulation CARLA.(12) Chapter 5 provides additional information.

Objective L. Measure and characterize statistical distribution of quality of service from CARMA Cloud for the work zone slowdown use case.(10)

Chapter 4 gives the segment-by-segment latency analysis of the CARMA TCR/TCM data flow pipeline.(3) A direct comparison was made between the CARMA TCR/TCM data flow of the live vehicle (which did not involve any TENA components (table 22) and the CARMA TCR/TCM data flow of the constructive vehicle at SRC, which used the VOICES TENA adapters (n/a = not applicable.

table 23).(5) The total TCR-to-TCM time for the live vehicle was approximately 136 ms, while the constructive vehicle at SRC averaged around 130 ms.

Objective M. Measure and characterize statistical distribution of work zone responses by the platoon of CARMA vehicles (live and constructed) hosted at physically distributed locations.(3)

Chapter 4, CDA Application Performance, details analysis, and results.(4) Findings from Objective M are like those from Objective K.

VOICES CORE REQUIREMENTS ACHIEVED

Table 2 contains the project team's self-assessment on how SIT-1 satisfied VOICES Core Requirements (CR). VOICES CR are a set of system requirements developed by the VOICES CoP during the conceptual design stage of the VOICES project.

If all or some of a specific requirement was achieved, it was noted in the "Satisfied by SIT-1" column. Remarks about whether the requirement was achieved—Yes, No, Partially, or N/A—are contained in the "Remarks" column.

Table 2. VOICES core requirements achieved by SIT-1.

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR1	Use Case Scenario Database	The system shall provide a database which contains the descriptions of use cases that the Cooperative Driving Automation community is interested in researching or testing.	High	Partially	A Scenario Publisher was included in SIT-1. The Scenario Publisher pulled a scenario run-time configuration file from the scenario database and published relevant scenario data.
CR2	Addition of New Component Adapters	The system shall provide a method for community users to submit adapter applications for system components, e.g., a simulator, to interact with other system components through a common middleware.	High	Yes	SIT-1 included a live intersection, a live vehicle, two constructive vehicles, and a manually driven virtual vehicle in CARLA.(12) Six VOICES adapters/ plugins were developed to enable data conversion and communication among four LVC simulation nodes.(6) SIT-1 Plan for more detailed block diagrams of VOICES adapters/plugins at each simulation node.

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR3	Addition of New Data Objects	The system shall provide a method for community users to submit OMs for data objects that will be shared within the execution environment.	High	Partially	TENA Design Language (TDL) is used to define OMs for VOICES.(5) The TENA TDL is in the process of being released to the public. Once released, the community users can use the OM libraries for the OMs to create their own adapters.
CR4	Protection of Repositories	The system shall provide a method to restrict access to the adapters, OMs, and use cases so that only authorized community users may access them.	High	Yes	All adapters and OMs are currently hosted on private repositories managed by TRMC, which requires user registration for access. Adapters and OMs can be permissioned on group basis by TRMC. Individual users can be given access to specific groups. Users will have access to all OMs contained within a group to which they have access.
CR5	Selective Privacy of Use Cases	The use case database shall allow use cases to be selectively shared with specified community users.	Medium	N/A	N/A
CR6	Shared Execution Environment	The system shall provide a network environment which facilitates the participation of multiple physical and simulated systems in a single execution of a test scenario.	High	Yes	SIT-1 was performed on a distributed network with three sites hosting four LVC simulation nodes.(6) TFHRC hosted a live and a virtual simulation node. SRC and MITRE each hosted a constructive simulation node.

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR7	Control of Participation in Shared Execution Environment	The system shall be designed such that participants' access to a cooperative execution can be explicitly controlled.	High	Partially	Access to the SIT-1 network was explicitly granted using secured IPsec tunnels between sites.(16)
CR8	Control of Access to Measured Data	The system shall be designed such that data captured during an event can be selectively exposed to users.	High	N/A	N/A
CR9	Functional Architecture	The system shall provide the ability to integrate different types of functional components to the system, e.g., different numbers of vehicles and pedestrians, different live, virtual, and constructive components.	High	Yes	SIT-1 involved a live traffic signal, a live vehicle, two constructive vehicles, and a manually driven virtual vehicle. The components and actors for each test can be easily added, removed, or modified in the scenario file.
CR10	Persistent Connectivity	The system shall provide persistent connectivity to participants.	Medium	Partially	Access to the SIT-1 network was persistent using secured IPsec tunnels between sites.(16) The network configuration used was specifically and temporarily stood up for SIT-1, then subsequently dismantled.

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR11	Virtual Environment Generation	When connecting live and virtual components, the system shall create and maintain a shared cyberphysical environment that allows the components to interact as if they were in the same environment.	Medium	Yes	SIT-1 created a shared cyberphysical environment for all four simulation nodes using a 3D map of the TFHRC campus and relevant VOICES adapters/plugins. The live vehicle and the two constructive vehicles successfully interacted with each other to form and maintain a platoon in the shared cyberphysical environment.
CR12	Compatibility With Co-Simulation Environments	The VOICES test bed shall seamlessly provide multiple modes of operation to assist in testing algorithms ranging from immersion in a fully simulated environment with thousands of entities from multiple simulations to embedding in a real-time system for testing in a live, virtually, constructive distributed system.	Medium	Yes	SIT-1 included a live intersection, a live vehicle, two constructive vehicles, and a manually driven virtual vehicle in CARLA. Data from all four LVC simulation nodes was integrated into a shared cyberphysical environment in real time.(6)
CR13	Live Observation of Parameters	The system shall provide the capability to observe desired parameters of the test execution in real or near-real time.	Medium	Yes	SIT-1 showed that the TDCS and the TENA DataView tool can be used to observe data in real time.(5)
CR14	Storage and Access of Parameters for Future Observation	The system shall provide the capability to observe desired parameters of a test after the test has been executed.	Medium	Yes	SIT-1 showed that the TDCS can collect the data from the demonstrations and store it in a file that can be observed after the test.

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR15	Authentication of System Users	The system shall provide a mechanism to identify users of the system so that access to test events and data can be controlled.	High	Partially	Access to the SIT-1 network was explicitly granted using secured IPsec tunnels between sites and all computers are password-protected. An authentication mechanism to accommodate remote and distributed users has not been developed.(16)
CR16	Representation of CDA Environments(4)	The system shall be capable of representing the roadway and infrastructure environments relevant to CDA use cases with associated traffic and vulnerable road users.(4)	High	Partially	SIT-1 showed the integration of CARMA Streets/V2X Hub and CARMA Messenger (for BSM broadcasting) into VOICES.(9,11,8)
CR17	Support Hardware Dedicated to Testing	The system shall support the addition of hardware to measure test data for the systems under test when the operational components of the systems are not sufficient to meet the needs of the test.	Low	N/A	N/A
CR18	Large-Scale Traffic Simulation	The system shall be capable of executing tests involving thousands of entities on multiple simulation systems as well as actual hardware entities.	Low	N/A	N/A

Req	Name	Description	Importance	Satisfied by SIT-1	Remarks
CR19	Data Analytics	The system shall support the collection and aggregation of data from test scenarios such that it can be analyzed by either: 1) analysis systems that are connected to the VOICES DSTE or 2) external systems.	Low	Yes	All SDO data from SIT-1 was gathered by the TDCS and can be ingested into an analytics framework.
CR20	Environmental Conditions	The system shall support the simulation of operational conditions such as weather, road type, time of day, and geographic location.	Medium	Yes	CARLA was chosen as the simulator for SIT-1.(12) CARLA supports simulated sensing environment, including weather, lighting, and terrain.
CR21	Big Data Analytics	The system shall support the capturing and logging of unstructured data such as the type typically logged by automated vehicle systems.	Low	Partially	Scripts were developed and used during SIT-1 to capture appropriate CARMA Platform logs and rosbags, but this functionality was not directly implemented into TENA data collection system (TDCS)SIT-1.(7,25,5)

DSTE = dynamic synthetic test environment; N/A = not applicable; — = no information; TDL =TENA design language.

CHAPTER 4. SIT-1 DETAILED ASSESSMENT

This chapter presents detailed assessment methodologies and results on the following:

- Baseline SIT-1 network performance. (table 4).
- Network performance under load. (Network Performance Results under Load: Throughput and Network Performance Results under Load: Packet Loss and Latency.)
- CDA application performance during distributed testing. (chapter 4, CDA Application Performance).(4)

BASELINE NETWORK PERFORMANCE RESULTS

This document reports four metrics for the baseline network performance: throughput, latency, jitter, and packet loss.

Throughput

Network throughput is the actual amount of data traffic flowing from a specific individual (or group of) source(s) to another specific individual (or group of) destination(s) at any given time. This differs from bandwidth, which is the theoretical maximum amount of data traffic a network can support at any given time. Peak throughput for the baseline network was measured by transferring large files from each SIT-1 site to every other SIT-1 site. Performed in August 2022, this experiment transferred rosbag files that ranged from 10–30 gigabytes. 25) table 3. shows the maximum recorded transfer speed between sites. All sites used internet service with an advertised bandwidth of 1 gigabit upload and download from their providers.

Table 3. SIT-1 network maximum throughput results.

	Destination		
Source	**TFHRC**	**SRC**	**MITRE**
TFHRC	N/A	30 MB/s	11.1 MB/s
Augusta	26 MB/s	NA	10.2 MB/s
MITRE	12.8 MB/s	12 MB/s	N/A

Latency, Jitter, and Packet Loss

The baseline network latency, jitter, and packet loss were evaluated using the TENA ping test, which generates round-trip latency values.(24) The TENA ping test is superior to the standard ping test because it uses the actual TENA protocol to perform communication exchanges with the TENA Execution participants using either reliable (i.e., TCP) or best effort (i.e., User Datagram Protocol Multicast) communication mechanisms. These ping tests were initiated from the TENA Console, but other applications can perform ping tests. The ping tests generate a matrix of results from every application to every other application. Tests can occur with varying numbers of packets and delays. Results reported in this document were calculated based on data obtained from the TENA ping test using TENA Canary as the test application on September 14, 2022.

Table 4 reports the SIT-1 baseline network average round-trip latency (*L*) and jitter (*J*) results, both in milliseconds, based on the number of successful pings (*P*). Let l_i denote the latency value of the *ith* successful ping, where *i* belongs to set N then jitter is calculated as seen in figure 6:

$$J = \frac{1}{P}\sum_{i=2}^{P}|l_i - l_{i-1}|$$

Figure 6. Equation. TENA ping test equation.[(24)]

Table 4. SIT-1 network baseline round-trip latency and jitter (in milliseconds).

	Destination			
Source	**TFHRC V2X Hub Computer**	**TFHRC Virtual Simulation Computer**	**MITRE (McLean, VA)**	**SRC (Augusta, GA)**
TFHRC V2X Hub computer[(11)]	N/A	P: 378 L: 0.613757 J: 0.416446	P: 378 L: 4.878307 J: 0.305040	P: 377 L: 21.397878 J: 1.417553
TFHRC virtual simulation computer	P: 378 L: 0.947090 J: 0.095491	N/A	P: 378 L: 4.706349 J: 0.135279	P: 377 L: 20.655172 J: 0.031915
MITRE	P: 378 L: 4.777778 J: 0.278515	P: 378 L: 4.616402 J: 0.145889	N/A	P: 378 L: 17.989418 J: 0.021220
SRC	P: 377 L: 20.798408 J: 0.236702	P: 377 L: 20.687003 J: 0.109043	P: 377 L: 18.002653 J: 0.005319	N/A

J = jitter (in milliseconds); L = latency (in milliseconds); N/A = not applicable; P = number of pings.

Testers observed that the average network round-trip latency was less than 1 ms between the two TFHRC computers; less than 5 ms between the two test sites in McLean, VA; and around 20 ms between Virginia and Georgia. All jitter values, except one (between TFHRC V2X Hub computer and SRC in Augusta, GA) were less than 1 ms.

Among the 4,531 pings performed, none of the pings was dropped, which was expected as the reliable (i.e., TCP) communication protocol was selected for SIT-1.

NETWORK PERFORMANCE RESULTS UNDER LOAD: THROUGHPUT

To capture throughput during CDA applications, testers replayed the TENA data recorded by a TDCS instance during SIT-1.(4,5) Testers loaded a captured database file into the TENA Playback tool, started a new TDCS instance to receive the data, then started the replay. They repeated the process once for the full dataset using all message types, then replayed each message type individually. The NetHogs tool captured throughput values over time and compiled the values in comma-separated values (CSV) format.(26)

STOL analysis results showed that the throughput hovered around 120 KB/s while running the two SIT-1 CDA applications.(4) This throughput value is an order of magnitude smaller compared to the baseline throughput achieved by transferring large data files. The small throughput captured while running SIT-1 CDA applications verifies that the TENA data being transferred is lightweight, which is promising for distributed testing.(5)

Figure 7 through figure 10 show the throughput analysis results for experiment 3, runs 1–3. These figures were produced by replaying the TENA data recorded by the TDCS instance run at the virtual simulation node at TFHRC during SIT-1.(5)

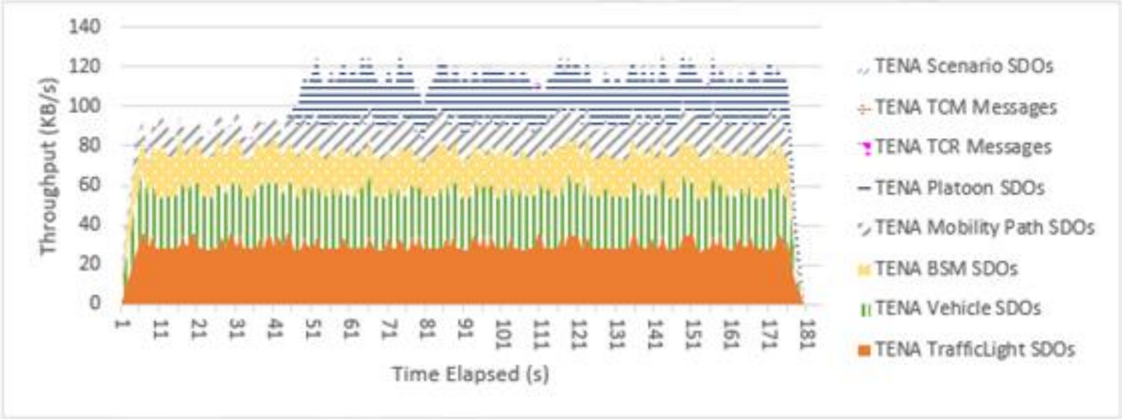

Source: FHWA.

Figure 7. Graph. CDA throughput per message type for experiment 3, run 1.

Before the CARMA Platform instances at the live and constructive sites engaged in automated driving mode, TENA TrafficLight, Vehicle, BSM, and Mobility Path SDOs comprised the TENA data traffic in the SIT-1 network.(7,5) Only one TENA Scenario SDO was published by the TENA Scenario Publisher; and all simulation nodes received the TENA Scenario SDO when they first joined the TENA execution. Figure 7 (experiment 3, run 1), shows the TENA TrafficLight SDOs hovered around 30–35 KB/s, TENA Vehicle SDOs were approximately 25–30 KB/s, TENA BSM SDOs approximately 20–25 KB/s, and TENA Mobility Path SDOs 10–20 KB/s. When the CARMA Platform instances at the live and constructive sites engaged in automated driving mode, TENA Platoon SDOs and TENA TCR and TENA message recording started. For experiment 3 run 1, the live vehicle's CARMA Platform started to engage around 44 s into the test (figure 7). TENA Platoon SDOs were approximately 20–30 KB/s, whereas

TENA TCR and TCM messages were negligible. Figure 8 shows similar observations made for experiment 3, run 2.

Around 83 s into the test for experiment 3, run 1 (figure 7 and figure 9), a significant drop in throughput occurred. Similarly, significant throughput dropped for experiment 3, run 2 about 43 s and 133 s into the test (figure 8 and figure 10).

Figure 9 and figure 10 plot the total throughput obtained through replaying the TENA data recorded by each TDCS instance at each of the four simulation nodes during experiment 3, runs 1 and 2, respectively.(5) This comparison shows that all sites received TENA data at approximately the same rate throughout each test.

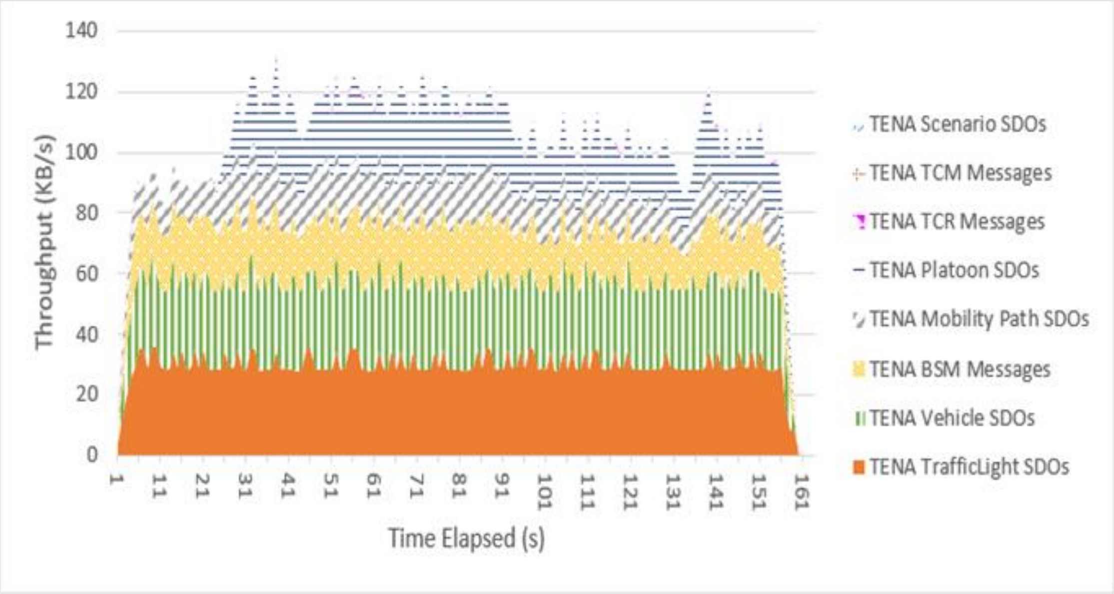

Source: FHWA.

Figure 8. Graph. CDA throughput per message type for experiment 3, run 2.

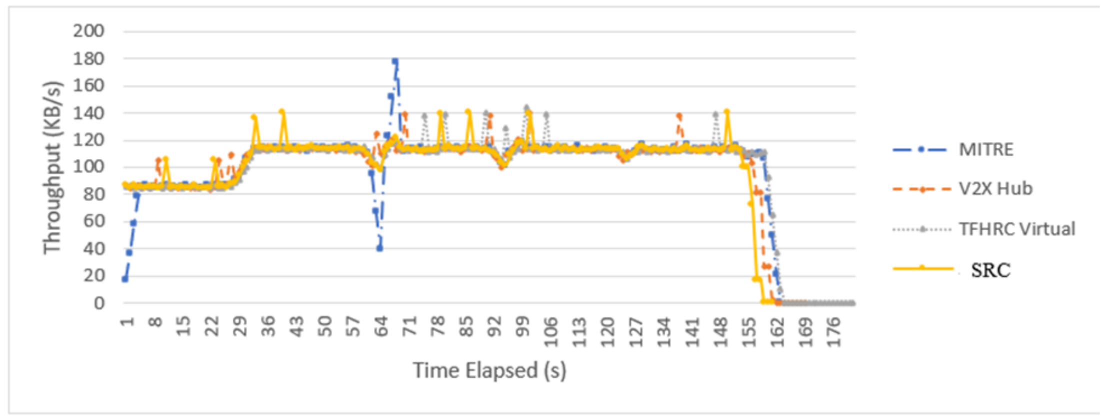

Source: FHWA.

Figure 9. Graph. Throughput for all sites during experiment 3, run 1.

Testers observed a large spike in data recorded by the MITRE TDCS instance (figure 9), which testers attributed to processing demands temporarily exceeding laptop computational capabilities. The additional small sharp spikes shown in figure 9 do not correlate among sites or to any known events at those periods. The spikes could be noise intrinsic to the data collection process, but this hypothesis must be verified during future testing.

Figure 10 shows similar data spikes for experiment 3, run 2 as occurred for experiment 3, run 1. Interestingly, the spikes were more pronounced in experiment 3, run 1 that was completed successfully. On the other hand, experiment 3, run 2 resulted in the constructive vehicle from MITRE running off the road at around 41 seconds due to an inability to generate and/or respond to steering controls. The observed data spikes were corroborated by the package loss analysis addressed in chapter 4, Latency.

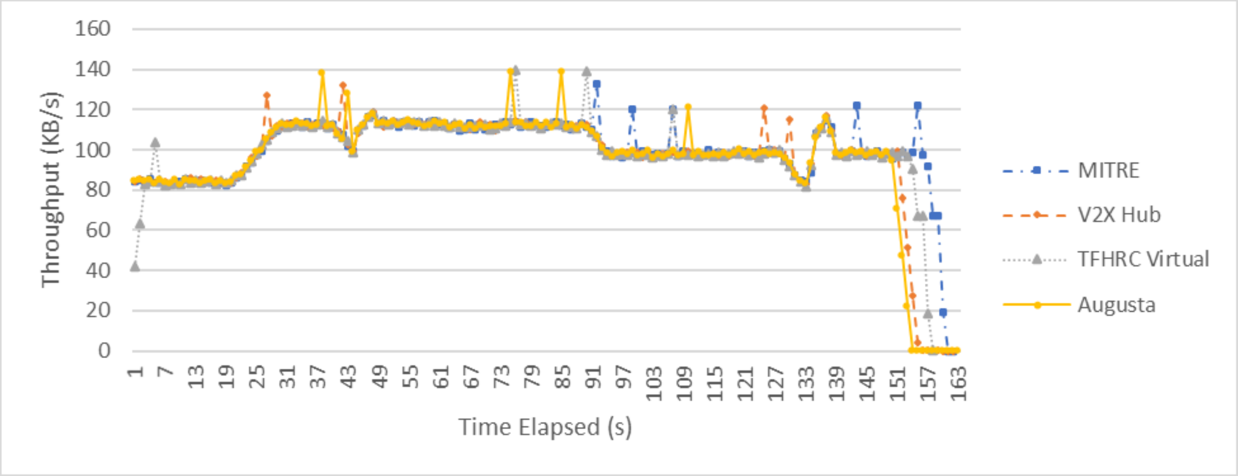

Source: FHWA.

Figure 10. Graph. Throughput for all sites during experiment 3, run 2.

NETWORK PERFORMANCE RESULTS UNDER LOAD: PACKET LOSS AND LATENCY

Methodology

Calculating the network performance results (latency, jitter, and packet loss) of the SIT-1 CDA applications (VOICE SIT-1 secondary objectives J and L) required in-depth analysis of the performance of each segment along the data pipeline.(4) The necessity for this degree of analysis is due to the inability to obtain end-to-end measures because data are transformed multiple times and transmitted between multiple systems along the pipeline, often without a persisting unique ID. One exception is the end-to-end latency and drop rate of the CARMA Mobility Request or Response Messages.(20) Chapter 4, Methodology, describes how end-to-end latency was computed for CARMA Mobility Request/Response messages computed end-to-end latency using rosbag data from a single vehicle.(25) The section also explains in detail the process of aligning data to compute segment-by-segment performance metrics for other types of messages used by the two SIT-1 CDA applications. The end of the section lists assumptions employed in the segment-by-segment analysis.

End-to-End Latency Calculation of CARMA Mobility Request/Response Messages from rosbag Data(20,25)

The CARMA Platform platooning protocol requires the exchange of two sets of CARMA Mobility Request and Mobility Response messages during the initial handshake between two vehicles.(7,20) After the two successful exchanges, the two vehicles are part of the same platoon.

A trailing vehicle generates the CARMA Mobility Request Message.(20) A preceding vehicle generates a CARMA Mobility Response Message for the trailing vehicle upon receiving its Mobility Request Message. The Mobility Response Message and Mobility Request Message include data fields that can be uniquely traced to the trailing vehicle. Therefore, using the trailing vehicle's rosbag data alone, calculating end-to-end latency of the Mobility Response/Request message combination is possible.(25) This latency measures the round-trip time from when a CARMA Mobility Request Message is generated to when a corresponding CARMA Response Message is received.

The end-to-end latency analysis of the CARMA Mobility Request/Response message combination led to very reliable results for two main reasons: 1) the uniquely identifiable traceability in the messages; and 2) the fact that data needed for this analysis came from a single system (the following vehicle), thus eliminating the need to align data.(20)

Data Alignment for Segment-by-Segment Performance Analysis

Network performance for all messages other than the CARMA mobility response or request message combination required segment-by-segment performance analysis.(20) To achieve this analysis, data was collected at each data transmission and transformation interface, wherever possible, leading to multiple datasets for each test run. Due to limited data logging capabilities at certain interfaces, the analysis adopted several approximations. The following subsection discusses these approximations. The datasets from a particular test run were processed and aligned to calculate the segment-by-segment network performance.

The system clocks of all computers used in SIT-1 needed to be synchronized to produce usable data. However, not all system clocks were correctly synchronized during experiment 1 and experiment 2 runs. As a result, the segment-by-segment packet loss and latency analysis results reported in this section contains only results from experiment 3.

All data collected was decoded and separated by message type. For transmission of J2735 encoded messages via dedicated short-range communications (DSRC) or TCP, PCAPs were taken at each source and destination.(19,21) A decoder script was utilized to decode the raw packets, extract relevant data, and generate separate CSV files for each J2735 message type. For TENA data creation and transmission, multiple instances of the TDCS were utilized at each site to capture all TENA data.(5) While all TDCS instances in SIT-1 subscribed to the same data, the TDCS instances were deployed at each site specifically to capture the receipt times of TENA data at each site. This TENA data was exported to CSV using the built-in TDCS export function.

The deployment of multiple TDCS instances increased the amount of TENA data traffic across the SIT-1 network.(5) However, since TENA OMs are designed to be lightweight, the additional data traffic resulting from multiple TDCS instances did not burden the SIT-1 network. As discussed in chapter 4, Network Performance Results Under Load Throughput, the peak throughput observed under load was around 120KB/s, a magnitude smaller than the peak throughput achieved in baseline conditions when transferring very large files.

Once all data was in CSV format, each packet was aligned for each step of its path from one entity to another. For example, a BSM sent from the live vehicle to a constructive vehicle takes the following path: the live vehicle equipped with a DSRC radio generates and broadcasts a J2735 BSM; V2X Hub receives that message; the V2XHub-TENA-BSM plugin converts that J2735 BSM into a TENA BSM SDO; that TENA BSM SDO data was transmitted to the constructive vehicle's CARMA-Platform-TENA adapter. (See references 21, 19, 11, 5, and 7.) Finally, the constructive vehicle's CARMA-Platform-TENA adapter converts the TENA BSM SDO back into a J2735 BSM and sends it to the corresponding CARMA Platform instance. Every step in that process contains its own dataset, and every message must be traced from source to destination to calculate performance metrics such as latency, jitter, and packet loss. This process must be repeated for every message type and every source and destination combination.

To analyze network test performance, an analysis script (calculate_e2e_perf.py) was developed.(27) The script performed the following major tasks:

- Loading the data.
- Filtering the datasets by desired message and datatype.
- Aligning the data.
- Checking for dropped packets.
- Calculating performance metrics.
- Generating summary metrics.

This process is not trivial due to the complexity of the datasets. Refer to appendix C for details on considerations, challenges, and approaches adopted in analyzing SIT-1 data.

This performance analysis script was designed for a single message type flowing from one entity to another.(27) To complete the analysis of data from a full run of VOICE SIT-1 scenario, the process needs to be repeated for every other message type and for each origin-destination combination. To streamline this process, a batch analysis script (batch calculate_e2e_perf.py) was developed to quickly locate all required input data files, generate runtime parameters, and execute the analysis for every combination of a single run.

Testers scrutinized the data processing and alignment scripts and outcomes to make sure that no error was introduced during this process.

Assumptions Employed in Segment-by-Segment Performance Analysis

The SIT-1 segment-by-segment performance analysis employed several assumptions due to the limited data logging capabilities of VOICES adapters. The adapters at each simulation node were designed to be lightweight and currently do not have data logging features that allow for traceability between data inputs and outputs. As a result, exact timestamps of when relevant data entered and left a specific adapter were not available. The TDCS instance at each simulation node was set up to subscribe to all TENA SDOs and TENA messages, which provided a close approximation to some exact timestamps of interest.(5) Additionally, certain J2735 messages as inputs to some adapters could not be captured at the transmission interface.(19) Therefore, testers made necessary approximations.

Assumption 1: For any TENA SDO or TENA Message input data to any VOICES adapter, the receipt time of the same TENA data at the TENA TDSC instance on the same computer as the adapter was used in this performance analysis.(5) Since the adapter and the TDCS instance were both hosted on the same computer and subscribed to the same data, testers assumed that both would receive the TENA data of interest at approximately the same time. While further analysis indicated that this assumption might not be ideal, it was the only approximation that was possible during SIT-1. Chapter 4, Latency, provides more details.

Assumption 2: For any TENA SDO or TENA Message output data from any adapter, the receipt time of the same TENA data at the TDSC instance on the same computer as the adapter was used in this analysis.(5) Since the adapter and the TDCS instance were both hosted on the same computer, testers assumed that the TDCS instance received the relevant data at approximately the same time the data was published by the VOICES adapter. The project team believes this assumption is reasonable, based on knowledge about TENA middleware. The analysis results seemed normal and did not suggest otherwise.

Assumption 3: The receipt time of J2735 signal phase and timing (SPaT) as an input into the V2X Hub-TENA-SPaT plugin (Figure 26 in appendix A) was approximated by the transmit time of the same J2735 SPaT message sent from the CARMA Streets/V2X Hub computer to the RSU. (See references 19, 24, 9, and 11.) This approximation is far from ideal. Given more time and resources, a more accurate approximation could be achieved by examining the NTCIP 1202 SPaT data into the CARMA Streets/V2X Hub computer (figure 26 in appendix A). Further analysis supported that this approximation should be reconsidered in the future. Chapter 4, Latency, provides more details.

These assumptions were needed only because the adapters currently do not have any data logging capabilities. Such capabilities could be developed for adapters in the future to provide accurate data to support network and CDA application performance analysis purposes.(4)

Packet Loss

Analysis examined the total number of dropped packets and the number of dropped packets over time for each vehicle and for each message type. As expected, overall, no TENA data dropped when transmitting across the SIT-1 network because TENA was configured to transmit data via TCP (reliable) for SIT-1.(5) Very occasional packet drops occurred when J2735 messages were transmitted from the source to the relevant adapters.(19) This behavior was also expected because transmission was via UDP (best effort). The drop rate of J2735 message transmission was about 1 or 2 in over 1,000 messages. Additionally, the project team saw a relatively high message drop rate of about 5–10 percent during experiment 3, run 2, via the TDCS, for the CARMA-platform-TENA adapter instance run at the MITRE simulation node. Testers suspect the elevated drop rate is the result of the computing demand exceeding the hardware specification.

The following sections present a subset of representative results from experiment 3, run 3 (successful), as well as comparable results from experiment 3, run 2 (unsuccessful).

Total Number of Dropped Packets for Experiment 3, Run 3

V2I Communications

The work zone application employed V2X communications.(28) CARMA Platform sent CARMA TCR messages to the CARMA Cloud, and CARMA Cloud responded with CARMA TCMs.(7,10) CARMA TCR/TCM were handled together in the dropped packet and latency calculation.

Table 5 shows the dropped packets for V2I communications between the live vehicle and the CARMA Cloud instance at TFHRC.(17,10) The table also shows that no packet was dropped between the live vehicle and the onsite CARMA Cloud instance.

Table 5. TCR/TCM packet loss results for experiment 3, run 3 Live.

Metric	Live Simulation—CARMA Cloud(10)
Total packets analyzed	18
TCR DSRC message transmission (live vehicle to V2X Hub)(21,11)	0
J2735 TCR receipt to J2735 TCM sent (TFHRC V2X Hub computer)(19,11)	0
TCM DSRC message transmission (V2X Hub to live vehicle)(21,11)	0

Table 6 presents the dropped packet results for the V2I communication between the constructive vehicle at Augusta, GA, and the CARMA Cloud instance at TFHRC.(10) The 0 value in the table represents no packet loss for those messages. The V2I communication pipeline was functional except for the last component; the CARMA-Platform-TENA adapter did not convert the TENA TCM Message into a J2735 TCM Message, which was to be used by the constructive vehicle's CARMA Platform.(See references 17, 5, 7, 19.)

Table 6. TCR/TCM packet loss results for experiment 3, run 3 SRC.

Metric	TCR/TCM
Total Packets Analyzed	11
TCR SDO message creation from J2735 (MITRE)(19)	0
TCR SDO network transmission (from SRC to TFHRC)	0
TCM SDO creation from TCR SDO receipt (V2X Hub)(11)	0
TCM SDO network transmission (from TFHRC to SRC)	0
J2735 message creation from SDO (SRC)(19)	—*

*TCM SDOs were not converted into J2735 messages for this testing and, therefore, no dropped packet data was available.(19)

V2V Communications

Table 7–table 9 show some excerpts of analysis on the total number of dropped packets from experiment 3, run 3.

Note that if a packet is dropped in an earlier dataflow step, it was also recorded as dropped in the subsequent steps. For example, in table 7 a BSM packet was dropped at the DSRC message transmission step along the data pipeline from the live lead vehicle to the constructive vehicle at SRC; and the same dropped packet was marked in all subsequent steps.(21) Similar situations were observed for the CARMA Mobility Path Message from the live lead vehicle to the constructive vehicle at SRC (Mobility Path Message column in table 7).(20) A single dropped packet was observed for the BSM and for the CARMA Mobility Path Message from the live lead vehicle to the constructive vehicle at MITRE (BSM and the Mobility Path Message columns in table 8).

In the current platooning logic, CARMA Mobility Request Messages are sent only from joining vehicles, and Mobility Response Messages are sent only from leading vehicles.(20) Therefore, table 7 for the live vehicle does not contain any Mobility Request Messages. Similarly, table 8 and table 9 do not contain any Mobility Response Messages.

In addition, some dropped packet values were marked as 0 even though the SDO data did not exist in the TDCS database (Mobility Response column in table 7 and table 8, and the Mobility Request column in table 9).(20) The packets are not considered dropped because the J2735 data still reached the destination even though the TDCS database was not updated.(19) This situation occurred because in the current implementation of relevant adapters, the platoon SDO data exchanges between adapters during the platoon initiation process were not recorded in the TDCS database. Future plans are to ensure all data exchanges are properly recorded and easily traceable in all datasets.

From table 7–table 9, the total number of dropped packets for most datasets was relatively low (fewer than 3 out of 1,000, or 0.3 percent). However, the older laptop used at MITRE dropped 5–10 percent of packets at the CARMA-Platform-TENA adapter step during the SDO to J2735 conversion in 1 of 8 test runs during experiment 3, run 3 (last row in table 8 and table 9).(7,5,19)

Table 7. Packet loss results for experiment 3 run 3 live to SRC.

Metric	BSM	Mobility Path	Mobility Request	Mobility Response*	Mobility Operations (INFO)	Mobility Operations (STATUS)
Total packets analyzed	1,257	1,830	0	4	591	1,241
DSRC message transmission (live vehicle to V2X Hub)(21,11)	1	1	N/A	0	0	0

Metric	BSM	Mobility Path	Mobility Request	Mobility Response*	Mobility Operations (INFO)	Mobility Operations (STATUS)
SDO creation (TFHRC V2X Hub computer)(11)	1	1	N/A	0*	0	0
SDO network transmission (from TFHRC to SRC)	1	1	N/A	0*	0	0
J2735 message creation from SDO (SRC)(19)	1	1	N/A	0	2	1

N/A = not applicable.

*Not all messages were reflected as SDO updates in the TDCS database, but all messages arrived at the intended destination as J2735 messages.(19)

Table 8. Packet loss results for experiment 3, run 3 live to MITRE.

Metric	BSM	Mobility Path	Mobility Request	Mobility Response	Mobility Operations (INFO)	Mobility Operations (STATUS)
Total packets analyzed	1,188	1,196	0	4	2,339	1,160
DSRC message transmission (live vehicle to V2X Hub)(21,11)	1	1	N/A	0	0	0
SDO creation (TFHRC V2X Hub computer)(11)	1	1	N/A	See asterisk footnote	0	0
SDO network transmission (from TFHRC to MITRE)	1	1	N/A	See asterisk footnote	0	0
J2735 message creation from SDO (MITRE)(19)	121	55	N/A	0	61	132

N/A = not applicable.

*Not all messages are reflected as SDO updates in the TDCS database, but all messages arrived at the intended destination as J2735 messages.(19)

Table 9. Packet loss results for experiment 3, run 3 MITRE to SRC.

Metric	BSM	Mobility Path	Mobility Request	Mobility Response	Mobility Operations (INFO)	Mobility Operations (STATUS)
Total packets analyzed	1,214	1,228	4	0	0**	1,160
SDO message creation from J2735 (SRC)(19)	0	0	See asterisk footnote	N/A	N/A	0
SDO network transmission (from SRC to MITRE)	0	0	See asterisk footnote	N/A	N/A	0
J2735 message creation from SDO (MITRE)(19)	19	19	0	N/A	N/A	0

N/A = not applicable.

*Not all messages are reflected as SDO updates in the TDCS database, but all messages arrived at the intended destination as J2735 messages.(19)

**Vehicles only broadcast CARMA MOM (INFO) messages if they were the leader of their platoon or if they were not currently platooning.

Dropped Packets Over Time for Experiment 3, Run 3

To further investigate the dropped packets observed during the successfully completed experiment 3, run 3 (table 7–table 9), testers observed the dropped packets over time and plotted them (figure 12, figure 13, and figure 14).

Figure 11 reflects the very few dropped packets in table 7 for the BSM, Mobility Path, Mobility Operations (Info), and Mobility Operations (Status) messages from the live lead vehicle at TFHRC to the constructive vehicle at SRC.(20)

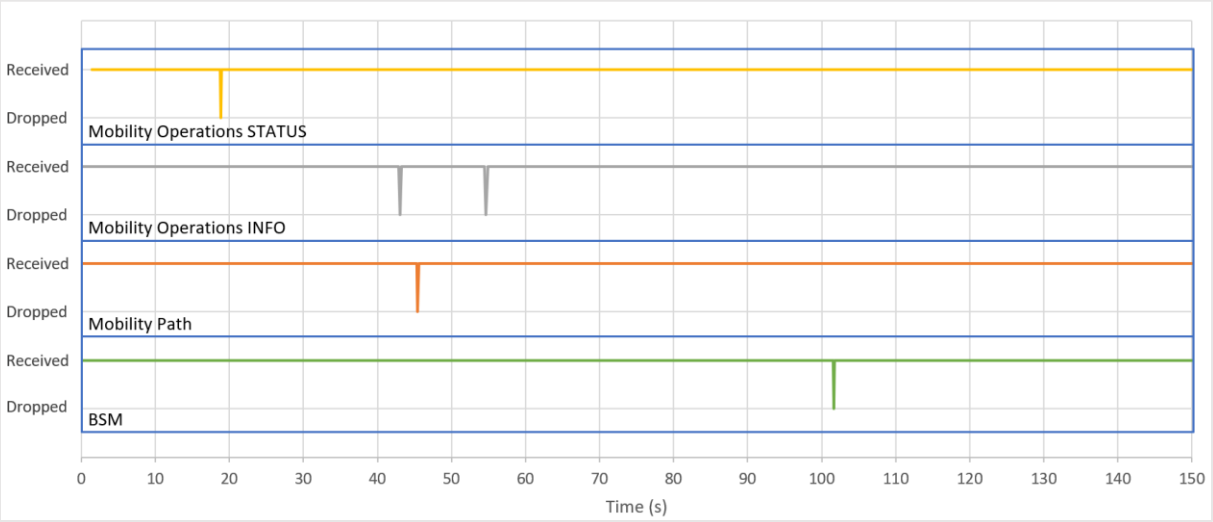

Source: FHWA.

Figure 11. Graph. Packet loss results over time for experiment 3, run 3, live to SRC.

Figure 12 shows the dropped packets plotted over time from the live simulation node at TFHRC to the constructive simulation node at MITRE. The packets dropped around the same time for all message types.

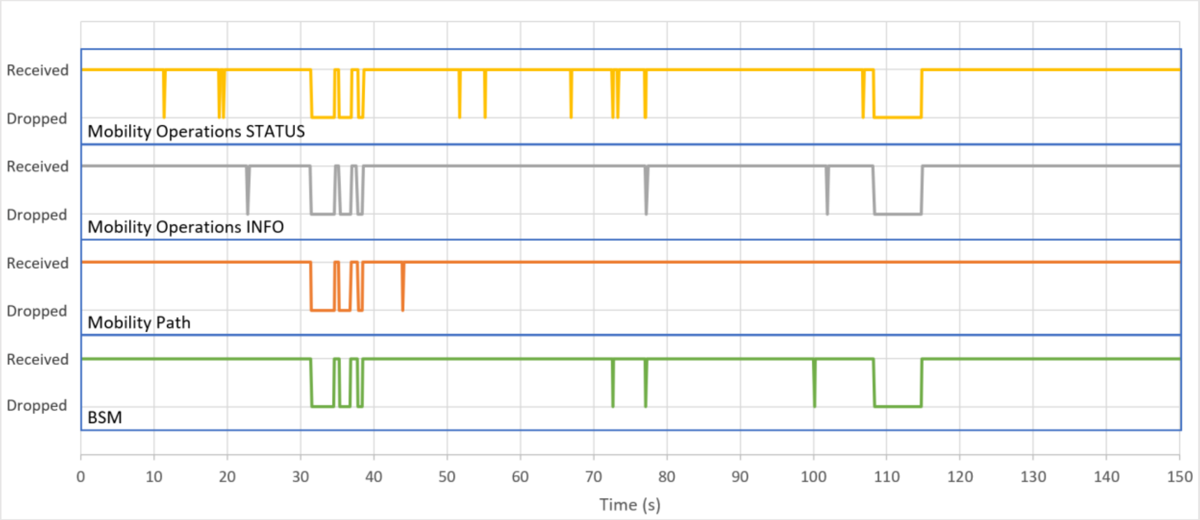

Source: FHWA.

Figure 12. Graph. Packet loss results over time for experiment 3 run, 3 live to MITRE.

Figure 13 plots relevant packet loss data over time for messages from the constructive vehicles at MITRE to the constructive vehicle at SRC. Figure 11–figure 13 show the dropped packets. Note that the gaps in figure 13 represent a lack of data for that period, not dropped packets. This observation supports the hypothesis that computing demands exceeded the processing capabilities of the MITRE laptop, which became unable to perform any calculations. The degree of impact on CDA functionalities depends on when packet loss occurs. For example, the packet loss observed in experiment 3, run 3 did not affect the successful completion of the SIT-1 scenario with two CDA applications.(4) Similar packet losses occurred during experiment 3, run 2, in which the constructive vehicle at MITRE ran off the road.

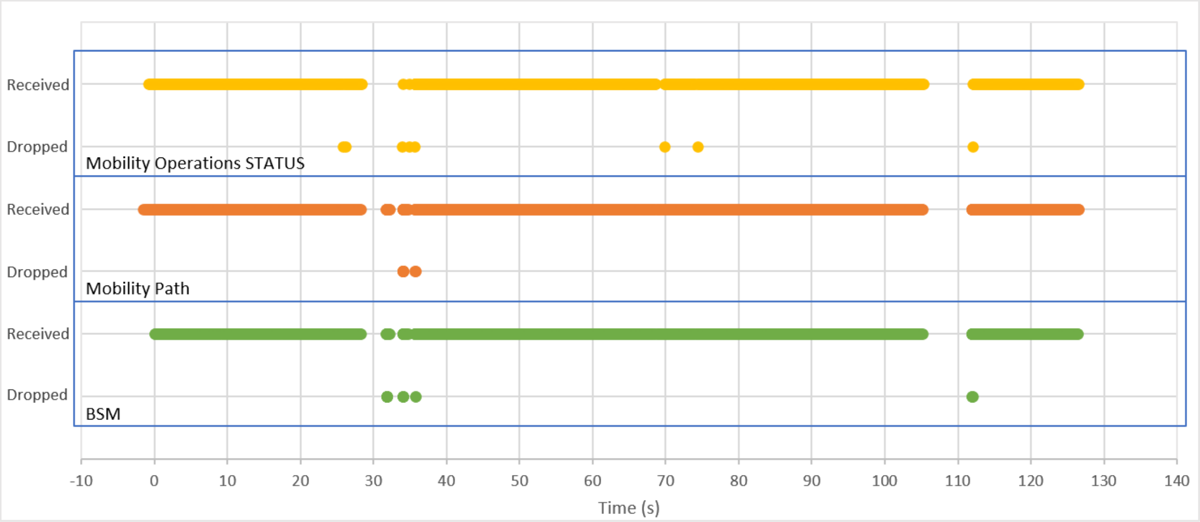

Source: FHWA.

Figure 13. Graph. Packet loss results over time for experiment 3, run 3 MITRE to SRC.

Dropped packets Over time for experiment 3, run 1

This section shows another set of dropped packet analysis results from the successfully completed experiment 3, run 1. Figure 14–figure 16 show similar trends as the results from experiment 3, run 3 (figure 11–figure 13). While testers observed minimal dropped packets for the communication between the live simulation node and the constructive simulation node at SRC (figure 14), testers also observed packet losses during communication with the constructive simulation node at MITRE (figure 15 and figure 16).

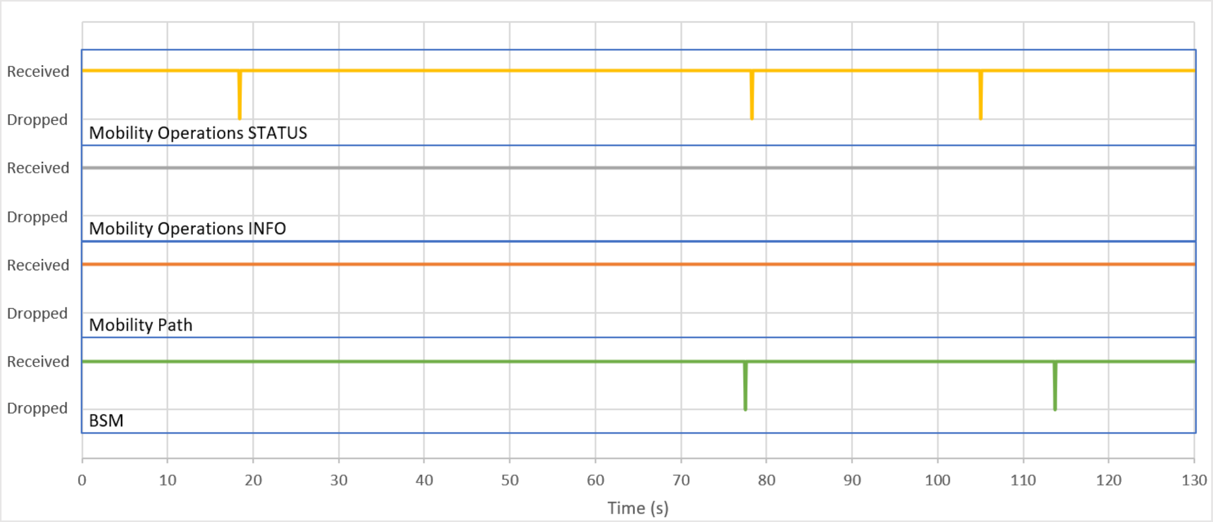

Source: FHWA.

Figure 14. Graph. Packet loss results over time for experiment 3, run 1 live to SRC.

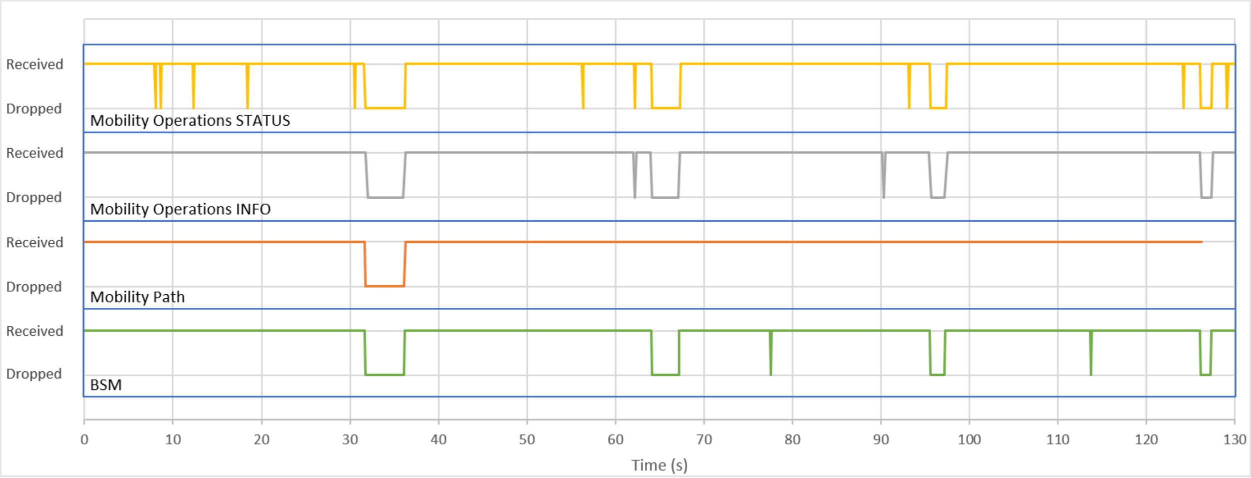

Source: FHWA.

Figure 15. Graph. Packet loss results over time for experiment 3, run 1 live to MITRE.

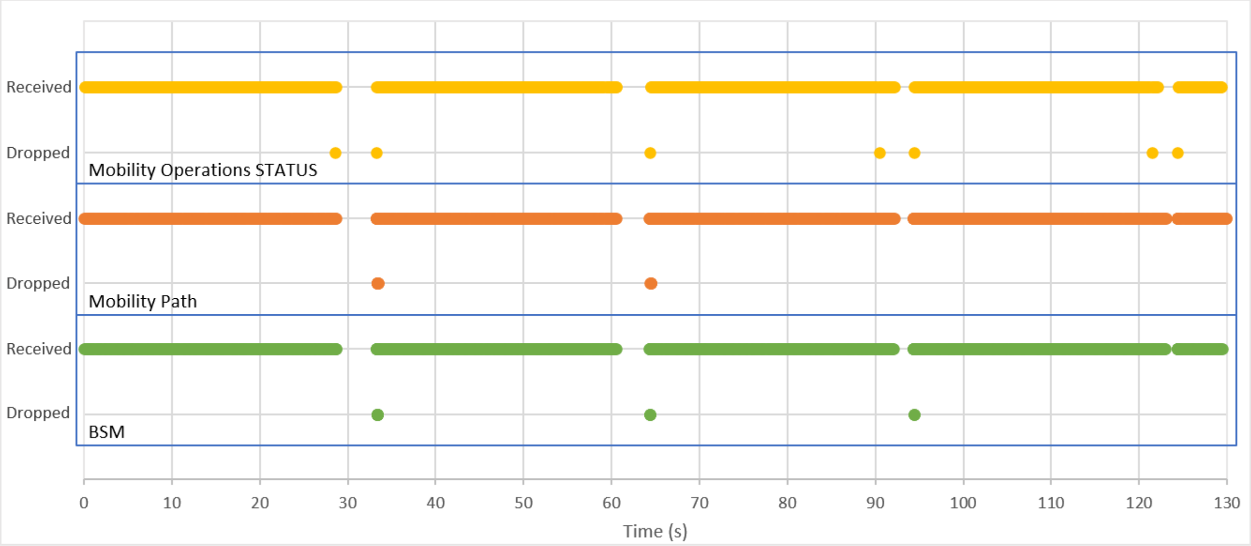

Source: FHWA.

Figure 16. Graph. Packet loss results over time for experiment 3, run 1 MITRE to SRC.

Latency

This section reports both the end-to-end and segment-by-segment latency analysis results.

End-to-end results are available for the CARMA mobility request/response messages based on the rosbag data for all eight runs across the three experiments.(20,25) Compared to the segment-to-segment analysis, this end-to-end latency analysis from rosbag data is more straightforward because each analysis required only rosbag data from a single vehicle. The results from this end-to-end analysis are hence more robust than segment-to-segment. Segment-by-segment analysis performed on the same data pipeline for experiment 3 runs was compared to the end-to-end results for the CARMA Mobility Request/Response Messages.(20) This comparison sheds light on the additional computation and communication latency introduced by the SIT-1 system.

Additional segment-by-segment results are included for CARMA MOMs, CARMA TCR and TCM messages, BSM, and SpaT messages.(20) The segment-by-segment analysis results in this section contain summary data for latency analysis from experiment 3, run 3. Similar data was collected for all other dataflows and message types and for experiment 3, run 1 and experiment 2, run 2. Summary data for each table includes average, minimum, and maximum latency for each step in the data flow, as well as jitter for data transmission steps and standard deviation for computational steps. Relevant data flow directions for each message type are highlighted in this section to adequately summarize the entire dataset.

Mobility Request and Mobility Response Messages

End-to-End Latency From rosbag Data(25)

Testers first analyzed the end-to-end latency of the initial handshake when two vehicles attempt to platoon using rosbag data from a follower.(25) For each test run, a rosbag containing internal CARMA Platform data was collected from each trailing constructive vehicle and used for analysis.(7) From these rosbags, calculations were made to measure the end-to-end time between a trailing constructive vehicle broadcasting a CARMA mobility request message to the live lead vehicle and to the same trailing constructive vehicle receiving a CARMA mobility response message from the live lead vehicle.(20) For each constructive vehicle joining the platoon in the SIT-1 experiment, this communication flow occurred twice. Because each test run involved two constructive vehicles, four end-to-end latency values were measured per test run. Analysis was performed for seven out of the eight test runs across the three sets of experiments. The only test run that was not analyzed using rosbag data was experiment 3, run 3 due to corrupted rosbag data. A total of 28 data points for the 7 test runs across the 3 sets of experiments were analyzed. table 10 reports the descriptive statistics for these 28 end-to-end latency measures related to the SIT-1 experiment and for a traditional field test with 22 data points.

Table 10. End-to-end latency statistics of the initial exchange of mobility request and mobility response messages.

Statistics (ms)	SIT-1 Experiment	Traditional Field Test
Number of data points	28	22
Latency average	129	21.7
Latency standard deviation	30.9	9.77
Latency minimum	57.4	9.40
Latency maximum	186	46.7

As a benchmark, table 10 also reports the end-to-end latency of the initial platooning handshake between two CARMA Platform vehicles in traditional field testing. In traditional field testing, the platooning use case relies purely on V2V communications with no infrastructure component.(7,18) CARMA Platform vehicles were equipped with DSRC radios during the traditional field test.(21) The traditional field test data shown in table 10 was from the CARMA Platform 4.2.0 release verification testing conducted in Auburndale, FL, in July 2022. The CARMA Platform 4.2.0 release verification testing involved two vehicles joining a front vehicle's platoon from the rear in the same lane, like the SIT-1 test objective. With regard to the

CARMA mobility request/response message exchange, CARMA Platform 4.2.0 (used in the traditional field tests reported) slightly differs from CARMA Platform 3.11.0 (used for SIT-1).(20) Minor differences exist in the CARMA Platform's logical checks when processing a received CARMA mobility response message from the platoon leader vehicle. The team does not believe that these minor modifications in the CARMA Platform logic would contribute significantly to the end-to-end latency differences between the traditional field test results and the SIT-1 test results. The end-to-end latency for the initial platooning handshake communication in SIT-1 (between a constructive and a live vehicle) is about 100 ms longer, on average, than that between two live vehicles.

Because a different live vehicle was used for SIT-1 experiment 1 than for experiments 2 and 3, and each experiment was conducted on a different day, testers compared the results from the three experiments. Table 11 reports the comparison statistics for these latency measures between the two experiments. These results led testers to conclude that the change in the live vehicle between experiment 1 and experiments 2 and 3 did not contribute significantly to the latency experienced in this end-to-end communication flow.

Table 11. End-to-end latency statistics of the initial exchange of mobility request and mobility response messages for SIT-1 experiments 1, 2, and 3.

Statistics (ms)	SIT-1 Experiment 1	SIT-1 Experiment 2	SIT-1 Experiment 3
Number of data points	12	8	8
Latency average	130	132	124
Latency standard deviation	33.1	28.9	32.8
Latency minimum	98.8	95.3	57.4
Latency maximum	186	180	169

Table 17 shows a histogram representation of the end-to-end latencies for these examined message exchanges from the three SIT-1 experiments and traditional field tests.

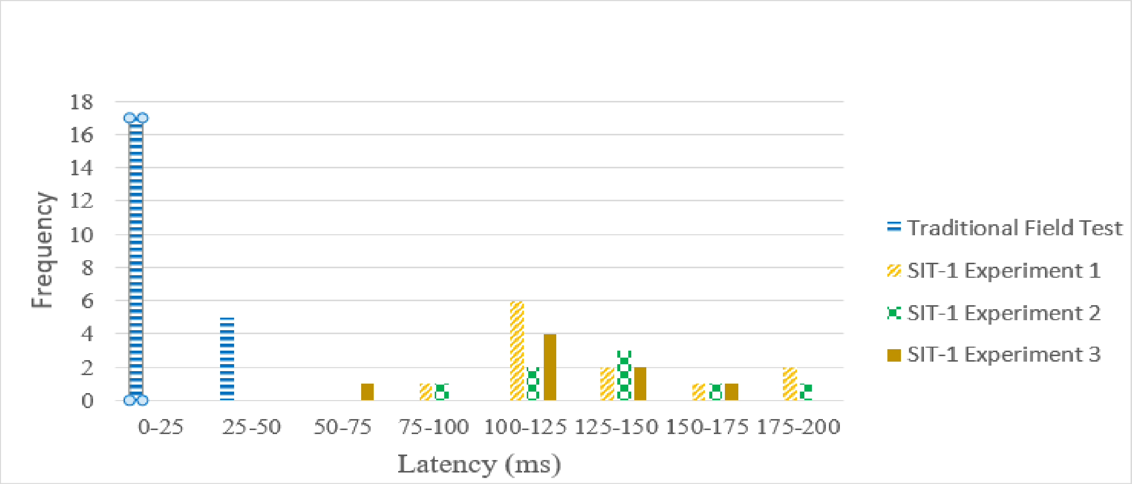

Source: FHWA.

Figure 17. Graph. End-to-end latency histogram of the initial exchange of mobility request and mobility response messages for SIT-1 experiments and traditional field tests.

Segment-by-Segment Latency Analysis

To understand the main sources contributing to the additional latency of about 100 ms, segment-by-segment analysis was performed to examine the latency (computation and communication) of each component along the data flow pipeline. The analysis follows the segment-by-segment data alignment process discussed in the Latency section of this chapter.

The following describes actions that occur at each step in the CARMA mobility request and response message data pipeline:(20)

> Step 1: The mobility request message sent by a constructive trailing vehicle is converted into a TENA platooning SDO by the TENA-CARMA-Platform adapter (figure 30) running on the same computer as the CARMA Platform instance associated with the constructive vehicle.(20,5,7)
>
> Step 2: The same TENA-CARMA-Platform adapter in step 1 (figure 30) broadcasts the TENA platooning SDO to the SIT-1 network.(5,7) The V2XHub-TENA-Mobility plugin (figure 28) receives the TENA platooning SDO through subscription.
>
> Step 3: The V2X Hub-TENA-Mobility plugin (figure 28) converts the TENA platooning SDO back into a CARMA mobility request message.(11,5,20)
>
> Step 4: The V2X Hub forwards the CARMA mobility request message to RSU, which broadcasts the message through DSRC radio.(11,20,21)

Step 5: The live lead vehicle receives the CARMA mobility request message originated from a constructive trailing vehicle, processes the message, and generates a CARMA mobility response message.(20)

Step 6: The live lead vehicle broadcasts the CARMA mobility response message through DSRC radio to the V2X Hub-TENA-Mobility plugin. (See references 20, 21, 11, and 5.)

Step 7: The V2X Hub-TENA-Mobility plugin (figure 28) converts the CARMA Mobility Response Message into a TENA platooning SDO.(11,5,20)

Step 8: The V2X Hub-TENA-Mobility plugin (figure 28) broadcasts the TENA platooning SDO to the SIT-1 network.(11,5) The same TENA-CARMA-Platform adapter used in step 1 (figure 30) receives the TEN Platooning SDO.

Step 9: The same TENA-CARMA-Platform adapter used in step 1 (figure 30) converts the TENA Platooning SDO back to a CARMA mobility response message and sends the message to the CARMA Platform instance associated with the constructive vehicle running on the same computer.(5,7,20)

Latencies of all the steps in the CARMA mobility request and response message data pipeline, except step 5, were calculated. Step 5 was not calculated because there was no instrumentation inside CARMA Platform to collect the data. IMPORTANT: The data input and output times of any adapters were approximate due to the limited data logging capabilities discussed in the Latency section of this chapter.(20)

Table 12 shows the mobility request latency data from the constructive vehicle in SRC (Augusta, GA) to the live vehicle at TFHRC (McLean, VA) table 12, and table 13 shows the associated mobility response latency data. For a single run, fewer than four requests and responses were sent between vehicles. This means that some statistics may not have been calculated and are potentially less precise due to the small sample size.

Table 12 and table 13 show that the transmission time of TENA Platooning SDOs (step 2 and step 8) are on par with the baseline latency values presented in the Latency section of this chapter.(5) The computational latencies of relevant adapters (steps 1, 3, 7, and 9) vary from sub-milliseconds to about 26 ms. Step 9's computation is likely inaccurate due to the approximation in data collection and assumption 1 discussed in in the Latency section of this chapter. Negative latency values are consistently observed in the last step of a data pipeline where a TENA SDO/TENA message is converted back to a J2735 test message.(19) This negative latency value issue is further discussed in the BSM latency analysis in which larger data sets are available to produce more statistically significant results.

Table 12. Mobility request latency results for experiment 3, run 3 SRC to live.(20)

Step	Description	Average (ms)	Minimum (ms)	Maximum (ms)
Step 1	SDO message creation from J2735 (SRC)(19)	13.4510994*	N/A	N/A
Step 2	SDO network transmission (from SRC to TFHRC)	2.360105515*	N/A	N/A
Step 3	J2735 message creation from SDO (V2XHub)(19,11)	11.62600517*	N/A	N/A
Step 4	DSRC message transmission (V2X Hub to live vehicle)(21,11)	6.773312887**	5.542039871**	7.543087006**

N/A = not applicable.
*only one data point.
**only three data points.

Table 13. Mobility response latency results for experiment 3, run 3 live to SRC.(20)

Step	Description	Average (ms)	Minimum (ms)	Maximum (ms)
Step 5	Mobility request processing and response generation (live vehicle)	—	—	—
Step 6	DSRC message transmission (live vehicle to V2X Hub)(21,11)	7.42226839*	3.48997116*	15.23089409*
Step 7	SDO creation (TFHRC V2X Hub computer) (11)	25.55251122**	20.43104172**	30.67398071*
Step 8	SDO network transmission (from TFHRC to SRC)	13.46004009**	10.68997383**	16.23010635**
Step 9	J2735 message creation from SDO (SRC)(19)	0.61953068*	−0.04506111*	1.28412247*

N/A = not applicable.
*only four data points.
**only two data points.

Results from experiment 3, run 1 are also show (table 14 and table 15) in addition to those from experiment 3, run 3. Experiment 3, run 1 shows similar trends.

Table 14. Mobility request latency results for experiment 3, run 1 SRC to live.(20)

Step	Description	Average (ms)	Minimum (ms)	Maximum (ms)
Step 1	SDO message creation from J2735 (SRC) (19)	10.82491874*	N/A	N/A
Step 2	SDO network transmission (from SRC to TFHRC)	3.50999832*	N/A	N/A
Step 3	J2735 message creation from SDO (V2X Hub)(19,11)	0.99396705*	N/A	N/A
Step 4	DSRC message transmission (V2X Hub to live vehicle)(21,11)	5.54760297**	3.99088859**	6.68597221**

N/A = not applicable; in this case the min and max is the same as the average
*only 1 data point
**only 3 data points

Table 15. Mobility response latency results for experiment 3 run, 1 live to SRC.(20)

Step	Description	Average (ms)	Minimum (ms)	Maximum (ms)
Step 5	Mobility request processing and response generation (live vehicle)	—	–	—
Step 6	DSRC message transmission (live vehicle to V2X Hub)(27,11)	5.63621520*	13.38005065*	6.41584396*
Step 7	SDO creation (TFHRC V2X Hub computer)(11)	42.86253452**	41.32604598**	44.39902305**
Step 8	SDO network transmission (from TFHRC to SRC)	14.06002044**	10.68997383**	13.46004009**
Step 9	J2735 message creation from SDO (SRC)(19)	−1.37495994*	−2.48599052*	−0.26392936*

*only 4 data points
**only 2 data points

Averaging the data between these two runs, the mobility request message took approximately 27.54 ms to go from the constructive vehicle at SRC (Augusta, GA) to the live lead vehicle at TFHRC (McLean, VA). The matching response took on average approximately 54.12 ms. This leads to an average total latency of 81.66 ms round-trip for these two runs in experiment 3.

Comparing to the end-to-end latency computed from rosbag data (table 16 and figure 17), the results from segment-by-segment analysis could be considered reasonable, despite the approximations employed in data collection.(25) Note that the segment-by-segment latency results do not include the computation time of the live vehicle processing the CARMA mobility request message (step 5), while the rosbag end-to-end latency analysis includes the computation time.(20)

BSM

BSMs were used in SIT-1 to create digital twins and were not used by any of the two CDA applications tested.(4) The following subsection presents segment-to-segment BSM latency analysis results for three representative data flows:

- From the live simulation node to the constructive simulation node at SRC (Augusta, GA).
- From the constructive simulation node at MITRE (McLean, VA) to the live simulation node and from the constructive simulation node at SRC (Augusta, GA) to the constructive simulation node at MITRE (McLean, VA).

From Live Simulation Node to Constructive Simulation Node at SRC (Augusta, GA)

The BSM dataset from Live to SRC shown in table 16 is one of the most complete and exemplary datasets collected during SIT-1. The sample size is large (*n* = 1257); data was translated completely from end to end; the dataset encompasses a live, infrastructure, and constructive element; and there are minimal dropped packets (table 7.).

Table 16. BSM latency results for experiment 3, run 3 live to SRC.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
DSRC message transmission (live vehicle to V2X Hub)(27,11)	5.006482	2.630949	18.825054	0.737428	N/A
SDO creation (TFHRC V2X Hub computer) (24,11)	21.966581	5.837202	72.128773	N/A	10.076960
SDO network transmission (from TFHRC to SRC)	13.814698	11.259079	66.103935	2.867804	N/A
J2735 message creation from SDO (SRC) (18)	-0.121645	-13.755083	6.36005	N/A	1.189280

From Constructive Simulation Node at MITRE to Live Simulation Node

Starting from the beginning of the message data pipeline to the receipt of the message, the DSRC message transmission latency (the time from when the live vehicle broadcasts the message via DSRC to when the V2X Hub receives it) averages 5 ms.(21,11) This value matches the expected duration for this data transmission.

SDO creation latency (time to convert a J2735 message to TENA SDO data) at the V2X Hub averaged approximately 22 ms, which is nearly 20 times higher than the time required for the SRC computer to complete the same task (table 18).(19,5,11) This increased computation time is caused by two main factors—1) the hardware for the V2X Hub computer is significantly less powerful than the SRC computer, and 2) messages must be received and decoded by the V2X Hub Message Receiver plugin before they are available to the V2X Hub-TENA-BSM plugin (figure 28)—caused the increased computation time. The individual effect of each of these factors could be further investigated in the future. Additionally, more powerful hardware could be used for the V2X Hub computer. Chapter 5 discusses more information on implications and lessons learned about system hardware.

Next, the SDO network transmission latency (one-way transmission latency of SDO data from TFHRC to SRC) from TFHRC to SRC averaged around 13.8 ms, which aligns with the baseline round-trip latency from TFHRC to SRC of 20 ms shown in table 4.

Finally, the J2735 message creation from SDO latency (time to convert SDO data to a J2735 message) fluctuated from −13.75 ms, the minimum latency, to 6.36 ms, the maximum latency, with an average −0.12 ms.(19) The Latency section of this chapter discusses how these negative latencies are likely due to assumption 1—using a TDCS instance on the same computer instead of actual application receipt data. Negative values could occur when the CARMA-Platform-TENA adapter received the SDO and generated a corresponding J2735 message before the adjacent TDCS received the SDO.(7,5) Negative values are plausible because of the imprecise data collection method and the speed at which the SRC computer was able to complete the conversion (J2735 to SDO conversion occurred in 0.75 ms. (table 18). Further, investigation the reason for negative latency will occur in the future.

Table 17 presents the BSM segment-by-segment latency results from the constructive vehicle at MITRE (McLean, VA) to the live vehicle at TFHRC (McLean, VA) and that the BSM data pipeline from a constructive vehicle to a live vehicle is not complete, which is the reason some cells indicate that there is no data. This incomplete BSM data pipeline, however, did not affect the two CDA applications tested in SIT-1.(4) While vehicles would receive BSMs from all other vehicles in a field-testing setup, BSMs are currently not used by either the CARMA implementation of platoon or work zone uses cases. The V2X Hub-TENA-BSM adapter is likely the broken link in the pipeline, and further investigation is needed to draw a definitive conclusion.(11,5) Table 17 also shows large maximum and average latency values for SDO message creation from J2735 BSMs on the MITRE laptop.(19) These values can be examined by plotting the SDO message creation latency over time.

Table 17. BSM latency results for experiment 3, run 3 MITRE to live.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard. Deviation (ms)
SDO message creation from J2735 (MITRE)(19)	250.450797	0.365019	6663.383007	N/A	947.239164
SDO network transmission (from MITRE to TFHRC)	3.265361	2.267122	25.866985	0.688149	N/A
J2735 message creation from SDO (V2X Hub)(19)	—	—	—	—	—
DSRC message transmission (V2X Hub to live vehicle)(21,11)	—	—	—	—	—

N/A = not applicable; —No data.

Figure 18 shows gaps in the SDO creation latencies at around the 30- and 110-s marks. These gaps align exactly with the packet loss presented in figure 12.

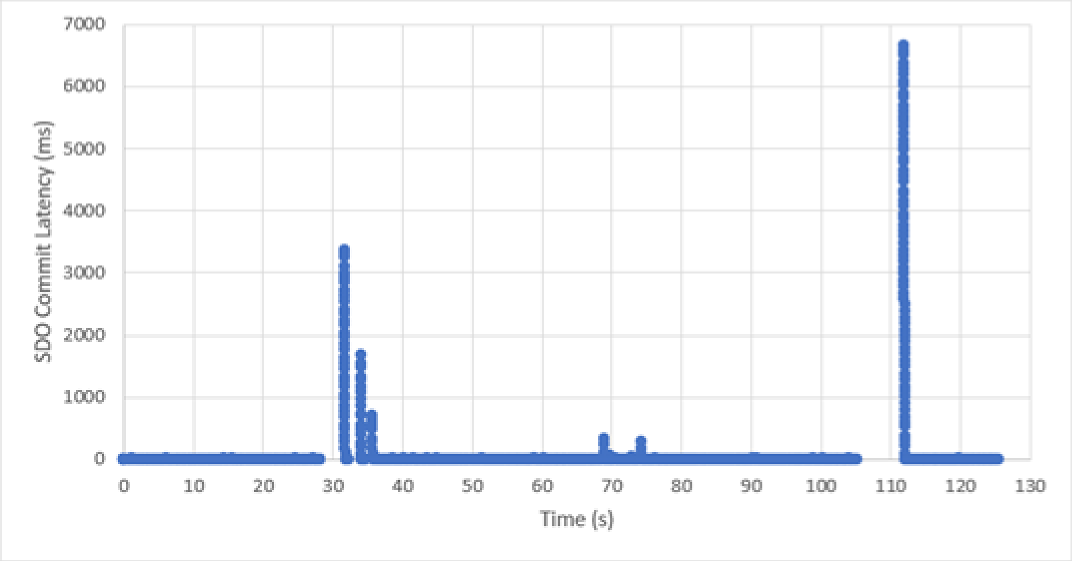

Source: FHWA.

Figure 18. Graph. MITRE BSM SDO creation latency over time for experiment 3, run 3.

Figure 19 shows similar behavior during experiment 3, run 1 that aligns with figure 15.

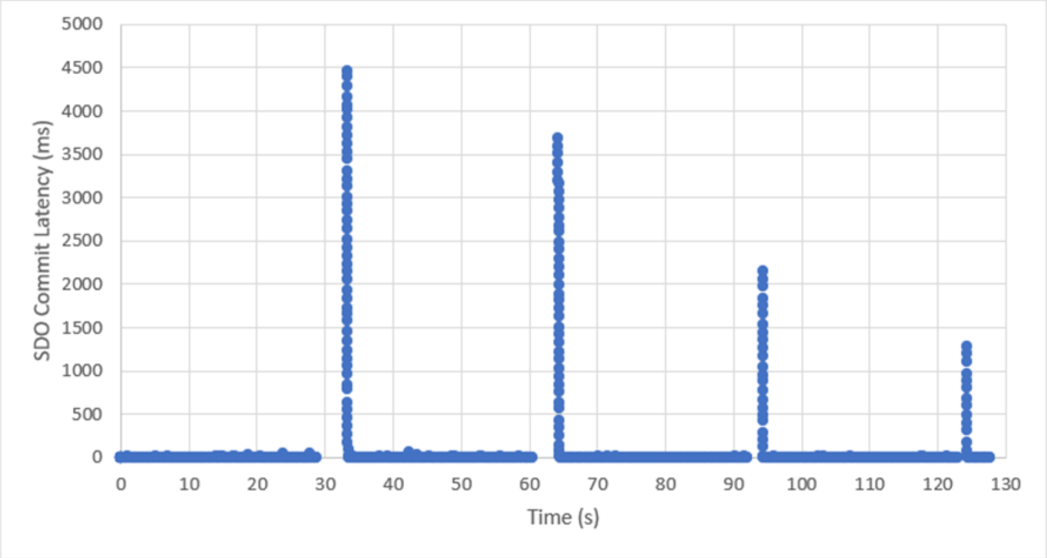

Source: FHWA.

Figure 19. Graph. MITRE BSM SDO creation latency over time for experiment 3, run 1.

From Constructive Simulation Node at SRC to Constructive Simulation Node at MITRE

Table 18 presents the BSM segment-by-segment latency results from the constructive vehicle at SRC (August, GA) to the constructive vehicle at MITRE (McLean, VA).

The BSM SDO creation from J2735 at SRC took 0.75 ms on average.(19) Compared to the 21.97 ms on average at the V2X Hub computer and the 250.45 ms on average at the MITRE constructive node, the low latency observed at SRC indicates that the efficiency of the relevant adapters is the likely cause. The discrepancies in the computation latencies of the same conversion observed at different simulation nodes further highlight the impacts of hardware selection on latency times.

The last step in this BSM pipeline from one constructive vehicle to another constructive vehicle shows a negative minimum latency value, consistent with observations shown in table 13, table 15, and table 16. Similarly, the team believes that the negative latency observations are caused by the approximate values adopted in data collection (more specifically for this case, assumption 1 discussed in the Methodology section of chapter 4).

Moreover, the last step in this BSM pipeline also shows a very large maximum latency value (1,671 ms). After further investigation of the J2735 message creation from SDO at the constructive simulation node at MITRE, testers determined that this large value was a sole outlier in the dataset.(19) The average of 3.46 ms shows the latency was consistently low, but computationally strained hardware may make a system prone to inconsistent performance, resulting in a latency value anomaly.

Table 18. BSM latency results for experiment 3, run 3 SRC to MITRE.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
SDO message creation from J2735 (SRC)(19)	0.753552	0.185013	16.293049	N/A	1.546494
SDO network transmission (from SRC to MITRE)	10.411779	9.355068	33.188820	1.157521	N/A
J2735 message creation from SDO (MITRE)(19)	3.463692	−8.454084	1671.032906	N/A	67.930768

N/A = not applicable.

SPaT

SPaT messages were used in SIT-1 to update the digital twin of the physical traffic signal in the simulated world. SPaT was not used by either of the two CDA applications tested in SIT-1.(4)

SPaT messages originated from the physical traffic signal controller as NTCIP messages. They are converted to J2735 SPaT messages by the V2X Hub SPaT plugin (figure 28).(19,11) Then the J2735 SPaT messages enter the VOICES SPaT data pipeline: first being converted to a TENA Traffic Light SDO, then transmitted through the SIT-1 network, and finally consumed by the CARLA-TENA adapter (figure 30) for updating the digital twin of the live traffic signal in CARLA.(5,12) TrafficLight SDOs were not converted from SDO to J2735 messages for constructive vehicles because that functionality was not required for the CDA applications tested in SIT-1.(4)

Note that computation latency results of SDO creation shown in table 19 are likely inaccurate due to the limitations and approximations adopted in data collection. (chapter 4, Methodology). In this case, assumption 3 is the likely contributor to the negative latency observations. The – in this and following tables represents missing data. Data is not present for those instances due to error, data corruption or missing data.

The timestamp in an NTCIP SPaT message would have been a better approximation of when the V2X Hub-TENA-SPaT plugin received a J2735 SPaT message.(11,5,19) However, STOL and the project teams currently do not have an efficient way to decode the NTCIP SPaT messages for analysis. Moreover, the computation latency of SDO creation varies widely, indicating the need for future investigation. The approximation adopted in data collection, the unstable performance of the V2X Hub computer, and/or the V2X Hub software could contribute to this wide range of latency values.(11)

Table 19. SPaT latency results for experiment 3, run 3 live simulation node to constructive simulation node at SRC (Augusta, GA).

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
NTCIP message transmission (TSC to V2X Hub)(11)	—	—	—	—	—
SDO creation (TFHRC V2X Hub computer)(11)	−1.459694	−67.244053	123.282194	N/A	14.474051
SDO network transmission (from TFHRC to SRC)	14.792430	10.411978	78.141928	4.147593	N/A

—No data; N/A = not applicable.

Mobility Operations Messages

CARMA MOM has two subtypes: CARMA mobility operations INFO message and CARMA mobility operations STATUS message.(20) CARMA vehicles broadcast only mobility operations INFO messages when they are searching for a platoon. Once the initial handshake to form a platoon completes (after the successful exchange of the two sets of CARMA mobility request/response messages), only the lead vehicle broadcasts the CARMA MOM INFO message; the other platoon vehicles broadcast only CARMA MOM STATUS messages. After the initial handshake to form a platoon completes, the vehicle following the lead vehicle tries to close the gap between it and the lead vehicle to maintain the desired gap. During this process, the vehicles continuously broadcast their statuses via CARMA mobility operations STATUS messages so platoon members can adjust their desired/actual gap and vehicle control.

The broadcast rate for CARMA MOMs (both INFO and STATUS subtypes) is 10 Hz, so each site had sufficient data points (over 1,000). The segment-by-segment latency results for CARMA mobility operations were as expected, very similar to the BSM latency results.(20)

Consistent with results observed in table 13, table 15, table 16, and table 18, negative latency values were seen during the J2735 creation from SDO in the CARMA MOM data pipeline (table 20 and n/a = not applicable.

table 21).(19,20) Again, negative latency was likely caused by assumption 1 discussed in chapter 4, Methodology section.

The computation latency of SDO creation from J2735 was higher and varied more widely on the V2X Hub computer (table 20) compared to latency computations on the SRC computer (n/a = not applicable.

table 21).(19,11) That comparison is consistent with previous observations and likely is due to hardware performance.

SDO transmission times across the SIT-1 network were on par with the baseline performances as seen in table 4.

Table 20. Mobility operations INFO latency results for experiment 3, run 3 live to SRC.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
DSRC message transmission (live vehicle to V2X Hub)(21,11)	6.064314	3.373861	26.346922	1.039802	N/A
SDO creation (TFHRC V2X Hub computer)(11)	25.552511	7.747173	111.013889	N/A	16.261043
SDO network transmission (from TFHRC to SRC)	13.198962	10.689974	16.230106	3.189905	N/A
J2735 message creation from SDO (SRC)(19)	−0.542413	−33.752918	6.509781	N/A	2.338668

N/A = not applicable.

Table 21. Mobility operations STATUS latency results for experiment 3, run 3 SRC to live.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
SDO message creation from J2735 (SRC)(19)	16.713120	11.240959	63.117981	N/A	6.409689
SDO network transmission (from SRC to TFHRC)	3.242215	1.339912	29.070139	2.346168	N/A
J2735 message creation from SDO (V2X Hub)(19,11)	5.825394	−28.946877	96.174002	N/A	15.261967
DSRC message transmission (V2X Hub to live vehicle)(21,11)	6.801305	3.164768	41.008949	2.151647	N/A

N/A = not applicable.

TCR and TCM

The work zone use case includes use of CARMA TCRs and TCMs. These two messages were traced together from the request source vehicle as J2735 TCR, to V2X Hub as TENA TCR Message, out of V2X Hub as TENA TCM Message, and back to the source vehicle as J2735 TCM.(19,11,5)

This subsection presents segment-by-segment analysis results of the CARMA TCR/TCM data flow for the live vehicle (table 22) and the constructive vehicle at SRC (n/a = not applicable.

table 23). The CARMA-Platform-TENA adapter run at the constructive node at MITRE (figure 30) did not produce any TENA TCR messages; therefore, no analysis was conducted.(7) Also, TCMs were not converted from TENA messages to J2735 by the CARMA-Platform-TENA adapter instance run at the constructive simulation node at SRC (as seen in the last row of n/a = not applicable.

table 23).(19) While the functionality did not operate as designed during VOICE SIT-1, it did not affect the work zone application. Because all vehicles were in a platoon, the constructive following vehicles slowed down in response to the lead vehicle's speed change upon entering the work zone. It is unfortunate that thorough testing of the CARMA-Platform-TENA adapter was not possible due to time constraints, but future work is planned to perform this testing and resolve any issues.

A direct comparison can be made between the CARMA TCR/TCM data flow of the live vehicle, which did not involve any TENA components (table 22) and the CARMA TCR/TCM data flow of the constructive vehicle at SRC, which used the VOICES TENA adapters (n/a = not applicable.

table 23).(3,5) The total TCR-to-TCM time for the live vehicle was approximately 135.83 ms, while the constructive vehicle at SRC averaged 127.96 ms. The total TCR/TCM latency for the constructive vehicle at SRC did not include the J2735 message creation from SDO, but data from similar conversions (e.g., the TCR SDO message creation from J2735) suggest this value should be less than 5 ms.(19) Therefore, using the VOICES system for the exchange of TCR and TCM does not add any significant latency to messaging and may even be faster than the speed a live vehicle could experience.

Table 22. TCR/TCM round trip latency results for Experiment 3 Run 3, live vehicle.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Standard Deviation (ms)
TCR DSRC message transmission (live vehicle to V2X Hub)(21,11)	5.046341	3.261089	7.892132	0.887913	N/A
J2735 TCR receipt to J2735 TCM sent	124.649339	64.167976	179.322004	N/A	32.323706

(TFHRC V2X Hub computer)(19,11)					
TCM DSRC message transmission (V2X Hub to live vehicle)(21,11)	6.136378	3.240108	11.920929	2.581330	N/A

N/A = not applicable.

Table 23. TCR/TCM round trip latency results for Experiment 3, Run 3 SRC.

Message Flow	Average (ms)	Minimum (ms)	Maximum (ms)	Jitter (ms)	Std Dev (ms)
TENA TCR message creation from J2735 (MITRE)(5,19)	2.210075	0.165939	5.302191	N/A	2.339234
TENA TCR message network transmission (from SRC to TFHRC)(5)	9.876208	9.660959	10.138273	0.168705	N/A
TENA TCM message creation from TCR SDO receipt (V2XHub)(5,11)	106.595906	64.719915	152.184010	N/A	36.137840
TENA TCM message network transmission (from TFHRC to SRC)(5)	9.277972	8.720875	9.763956	0.386882	—
J2735 TCM creation from SDO (SRC)(19)	—	—	—	—	—

—No data; N/A = not applicable.

CDA APPLICATION PERFORMANCE

Given the additional latency introduced by computations required (conversions by VOICES adapters) and communication across longer distances through the SIT-1 network, this project should investigate how the added latency may affect CDA.(4) Additionally, the simulated physics in CARLA is expected to differ from actual physics and may affect CARMA Platform's performances on vehicle control, which are individually tuned for different vehicles at different testing sites.(12,7) This project uses the following performance metrics were analyzed:

- Speed (milliseconds) of each individual vehicle.
- Vehicle response times (seconds) to commanded vehicle control inputs.
- Space headway (meters) between two consecutive vehicles in the platoon.

To enable CDA application performance analysis for specific vehicles, a rosbag was logged from each vehicle for each SIT-1 test run.(4,25) Each rosbag contained all published data from a vehicle's CARMA Platform for a corresponding test run. An analysis script was developed to parse each rosbag for a given test run to analyze and plot relevant results (vehicle speed, platoon follower actual space headway versus desired space headway, etc.). To process the data contained within each rosbag, this project used the Python rosbag library.

Of note is that rosbag data from all three sets of SIT-1 experiments were available and valid for analyzing the CDA application performances.(4,25) The segment-by-segment latency and packet loss analysis were only experiment 3 results were produced due to a clock synchronization issue in earlier experiments.

Vehicle Speed Profiles

Since the SIT-1 test was the first type of test that STOL has conducted that includes both live and constructive vehicles, the team is interested in characterizing and comparing the ability of each vehicle to follow the speed values commanded by the CARMA Platform.(7)

Because two different live vehicles were used across the three sets of experiments, figure 20 and figure 21 were produced to show the commanded and actual speed profiles for both live vehicles. These plots clearly show that the 2019 Lexus RX 450h performed better than the 2019 Chrysler Pacifica in following the commanded acceleration at the start of the test run as well as the commanded deceleration at the entrance of the work zone. The 2019 Chrysler Pacifica experienced greater delay in reaching the commanded speed when not travelling at steady state. This difference is likely due to the different low-level vehicle controllers installed in the two vehicles. The live vehicles' actual speed profiles were smooth when travelling with a near-constant commanded speed.

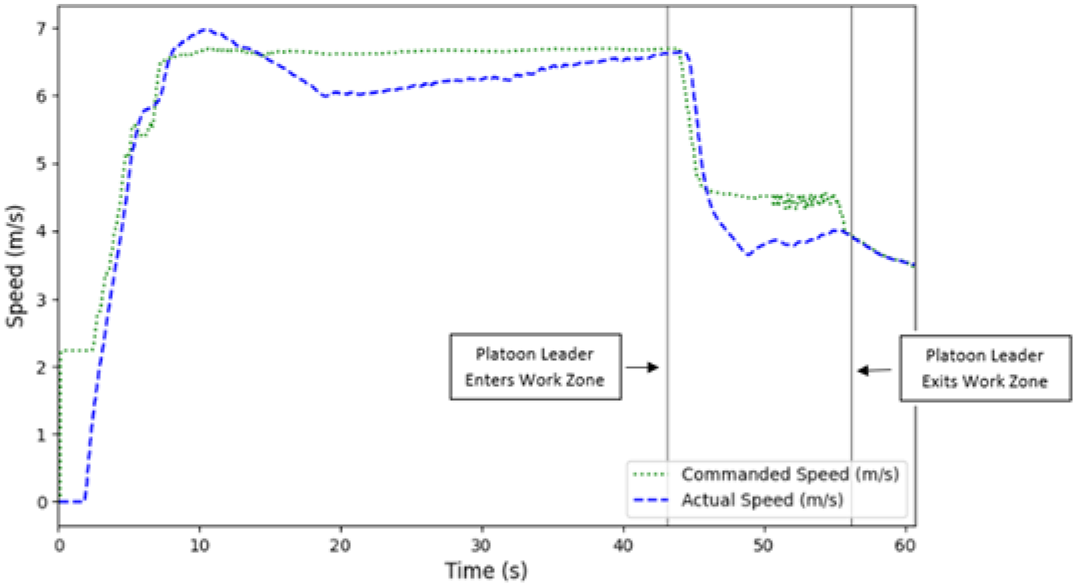

Source: FHWA.

Figure 20. Graph. Commanded and actual speed for experiment 1, run 3, vehicle 1.

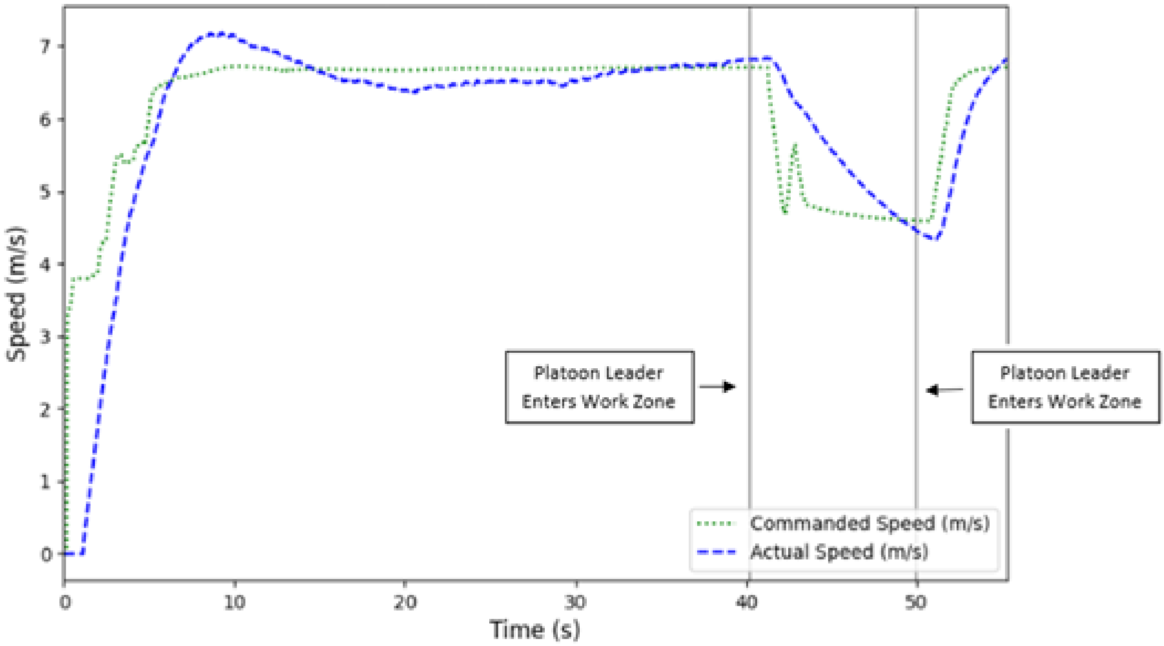

Source: FHWA.

Figure 21. Graph. Commanded and actual speed for experiment 2, run 3, vehicle 3 (2019 Chrysler Pacifica).

Figure 22 and figure 23 display both the commanded and actual speed profiles for vehicle 2 (constructive vehicle at SRC) and vehicle 3 (constructive vehicle at MITRE) in a typical SIT-1 test run. In both figures, the x-axis represents time since vehicle 1 engaged in automated driving. The vehicle speed plots for vehicle 2 and vehicle 3 indicate some acceleration and deceleration lags during the SIT-1 test, as well as some oscillatory behavior during steady state (approximately from $t = 25$ to $t = 40$ s). The lag in acceleration and deceleration could be due to the additional latency associated with the distributed LVC test, as well as to the simulation physics in CARLA.(6,12) Additionally, the oscillatory speed behavior during steady-state instances at constant speed could be due to limited simulated vehicle dynamics in the CARLA simulation platform.

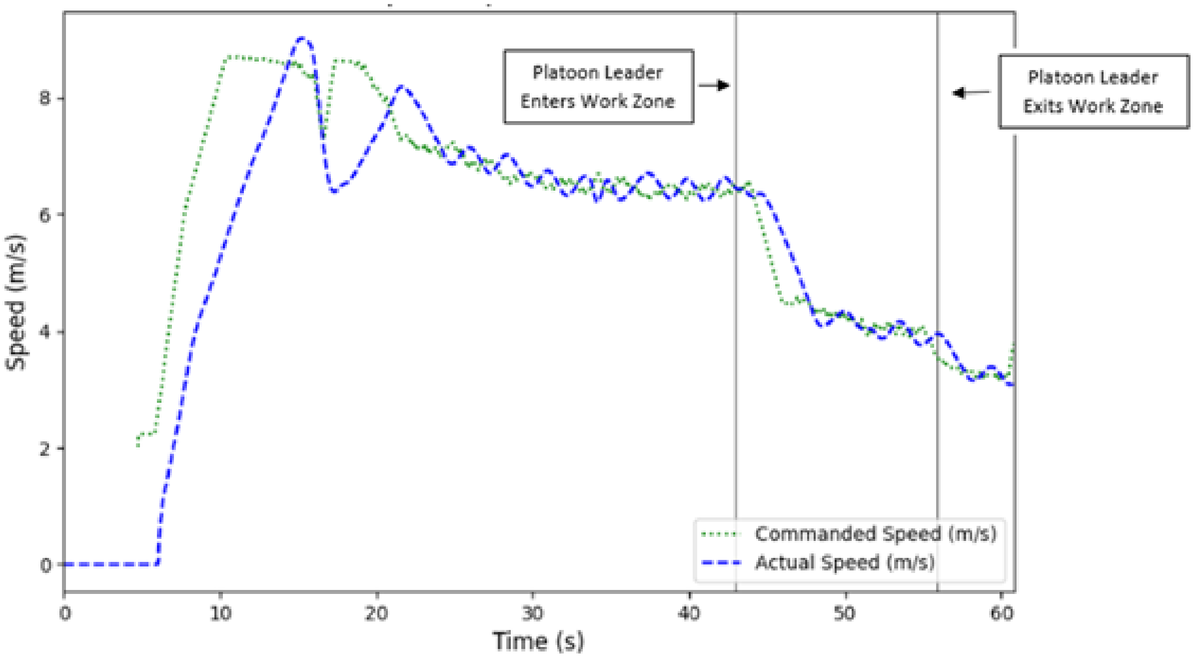

Source: FHWA.

Figure 22. Graph. Commanded and actual speed for experiment 1, run 3, vehicle 2.

Although there were differences in the speed performance for the constructive vehicles and both live vehicles, each vehicle was able to platoon effectively using CARMA platform.(7) Figure 24 displays the speed profiles of all three vehicles for a typical SIT-1 test run. This plot shows that both vehicle 2 and vehicle 3 were able to increase their speeds above the speed of their preceding platoon member to achieve their desired space headway with that member. Moreover, vehicle 2 was able to lower its speed to match the speed of vehicle 1 once the desired space headway was achieved. On the other hand, the speed of vehicle 3 was greater than the speeds of vehicle 1 and vehicle 2 until approximately $t = 50$ s, when vehicle 3 finally achieved the desired space headway (figure 27). This is likely due to the delayed start of vehicle 3 and the subsequent short amount of time spent at steady state.

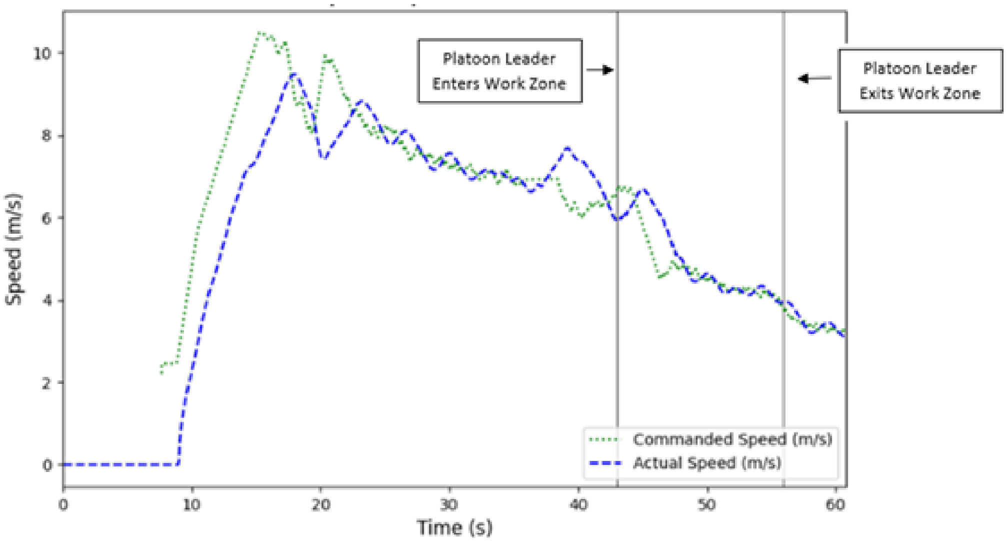

Source: FHWA.

Figure 23. Graph. Commanded and actual speed for experiment 1, run 3, vehicle 3.

Figure 24 is a typical example of the desired versus actual space headway between vehicle 2 (constructive vehicle at SRC) and vehicle 1 (live vehicle) in experiment 1. In the figure, the x-axis represents time since vehicle 2 joined the platoon. For all SIT-1 test runs, each follower vehicle officially joined the platoon (completed the initial handshake by exchanging CARMA mobility request response messages) within 1 s of engaging in automated driving.(20) For about the first 20 s that vehicle 2 engaged in automated driving, the vehicle was trying to close its gap with the lead vehicle. To determine the time at which a vehicle reached platooning steady state, the test team decided that time would represent the time the follower vehicle's actual space headway became smaller than its desired space headway. The team also decided on this approach for defining the start of platooning steady state because of the complexity involved in creating a more robust definition and in developing an automated script to identify it. Additionally, this simpler approach was sufficient to enable the team to identify patterns and make comparisons across experiments. At around $t = 22$ s, vehicle 2 was considered to have achieved steady state, which lasted until around $t = 38$ s, at which time the lead vehicle entered the work zone and started to reduce its speed. The desired space headway fluctuated in steady state because of the fluctuation in the ego vehicle speed. Ego speed is used to calculate the desired space headway based on a desired constant time headway. More specifically, the CARMA Platform 3.11.0 platoon logic calculates the desired space headway as the ego vehicle's speed multiplied by the predetermined desired time headway.(7) The desired constant time headway (configurable parameter) was set as 2.5 s for all CARMA Platform instances during SIT-1.

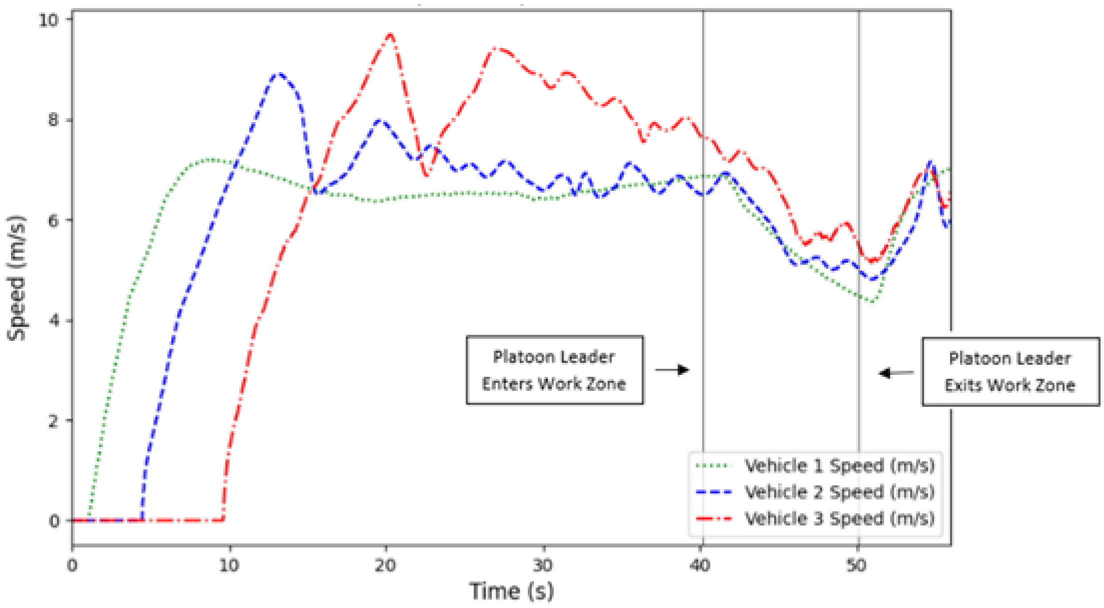

Source: FHWA.

Figure 24. Graph. Vehicle speeds for experiment 2, run 2.

Space Headway

Figure 25 shows similar results for a typical run-in experiment 2. One section that differs significantly between figure 24 and figure 25 is the section spanning from approximately *t* = 40 s to *t* = 50 s. During this time span, the desired space headway dropped sharply in both experiments due to the following vehicle slowing down as the platoon lead vehicle entered the work zone. The decrease in vehicle 2's speed resulted in a decrease in its desired space headway. Although a decrease in ego vehicle speed causes a decrease in desired space headway, the desired space headway cannot decrease below 12 m, which was the configuration parameter setting for all SIT-1 test runs.

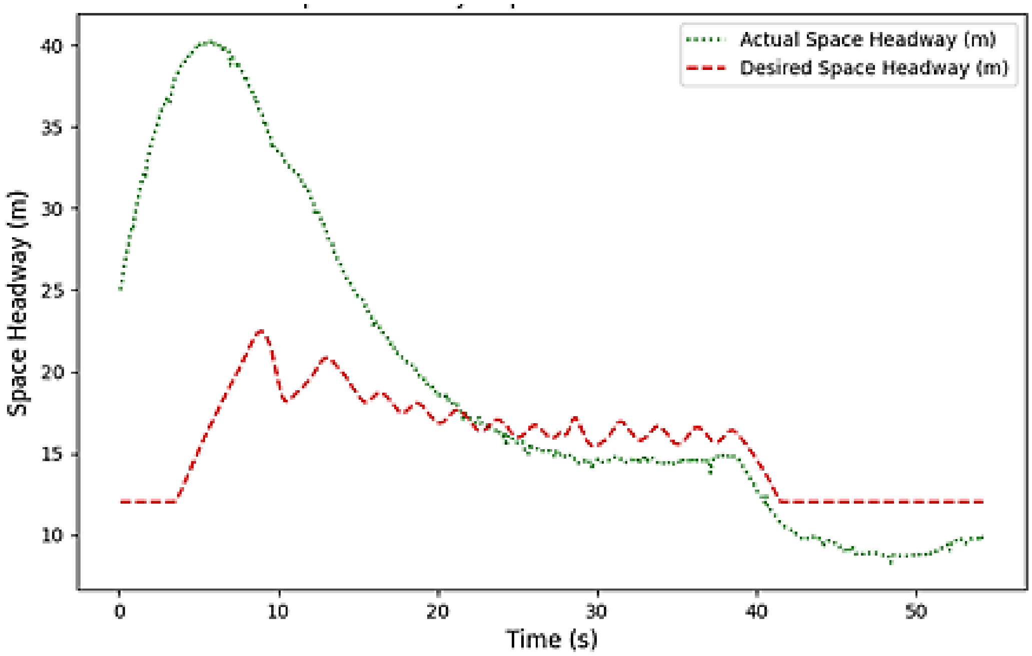

Source: FHWA.

Figure 25. Graph. Example space headway: experiment 1, run 2, vehicle 2.

In experiment 1 (figure 25), the actual space headway remained below the desired space headway after the lead vehicle entered the work zone. The behavior in experiment 1 was due to the live vehicle having a larger deceleration rate than the constructive vehicle, causing the actual space headway to decrease rapidly during the deceleration phase. On the other hand, in experiment 2, the live vehicle had a deceleration rate close to that of the constructive vehicle, so the actual space headway did not decrease during the deceleration phase. If the work zone segment had been longer, testers expected vehicle 2's speed would have eventually decreased and achieved its desired space headway with the lead vehicle.

Testers analyzed the difference between the desired and the actual space headways between vehicle 2 and vehicle 1 during steady state (the first instance where actual space headway is less than the desired value to when the lead vehicle enters the work zone) performed, and table 24 shows the summary statistics provided in table 24 of that analysis. In general, vehicle 2 maintained a space headway within three meters of its desired space headway throughout steady-state platooning operations across all test runs. On average, the difference between the desired and the actual space headways across all runs for experiment 1 trended negative, while the same value trended slightly positive for experiment 2 runs, as corroborated by the example headway plots shown in figure 22 and figure 23. Both experiment 3 test runs recorded slightly better average results.

Throughout each SIT-1 test run across all three sets of experiments, vehicle 3 consistently controlled its desired headway and control outputs based on the position and speed information received from vehicle 2. From all test runs, vehicle 3 was unable to reach its desired steady-state space headway consistently due to its late start and the short length of the designated route at the test facility for SIT-1 test runs. Figure 26 shows a typical example of the desired and actual space headway between vehicle 3 and vehicle 2. Figure 26 also indicates that vehicle 3 eventually achieved its desired headway space with vehicle 2 within the final seconds of the test run. As a result, no steady-state performance data is available for vehicle 3.

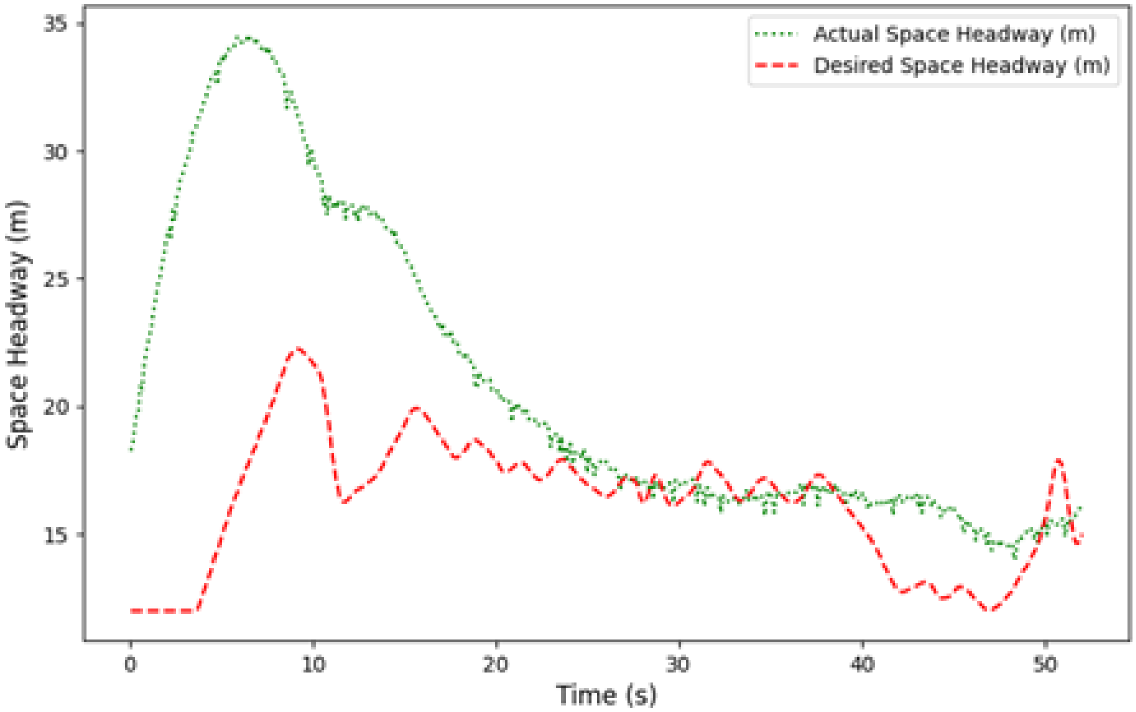

Source: FHWA.

Figure 26. Graph. Example space headway: experiment 2, run 2 vehicle 2.

Table 24. Average difference between desired and actual space headway during steady state: vehicle 2.

Experiment	Run No.	Desired Space Headway – Actual Space Headway (meters)			
		Average	Standard Deviation	Minimum	Maximum
1	1	−0.883	0.493	−1.97	0.240
	2	−1.19	0.766	−2.55	0.802
	3	−1.29	0.714	−2.72	0.191
2	1	0.271	0.458	−0.564	1.60
	2	0.0941	0.708	−1.66	1.44
3	1	−0.00870	0.662	−1.31	1.37
	2	0.0259	0.742	−1.43	1.67

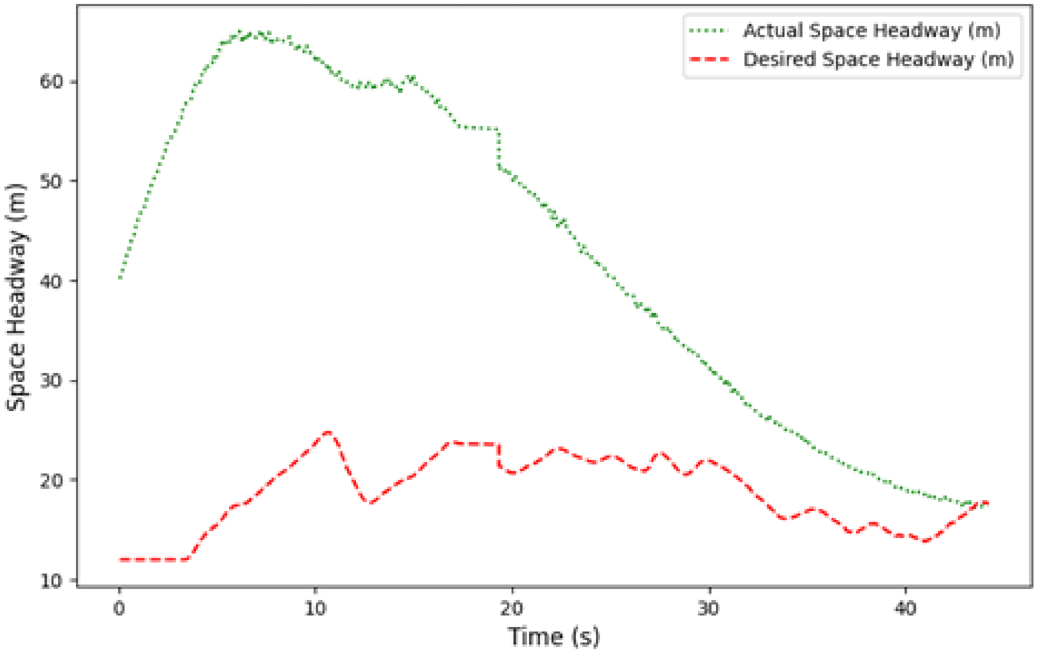

Source: FHWA.

Figure 27. Graph. Example space headway: experiment 2, run 1 vehicle 3.

CHAPTER 5. DISCUSSION AND CONCLUSION

LESSONS LEARNED

Logging Capabilities of VOICES Adapters

As discussed in chapter 4, Network Performance and Results under Load, the VOICES adapters used for SIT-1 at each simulation node, were designed to be lightweight and currently do not have data logging features that allow for traceability between data inputs and outputs. As a result, exact timestamps of when relevant data entered and left specific adapters were not available. This limited logging capability proved to be challenging for the network performance analysis under load. Three assumptions (chapter 4) were employed to approximate relevant timestamps. Two of the three assumptions are not ideal (albeit the best approximations the team was able to obtain), as shown in various tables throughout chapter 4. More specifically, the performance of adapters in creating J2735 messages from relevant TENA SDOs is largely still unknown, due to the limited logging capabilities of adapters.(19,5) In the future, logging capabilities could be developed for adapters to enable data traceability to better understand the adapters' performances. Including in the development a logging feature that could be toggled on or off depending on testing and research needs would be helpful.

Hardware Requirements

During the data analysis, the importance of computational hardware specifications became very clear. Specifically, the V2X Hub computer and the older MITRE laptop both struggled to process and distribute data underload.(11) Exact hardware specification used for SIT-1 are addressed in chapter 2. No minimum hardware computational requirements were identified prior to SIT-1 participation, but it became apparent during testing that, in the future, it would be useful to establish and refine understanding of such criteria for each node.

Although considered powerful upon its release in 2016, the older MITRE laptop appeared unable to keep up with all processing demands multiple times during most test runs (2–5 min/run), which resulted in dropped packets and gaps in data transmission. The hypothesis is that the combination of CARMA Platform, the CARLA simulator, the CARMA-CARLA integration tool, and multiple TENA adapters accumulated in processing requirements that occasionally exceeded the test laptop's capabilities. (See references 7, 12, 15, and 5.) In the future, different hardware configurations could be tested to define the minimum hardware requirements for each node type.

The V2X Hub computer was running much less demanding software (V2X Hub and four adapters).(11) While it did not appear to be computationally overloaded, the TENA adapters running on it consistently performed up to 20 times slower converting J2735 to TENA SDO and vice versa than did the SRC computer running similar operations (table 16).(19,5) The computer used for V2X Hub was an older edge computer with a processor released in 2013. Identical machines at STOL have proved sufficient to run previous applications of V2X Hub for years; but, given V2X Hub computer's underperformance compared to other computers used during SIT-1, future distributed tests using V2X Hub may warrant a more powerful edge computer to support the computational needs of a live simulation node.

The SRC computer (procured specifically for the VOICES project) did not present any observable computational or data communication issues. Similarly, the computer used for the virtual simulation node at TFHRC did not present any performance issues. These devices' specifications, therefore, can be used as baseline examples of hardware specifications for future constructive and virtual simulation nodes, and should be documented and validated during study preparations.

Software Compatibility

The project team encountered a software compatibility issue during preparation of VOICE SIT-1. More specifically, there was a mismatch of the ASN.1 library files use by CARMA Platform 3.11.0 and the CARMA Cloud master branch as of April 2022.(7,10) The ASN.1 file contains the structure and definitions of each J2735 message, including custom defined J2735 test messages.(19) The CARMA ecosystem software used custom J2735 test messages to activate prototype CDA applications.(3,4) Due to the rapid prototyping nature of the J2735 test messages, relevant J2735 test messages used by CARMA software are constantly being modified and refined. Since CARMA Platform 3.11.0 is an older version of the software, the ASN.1 file that was packaged with it contained older definitions of CARMA TCR/TCM messages, compared to the ASN.1 file that was packaged with the CARMA Cloud master branch as of April 2022. The team resolved the mismatch by creating a SIT-1 ASN.1 file that reconciled the differences between the two ASN.1 files packaged with the version of the two CARMA software versions.

As mentioned in the software configuration management section of chapter 2, the CARMA Platform 3.11.0 was used instead of the most recent CARMA Platform release because the 3.11.0 version is the most up-to-date version the CARMA-CARLA integration tool (as part of the CARMA XiL cosimulation tool) currently supports.(7,15,3) The CARMA XiL project is currently upgrading the CARMA XiL cosimulation tool to support ROS2, which will allow CARMA Platform and above to be integrated into the cosimulation tool. This upgrade will be able to eliminate the ASN.1 file mismatch issue encountered during preparation of SIT-1. In the future, definitions of custom J2735 test messages should be clearly tracked as they evolve.(19) Relevant ASN.1 files could also be provided to all VOICES participants joining the same test.

Differences in Controlling Live and Constructive Vehicles

To prepare for SIT-1 test execution using a live vehicle at TFHRC and two off-site constructive vehicles, the team first began testing with a live vehicle and two on-site constructive vehicles. During this initial testing phase, the team sought to identify both CARMA Platform software implementation updates and software tuning updates required for the constructive vehicles to platoon smoothly and successfully with the live vehicle.(7)

During this initial testing phase, the project team discovered that the constructive vehicles did not accurately follow the control commands output from CARMA Platform while platooning.(7) The team's investigation found that the CARLA simulation requires vehicle acceleration command inputs to simulate physics and vehicle dynamics more accurately than if it did not receive vehicle acceleration inputs.(12) The CARMA Platform platooning control plugin—which is responsible for outputting speed and steering commands during platooning—did not include an outputted acceleration command, as it was not directly required from CARMA Platform (but instead handled by a lower-level vehicle controller) to control a physical vehicle. A software update was made within the CARMA Platform to add a direct acceleration command to the output of the platooning control plugin.

In addition to the software implementation update required to enable constructive vehicles to closely follow the outputted control commands from CARMA Platform in simulation, there were also several configuration parameter updates that were required to improve constructive vehicle platooning performance.(7) In general, these configuration parameter updates were required because the original parameters were tuned for live vehicles, and the team discovered that the original parameters caused constructive vehicles to accelerate and decelerate at higher rates than desired. For both the CARMA Platform platooning trajectory-generation plugin and the platooning control plugin, maximum acceleration and deceleration limits were reduced from 1.5 m/s^2 to 1.0 m/s^2. The proportional gain for the CARMA Platform platooning proportional-integral-derivative controller, which outputs a commanded speed for a platoon follower to maintain a configured constant time headway with a platoon leader, was updated as well, with the final value reduced from 0.4 to 0.15.

CONCLUSION AND NEXT STEPS

VOICES SIT-1 included four LVC simulation nodes hosted at three geographically distributed sites.(6) SIT-1 featured two prototype CDA applications: work zone and platooning, both implemented using FHWA's CARMA suite of CDA tools.(4) In all the test runs except one across all three sets of experiments, the two CDA applications were successfully executed. The unsuccessful run was attributed to insufficient local hardware computing power (no minimum computing performance had been defined prior to the test) and resulted in the CARMA Platform losing control of the constructive vehicle.(7) In the future, the project team recommends analyzing and defining hardware computing requirements for each node in advance of testing.

The results from SIT-1 verified that VOICES can be used for distributed testing of cooperative and connected transportation applications, including CDA prototypes and simulations over a secure common network.(4)

Transmission latencies across the SIT-1 network were low and did not differ significantly from baseline unloaded network conditions to loaded conditions. The round-trip transmission latency between the two simulation sites in McLean, VA, was under 5 ms on average and under 20 ms on average between any of the three sites.

With the two CDA applications running, SIT-1 observed a total network throughput of around 120 KB/s.(4) The baseline throughput achieved by transferring large files was at least 10 MB/s. The throughput results show that TENA is a promising technology for VOICES by keeping TENA SDOs and TENA Messages lightweight.(5)

Results from the segment-by-segment latency analysis (wherever applicable) show that relevant VOICES adapters themselves are effective and efficient in converting J2735 messages to TENA SDOs and TENA Messages (and vice versa).(19,5) The performances of TENA adapters were shown to vary, depending on the overall computational load a computer is under in relation to the computing power it has.

For the platooning application, a direct comparison was made between the end-to-end latency values observed during similar past field testing and those observed during SIT-1 for the initial platooning handshake (where vehicles exchanged CARMA Mobility Request/Response Messages).(20) It was revealed that the VOICES SIT-1 platform on average added about 100 ms to this initial handshake. The additional latency was introduced by the conversion from J2735 messages to TENA SDOs and vice versa, as well as by the network transmission latency.(19,5) As noted previously, the computational latency varied based on computational load. The additional 100 ms latency introduced by VOICES in SIT-1 can be significantly reduced by using more powerful computers. The network latency could also be reduced by establishing a more robust network with more efficient protocols.

Future research needs include better understanding of local computational requirements for each node type. Additionally, testing could be expanded to other CDA applications, such as eco approach and departure at signalized intersections and cooperative perception.(4) There is also a desire to demonstrate collaborative testing using multiple cosimulation platforms and CDA prototype systems beyond CARLA and CARMA, respectively, recognizing that there is a diversity of simulation platforms and models used throughout the industry.(12,3)

APPENDIX A. VOICES SIT-1 DETAILED FUNCTIONAL ARCHITECTURE DIAGRAMS

Appendix A presents detailed functional architecture diagrams for each simulation node in SIT-1. The narrative and the diagrams in appendix A originally appeared in the SIT-1 Test Plan.

TFHRC (MCLEAN, VA)

TFHRC hosted two simulation nodes: the live and virtual simulation nodes.

Live Simulation Node

Figure 28 illustrates the detailed functional architecture of various hardware and software components at the live simulation node. The CARMA Streets-V2X Hub box contains small boxes in two columns.(9,11) The left column boxes represent native components of V2X Hub. The right column boxes represent TENA adapters that operate as V2X Hub plugins.(5) Arrows going to and from all TENA adapters indicate that the adapters subscribe to the TENA scenario SDO, which specifies the runtime description of the testing scenario, including the LVC nature of various entities.(6) The TENA scenario SDO enables the TENA adapters to determine the relevant messages to select and transmit to the V2X Hub SPaT plugin. The V2X Hub-TENA SPaT plugin work together to convert traffic signal status information from an NTCIP 1202 object to TENA Traffic Light SDO. In the SIT-1 implementation, each physical traffic signal needs a V2X Hub SPaT plugin and V2X Hub-TENA SPaT plugin.

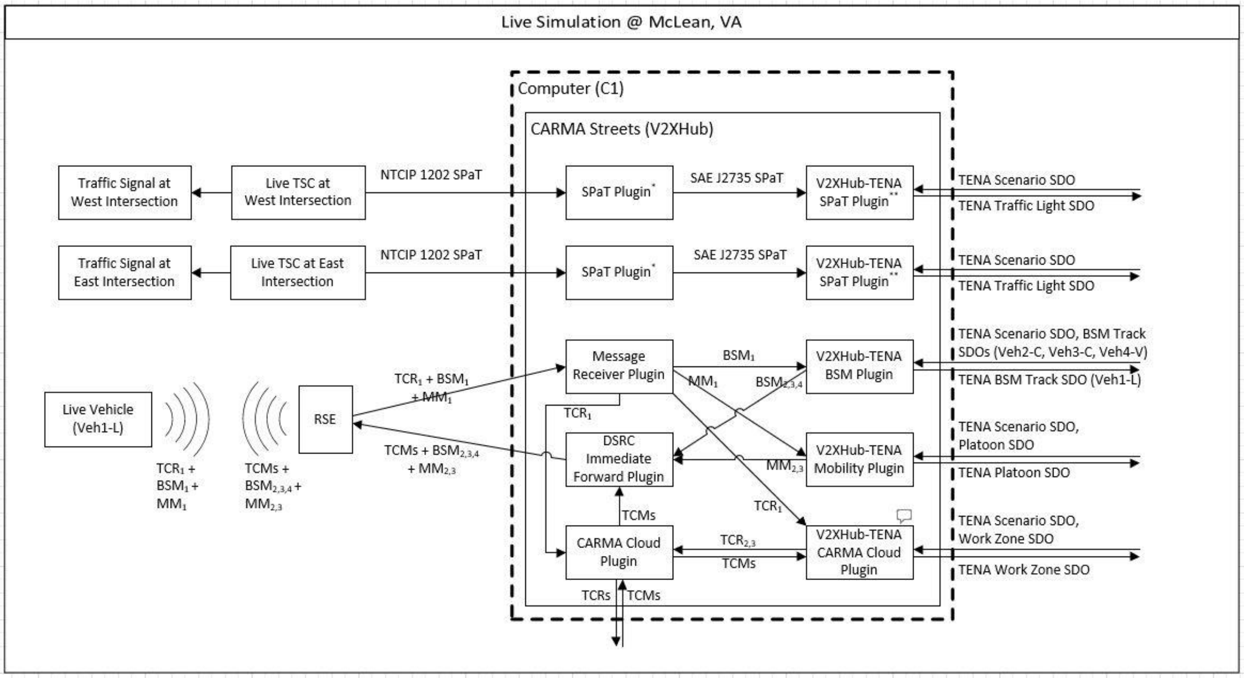

Source: FHWA.
See references 5, 9,10, 11, 21, 28.
*Multiple instances of the SPaT plugins are needed when multiple TSCs are involved, as the TCIP SPaT data from different TSCs and come in on different ports. The manifest file needs to be updated to assign each instance a unique name.
**In the current implementation as of August 2022, multiple instances of the V2X Hub-TENA SPaT plugins are needed when multiple TSCs are involved. The manifest file needs to be updated to assign each instance a unique name.

Figure 28. Data flow. SIT-1 detailed functional architecture.

The V2X Hub message receiver plugin handles all messages received from the live vehicle.(11) In SIT-1, such messages include the BSMs, CARMA TCRs, CARMA mobility request, mobility response, and MOMs (together referred to as mobility messages (MMs) in figure 28.(20) The V2X Hub message receiver plugin decodes the messages and publishes them on a message bus.

The V2X Hub DSRC Immediate Forward plugin handles all messages broadcasted by the RSU to roadway entities.(11) In SIT-1, such messages include CARMA TCMs from CARMA Cloud and BSMs and CARMA MMs from the two constructive vehicles.(10,20)

The V2X Hub-CARMA Cloud plugin listens to CARMA TCRs from the message bus and passes them on to CARMA Cloud.(11,10) When CARMA TCMs are received back from CARMA Cloud, they are published to the message bus by the V2X HubCARMA Cloud plugin.

Three TENA plugins (V2X Hub-TENA BSM plugin, V2X Hub-TENA mobility plugin, and V2X Hub-TENA traffic control plugin) subscribe to the message bus and listen for BSMs, CARMA MMs, and CARMA TCRs from the live vehicle, respectively.(5,11,20) The TENA plugins convert the messages into relevant TENA SDOs and TENA messages and publish them to the network.(6) VOICES adapters at the two constructive simulation nodes pick up the TENA BSM SDOs and TENA platoon SDOs messages, which encompass MMs, from the live simulation node, process them, and pass them on to the constructive vehicles. The TENA message TCR from the live simulation node will be captured by the TDCS. These TENA plugins also listen for relevant TENA SDOs on the network from other simulation nodes and convert them back into appropriate forms used by native V2X Hub plugins to enable communication between LVC simulation nodes.

Virtual Simulation Node

Figure 29 shows the detailed functional diagram of the virtual simulation node hosted at TFHRC. A CARLA-TENA adapter works with the CARLA simulator to construct digital twins of the live and constructive vehicles as well as the live traffic signals.(12,5) The CARLA-TENA adapter at the virtual simulation site also obtains vehicle status of the manually controlled virtual vehicle from CARLA, converts the information into relevant SDOs, and publishes the SDOs to the network.

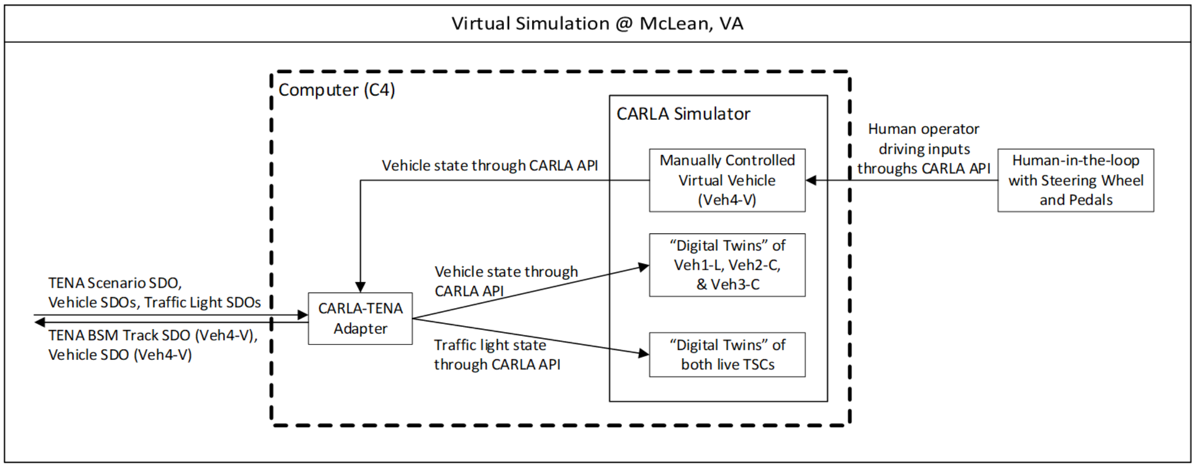

Source: FHWA.
See references 5 and 12.
API = application programming interface; BSM = basic safety message; SDO = stateful distributed object; TSC = traffic signal controller.

Figure 29. Illustration. SIT-1 detailed functional architecture: Virtual simulation node.

A constructive simulation node was hosted at SRC in Augusta. The site was chosen because of the network expertise of the onsite staff members and their familiarity with distributed network systems. The SRC team was a vital contributor in network design and configuration through every stage of design and configuration. Due to the site's distance from TFHRC, SRC received a computer with all necessary software installed and configured was shipped to SRC for SIT-1.

Figure 28 shows the detailed functional diagram of the constructive simulation node. Like the virtual simulation node, a CARLA-TENA adapter handles the information exchange of TENA Vehicle and Traffic Light SDOs and constructs digital twins in the CARLA simulator at the constructive simulation node.(12,5) In addition to the CARLA-TENA adapter, a CARMA-Platform-TENA adapter operates at the constructive simulation node and directly interacts with CARMA platform to translate BSMs, MMs, TCRs and TCMs.(7)

MITRE

MITRE Corporation participated in SIT-1 as the third site running a constructive simulation node identical to that at SRC (figure 30). Different from SRC, MITRE compiled the constructive simulation node from scratch and installed all necessary software.

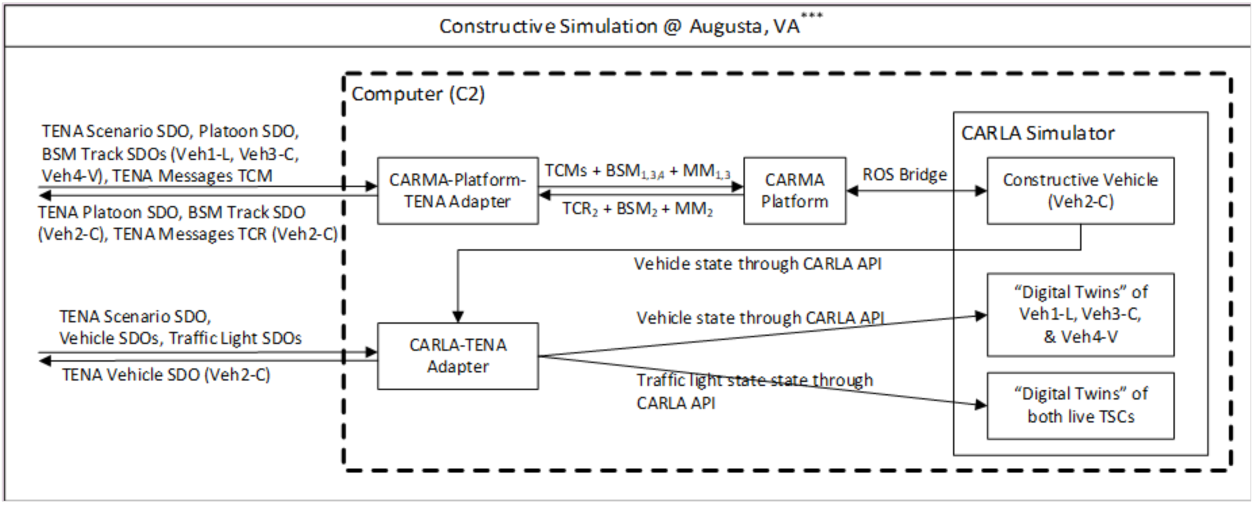

Source: FHWA.
See references 5, 7, 12, 15.
***An identical constructive simulation was hosted by MITRE in McLean, VA.
API = application programming interface; BSM = basic safety message; MM = mobility message; ROS = robot operating system; SDO = stated distributive object; TSC = traffic signal controller; TSR = traffic sign recognition.

Figure 30. Illustration. SIT-1 detailed functional architecture.

APPENDIX B. VOICES SIT-1 TEST PROCEDURES

Appendix B presents table 25, through table 28. SIT-1 test procedures as originally included in the SIT-1 Test Plan.

Table 25. SIT-1 test setup: Three-vehicle platooning.

Test Case #	SIT-1.1
Test Case	3-vehicle platooning
Objective	Provide proof of concept of VOICES as a common interface to connect three geographically distributed LVC nodes for CDA platooning.(6)
Entrance Criteria	TFHRC Live and Virtual Simulation Nodes • An RSU that is connected to the VOICES network and is configured to receive BSMs via DSRC and forward them to CARMA Streets.(21,9) • A TSC that is connected to the VOICES network and is configured to send NTCIP1202 messages to CARMA Streets.(9) • A live vehicle equipped with CARMA Platform parked within range of the RSU. A CARMA Platform instance is started and running, and is broadcasting BSMs with a constant BSM ID.(7) • A roadside computer (computer 1) is connected to the VOICES network running CARMA Streets with VOICES adapters configured and built.(9) The V2XHub Message Receiver, DSRC Immediate Forward, and SPaT plugins are configured and running.(9,11,21) • A computer (computer 4) that is connected to the VOICES network with TENA middleware, the TENA console, Scenario Publisher and VOICES DCS built and configured.(5) This live vehicle will be the first (lead) vehicle in the platoon. • A computer (computer 4) also has a steering wheel and the necessary script to start and control the virtual vehicle. • Constructive Simulation Node at SRC.

Entrance Criteria	- A computer (computer 2) that is connected to the VOICES network with TENA middleware, CARMA Platform, CARMA-CARLA integration), and the relevant VOICES adapters built and configured.(5, 7, 15) The constructive vehicle from this constructive simulation site will be the second (middle) vehicle in the platoon. Constructive Simulation Node at MITRE - A computer (computer 3) that is connected to the VOICES network with TENA middleware, CARMA Platform, CARMA-CARLA integration), CARLA simulator, and the relevant VOICES adapters built and configured. (See reference 5, 7, 15, 12.) The constructive vehicle from this constructive simulation site will be the third (last) vehicle in the platoon.
Data Inputs	J2735 BSMs, CARMA Mobility Messages (conforming to J2735 Test Messages), NTCIP1202 messages, Scenario File.(19,20)
Data Outputs	CARLA simulation of the digital twins of the live and constructive vehicles and the live TSC, VOICES DCS SDO capture database for BSM, Vehicle, Platoon, TrafficLight, and Scenario SDOs.(12)
Exit Criteria	The live vehicle successfully leads a platoon of itself and the two constructive vehicles in simulation. The simulated traffic signals update to reflect the state of the live traffic signal.

Table 26. SIT-1 test execution: Three-vehicle platooning.

Test Procedures	Step Description	Expected Outcome
1	On Computer 4, start the TENA Console.(5)	The TENA Console is started and displays a graphical user interface (GUI) to start and monitor a TENA execution.(5)
2	On the TENA Console, start the Execution Manager (EM).(5)	The TENA EM is started and there are no TENA applications in the current execution. (5)
3	On Computer 4, start the TENA Data Collection System (TDCS) and join the execution.(5)	The TDCS is started and the TDCS is shown in the TENA Console.(5)
4	On Computer 4, start the TENA Scenario Publisher.(5)	The Scenario Publisher joins the TENA execution, appears in the TENA Console.(5)
5	On CARMA Streets, start the TENA BSM plugin, TENA SPaT plugin, and the TENA Mobility plugin.(9,5)	The plugins TENA adapters join the TENA execution and appear in the TENA console.(5)
6	On Computer 4, start the CARLA Simulator, TENA CARLA adapter, Entity Generator, and the virtual vehicle control script.(12,5)	The CARLA Simulator viewing window appears, the TENA CARLA adapter and Entity Generator join the execution and appear in the TENA console, and the control window for the virtual vehicle is appears.(12,5)
7	On Computer 2, start CARMA Platform, CARMA-CARLA integration, the TENA CARLA adapter, and the TENA CARMA Platform adapter.(7,15,5)	The CARLA Simulator viewing window appears, the TENA CARLA adapter and TENA CARMA Platform adapter join the execution and appear in the TENA console.(7,15,5)
8	On Computer 3, start CARMA Platform, CARMA-CARLA integration), the TENA CARLA adapter, and the TENA CARMA Platform adapter.(7,15,5)	The CARLA Simulator viewing window appears, the TENA CARLA adapter and TENA CARMA Platform adapter join the execution and appear in the TENA console.(7,15)
9	Verify all TENA adapters have received Scenario SDOs.(5)	All adapters show they received a Scenario SDOs in their OM Stats tab on the TENA Console.(5)
10	Verify the TENA BSM plugin and TENA CARMA Platform adapters are publishing live BSM Track SDOs.(5,7)	The adapters show increasing numbers of BSM SDOs in their OM Stats tab on the TENA Console.(5,7)

Test Procedures	Step Description	Expected Outcome
11	Verify the Entity Generator and TENA CARLA adapters are publishing Vehicle SDOs.(5)	The adapters show increasing numbers of Vehicle SDOs in their OM Stats tab on the TENA Console.(5)
12	Verify the TENA Mobility plugin and TENA CARMA Platform adapters are publishing Mobility Path SDOs.(5,7)	The adapters show increasing numbers of Mobility Path SDOs in their OM Stats tab on the TENA Console.(5,7)
13	Verify TENA SPaT plugin is publishing TrafficLight SDOs.(5)	The TENA SPaT plugin shows increasing numbers of TrafficLight SDOs in its OM Stats tab on the TENA Console.(5)
14	Verify the live vehicle, constructive vehicles, and virtual vehicles are all shown in their proper locations and snapped to the roadway on all instances of the CARLA simulator.(12)	All vehicles are located in the proper location on the roadway on all instances of CARLA.(12)
15	Verify the traffic lights at the West Intersection in all CARLA simulators matches the state of the live traffic signal.(12)	The state of the traffic light at the West Intersection in all CARLA simulators matches that of the real traffic light.(12)
16	Open the CARMA Platform web interface on the live and constructive vehicles and select the appropriate route.(7)	The web interface for each vehicle is opened and the route is selected.
17	Engage CARMA on the lead (live) vehicle.(3)	CARMA Platform on the live vehicle is engaged and it begins to drive its route.(7)
18	When the live vehicle begins to move on the CARLA simulator for the second vehicle (Computer 2), engage CARMA for that vehicle (Constructive Vehicle 1).(12,3)	CARMA Platform on Computer 2 is engaged and begins to drive its route.(7)
19	Verify that the live vehicle begins platooning with Constructive Vehicle 1.	The states for the live and joining constructive vehicle progress from Searching, to Connecting to new leader/follower, and finally leading/following.
20	Engage CARMA Platform for Constructive Vehicle 2 (on Computer 3).(7)	CARMA Platform on Computer 3 is engaged and begins to drive its route.(7)

Test Procedures	Step Description	Expected Outcome
21	Verify that the live vehicle begins platooning with Constructive Vehicle 2.	The states for the live and joining constructive vehicle progress from Searching, to Connecting to new leader/follower, and finally leading/following.
22	Once the live and constructive vehicles form a platoon, begin manually driving the virtual vehicle along the desired route.	The virtual vehicle is manually driven along the desired route and is updated on all CARLA instances.(12)
23	Verify the live and constructive vehicles close the gap in the platoon and maintain formation throughout the duration of the route.	The following vehicles close the platoon gap to the desired distance and maintain that distance for the duration of the route.
24	When the live vehicle reaches the end of its route, disengage CARMA Platform by pressing the brake pedal. Do not stop the CARMA Platform process.(7)	CARMA Platform on the live vehicle is disengaged.(7)
25	Manually drive the live vehicle along the duration of the route and stop at the red light at the West Intersection.	The live vehicle is driven to the West Intersection and stops at the red light. The live vehicle's location is still updated on all CARMA simulators.(3)
26	Wait for the traffic signal at the West Intersection to turn green and verify signal states in CARLA match with the live traffic signal.(12)	The West Intersection traffic light turns green for the live vehicle and the signal states are appropriately updated in CARLA.(12)
27	Turn left in the live vehicle and manually drive it back to the starting location.	The live vehicle is manually driven back to the starting location.
28	Stop all CARMA Platforms, CARLA simulators, TENA adapters, and data collections.(7,5)	All CARMA Platforms, CARMA Platforms, CARLA simulators, TENA adapters, and data collections are shut down. No applications are shown in the TENA Console.(7,5)
29	Stop the TENA EM.(5)	The TENA EM is shut down.(5)

Table 27. SIT-1 work zone test setup.

Test Case No.	SIT-1.2
Test Case	Work zone
Objective	Provide PoC of VOICES as a common interface facilitate CDA work zone functionality.(4)
Entrance Criteria	TFHRC Live and Virtual Simulation Nodes • An RSU that is connected to the VOICES network and is configured to receive BSMs via DSRC and forward the messages to CARMA Streets.(21,9) • A TSC that is connected to the VOICES network and is configured to send NTCIP1202 messages to CARMA Streets.(9) • A live vehicle equipped with CARMA Platform parked within range of the RSU. A CARMA Platform instance is started and running, and is broadcasting BSMs with a constant BSM ID.(7) • A roadside computer (computer 1) is connected to the VOICES Network running CARMA Streets with VOICES adapters built and configured and built.(9) The V2X Hub Message Receiver, DSRC Immediate Forward, and SPaT plugins are configured and running.(11,21) • A computer (computer 4) that is connected to the VOICES network with TENA middleware, the TENA Console, Scenario Publisher, and VOICES DSC built and configured.(5) This live vehicle will be the first (lead) vehicle in the platoon.
Data Inputs	J2735 BSMs, NTCIP1202 messages, scenario file.(19)
Data Outputs	CARLA simulation of the digital twins of the live and constructive vehicles and the live TSC, VOICES DCS SDO capture database for BSM, Vehicle, TrafficLight, and Scenario SDOs.(12)
Exit Criteria	The live vehicle successfully leads a platoon of itself and the two constructive vehicles in simulation. The platoon of vehicles successfully adjust speed according to work zone speed limit when traversing through the geofenced work zone area. The simulated traffic signals update to reflect the state of the live traffic signal.

Table 28. SIT-1 work zone test execution.

Test Procedures	Step Description	Expected Outcome
1	On computer 4, start the TENA Console.(5)	The TENA console is started and displays a GUI to start and monitor a TENA execution.(5)
2	On the TENA console, start the TENA EM.(5)	The TENA EM is started and there are no TENA applications in the current execution.(5)
3	On computer 4, start the TDCS and join the execution.	The TDCS is started and the TDCS is shown in the TENA console. (5)
4	On computer 4, start the TENA Scenario Publisher and Entity Generator.(5)	The Scenario Publisher and entity generator join the TENA execution, appear in the TENA Console.(5)
5	On CARMA Streets, start the TENA BSM plugin, and TENA SPaT plugin.(9,5)	The plugins TENA adapters join the TENA execution and appear in the TENA console.(5)
6	Verify all TENA adapters have received Scenario SDOs.(5)	All adapters show they received a Scenario SDOs in their OM Stats tab on the TENA console.(5)
7	Verify the TENA BSM plugin is publishing live BSM Track SDOs. (5)	The plugin shows increasing numbers of BSM SDOs in the OM Stats tab on the TENA console.(5)
8	Verify the entity generator is publishing Vehicle SDOs.	The entity generator shows increasing numbers of Vehicle SDOs in the OM Stats tab on the TENA console.(5)
9	Verify TENA SPaT plugin is publishing TrafficLight SDOs. (5)	The TENA SPaT plugin shows increasing numbers of TrafficLight SDOs in its OM Stats tab on the TENA console. (5)
10	Verify the live vehicle is shown in its proper location and snapped to the roadway.	The live vehicle is located in the proper location on the roadway.
11	Verify the traffic lights at the West Intersection in all CARLA simulators matches the state of the live traffic signal.(12)	The state of the traffic light at the West Intersection in all CARLA simulators matches that of the real traffic light.(12)
12	Open the CARMA Platform web interface on the live vehicle and select the appropriate route.(7)	The web interface is open, and the route is selected.
13	Engage CARMA Platform on the live vehicle.(7)	CARMA Platform on the live vehicle is engaged and it begins to drive its route.(7)

Test Procedures	Step Description	Expected Outcome
14	When the live vehicle enters the speed reduction work zone, verify an active event begin alert shows in the CARMA web user interface, and its targeted speed reduces to the specified value.(3)	The live vehicle shows the active event begin alert in the web UI and reduces its speed inside the work zone.
15	When the live vehicle leaves the speed reduction work zone, verify an active event end alert shows in the CARMA web user interface, and its targeted speed returns to its previous value.(3)	The live vehicle shows the active event end alert in the web UI and its targeted speed is returned to its previous value.
16	When the live vehicle reaches the end of its route, disengage CARMA Platform by pressing the brake pedal. Do not stop the CARMA Platform process.(7)	CARMA Platform on the live vehicle is disengaged.(7)
17	Manually drive the live vehicle along the duration of the route and stop at the red light at the West Intersection.	The live vehicle is driven to the West Intersection and stops at the red light. The live vehicle's location is still updated on all CARMA simulators.(3)
18	Wait for the traffic signal at the West Intersection to turn green and verify signal states in CARLA match with the live traffic signal.(12)	The West Intersection traffic light turns green for the live vehicle and the signal states are appropriately updated in CARLA.(12)
19	Turn left in the live vehicle and manually drive it back to the starting location.	The live vehicle is manually driven back to the starting location.
20	Stop all CARMA Platforms, CARLA simulators, TENA adapters, and data collections.(7,12,5)	All CARMA Platforms, CARMA Platforms, CARLA simulators, TENA adapters, and data collections are shut down.(7, 12,5)
21	Stop the TENA EM.(5)	The TENA EM is shut down.(5)

APPENDIX C. DETAILED DISCUSSION OF VOICES SIT-1 SEGMENT-BY-SEGMENT DATA ANALYSIS METHODOLOGY

Appendix C expands on the segment-by-segment data analysis methodology discussed in chapter 4. The data referred to in appendix C comprised multiple datasets collected at various interfaces along a data flow pipeline. The datasets include multiple PCAP files and recorded TENA data.(5)

LOADING DATA

The data was loaded into the script and assigned dynamic tags for organization and ease of access. These tags defined the name, dataflow order, format, message type, and other properties of the dataset to be used during processing.

FILTERING DATA

Once all datasets were loaded, they were filtered down to only the desired entries. For example, when analyzing CARMA MOMs sent from the live vehicle to a constructive vehicle, all CARMA MOMs from other vehicles must be filtered out.(20) For some datasets, namely the TENA Platoon SDO data, filtering out other vehicles was difficult because the current design of the TENA platoon OM associates data from all platooned vehicles to the platoon leader; and the corresponding TENA adapter of the platoon leader publishes that data.(5) Analysis of the filtered data interpreted all data points for Platoon SDOs as originating from the platoon leader. To distinguish the actual source vehicle, analysts used the source vehicle's simulation time, which is a persistent field from each vehicle's originally published data. Simulation time is simply an increasing counter starting when a node's simulation began. Since all simulation nodes were started at slightly different times in SIT-1, the origin of a Platoon SDO data update could be identified by its timestamp relative to the desired vehicle's simulation timestamp. Moreover, CARMA mobility operations, mobility request, and mobility response messages are all contained in the Platoon SDO but are not explicitly distinguished in the data. Therefore, the type of message had to be inferred based on the values of a combination of different data fields. Fortunately, these levels of extrapolation were not often required and will be remedied or streamlined in future releases of the relevant TENA OMs.

ALIGNING DATA

Once the datasets were filtered, they had to be aligned with the same initial packet since the data captures at each site were started at slightly different times. First, the script searched for the first packet from the filtered source dataset in all other datasets. If not found at any point in one or more datasets, the script moved to the second packet in the source dataset up to the first 30 source packets to prevent excessive processing loops. The script verified packets that matched across datasets using configured match keys. For example, BSMs were matched using the latitude, longitude, sec mark, and velocity values. Match keys (imported from columns in the Mcirosoft® Excel™ data) are different for PCAPs and TDCS data (speed(m/s) versus tspi.velocity.ltpENU asTransmitted.vxInMetersPerSecond,Float32 (optional), respectively) and were configured accordingly.

In addition to the match key name, there were other slight differences between the packet capture and TDCS data. For some types of data, TENA uses built-in functions and subsystems to manage and translate data, including timestamps and location data.(5) A simple example is timestamp data, which is recorded in TDCS in nanoseconds, which had to be converted to match the PCAP's milliseconds. A more complicated example is location data, which is received by the TSPI module (in the form of latitude, longitude, elevation) and immediately is translated into many various coordinate systems such as XYZ and other common projections for easy access by other systems. The result of this process was a value that was slightly different from the original and has the potential for a difference in precision (as seen in table 29). To account for this, modifiers were added to the match key values such as decimal place rounding and match buffers. Values for these modifiers were carefully chosen and tuned to allow for appropriate matching and prevent over-rounding shows an example of this rounding. For this speed example, TDCS values were rounded to two decimal places with a buffer of 0.03. Additional modifiers were added as needed, including converting radians to degrees, and extracting values from compound fields (e.g., platooning operations parameters).

While these approximations in a single key might cause uncertainty on their own, the use of multiple keys, including those with more cross-dataset precision, allow for a high certainty of data alignment. More or fewer match keys were used depending on the cross-dataset precision and detail within the key itself. table 30 shows a complete set of all match keys used in the data analysis by message type.

Table 29. Match key rounding example.

PCAP Velocity Value (m/s)	TDCS Velocity Value(m/s)
0.16	0.1599999964237210
0.42	0.4199999868869780
0.62	0.6200000047683720
0.78	0.7799999713897710
0.98	0.9800000190734860
1.16	1.1599999666214000
1.36	1.3600000143051100
1.56	1.5599999427795400
1.76	1.7599999904632600
1.96	1.9600000381469700
2.14	2.1400001049041700
2.28	2.2799999713897700
2.42	2.4200000762939500

m/s= meters/second

Table 30. Match keys per message type.

Data Type	Match Key	PCAP	TDCS	Modifiers
BSM	1	latitude	tspi.position.geodetic_asTransmitted.latitudeInDegrees,Float64 (optional)	Round: 6 decimals Buffer: 0.000002
	2	longitude	tspi.position.geodetic_asTransmitted.longitudeInDegrees,Float64 (optional)	Round: 6modi decimals Buffer: 0.000002
	3	secMark	msWithinMinute,UInt16	—
	4	speed(m/s)	tspi.velocity.ltpENU_asTransmitted.vxInMetersPerSecond,Float32 (optional)	Round: 2 decimals Buffer: 0.03
SPAT	1	phase2_eventState	Enum,VUG::Entities::Signals::TrafficLightState, currentState	SPAT state comparison
Mobility path	1	hostStaticId	header.hostStaticId,String	—
	2	hostBSMId	header.hostBSMId,String	—
	3	planId	header.planId,String	—
Mobility operations	1	headerTimestamp	timestamp.nanosecondsSince1970, Int64	Convert ms to ns
	2	operationParams	downtrackDistanceInMeters,Float32 or joinedVehicles^strategyParameters,String or joinedVehicles^strategyParameters,String	Extract DTD from string Round: 3 decimals Buffer: 0.002
Mobility request	1	headerTimestamp	timestamp.nanosecondsSince1970,Int64	Convert ms to ns
	2	strategyParams	requestedVehicles^strategyParameters,String	N/A
Mobility response	1	headerTimestamp	timestamp.nanosecondsSince1970,Int64	Convert ms to ns
Traffic control request/message	1	reqid_hex	requestID, String	—

— No modifier used.

Once all datasets contained an initial packet, the script iterated through each source packet and checked for the same packet in every other dataset. If the script found no initial packet in a dataset, it assumed the inserted script dropped packet rows until the data was realigned. This process both identified dropped packets and produced a complete set of aligned data that could be processed easily for performance data. Table 31 shows an example of the dropped packet alignment process.

Table 31. Dropped packet example excerpt.

Source Data	Downstream Data—BEFORE	Downstream Data—AFTER
38.9556557	38.9556557	38.9556557
38.9556509	38.9556509	38.9556509
38.9556469	38.9556469	38.9556469
38.955641	38.955636	DROPPED PACKET
38.955636	38.955632	38.955636
38.955632	38.9556261	38.955632
38.9556261	38.955622	38.9556261
38.955622	38.9555841	38.955622

Once all datasets were filtered and aligned, performance metrics were calculated, using the performance analysis script, for each packet and each step of the process in the VOICES system. This process outputs a CSV file containing the following for each step in the data flow:

- Packet index (relative to the input data, for traceability).
- Packet timestamp.
- Total latency (time from message generation to current step).
- Incremental latency (time for current step).

All computers were time-synced to GPS time, so latencies were calculated by subtracting the current step timestamp from the previous step. Datasets containing dropped packets were marked with "DROPPED PACKET," and subsequent incremental latencies were labeled as "NO PREV PACKET." If one dataset ended before another, the remaining columns for the ended dataset would contain EOF (end of file). Finally, summary metrics calculated for each dataset in the generated results are the following:

- Total input packets.
- Total filtered packets.
- Dropped packets.
- Minimum latency.
- Maximum latency.
- Mean latency.
- Jitter or standard deviation (for network transmission and computation, respectively).

All scripts described in this section can be found on the VOICES-POC GitHub.(22) The single analysis script (calculate_e2e_perf.py) can be executed using command line options or prompted user input. Data can be loaded by user input or from data input files containing the location and tag definition of each dataset. The batch analysis script (batch_calculate_e2e_perf.py) is configured using the locations of the input data files for each directory and executed from the command line with no additional arguments required. The GitHub readme file contains more information on the execution of these scripts.

ACKNOWLEDGMENTS

This proof-of-concept project and report would not have been possible without the support and hard work of many talented individuals. Developing SIT-1 and testing was truly a collaborative effort deserving of acknowledgments and thanks to a large group of contributors.

The authors would like to thank the leaders at the Office of the Secretary of Transportation: Robert Heilman, Taylor Lochrane, Santiago Navarro, Ryan Steinbach, and Dr. Robert Hampshire for their steady guidance and oversite during the project. Much of what has been accomplished is due to their unwavering support.

Another group that has made this project possible is the incredible public servants at USDOT FHWA: Dale Thompson, Randall VanGorder, Govind Vadakpat, Craig Thor, and all the other FHWA teams that support contracts, public affairs, and operations.

From the USDOT VOLPE Center, we wish to thank Samuel Toma, Wassim Najm, Faroog Ibrahim, and Sharon Chan Edmiston for their support.

Much of the technical work for this project was supported by the DoD and their contractors. The authors would like to thank Ryan Norman, Scott Wales, and Nicolas Roca, for all the hard work and expertise they put into accomplishing all that CARMA has done.

Many members from the VOICES CoP contributed to VOICES. To all who are a part of the CoP, thank you so much for your time, contributions, expertise, and willingness to invest in this project to help shape the future of transportation and increase safety.

The original map shown in figure 3 is the copyright property of Google® Earth and can be accessed from https://www.google.com/earth.(23) Project engineers created the overlay that depicts the vehicle test route.

REFERENCES

1. U.S. Department of Transportation. n.d. "Virtual Open Innovation Collaborative Environment for Safety (VOICES) Proof of Concept (PoC)" (web page). https://www.transportation.gov/hasscoe/voices, last accessed January 29, 2024.

2. FHWA. n.d. "VOICES" (web page). https://highways.dot.gov/research/projects/voices, last accessed January 29, 2024.

3. USDOT. n.d. "CARMA" (web page). https://its.dot.gov/cda, last accessed May 30, 2023

4. FHWA. 2022. *Cooperative Driving Automation Transportation Systems Management Operations Strategies and Use Cases—An Overview*. Publication No. FHWA-HRT-22-067. Washington DC: Federal Highway Administration. https://doi.org/10.21949/1521869, last accessed June 16, 2023

5. The Foundation for DoD Range Systems Interoperability. 2020. "Test and Training Enabling Architecture" (web page). https://www.tena-sda.org/attachments/TENA-OverviewFS-2020-03-16-DistA.pdf, last accessed September 14, 2021.

6. Demers, A. 2019. "Live-Virtual-Constructive (LVC) Training" (web page). https://www.presagis.com/en/blog/detail/live-virtual-constructive-lvc-training-then-and-now/, last accessed June 8, 2023.

7. FHWA. n.d. *CARMA Platform* (software). Version 4.3.0

8. FHWA. 2023. *CARMA Messenger* (software). Version 4.4.0.

9. FHWA. 2023. *CARMA Streets* (software). Version 4.3.0.

10. FHWA. 2023. *CARMA Cloud* (software). Version 4.3.0

11. FHWA. 2023. *V2X Hub* (software). Version 7.4.0.

12. CARLA. 2022. *Cars Learning to Act* (software). Version 0.9.14. https://carla.org/, last accessed February 7, 2024.

13. Eclipse. 2023. Simulation of Urban MObility (software). Version 1.17.0.

14. NSNAM. 2024. *NS-3* (software). https://www.nsnam.org/, last accessed February 7, 2024.

15. FHWA. 2024. "CARMA-CARLA Integration Tool" (Github repository). https://github.com/usdot-fhwa-stol/carma-carla-integration, last accessed May 31, 2023.

16. Loshin, P. 2021. "IPsec (Internet Protocol Security)" (web page). https://www.techtarget.com/searchsecurity/definition/IPsec-Internet-Protocol-Security, last accessed June 5, 2023.

17. U.S. Department of Transportation. n.d. *Vehicle-to-Infrastructure (V2I) Communications for Safety.* Washington, DC: U.S. Department of Transportation. https://www.its.dot.gov/research_archives/safety/v2i_comm_plan.htm, last accessed January 30, 2024.

18. Hoekstra-Atwood, L., C. M. Richard, and V. Venkatraman. 2022. *Multiple Sources of Safety Information from V2V and V2I: Phase II Final Safety Message Report*. Report No. FHWA-HRT-22-013. Washington, DC: Federal Highway Administration. https://www.fhwa.dot.gov/publications/research/safety/22013/22013.pdf, last accessed January 31, 2024.

19. SAE International. 2022. *V2X Communications Message Set Dictionary*. J2735_202211. 2022-11-14 revision. Warrendale, PA: SAE International. https://www.sae.org/standards/content/j2735_202211/, last accessed May 31, 2023

20. FHWA. 2021. "CARMA Adds to Cooperative Driving Automation (CDA) Message Set to Improve Intelligent Transportation Systems (ITS) Communications" (web page). https://www.fhwa.dot.gov/publications/research/operations/21104/21104.pdf, last accessed January 30, 2024.

21. SAE International. 2017. *Dedicated Short-Range Communication (DSRC) Systems Engineering Process Guidance for SAE J2945/X Documents and Common Design Concepts.* J2945_201712, 2017-12-07 revision. Warrendale, PA: SAE International. https://www.sae.org/standards/content/j2945_201712/, last accessed May 24, 2023

22. FHWA. 2024. "voices poc" (Voices software and configuration files in GitHub repository). https://github.com/usdot-fhwa-stol/voices-poc, last accessed February 7, 2024.

23. Google. 2023. *Google Earth*. Mountain View, CA. https://www.google.com/earth, last accessed January 19, 2024.

24. Test and Training Range Community. n.d. "TENA Ping Test" (web page). https://www.trmc.osd.mil/wiki/display/CONSOd.LE/TENA+Console+User+Guide#TENAConsoleUserGuide-BasicApplication-to-ApplicationTests, last accessed May 17, 2023

25. FHWA. 2022. "VOICES SIT-1 rosbag Analysis Script" (Github repository). https://github.com/usdot-fhwa-stol/carma-analytics-fotda/blob/develop/src/rosbag_analysis_tools/scripts/analyze_voices_sit1_rosbags.py, last accessed December 8, 2022.

26. Engelen, A. 2022. *Nethogs* (software). Version 0.87. https://github.com/raboof/nethogs, last accessed December 8, 2022.

27. FHWA. 2022. "VOICES Performance Analysis Script" (GitHub repository). https://github.com/usdot-fhwa-stol/voices-poc/tree/develop/scripts/performance_analysis, last accessed February 7, 2024.

28. FHWA. 2021. “Vehicle-to-Everything (V2X) Hub: Open-Source Connected Vehicle (CV) Software” (web page). https://www.fhwa.dot.gov/publications/research/operations/22047/22047.pdf, last accessed February 2, 2024.

Recommended citation: Federal Highway Administration,
VOICES Cooperative Driving Automation (CDA)
Proof-of-Concept Systems Integration Test 1
(Washington, DC: 2024) https://doi.org/10.21949/1521742

HRSO-40/04-24(WEB)E