

**National Aeronautics and
Space Administration**

UAS-NAS

Live Virtual Constructive
Distributed Test
Environment

Characterization Report

Document No.
UAS-04.05.0001.01

Release Date: September 2013

Signatures Page

Prepared By:

Jim Murphy – Project Engineer, NASA ARC

Sam Kim – Project Engineer / Deputy Project Manager, NASA DFRC

Concurred By:

Debra Randall – Project Chief Engineer, NASA DFRC

Approved By:

Chuck Johnson – Project Manager, NASA DFRC

REVISION SHEET

REVISION	AUTHOR	DATE	LIST OF CHANGES
Baseline	J. Murphy	XXXX	Initial Release

Executive Summary

This report documents message latencies observed over various Live, Virtual, Constructive, (LVC)¹ simulation environment configurations designed to emulate possible system architectures for the Unmanned Aircraft Systems (UAS) Integration in the National Airspace System (NAS) Project integrated tests. For each configuration, four scenarios with progressively increasing air traffic loads were used to determine system throughput and bandwidth impacts on message latency.

This report serves as the UAS in the NAS Fiscal Year 2013 Annual Performance Goal submission. The FY13APG (APG 4.2.1.1: AR-13-7) is defined as: “Complete flight evaluations to assess the capabilities of the Live, Virtual, Constructive (LVC) distributed simulation environment.” The analyses in this report cover the observed latencies for messages sent from a live aircraft during flight, augmented with analyses of observed latencies among virtual and constructive aircraft data sources to air traffic control (ATC) displays across the LVC system. Together this provides a comprehensive data set that will inform the LVC developers and researchers as a distributed environment is designed to meet the integrated event requirements.

The Integrated Test and Evaluation (IT&E) subproject conducted the LVC characterization testing working closely with members of the Communications subproject. IT&E members from NASA Ames took the lead role in the testing, hosting the core LVC infrastructure. IT&E members from NASA Dryden supported one test configuration utilizing the Ikhana Simulator and Communication members supported two test configurations, the primary flight test as well as a infrastructure test between NASA Ames and NASA Glenn.

The LVC is being developed in support of the UAS Integration in the NAS Project, which is investigating and integrating technologies that are intended to reduce technical barriers related to the safety and operational challenges associated with enabling routine UAS access to the NAS. To support this goal, the Integrated Test and Evaluation (IT&E) subproject is developing a distributed LVC test environment to enable human-in-the-loop (HITL) simulation and flight test activities. The LVC test environment for the UAS Integration in the NAS Project is comprised of ATC, constructive and virtual aircraft simulators, and UAS ground control stations (GCS) that together provide researchers with a relevant unmanned environment. For each specific flight test or simulation a subset of available live assets and software components are integrated to form an instance of the LVC. To maximize the use of available resources, the LVC test environment is designed to be distributed such that technologies and concepts developed by the Project as well as its external partners can be easily integrated with the simulation or flight environment.

The data captured in this report is critical to support building the appropriate LVC configuration in support of the Project integrated events. Due to the distributed nature of the LVC test environment, the latencies of messages passed between the LVC components observed in standalone simulations must be characterized and clearly understood to assess the effect of latency on the overall simulation. In addition, to properly synchronize live, virtual, and constructive data, it is critical to understand the latency inherent between the various possible distributed components of the LVC test environment. Understanding the LVC capabilities and performance characteristics will allow developers to account for and mitigate known system delays in order to define LVC requirements with the researchers that work within the LVC capabilities, or levy new requirements on LVC development creating a realistic test environment.

¹ The term LVC is a broadly used name for classifying modeling and simulation (M&S). It is recognized that categorizing a simulation as a live, virtual, or constructive is problematic since there is no clear division between these categories. Also, the degree of human participation in a simulation is infinitely variable, as is the degree of equipment realism. Generally live M&S involve real actors operating real systems. Virtual M&S involve real actors operating simulated systems. Constructive M&S involve simulated people operating simulated systems.

Since an LVC instantiation can be constructed in many different configurations depending on the specific requirements of a simulation event, the tests focused on determining the latencies between key components of the distributed simulation environment as characterized by four primary categories:

- 1.) Latencies to publish live aircraft state data for distribution to the rest of the LVC
- 2.) Latencies to publish virtual aircraft state data for distribution to the rest of the LVC
- 3.) Latencies to publish constructive aircraft state data for distribution to the rest of the LVC
- 4.) Latencies between distributed facilities.

In order to capture the data required to inform the latencies measurements of our four focus categories, the LVC Characterization test had two primary objectives:

- 1.) Determine the time differential for each aircraft state data message produced by the aircraft data sources; measured as the time between when the message originated and when it was received at specific points within the LVC network.
- 2.) Measure the throughput of the aircraft state data messages at specific points on the LVC system.

To meet those primary objectives eight LVC configurations were tested, namely:

- 1.) Simulated traffic to ATC on internal network
- 2.) Simulated traffic to ATC via middleware on internal network
- 3.) Simulated traffic to ATC via HLA between co-located NASA Ames facilities
- 4.) B747 flight simulator data to co-located NASA Ames facility
- 5.) Simulated traffic to ATC via LVC Gateway between co-located NASA Ames facilities
- 6.) Simulated traffic to GCS and Ikhana simulator data between NASA Ames and NASA Dryden
- 7.) Simulated traffic to GCS between NASA Ames and NASA Glenn via Internet
- 8.) Live aircraft data between S-3B Viking and LVC Gateway at NASA Glenn.

Since the Ikhana aircraft was taken out of service for a planned upgrade, the timing of the characterization tests was set to coincide with existing prototype radio testing conducted by the UAS Communication subproject. The NASA S-3B Viking aircraft, which is a candidate participating aircraft for Flight Test 4, provided live telemetry through a 3G cellular connection during a test flight in the Cleveland Center airspace. The B747 FAA Level D flight simulator located at NASA Ames and the Ikhana Predator-B simulator located at NASA Dryden provided virtual traffic data, while constructive traffic was supplied by the Multi-Aircraft Control System (MACS) Simulation Manager capability housed at NASA Ames. NASA Ames provided the facilities that ran the LVC messaging communication hub utilizing a High Level Architecture format model. Due to the location of the assets and facilities in the LVC, messaging latency was measured between facilities at NASA Ames, and between NASA Ames and facilities at NASA Dryden and NASA Glenn. The Cockpit Situational Display developed at NASA Ames was used to emulate a UAS ground control station.

The results from the data captured during the flight of the NASA S-3B proved to be significant. The time to publish the S-3B data under the best case scenario was at or near the latency requirement for data in the Terminal airspace and approached the threshold for en-route airspace even after accounting for the fact that live aircraft state data time stamp was truncated to the nearest second. While the transmission mechanism for ingesting the data into the LVC network via 3G cellular technologies may have too great of an inherent latency to be truly effective for our purposes, it is not the only transmission option. Once the Ikhana aircraft is available for testing, data will be collected using its existing Ethernet connection between the GCS and the LVC.

The B747 and Ikhana virtual simulators performed with minimal latencies observed (in the tens of milliseconds range) when publishing data to the LVC messaging components. In addition, data gathered to calculate the time it takes to transmit the state message between NASA Ames and NASA facilities at Glenn and Dryden indicate no issues with running those assets remotely.

While the MACS Simulation Manager was able to publish constructive state updates well under the required operational threshold, its latency was significantly greater than the virtual simulators. In addition, under the higher traffic loads, MACS dropped and duplicated state messages, and missed over 50% of the messages during the highest aircraft traffic loads. Because the high traffic loads tested are well beyond the designed limits for MACS (50-60 aircraft maximum), this issue is not a concern for the simulations. However, understanding this inherent behavior is critical.

Questions raised by these analyses will continue to be investigated as the project moves forward and exercises the LVC environment. In particular, tests should be conducted utilizing the Ikhana's or other aircrafts' transmission mechanisms for sending live state data to the LVC. As stated above, the nature of the MACS Simulation Manager state data latency should be investigated and mitigated if possible. The candidate air traffic control display should be instrumented to record the time state data as actually displayed; this would provide the missing latency data not covered in this report. Finally, since any changes to the LVC could impact overall latency, each instance of the LVC developed to support a simulation or flight test should be tested to determine whether the latencies are still within acceptable levels.

Table of Contents

Executive Summary	4
1 Introduction	10
1.1 Test Item Description.....	10
1.2 Overall Test Objectives.....	12
2 Method and Materials	13
2.1 Test Resources	13
2.1.1 Software Components	13
2.1.2 Live Resources	15
2.1.3 Virtual Resources	15
2.1.4 Constructive Resources	15
2.1.5 Test Facilities	15
2.1.6 System Time Synchronization	15
2.2 Test Design	16
2.2.1 Measures of Performance (MOP)	16
2.2.2 Characterization Configurations	16
3 Results.....	30
3.1 Configuration Results	30
3.1.1 Configuration 1: Simulated Traffic to ATC on DSRL Network.....	31
3.1.2 Configuration 2: Simulated Traffic to ATC via HLA on DSRL Network	33
3.1.3 Configuration 3: Simulated Traffic to ATC via HLA between DSRL and CVSRF	34
3.1.4 Configuration 4: B747 Flight Simulator Data to DSRL via HLA	35
3.1.5 Configuration 5: Simulated Traffic to ATC via LVC Gateway between DSRL and CVSRF.....	36
3.1.6 Configuration 6: Simulated Traffic to GCS and Ikhana Simulator Data between DSRL and RAIF via NISN	37
3.1.7 Configuration 7: Simulated Traffic to GCS between DSRL and NASA Glenn UAS Communication lab via Internet	38
3.1.8 Configuration 8: Live Aircraft Data between S-3B Viking and LVC Gateway at NASA Glenn UAS Communication Lab	39
3.2 Compilation Results	40
3.2.1 Publishing Live Aircraft State Data to the LVC	40
3.2.2 Publishing Virtual Aircraft State Data to the LVC	41
3.2.3 Publishing Constructive Aircraft State Data to the LVC	42
3.2.4 Documentation of Latencies Between Tested Facilities	43
3.2.5 Combined Analyses	44
4 Conclusion	46
Appendix A: Acronyms	48
References.....	50

List of Figures:

Figure 1. High-level view of the system under test	12
Figure 2. Diagram of all test components - test configurations are a subset of this set.	18
Figure 3. Configuration 1: LVC system configuration used to determine internal MACS/ADRS latencies	19
Figure 4. Configuration 2: LVC system configuration used to determine latency added by the use of HLA	20
Figure 5. Configuration 3: LVC system configuration used to determine latency added due to the distribution of the LVC across different networks	22
Figure 6. Configuration 4: LVC system configuration used to determine the latency of remote facilities sending aircraft data back to the LVC Hub	23
Figure 7. Configuration 5: LVC system configuration used to determine whether the LVC Gateway adds any additional latency beyond the use of the specific HLA Toolboxes	25
Figure 8. Configuration 6: LVC system configuration used to determine the latency when distributing across the NASA Integrated Services Network	26
Figure 9. Configuration 7: LVC system configuration used to determine the latency when distributing a remote site via a VPN over the Internet.....	28
Figure 10. GRC S-3B Channel Sounding Flight Test Communications Architecture	29
Figure 11. Configuration 8: LVC system configuration used to determine the latency of receiving live aircraft telemetry data	30
Figure 12. Observed latencies for direct MACS state data generation to display time	32
Figure 13. Observed latencies for a typical aircraft during the 200 aircraft traffic load	33
Figure 14. Observed latencies during the 400 Aircraft Scenario of the B747 state data including time to the HLA at the DSRL	36
Figure 15. Example of Discrete Buffering from Live Aircraft.....	40
Figure 16. State Data Publishing Latency, Ikhana vs B747	42
Figure 17. State Data Publishing Latency for the MACS SimMgr with Standard Deviations	43
Figure 18. HLA Transit Latency, CVSRF vs Dryden vs Glenn	44
Figure 19. End-to-End Latency Comparison based on Average Latencies	45
Figure 20. End-to-End Latency Comparison based on Average Latencies plus 2 Standard Deviations.....	46

List of Tables:

Table 1. Facilities used during LVC Characterization testing and the type of network connection they have to the LVC Hub at DSRL	15
Table 2. Data collected for Configuration 1	19
Table 3. Data collected for Configuration 2	20
Table 4. Data collected for Configuration 3	22
Table 5. Data collected for Configuration 4	23
Table 6. Data collected for Configuration 5	25
Table 7. Data collected for Configuration 6	26
Table 8. Data collected for Configuration 7	28
Table 9. Data collected for Configuration 8	30
Table 10. Observed Latencies between MACS SimMgr and MACS DSR.....	31
Table 11. MACS publishing time to LVC (HLA)	33
Table 12. HLA Transit Times (local network)	34
Table 13. HLA Transit Times between DSRL and CVSRF.....	35
Table 14. B747 Latency to HLA	35
Table 15. Transit Time between HLA in DSRL and LVC Gateway in CVSRF.....	37
Table 16. HLA Transit Times between DSRL at NASA Ames and RAIF at NASA Dryden	38
Table 17. Ikhana State Data Publishing Time to Dryden LVC Gateway	38
Table 18. Transit time between Ames HLA and Glenn LVC Gateway	39
Table 19. S-3B State Data Publishing Latency to Glenn LVC Gateway	39
Table 20. S-3B State Data Publishing Latency to Glenn LVC Gateway (Best Case).....	41

1 Introduction

The Unmanned Aircraft Systems (UAS) Integration in the National Airspace System (NAS) Project is investigating and integrating technologies that are intended to reduce technical barriers related to the safety and operational challenges associated with enabling routine UAS access to the NAS. To support this goal, the Integrated Test and Evaluation (IT&E) subproject is developing a distributed Live, Virtual, Constructive (LVC) test environment to enable human-in-the-loop (HITL) simulation and flight test activities. LVC test environments are not a new concept; they are widely used by the Department of Defense to provide a safe and relevant test environment.^{1,2} A constructive simulation generally has no interactive human involvement in simulated conditions. Instead, scenarios unfold using rule-based decisions that control the interactions between simulated actors. Virtual simulations involve real actors operating simulated systems (i.e., human/operator interaction in the use of the model or simulation). A live test environment involves real actors operating real systems. Categorizing a simulation as live, virtual, or constructive is problematic since there is no clear division between these categories. Also, the degree of human participation in a simulation is infinitely variable, as is the degree of equipment realism.

The LVC test environment for the UAS Integration in the NAS Project is comprised of air traffic control (ATC), constructive and virtual aircraft simulators, and UAS ground control stations (GCS) that together provide researchers with a relevant unmanned environment. In order to maximize the use of available resources, the LVC test environment is designed to be distributed in such a way that technologies developed by our research and external partners can be more easily integrated into the simulation or flight environment. Due to the distributed nature of the LVC test environment, the latencies of messages passed between the LVC components observed in standalone simulations must be characterized and clearly understood to assess the overall simulation. In addition, to properly synchronize live, virtual, and constructive data, it is critical to understand the latency inherent between distributed components of the LVC test environment, henceforth referred to as the LVC.

Utilizing an existing government off the shelf aircraft traffic generator and air traffic control display, a distributed LVC system was developed to test the message latencies between distributed LVC components. The system contains the core infrastructure components specifically developed to distribute and record the messages. This test version of the LVC system allows the LVC components to be distributed across different facilities providing for the testing of message latency across various network topologies with the understanding that message latency may be affected by the throughput of the message traffic and bandwidth of the network.

This report documents message latencies observed over eight LVC instantiations or configurations designed to address different possible system architectures. For each configuration, four scenarios with progressively increasing air traffic loads were used to determine system throughput and bandwidth impacts on message latency. The analyses cover the observed latencies for messages sent from the aircraft data sources (constructive state data generators, virtual aircraft simulators, and live aircraft telemetry data) to ATC displays across the LVC system. Understanding the LVC capabilities and performance characteristics will allow developers to account for and mitigate known system delays to create a more realistic test environment.

1.1 Test Item Description

The characterization tests were designed to evaluate the throughput and data latencies for specific communication paths between the core LVC system components, specifically the LVC Gateway and the High Level Architecture (HLA) middleware. Figure 1 depicts the high level LVC system architecture. In this diagram, the components representing live, virtual, and constructive systems (shown as ovals) send

position updates of the aircraft they support and receive position updates for all other aircraft in the system. The LVC Gateways and HLA Toolboxes (shown as rectangles) distribute these data to the components that subscribe to the data including the Cockpit Display of Traffic Information (CDTI) displays residing at the traffic sources and air traffic control (ATC) displays (also oval). Many factors can affect the flow of data from the aircraft data sources through the LVC system components, including the type and architecture of the network, speed of the processors running the systems, and the way the components have been implemented. The LVC is not a static simulation environment, but a dynamic system that provides the infrastructure for connecting different air traffic components in order to emulate the relevant environment required for a simulation. These tests were designed to provide data to measure the performance of the existing core LVC prototype software and hardware by measuring the latencies between specific LVC components across different network architectures. Since an LVC instantiation can be constructed in many different configurations, the tests focus on determining the latencies between several component connection options and fall into four categories:

1. Latencies to publish live aircraft state data for distribution to the rest of the LVC
2. Latencies to publish virtual aircraft state data for distribution to the rest of the LVC
3. Latencies to publish constructive aircraft state data for distribution to the rest of the LVC
4. Latencies between distributed facilities.

The outcomes of these tests inform the development of future LVC instantiations for upcoming tests (flight and human in the loop simulations) and help determine whether changes to the existing LVC will be necessary to meet required performance characteristics. Eight distinct test configurations have been designed featuring combinations of simulated and live data across various distributed facilities; each is described in detail in later sections.




Figure 1. High-level view of the system under test

1.2 Overall Test Objectives

The goal of the LVC is to provide a simulation infrastructure that emulates an operational air traffic control environment able to mix live and simulated air traffic data. This goal provides the LVC development team with timing data that can be used to bound the latency results. Operationally, the maximum allowable latency is based on a combination of the surveillance source timing and the required time for processing and display at the facility. ATC Terminal facilities have 1.0 second processing requirement, while En-route facilities have a 1.6 second processing requirement.³ When combined with the radar sensor and communication timing this allows 2.2 seconds and 3.0 seconds for detection to display time to the Terminal and En route facilities, respectively.^{3,4} The maximum allowed generation and transmission time for ADS-B data is 2.5 seconds allowing for a total of 5.0 seconds for display in the cockpit.^{5,6}

In order to capture the data required to inform the latencies measurements of our four focus categories, the LVC Characterization test had two primary objectives:

1. Determine the time differential for each aircraft state data message produced by the aircraft data sources between when the message originated and when it was received at specific points within the LVC network.
2. Measure the throughput of the aircraft state data messages at specific points on the LVC system.

These objectives are not mutually exclusive. The second (or throughput) objective supports the first by providing the opportunity to measure latencies while increasing the number of aircraft in the system. The

anticipation is that increased aircraft messages will also increase the observed latencies, however, can a threshold where the differences in latencies are significant be determined? In addition, the LVC is built upon components that have traditionally supported scenarios of 80 or less aircraft. By increasing the traffic load during each test configuration, the intention is to investigate whether a throughput capacity can be determined.

The first not only provides the data to understand how long it takes to send an aircraft message to remote systems, but more importantly can be used to understand the time it takes to send messages between two specific components (in Figure 1, refer the LVC Gateway at each aircraft data source and the LVC Gateway Toolbox at the ATC Hub). In this way partial latency contributions between intermediate components can be used to build a unique LVC instantiation for a given set of requirements.

These objectives, when applied together along different points and among different LVC configurations, provide a general understanding of the system in terms of its ability to transmit the appropriate data in a timely manner. Due to the anticipated need to synchronize data from live, virtual, and constructive aircraft during testing, precise measurement of the latencies for these different air traffic inputs is critical.

2 Method and Materials

2.1 Test Resources

2.1.1 Software Components

This section provides background information on the LVC components that were used during the characterization tests.

2.1.1.1 Multi-Aircraft Control System and the Aeronautical Data Link and Radar Simulator

The Multi-Aircraft Control System (MACS) is a software program that can be configured to emulate either a pseudo pilot control station[†] or an air traffic control display. The MACS Simulation Manager (SimMgr) reads in a simulation file that specifies the flight path, flight intent, and starting position for a set of aircraft. It then generates flight trajectories for these aircraft and provides the LVC with position updates. For this series of tests, the SimMgr was run as a constructive aircraft data source, providing simulated aircraft data without pilot input. On the ATC side, MACS was configured to run a Display System Replacement (DSR) or Host emulation for test monitoring and display time logging for some of the test configurations. The Aeronautical Data Link and Radar Simulator (ADRS) is a companion program to MACS. It translates, filters, and transmits messages to and from instances of both the MACS SimMgr and MACS DSR.⁷ MACS and ADRS were developed at NASA Ames Research Center (ARC) for the purpose of Air Traffic Control simulation and are treated as Government off the Shelf (GOTS) software.

2.1.1.2 Cockpit Situation Display

The Cockpit Situation Display (CSD) is a software platform developed by NASA Ames to research concepts related to the display of information to a pilot.⁸ Many of the Human System Integration Ground Control Station research technologies have been tested via the CSD. For this test, the CSD was used for test monitoring and to provide a destination for messages to be sent. The CSD is considered GOTS software.

[†] The MACS pseudo pilot station allows for a single person to control and interact with many constructive aircraft to perform typical flight maneuvers. During a simulation this “pseudo pilot” would be communicating with the air traffic controllers for each aircraft under his/her control and maneuvering the aircraft based on the ATC direction. Reference 7 has a good description of this functionality.

2.1.1.3 High Level Architecture Middleware

The LVC used a version of the IEEE 1516 standard Pitch portable Real Time Infrastructure (RTI) High Level Architecture (HLA) and Federation Object Model (FOM) middleware, modified at NASA Ames, to exchange information about the air traffic environment (aircraft state, flight plans, digital messaging) among the participants operating from distributed facilities.^{9,10} The use of an HLA FOM provides an interface to well-defined air traffic data structures and promotes interoperability with simulation architectures using other middleware solutions such as AviationSimNet, Distributed Interactive Simulation (DIS), or Test and Training Enabling Architecture (TENA) FOMs.^{11,12,13} The HLA RTI ran at NASA Ames and served as the backbone for the LVC control capability, determining which components were connected and routed traffic information as appropriate. Simulation components are connected to the HLA via a Toolbox, which formats the messages as defined by the HLA interface.

2.1.1.4 LVC Gateway

The LVC Gateway was developed to allow connectivity to an external software component where connection directly to the HLA environment is not desired. For example, an external facility may have multiple components, each requiring communication to the HLA or directly to each other. Instead of each connecting remotely to the HLA, an LVC Gateway was used to route local message traffic and provide a single connection from the remote facility to the HLA at Ames. Components connecting to the LVC Gateway conform to a published interface control document (ICD) that automatically maps messages to the format expected by the HLA FOM. Any number of sites can be added to the LVC system by connecting additional LVC Gateways to the architecture.

2.1.1.5 HLA Toolboxes

While the HLA has a well-defined message interface, each software component connecting to the LVC for a simulation may have its own, which may not be consistent. Toolboxes are used to translate the messages from a software component to comply with the defined HLA interface. There are two primary reasons for the use of Toolboxes instead of developing the interface directly in the component software: the software component may be GOTS or COTS (i.e. the development team does not control the software in order to implement the interface); and the software component may be used to connect to multiple different versions of middleware. The LVC system utilized Toolboxes to connect to ADRS, the LVC Gateway, and the B747 flight simulator. The Toolboxes were designed to record and output the times messages are received; thus providing another physical location within the LVC network where data are time tagged.

2.1.1.6 Gateway Data Logger

The Gateway Data Logger process was developed to collect message timing and throughput data. The Gateway Data Logger connected directly to the LVC Gateway and stored the time an aircraft position update message was created by the sending process and the time the message was received by the LVC Gateway. The purpose for creating a separate process was to minimize the work required by the LVC Gateway during real-time data processing and its software complexity. During post-simulation processing, the Gateway Data Logger files were converted into time-series data sets for specialized analyses by researchers.

2.1.1.7 NASA Glenn Aircraft Ground Station and Map Display

Two specialized programs were developed to support and utilize live flight data specifically for Glenn flight-testing. The Aircraft Ground Station is a bridge between data from a live aircraft and the LVC system. It receives formatted telemetry data from a live aircraft via the prototype UAS datalink radio and relays the data to the LVC Gateway. The Map Display is a web-based application that provides an interface to a world map. It receives aircraft position updates from the LVC Gateway and overlays them on the map. It supported ground personnel by providing the location of the live asset during Communication flight-testing.

2.1.2 Live Resources

NASA Glenn's S-3B Viking aircraft, used for the Communication subproject's prototype radio testing as well as earlier channel sounding experimentation, served as the live asset for the characterization test. Data were collected during a planned Communication test flight in Northern Ohio, with the ground station located at the NASA Glenn UAS Communications Lab. Two programs acquire and transmit telemetry data from the aircraft. The Flight Data Aggregator program composes flight state messages from data monitored on the aircraft's MIL-STD 1553 and ARINC 429 data busses. The LVC Interface program sends the messages to the Data Link radio for transmission to the Ground Station.

2.1.3 Virtual Resources

The Ikhana UAS Simulator at DFRC and the B747 flight simulator at ARC provided virtual aircraft state data during specific test configurations.

2.1.4 Constructive Resources

The MACS SimMgr provided constructive aircraft state data for all the characterization tests. Virtual aircraft operated in the Cleveland Center airspace for the live flight configuration.

2.1.5 Test Facilities

Table 1. Facilities used during LVC Characterization testing and the type of network connection they have to the LVC Hub at DSRL

Facility	Location	Connection Type to DSRL
Distributed System Research Laboratory (DSRL)	NASA Ames	Not Applicable
Crew Vehicle Simulation Research Facility (CVSRF)	NASA Ames	Firewalled SimLab Network (SimNet)
Research Aircraft Integration Facility (RAIF)	NASA Dryden	Encrypted VPN via the NASA Integrated Services Network (NISN)
Simulation Lab	NASA Dryden	Encrypted VPN via the NISN
UAS Communications Lab	NASA Glenn	Encrypted VPN via the Internet

2.1.6 System Time Synchronization

Network Time Protocol (NTP) was used for clock synchronization between computer systems. NTP provides Coordinated Universal Time (UTC) including scheduled leap second adjustments. The NTP uses a hierarchical, semi-layered system of levels of clock sources. Each level of this hierarchy is termed a stratum and is assigned a layer number starting with 0 (zero) at the top. The stratum level defines its distance from the reference clock and exists to prevent cyclical dependencies in the hierarchy. It is important to note that the stratum is not an indication of quality or reliability. All computers on the SimLabs network were synchronized to a Stratum 1 or Stratum 2 time-server. At NASA Ames and NASA Glenn, the UTC source for the Stratum 1 time-server was provided by a Stratum 0 GPS device. An inter-range instrumentation group (IRIG) distribution amplifier (fed from a GPS antenna) was used at NASA Dryden.

The programs “tquery” and “timesvr” were used to baseline the time offset between computers residing in two remote facilities of the LVC environment prior to or after recording state data for a given configuration test. This informed the data analysis process, providing information as to whether the time offset was a significant factor in system latency calculations. The tquery program calculates the time it takes to send a message to timesvr and then receive the reply. The offset is calculated by comparing the average of the time a message tquery send and received times to the time the message was originally received by the remote timesvr program. The tquery program has the ability to run large sample pools to determine an average and standard deviation on the final difference output by the program.

2.2 Test Design

2.2.1 Measures of Performance (MOP)

MOP 1: Latency: The latency of each message between components of the LVC is measured by comparing the time the message was received at each of the components. If one of the components is the air traffic data source, the time the message was created is used.

MOP 2: Throughput: The throughput of the system is measured as a function of system load by successively increasing the system load and measuring the number of data messages that are successfully sent between LVC components.

2.2.1.1 Success Criteria

The success criteria for both the Latency and Throughput MOPS were the same, namely: Latency measurements were collected for each traffic sample. The first minute (+) of data from a run was discounted to account for system spikes associated with the gradual build-up of flight plan and state information for each set of aircraft. The beginning of the data collection run (T_0) was defined as the time when all flight plans for the load test were input into the system and each aircraft had at least three state data points. Enough data was considered to be collected when, starting at T_0 , a total of five consecutive minutes of traffic data were recorded at each planned point in the test configuration. Five minutes corresponds to the amount of time available to complete testing during the planned flight test. Since the test are looking at the latency of data message based on changing throughput, which does not change over time, five minutes was considered sufficient.

2.2.2 Characterization Configurations

In order to determine the appropriate test configurations, a high-level example of how the LVC could be used for a simulation (or use case) was developed. This use case demonstrated LVC integration capability of the proposed architecture topology as well as its scalability in the event the UAS Integration in the NAS Project expands the LVC test environment in the future. The initial requirements of the UAS-NAS LVC test environment included a constructive air traffic simulation, a virtual aircraft simulation, an emulated ATC environment and a live UAS (or surrogate UAS) aircraft. The core components of the LVC architecture are the HLA middleware and the in-house developed LVC Gateway that enable data exchange between different participants integrated with the LVC. Figure 2 shows a diagram of the overall use case from which each of the eight configurations included in the characterization test were derived. The LVC components outlined in red denote programs where data files were collected to support characterization. It should be noted that no instantiation of the LVC would utilize the system as diagramed in Figure 2; this is simply provided as reference to help with the understanding of how the test configurations are related.

The configurations cover eight variants for connecting live, virtual, and constructive traffic to the LVC system for display in an ATC environment. Each test configuration contains four simulation input scenario files of 50, 100, 200, and 400 aircraft used by the MACS SimMgr to generate simulated aircraft state data at a one-second update rate. A 50 aircraft input file was chosen as the low traffic sample because it represents a fairly light traffic load. From 50, the traffic load increments by factors of two up to 400, which

represents a traffic load that is well beyond what is anticipated for the planned simulations. The distinct test configurations are listed below and will be discussed in detail in subsequent sections of this document:

- 1.) Simulated Traffic to ATC on DSRL Network
- 2.) Simulated Traffic to ATC via HLA on DSRL Network
- 3.) Simulated Traffic to ATC via HLA between DSRL and CVSRF
- 4.) B747 Flight Simulator Data to DSRL via HLA
- 5.) Simulated Traffic to ATC via LVC Gateway between DSRL and CVSRF
- 6.) Simulated Traffic to GCS and Ikhana Simulator Data between DSRL and RAIF via NISN
- 7.) Simulated Traffic to GCS between DSRL and NASA Glenn UAS Communication lab via Internet
- 8.) Live Aircraft Data between S-3B Viking and LVC Gateway at NASA Glenn UAS Communication Lab

The following provides a list of LVC system features applied during the characterization tests:

- The HLA middleware as a core component of the LVC system was run at NASA Ames
- All processes included in a test configuration were run on separate machines (with the exception of the Gateway Data Logger, see next bullet)
- LVC Gateways were configured to include a Gateway Data Logger that captured source and receive times for each aircraft state message in the system (run on the same machine)
- The Cleveland Center airspace was used for all test configurations
- The same four constructive aircraft scenarios were used for each test configuration
- The MACS SimMgr was run for each test configuration to provide a control of system message throughput, whether the data was required for the configuration or not
- Each aircraft state data source provided a timestamp indicating when the position update was created
- The creation timestamp was carried forward through each of the LVC components to provide a common message identifier
- Each LVC component that was able record data (see Figure 2) provided the creation time of the state data along with the time the message was received and the size of the message
- The MACS DSR did not have a mechanism for recording the time aircraft state data are displayed.
- The CSD did not have a mechanism for recording the time aircraft state data are displayed.

Since the focus of the characterization is on the flow of data through the LVC core components, the missing display times, while inconvenient, are not critical. The display time requirements from References 1, 2, and 3 will still be used to indicate latency thresholds, with the understanding that additional latency buffer is required to account for the yet untested display times.




Figure 2. Diagram of all test components - test configurations are a subset of this set.

2.2.2.1 Configuration 1: Simulated Traffic to ATC on DSRL Network

2.2.2.1.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator and an ATC display. This provides the minimal baseline architecture to evaluate against, without the introduction of any LVC specific components. The intent of this test configuration is to measure latency between the times when aircraft state data is generated and when it reaches MACS DSR.

2.2.2.1.2 Configuration Methodology

The first test configuration was conducted in the DSRL at NASA Ames. As seen in Figure 3, the system included the MACS SimMgr, MACS DSR, and ADRS on three separate computers.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the MACS DSR via the ADRS. It then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. The MACS DSR is configured to output a file that contains the time the aircraft state information is generated by the SimMgr and the time the MACS DSR receives the state information. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.




Figure 3. Configuration 1: LVC system configuration used to determine internal MACS/ADRS latencies

2.2.2.1.3 Data Collected

Table 2. Data collected for Configuration 1

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	MACS DSR Output File
Time message received	MACS DSR	MACS DSR Output File
Size of message	MACS DSR	MACS DSR Output File

2.2.2.2 Configuration 2: Simulated Traffic to ATC via HLA on DSRL Network

2.2.2.2.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator and an ATC display routed through the HLA middleware on a local network. This configuration introduces the HLA middleware and associated toolboxes to the simulation infrastructure. The intent of this test configuration is to establish latency between the times when aircraft state information is generated and when it is received by the LVC (time to publish). This data also supports investigation of the latency of sending data through the HLA.

2.2.2.2.2 Configuration Methodology

This test configuration was also conducted solely in the DSRL lab at NASA Ames. The system included an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox. The MACS DSR was also connected to the HLA via a separate instance of ADRS and the ADRS Toolbox. In order to collect data to be compared to other configurations, an instance of the Gateway Data Logger was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. The diagram for configuration 2 can be seen in Figure 4.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the MACS DSR via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second.

Each ADRS Toolbox and the LVC Gateway (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.




Figure 4. Configuration 2: LVC system configuration used to determine latency added by the use of HLA

2.2.2.2.3 Data Collected

Table 3. Data collected for Configuration 2

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	MACS DSR Output File
Time message received	ADRS Toolbox #1	ADRS Toolbox Output File
Size of message	ADRS Toolbox #1	ADRS Toolbox Output File
Time message received	ADRS Toolbox #2	ADRS Toolbox Output File
Size of message	ADRS Toolbox #2	ADRS Toolbox Output File
Time message received	LVC Gateway	Gateway Data Logger File
Size of message	LVC Gateway	Gateway Data Logger File
Time message received	MACS DSR	MACS DSR Output File
Size of message	MACS DSR	MACS DSR Output File

2.2.2.3 Configuration 3: Simulated Traffic to ATC via HLA between DSRL and CVSRF

2.2.2.3.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator and an ATC display routed through the HLA middleware on a distributed network within a NASA Center. This configuration distributes the LVC onto separate networks. The intent of this test configuration is to establish the contribution the NASA Ames network to the HLA transit latency.

2.2.2.3.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab and CVSRF at NASA Ames. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. An instance of the Gateway Data Logger running at DSRL was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. The MACS DSR was also connected to the HLA via a separate instance of ADRS and the ADRS Toolbox all running in the CVSRF. The diagram for configuration 3 can be seen in Figure 5 and shows the connection between the ADRS Toolbox #2 running at CVSRF and the HLA running at DSRL.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the MACS DSR via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. Each ADRS Toolbox and the LVC Gateway (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.

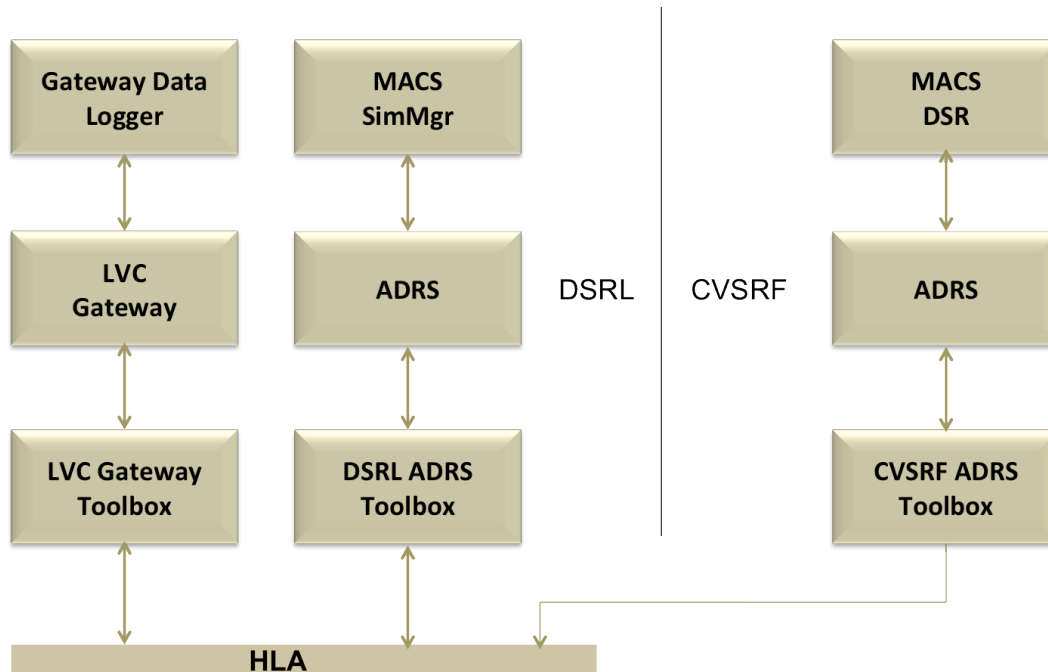


Figure 5. Configuration 3: LVC system configuration used to determine latency added due to the distribution of the LVC across different networks

2.2.2.3.3 Data Collected

Table 4. Data collected for Configuration 3

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	Gateway Data Logger File
Time message received	ADRS Toolbox #1	ADRS Toolbox #1 Output File
Size of message	ADRS Toolbox #1	ADRS Toolbox #1 Output File
Time message received	ADRS Toolbox #2	ADRS Toolbox #2 Output File
Size of message	ADRS Toolbox #2	ADRS Toolbox #2 Output File
Time message received	LVC Gateway	Gateway Data Logger File
Size of message	LVC Gateway	Gateway Data Logger File

2.2.2.4 Configuration 4: B747 Flight Simulator Data to DSRL via HLA

2.2.2.4.1 Configuration Objective

Measure and evaluate the latency of data messages sent from a flight simulator at an external facility through the HLA middleware running on a separate network at the same NASA Center. This configuration introduces the sending of traffic data from an external facility back to the virtual ATC system. The intent of this test configuration is to assess the latency of sending state data updates from a remote facility.

2.2.2.4.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab and CVSRF at NASA Ames. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. The MACS DSR was also running at DSRL and connected directly through the ADRS, in this configuration the MACS DSR was used only for test support. An instance of the Gateway Data Logger running at DSRL was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. The B747 simulator located in the CVSRF laboratory was connected to the HLA via the B747 Toolbox. The diagram for configuration 4 can be seen in Figure 6 and shows the connection between the B747 Toolbox running at CVSRF and the HLA running at DSRL.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the LVC via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. For this test configuration, the timing of the aircraft from the MACS SimMgr is not under test, but used to ensure the LVC Gateway and Toolboxes are processing a full complement of air traffic. The B747 flight simulator is flown for the duration of the test and sends position updates to the MACS DSR via the HLA. The ADRS Toolbox and the LVC Gateway (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. The B747 also receives the simulated traffic from the SimMgr, but for this test those data are not displayed. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.

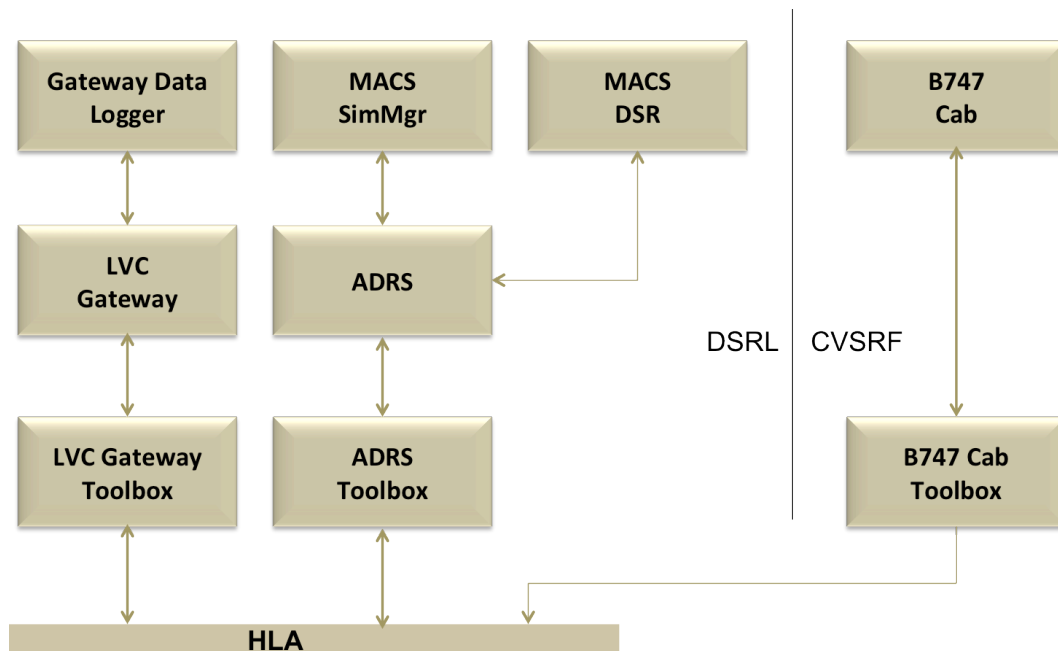


Figure 6. Configuration 4: LVC system configuration used to determine the latency of remote facilities sending aircraft data back to the LVC Hub

2.2.2.4.3 Data Collected

Table 5. Data collected for Configuration 4

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	Gateway Data Logger File
Time aircraft data is created	B747 Cab	Gateway Data Logger File

Time message received	ADRS Toolbox	ADRS Toolbox Output File
Size of message	ADRS Toolbox	ADRS Toolbox Output File
Time message received	LVC Gateway	Gateway Data Logger File
Size of message	LVC Gateway	Gateway Data Logger File

2.2.2.5 Configuration 5: Simulated Traffic to ATC via LVC Gateway between DSRL and CVSRF

2.2.2.5.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator routed through the HLA middleware on a local network and a remote LVC Gateway located at a distributed facility at the same NASA Center. The intent of this test configuration is to determine associated with the use of the LVC Gateway.

2.2.2.5.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab and CVSRF at NASA Ames. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. The MACS DSR was also running at DSRL and connected directly through the ADRS, in this configuration the MACS DSR was used only for test support. An instance of the Gateway Data Logger, also running at DSRL, was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. Figure 7 shows the connection between the LVC Gateway running at CVSRF and the LVC Gateway Toolbox running at DRSL.

During the test run, the MACS SimMgr reads simulation input files and sends the flight plan information for each aircraft specified in the file to the LVC via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. The ADRS Toolbox and the LVC Gateways (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.




Figure 7. Configuration 5: LVC system configuration used to determine whether the LVC Gateway adds any additional latency beyond the use of the specific HLA Toolboxes

2.2.2.5.3 Data Collected

Table 6. Data collected for Configuration 5

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	DSRL Gateway Data Logger File
Time message received	ADRS Toolbox	ADRS Toolbox Output File
Size of message	ADRS Toolbox	ADRS Toolbox Output File
Time message received	DSRL LVC Gateway	DSRL Gateway Data Logger File
Size of message	DSRL LVC Gateway	DSRL Gateway Data Logger File
Time message received	CVSRF LVC Gateway	CVSRF Gateway Data Logger File
Size of message	CVSRF LVC Gateway	CVSRF Gateway Data Logger File

2.2.2.6 Configuration 6: Simulated Traffic to GCS and Ikhana Simulator Data between DSRL and RAIF via NISN

2.2.2.6.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator routed through the HLA middleware on a local network and a flight simulator and cockpit traffic display located at an external facility connected via NISN. The intent of this test configuration is to establish HLA transit latency between laboratories at NASA Ames and NASA Dryden; it also introduces simultaneously sending state data for a single aircraft back to NASA Ames, further stressing the remote LVC Gateway.

2.2.2.6.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab at NASA Ames and the RAIF at NASA Dryden. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. The MACS DSR was also running at DSRL and connected to directly through the ADRS, in this configuration the MACS DSR was used only for test support. The CSD and the Ikhana UAS simulator were run at the RAIF, with their own instance of an LVC Gateway, and connected to the HLA via an LVC Gateway Toolbox running at DSRL. The diagram for configuration 6 can be seen in Figure 8 and shows the connection between the LVC Gateway running at the RAIF and the LVC Gateway Toolbox running at DSRL.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the LVC via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. The Ikhana UAS simulator is flown for the duration of the test and sends position updates to the MACS DSR via the HLA. The ADRS Toolbox and the LVC Gateway (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.

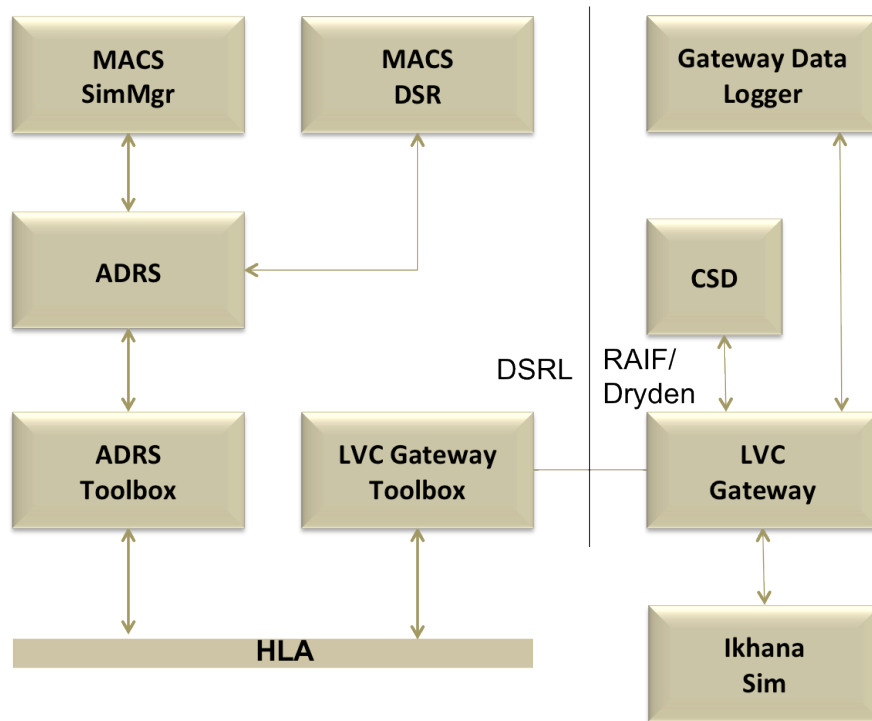


Figure 8. Configuration 6: LVC system configuration used to determine the latency when distributing across the NASA Integrated Services Network

2.2.2.6.3 Data Collected

Table 7. Data collected for Configuration 6

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	ADRS Toolbox Output File

Time message received	ADRS Toolbox	ADRS Toolbox Output File
Size of message	ADRS Toolbox	ADRS Toolbox Output File
Time message is sent	Ikhana UAS Simulator	Gateway Data Logger File
Size of message	Ikhana UAS Simulator	Gateway Data Logger File
Time message received	LVC Gateway	Gateway Data Logger File
Size of message	LVC Gateway	Gateway Data Logger File

2.2.2.7 Configuration 7: Simulated Traffic to GCS between DSRL and NASA Glenn UAS Communication lab via Internet

2.2.2.7.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator routed through the HLA middleware on a local network and a cockpit traffic display located at an external facility connected via the Internet. The intent of this test configuration is to determine latency pertaining to a distributed system running over a VPN via the Internet between NASA Ames and NASA Glenn.

2.2.2.7.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab at NASA Ames and UAS Comm lab at NASA Glenn. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. The MACS DSR was also running at DSRL and connected directly through the ADRS, in this configuration the MACS DSR was used only for test support. An instance of the Gateway Data Logger running at DSRL was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. The CSD was run at UAS Comm lab, with its own instance of an LVC Gateway, and connected to the HLA via an LVC Gateway Toolbox running at DSRL. The diagram for configuration 7 can be seen in Figure 9 and shows the connection between the LVC Gateway running at UAS Comm lab and the LVC Gateway Toolbox running at DSRL.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the LVC via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. The ADRS Toolbox and the LVC Gateways (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. It should be noted that since the CSD is GOTS, we are not able to modify the software to record when it received the data, instead the LVC Gateway running in UAS Comm lab records the time it receives the position messages. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.

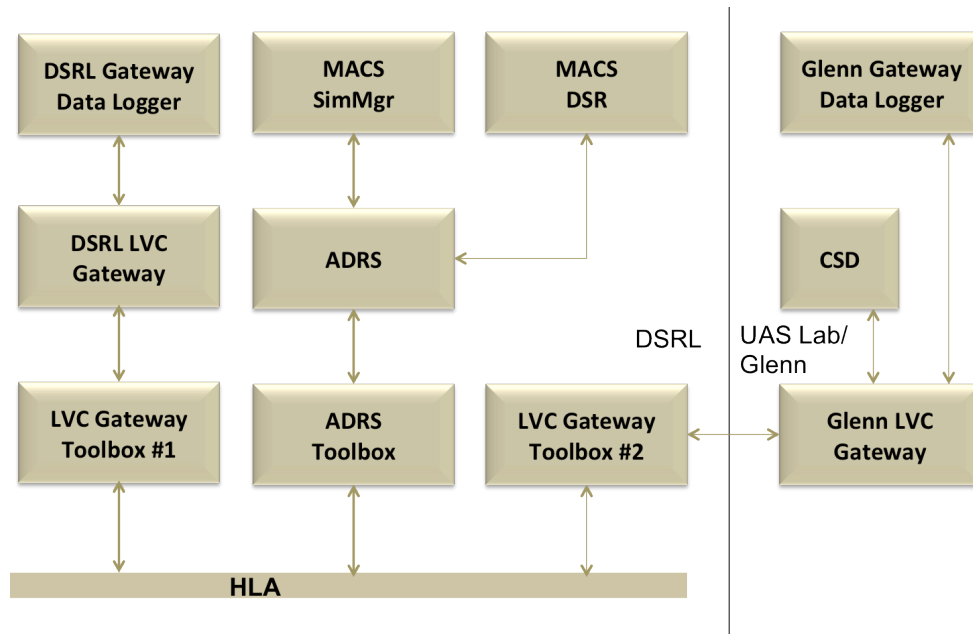


Figure 9. Configuration 7: LVC system configuration used to determine the latency when distributing a remote site via a VPN over the Internet

2.2.2.7.3 Data Collected

Table 8. Data collected for Configuration 7

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	DSRL Gateway Data Logger File
Time message received	DSRL LVC Gateway	DSRL Gateway Data Logger File
Size of message	DSRL LVC Gateway	DSRL Gateway Data Logger File
Time message received	UAS Comm LVC Gateway	Glenn Gateway Data Logger File
Size of message	UAS Comm LVC Gateway	Glenn Gateway Data Logger File
Time message received	ADRS Toolbox	ADRS Toolbox Output File
Size of message	ADRS Toolbox	ADRS Toolbox Output File

2.2.2.8 Configuration 8: Live Aircraft Data between S-3B Viking and LVC Gateway at NASA Glenn UAS Communication Lab

2.2.2.8.1 Configuration Objective

Measure and evaluate the latency of data messages under a series of four discrete throughputs between a constructive aircraft state data simulator routed through the HLA middleware on a local network and flight

data and traffic display located at an external facility connected via the Internet. The intent of this test configuration is to determine the latency of integrating live telemetry data from an aircraft into the LVC.

2.2.2.8.2 Configuration Methodology

The components of this test configuration were distributed between the DSRL lab at NASA Ames and UAS Comm lab at NASA Glenn. The system used an instance of the MACS SimMgr connected to HLA via the ADRS and ADRS Toolbox running at DSRL. The MACS DSR was also running at DSRL and connected directly through the ADRS, in this configuration the MACS DSR was used only for test support. An instance of the Gateway Data Logger running at DSRL was connected to the HLA via an LVC Gateway and LVC Gateway Toolbox. At NASA Glenn, the Map Display and Aircraft Ground Station programs were connected to an instance of the LVC Gateway running at the UAS Comm lab connected to the HLA via an LVC Gateway Toolbox running at DSRL. The diagram for configuration 8 can be seen in Figure 11 and shows the connection between the LVC Gateway running at UAS Comm lab and the LVC Gateway Toolbox running at DSRL.

During the test run, the MACS SimMgr reads in simulation input files and sends the flight plan information for each aircraft specified in the file to the LVC via the HLA. The MACS SimMgr then “flies” each simulated aircraft according to the flight plan information and outputs aircraft state data every second. For this test configuration, the timing of the aircraft from the MACS SimMgr is not under test, but used to ensure the LVC Gateway and ADRS Toolbox are processing a full complement of air traffic. The test is set up to correspond with a previously scheduled UAS Communication flight test. Glenn’s S-3B aircraft is flown for the duration of the test and sends position updates to the Aircraft Ground Station using a line-of-site R-C prototype radio (Figure 9).

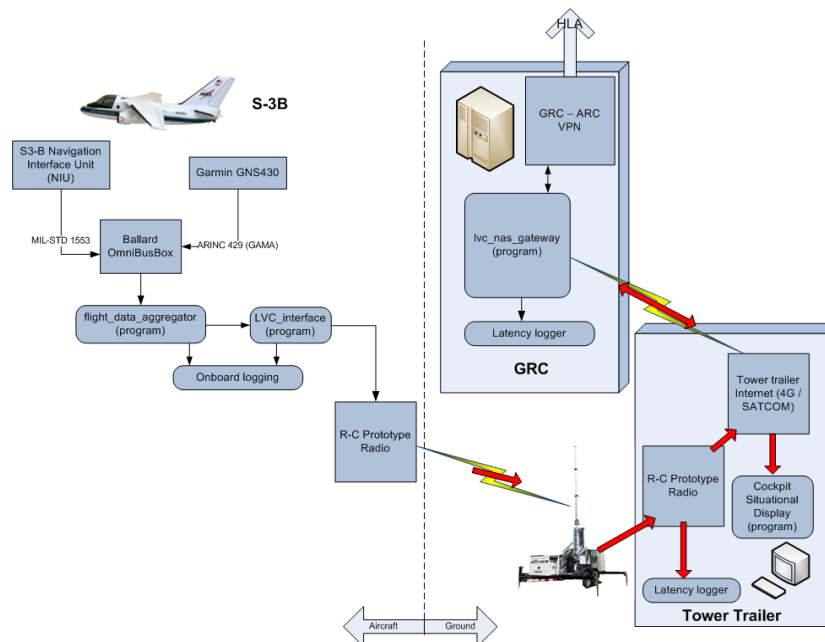


Figure 10. GRC S-3B Channel Sounding Flight Test Communications Architecture

These data are then sent to the MACS DSR via the HLA. The ADRS Toolbox and the LVC Gateways (via the Gateway Data Logger) record the time the aircraft state messages are received along with their associated creation time. In order to determine whether the traffic load has any impact on the system capacity or latency, the test was run with four separate simulation input files, each with an increasing number of active aircraft. The test load levels were 50, 100, 200, and 400 active aircraft.

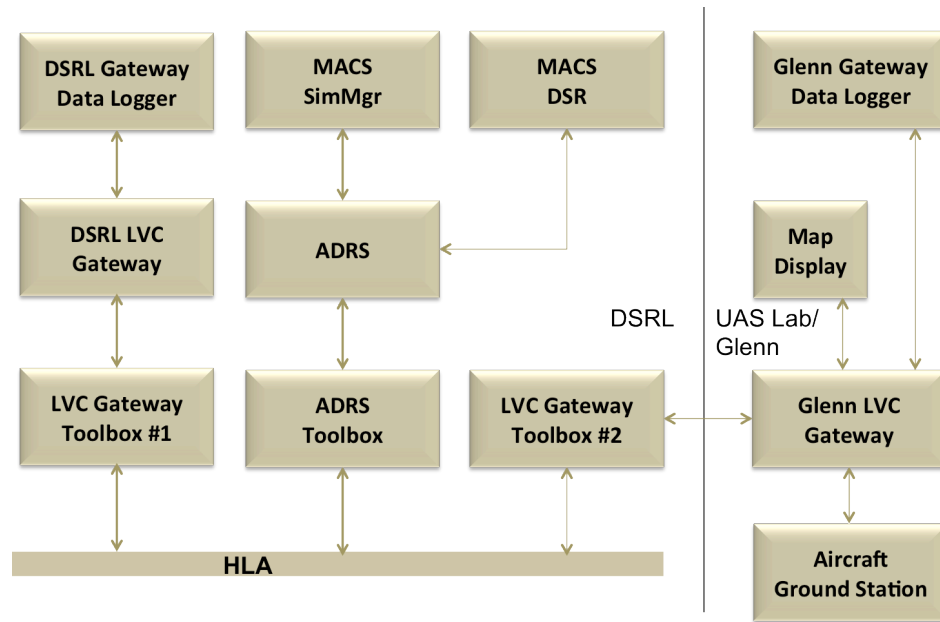


Figure 11. Configuration 8: LVC system configuration used to determine the latency of receiving live aircraft telemetry data

2.2.2.8.3 Data Collected

Table 9. Data collected for Configuration 8

Type of Data	Program Supplying Data	File Containing Data
Time aircraft data is created	MACS SimMgr	DSRL Gateway Data Logger File
Time message received	Live aircraft telemetry	DSRL Gateway Data Logger File
Size of message	Live aircraft telemetry	DSRL Gateway Data Logger File
Time message is sent	DSRL LVC Gateway	DSRL Gateway Data Logger File
Size of message	DSRL LVC Gateway	DSRL Gateway Data Logger File
Time message received	UAS Comm LVC Gateway	Glenn Gateway Data Logger File
Size of message	UAS Comm LVC Gateway	Glenn Gateway Data Logger File
Time message received	ADRS Toolbox	ADRS Toolbox Output File
Size of message	ADRS Toolbox	ADRS Toolbox Output File

3 Results

3.1 Configuration Results

The results presented in this section represent the analyses of interest for each of the specific test configurations. Comparisons of results across different test configurations are presented in Section 3.2.

3.1.1 Configuration 1: Simulated Traffic to ATC on DSRL Network

3.1.1.1 Test Conditions

This test configuration was the simplest (Figure 3) and provided the baseline latency times for sending data between the MACS SimMgr state data generator and the MACS DSR air traffic control display. Since HLA was not part of the configuration, native MACS output files were recorded for analysis. The computer system clocks were not explicitly checked for this test, primarily because the tquery process only runs on machines with the Linux operating system (while MACS runs on the Windows OS). However, daily testing of multiple similar machines (that are able to run the tquery process) on the same network have indicated a time difference below 1 millisecond on average. The time differences between machines in this LVC-DE configuration are assumed to be negligible.

3.1.1.2 Test Results

Table 10 and Figure 12 show the latencies observed between the times the state data were generated by the MACS SimMgr and when they were received by the MACS DSR. For each traffic load, five minutes of aircraft state data was analyzed (after the initial flight plan initialization period). For each data sample, the number of missing aircraft stat data messages is calculated by taking the number of expected stat data messages for the traffic sample and subtracted the actual number of messages received. As the traffic load increases, the average magnitudes of the latencies also increase only slightly between the 50 aircraft and 100 aircraft traffic loads, but more sharply for the 200 and 400 traffic loads. While on average the latencies are well below our emulation thresholds (2.2 seconds Terminal, 3.0 seconds en route), the maximum observed latency for the 400 aircraft traffic sample was greater than the required value. In addition, the number of state data updates that were dropped (as seen in Table 10) rises significantly for each traffic sample. In order to attempt to explain the missing state data messages and the relatively high data variability, the track histories of a few specific aircraft were analyzed.

Table 10. Observed Latencies between MACS SimMgr and MACS DSR

Traffic Load	SimMgr to DSR Average Latency (sec)	Standard Deviation (sec)	Max Latency (sec)	Missing State Updates Messages
50 Aircraft	0.694	0.214	1.55	385 (2.6%)
100 Aircraft	0.716	0.390	1.28	1231 (4.1%)
200 Aircraft	0.888	0.339	1.72	5200 (8.7%)
400 Aircraft	1.276	0.744	3.35	20135 (16.8%)




Figure 12. Observed latencies for direct MACS state data generation to display time

Figure 13 shows the track latencies for one of the simulated aircraft during the 200 aircraft traffic scenario. Although the average (0.83 seconds) and maximum (1.38 seconds) latencies for this aircraft fall within the acceptable range for surveillance data, its variability seems to follow a somewhat cyclic behavior. This behavior triggered additional analysis of how MACS handles the bundling and release of its state data. It was discovered that MACS generates state update messages at 4 Hertz, but not all of state updates are sent or logged. MACS sends state message to ADRS at a rate that is pre-set in the MACS configuration at start-up (default 1 second), however, MACS will not send a message until at least that amount of time has passed. It is believed this rule is responsible for the occasional major shift in latency that causes the large jump in the “sawtooth” pattern. However, the flight state sent is computed legitimately and falls well within the latency threshold for the display at the ATC controller’s position. It may be possible, if the MACS SimMgr is overburdened, that this same behavior explains the missing track information. The missing or dropped state data issue is further discussed under Configuration 2. It should be noted that this fluctuating track latency per aircraft is not anomalous, but seen throughout the aircraft in all traffic samples.




Figure 13. Observed latencies for a typical aircraft during the 200 aircraft traffic load

3.1.2 Configuration 2: Simulated Traffic to ATC via HLA on DSRL Network

3.1.2.1 Test Conditions

Configuration 2 connected a MACS SimMgr with a MACS DSR through HLA using two ADRS and associated ADRS Toolboxes. The test was conducted in the NASA Ames DSRL lab, with all processes run on local machines. Since all machines connect directly to the same NTP server, time differences are assumed to be negligible.

3.1.2.2 Test Results

Table 11, displays the mean latency times between state data generation by the MACS SimMgr and when the state data was published on HLA, as captured by the ADRS Toolbox. The average and maximum values are consistent for the 50 and 100 aircraft traffic loads. More variability is seen for the 200 aircraft sample, but the values continue to be well below the threshold display times (2.2 Terminal, 3.0 En-route). For the 400 aircraft traffic sample, a significant increase in the maximum observed latency is observed, which would be beyond the tolerance for Terminal airspace (even without adding the processing and display time). A significant number of missing or dropped state data messages is also observed.

Table 11. MACS publishing time to LVC (HLA)

Traffic Load	MACS publish to HLA Average Latency (sec)	Standard Deviation (sec)	Maximum Observed Latency	Missing State Updates Messages
50 Aircraft	0.812	0.021	0.851	0 (0%)
100 Aircraft	0.940	0.021	0.980	0 (0%)

200 Aircraft	1.025	0.215	1.377	1278 (2.13%)
400 Aircraft	0.9215	0.351	2.450	64758 (53.97%)

The issue with the missing or dropped state messages may be explained by analyzing the HLA transit time data. In order to calculate the results seen in Table 12, the time at the ADRS Toolbox handling the MACS SimMgr is compared to the time seen at the ADRS Toolbox handling the MACS DSR process for each message (refer to Figure 4). These will be called SimMgr Toolbox and DSR Toolbox, for convenience. As can be seen, the HLA transit times are on average well under 100 milliseconds and have relatively little variability. This is to be expected, since the SimMgr and DSR Toolboxes were run on the same network. However, Table 12 also provides the total number of messages recorded for the 5-minute test sample. For the 50 and 100 aircraft traffic loads, no missing or dropped messages were detected. The 200 aircraft sample has a total of 1278 missing messages of which 425 were missing from the SimMgr Toolbox and an additional 853 were missing from the DSR Toolbox. For the 400 aircraft traffic sample, the numbers go to 64,758 total missing messages, 6084 missing from the SimMgr Toolbox and an additional 58,674 missing from the DSR Toolbox.

The reason for the SimMgr Toolbox missing messages is under investigation. The MACS SimMgr and ADRS programs were not the focus of these analyses and not modified for the tests. The LVC development team is now instrumenting both software baselines to determine the cause. However, additional missing messages as detected by the DSR Toolbox can be partially explained. Each of those missing messages can be traced to a message in the SimMgr Toolbox that is a duplicate of a previous message. The SimMgr Toolbox detects the repeated message and does not forward to HLA. The cause for the duplicate messages from MACS SimMgr is also under investigation. Since the dropped and missing messages are directly attributed to the behavior of MACS SimMgr, the results of those analyses will no longer be presented, while the latencies of the messages received will be continued.

Table 12. HLA Transit Times (local network)

Run	HLA Transit Time (sec)	Std Dev (sec)	Number of Expected State Messages	Number of Messages at ADRS Toolbox #1	Number of Messages at ADRS Toolbox #2
50 Aircraft	0.019	0.008	15000	15000	15000
100 Aircraft	0.032	0.013	30000	30000	30000
200 Aircraft	0.049	0.024	60000	59575	58722
400 Aircraft	0.056	0.033	120000	113916	55242

3.1.3 Configuration 3: Simulated Traffic to ATC via HLA between DSRL and CVSRF

3.1.3.1 Test Conditions

Configuration 3 used an identical architecture to Configuration 2. However the MACS DSR station was run from a geographically separated ARC lab operating on the same network domain as the DSRL. The intention was to measure whether routing the network traffic through multiple routers, but still on the same local network, had any significant latency impact.

3.1.3.2 Test Results

Since the MACS publishing times have already been investigated in the previous two test configurations, this configuration focuses on the HLA transit time. As before, since HLA does not directly record message times, the message times as recorded in the DSRL ADRS Toolbox and CVSRF ADRS Toolbox serve as proxy. In addition, this allows us to conveniently analyze the transit time of messages across the separate facilities. The resulting transit times, presented in Table 13, show average latencies for each of the traffic samples that increase with increased traffic load along with the standard deviations and the maximum observed latency. Of interest are not the averages, but the maximums. Though below the total required latency thresholds, when building an instance of the LVC for a simulation the latencies are incremental.

Table 13. HLA Transit Times between DSRL and CVSRF

Traffic Load	HLA Transit Time (sec)	Std Dev (sec)	Maximum Observed Latency (sec)	Number of Messages
50 Aircraft	0.038	0.040	0.311	14882
100 Aircraft	0.062	0.067	0.600	29934
200 Aircraft	0.127	0.167	1.200	57277
400 Aircraft	0.143	0.134	1.231	55648

As the number of aircraft in the traffic sample increases, the latency also increases. However, an interesting condition can be seen in the 400 aircraft case; the number of messages actually passed between the facilities goes down as compared to the 200 aircraft traffic sample (because of duplicate tracks as described above). With fewer messages to handle, the expectation is for the transit times to decrease. The reason for the seemingly odd behavior may be attributed to the time the DSRL ADRS Toolbox is required to handle and determine whether messages should be distributed further. Since the DSRL ADRS Toolbox handles nearly twice as many messages in the 400 traffic sample as the 200 traffic sample, even though less are distributed, the processing takes more time.

3.1.4 Configuration 4: B747 Flight Simulator Data to DSRL via HLA

3.1.4.1 Test Conditions

Configuration 4 extended the previous configuration by adding the Boeing 747 Level D flight simulator to measure latency across the LVC from a high fidelity simulation sending data at a rate of greater than 1 Hertz. In this case, the B747 simulator published state data to HLA at a rate of 5 Hz.

3.1.4.2 Test Results

Table 14 shows the effect of increased MACS traffic loads on the latency of B747 state data transit times to HLA. It should be noted that data recording was not available in the B747 Toolbox at the time of the configuration test. As a result, the transit times include the time from CVSRF (location of the B747 simulator) and DSRL (location of HLA). Although some additional latency can be noted, the effect was not significant even at the highest MACS traffic load and increased transmission rate from the simulator. Also of note is the number of missing B747 state data messages; accounting for only 5 messages (one during the 200 aircraft traffic sample, four during the 400 aircraft traffic sample) out of 6000 total messages. The B747 Toolbox will be instrumented and tested again to determine if the reason for the few missing target can be found.

Table 14. B747 Latency to HLA

Run	747 Time to	Standard	Max Latency	Missing State Updates
-----	-------------	----------	-------------	-----------------------

	HLA (sec)	Deviation (sec)	(sec)	Messages
50 Aircraft	0.028	0.011	0.062	0 (0.0%)
100 Aircraft	0.038	0.013	0.121	0 (0.0%)
200 Aircraft	0.047	0.017	0.218	1 (0.1%)
400 Aircraft	0.057	0.018	0.279	4 (0.2%)

Figure 14 depicts the plot of the B747 state data to HLA during the 400 aircraft scenario run. The data points are highly clustered around the mean (0.057 seconds), with a few state data messages showing significantly higher latencies. This is expected, since there is a hard lower bound to the latency, while some messages may encounter network delays. Note as described above, the B747 HLA latencies include the time to send data from CVSRF to DSRL.

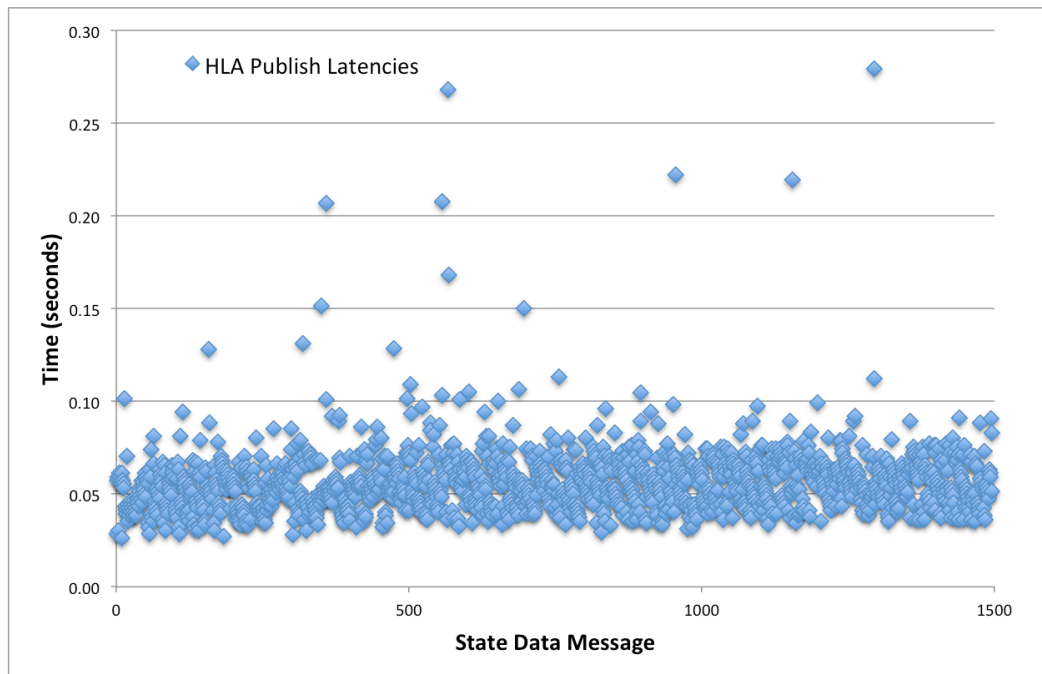


Figure 14. Observed latencies during the 400 Aircraft Scenario of the B747 state data including time to the HLA at the DSRL

3.1.5 Configuration 5: Simulated Traffic to ATC via LVC Gateway between DSRL and CVSRF

3.1.5.1 Test Conditions

Configuration 5 was the first to measure the latency of MACS data sent to a remote LVC Gateway. This was the first step to test latencies associated with connecting a remote LVC Gateway to LVC components, which allows for connection of multiple remote components. To simplify the configuration, there were no data producer or consumer components connected to the LVC Gateway other than the Gateway Data Logger, which recorded the arrival time of MACS data at the LVC Gateway.

3.1.5.2 Test Results

Table 15 lists the transit times to send data from the DSRL laboratory to CVSRF via HLA and an LVC Gateway. The latencies continue to follow the similar patterns established in earlier test configurations. The 400 traffic load scenario again has significant data drop-out associated with the SimMgr duplicated state data updates. Similar behavior is observed with the average LVC transit time for the 400 aircraft traffic sample greater than the 200 aircraft traffic sample, but fewer number of state data messages actually handled. As before, it is expected that the processing of the duplicated state data from the SimMgr that is dropped by the DSRL ADRS Toolbox adds to the increase in HLA transit time.

Table 15. Transit Time between HLA in DSRL and LVC Gateway in CVSRF

Traffic Load	HLA Transit Time (sec)	Standard Deviation (sec)	Max Latency (sec)	Missing State Updates Messages
50 Aircraft	0.061	0.027	0.137	0 (0.0%)
100 Aircraft	0.102	0.054	0.244	80 (0.3%)
200 Aircraft	0.179	0.105	0.604	1724 (2.9%)
400 Aircraft	0.208	0.131	0.802	64563 (53.8%)

3.1.6 Configuration 6: Simulated Traffic to GCS and Ikhana Simulator Data between DSRL and RAIF via NISN

3.1.6.1 Test Conditions

Configuration 6 was designed to test latency and throughput issues associated with data exchange between two geographically separated locations (NASA Ames and the NASA Dryden) over a wide area network—in this instance the NASA Integrated Services Network (NISN). It also is the first test conducted with a remote LVC Gateway that is serving both an information producer (Ikhana Sim) and consumer (CSD). The Ikhana Sim published data at a 1Hz rate. The NTP servers at both locations used GPS derived UTC time. The computed average offset between the uasgw3 computer running the LVC Gateway Toolbox at Ames and the LVC Gateway computer at Dryden was 0.016 seconds, which resulted in no significant impact on latency calculations between the two sites.

3.1.6.2 Test Results

Table 16 lists the mean times and standard deviations of sending data between the ADRS Toolbox at the DSRL at NASA Ames and the LVC Gateway at the RAIF lab at NASA Dryden. The values were calculated by subtracting the time each message reached the LVC Gateway at Dryden from the time the messages was first received by the ADRS Toolbox in the DSRL lab (factoring in the computer clock difference, which was calculated by the tquery process run at the beginning of the tests) and averaging over the 5-minute sample. The results of this test follow the similar pattern of monotonic increases in transit time as the traffic load increases. However, this test configuration did not have a lower number of messages passed between the ADRS Toolbox and the receiving LVC Gateway when comparing the 200 and 400 traffic samples. The average and maximum observed transit times continue to remain below the required operational detection to display latency. Direct comparison of transit time among the different facilities will be present in Section 3.2.

Table 16. HLA Transit Times between DSRL at NASA Ames and RAIF at NASA Dryden

Traffic Load	HLA Transit Time (sec)	Standard Deviation (sec)	Max Latency (sec)	Missing State Updates Messages
50 Aircraft	0.102	0.034	0.177	235 (1.6%)
100 Aircraft	0.157	0.064	0.286	862 (2.9%)
200 Aircraft	0.259	0.125	0.519	5143 (8.6%)
400 Aircraft	0.335	0.201	0.945	61560 (51.3%)

The observed latency times required to publish Ikhana state data to the Dryden LVC Gateway were extremely consistent across all four test runs. The values provided in Table 17 indicate that the increasing the total number of MACS state data had no effect on latency at the Dryden LVC Gateway.

Table 17. Ikhana State Data Publishing Time to Dryden LVC Gateway

Run	Ikhana to Dryden GW Latency (sec)	Standard Deviation (sec)	Max Latency (sec)	Number of Missing State Data Messages
50 Aircraft	0.0216	0.002	0.097	0 (0.0%)
100 Aircraft	0.0215	0.001	0.033	0 (0.0%)
200 Aircraft	0.0215	0.001	0.033	0 (0.0%)
400 Aircraft	0.0214	0.001	0.029	0 (0.0%)

3.1.7 Configuration 7: Simulated Traffic to GCS between DSRL and NASA Glenn UAS Communication lab via Internet

3.1.7.1 Test Conditions

The Configuration 7 tests were used as a baseline to test latency between the Ames and Glenn laboratories over an internet-based virtual private network (VPN). The NTP servers at both locations used GPS derived UTC time. The computed average offset between the systems was -0.297 seconds.

3.1.7.2 Test Results

Table 18 the mean times and standard deviations of sending data between the ADRS Toolbox at the DSRL at NASA Ames and the LVC Gateway at the UAS Communications lab at NASA Glenn. The values were calculated by subtracting the time each message reached the LVC Gateway at Glenn from the time the messages was first received by the ADRS Toolbox in the DSRL lab (factoring in the computer clock difference, which was calculated by the tquery process run at the beginning of the tests) and averaging over the 5-minute sample. The results of this test follow the similar pattern of monotonic increases in transit time as the traffic load increases. Direct comparison of transit time among the different facilities will be present in Section 3.2.

Table 18. Transit time between Ames HLA and Glenn LVC Gateway

Run	LVC Transit Time (sec)	Standard Deviation (sec)	Max Latency (sec)	Number of Missing State Data Messages
50 Aircraft	0.142	0.075	0.838	2 (0.01%)
100 Aircraft	0.204	0.106	2.045	6 (0.02%)
200 Aircraft	0.292	0.178	1.483	2456 (4.1%)
400 Aircraft	0.337	0.240	2.258	64574 (53.8%)

3.1.8 Configuration 8: Live Aircraft Data between S-3B Viking and LVC Gateway at NASA Glenn UAS Communication Lab

3.1.8.1 Test Conditions

Configuration 8 introduced a live S-3B Viking aircraft in flight sending state data to the Glenn Ground Station, which published it to the LVC Gateway also located at NASA Glenn. At the same time, the Glenn LVC Gateway was connected to the HLA running at NASA Ames, which received the state data from the live aircraft while simultaneously publishing simulated flight state data back to Glenn.

3.1.8.2 Test Results

The various latencies associated with the sending and receiving of the live aircraft data was the primary focus of this test configuration. The simulated traffic data generated at NASA Ames served to provide the LVC Gateways with varying traffic loads in order to evaluate any impact on the latency of the live data. Latencies for sending simulated data between NASA Ames and Glenn are documented in Configuration 7.

Table 20 shows the timing of sending the data from the S-3B aircraft to the LVC Gateway running at NASA Glenn. Notice that the average latencies from the live aircraft to the LVC are much greater than the corresponding latencies seen in the B747 Flight Simulator and the Ikhana Simulator from previous configurations (see Figure 23). This is as expected, since the live aircraft must transmit the position update via line-of-sight 3G Cellular transmission. However another feature of the data also played an important role, as the analysis showed that the creation time of the position message from the S-3B was truncated to the nearest second. This means the actual latency may be as much as 0.999 seconds less than the calculated value.

Table 19. S-3B State Data Publishing Latency to Glenn LVC Gateway

Simulated Traffic Load	S-3B to Glenn LVC Gateway (sec)	Standard Deviation (sec)	Min Latency (sec)	Max Latency (sec)	Number of Missing State Data Messages
50 Aircraft	2.450	0.091	2.403	3.234	2 (0.6%)
100 Aircraft	2.592	0.097	2.497	3.324	3 (1.0%)
200 Aircraft	2.111	0.169	2.003	3.440	2 (0.6%)
400 Aircraft	2.062	0.124	1.998	2.963	1 (0.3%)

400 Aircraft (During buffering)	2.242	0.886	1.923	9.123	2 (0.6%)
------------------------------------	-------	-------	-------	-------	----------

Another interesting feature seen during the live flight data collection was a period of message buffering observed in the data during collection. The last row of Table 19 provides data that was collected after the primary data collection for the 400 aircraft traffic sample. During this time period, position messages from the aircraft were buffered before being recorded in the Glenn LVC Gateway. This is seen in the greater standard deviation and unusually high maximum observed latency. The location of the buffering detection indicates it occurs either in the transmission of the data to the ground antenna from the aircraft or during the transmission of the data over the 3G cellular connection to the LVC Gateway. The data recording does not provide enough detail to determine the exact source of the buffering. Figure 15 shows how this buffering happens at discrete times, and then the latencies go back to a nominal magnitude.

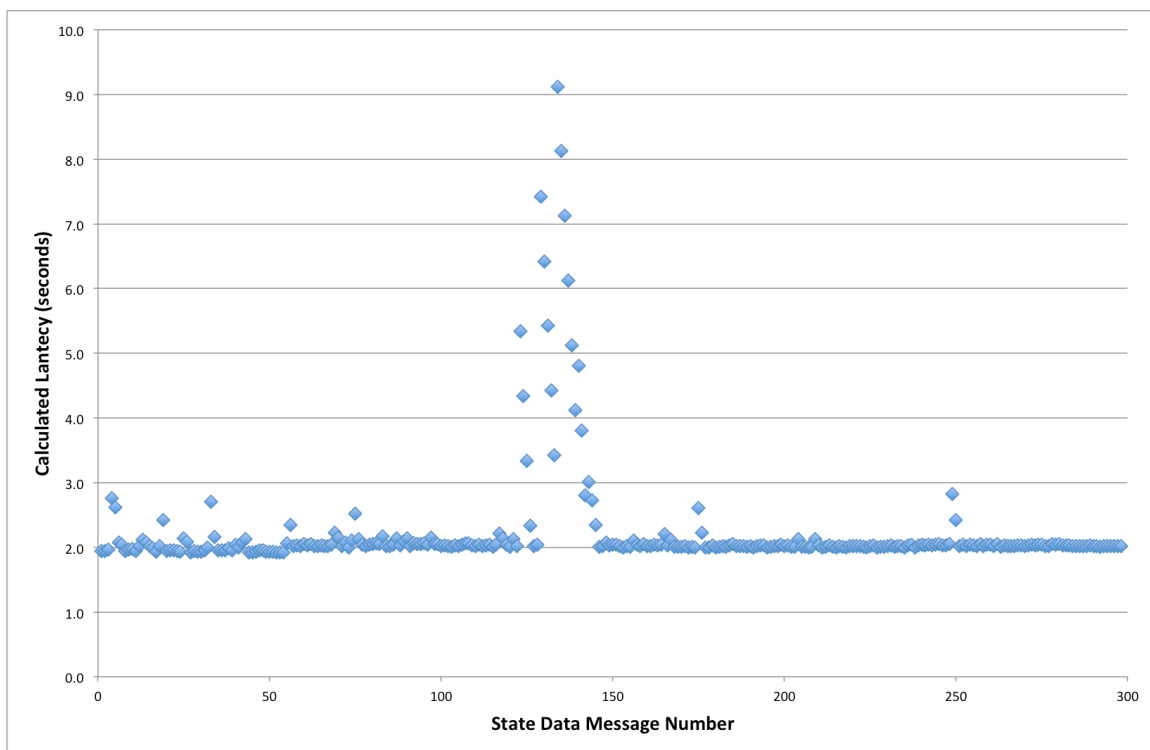


Figure 15. Example of Discrete Buffering from Live Aircraft

3.2 Compilation Results

The results detailed in this section provide further insight by comparing latencies across selected test configurations. These results cover our four primary focus areas:

- 1.) Latencies to publish live aircraft state data for distribution to the rest of the LVC
- 2.) Latencies to publish virtual aircraft state data for distribution to the rest of the LVC
- 3.) Latencies to publish constructive aircraft state data for distribution to the rest of the LVC
- 4.) Latencies between distributed facilities.

3.2.1 Publishing Live Aircraft State Data to the LVC

Since we only had access to the S-3B during our data collection the live data results do not differ from those in section 3.1.8. As described above, due to the truncation of the state data creation time, these values may be greater than the actual value by as much as 1 second. Table 20 shows the live flight data accounting for the maximum possible error. The purpose of these numbers isn't to suggest that obtaining these observations is achievable, but to show that even with the best case, the latencies would be very large. Even without accounting for transit to the ATC machines, processing and display times, they are already greater than the required maximum for the Terminal airspace (2.2 seconds for radar data).

Table 20. S-3B State Data Publishing Latency to Glenn LVC Gateway (Best Case)

Simulated Traffic Load	Observed Min Latency (sec)	Best Possible Minimum Latency (sec)	Observed Max Latency (sec)	Best Possible Max Latency (sec)
50 Aircraft	2.403	1.404	3.234	2.235
100 Aircraft	2.497	1.498	3.324	2.325
200 Aircraft	2.003	1.004	3.440	2.441
400 Aircraft	1.998	0.999	2.963	1.964

It should be noted that the transmission of the S-3B telemetry data including routing the data from the computers receiving the messages on the ground to the LVC via a 3G cellular connection. Cellular transmissions can be prone to high latencies, depending on bandwidth usage.

3.2.2 Publishing Virtual Aircraft State Data to the LVC

Test Configurations 4 and 6 each included the publishing of virtual traffic data for distribution to the rest of the LVC. Figure 16 provides a comparison between the local publishing latencies of the Ikhana Sim (red) and the B747 (blue). This shows that the latency is monotonically increasing as the traffic load doubles. While the data from the Ikhana Simulator is incredibly precise, with almost no variability. However, recall that the B747 publish data includes time to transit between CVSRF and DSRL due to data not being recorded by the B747 Toolbox. A likely explanation under investigation is that the ADRS Toolbox recording the B747 data once it reaches the DSRL slows as it processes the simulated data for distribution. The implication of these results is that while increased traffic loads may not impact remote facilities connected to the LVC, latencies of the data from the remote facilities may have different observed latencies based on the underlying simulation traffic. From these results, however, the difference is almost negligible.




Figure 16. State Data Publishing Latency, Ikhana vs B747

3.2.3 Publishing Constructive Aircraft State Data to the LVC

The analyses uncovered two critical features of the MACS SimMgr state data generation capabilities that will need to be further investigated and understood before using the software for simulation:

- 1.) The latency of publishing data to the MACS DSR through the ADRS and to the HLA via the ADRS Toolbox is highly variable.
- 2.) Missed and or duplicated state messages increase as the traffic load increases.

Both of these issues may be explained by the same software mechanism. The MACS SimMgr generates state data at a specified update rate (default is 4 times per second) to be used by the process as needed. A separate functionality in the software tracks how often the position of the aircraft needs to be updated (e.g. every 12 seconds to emulate en route radar data, or every 1 second to emulate ADS-B data). If the MACS SimMgr is unable to generate the state data for all aircraft within the time limit of the update rate, it could fall behind. The result would be dropped or duplicate tracks, and a greater time difference between the when the state data was generated and when it was sent out. The LVC development team is investigating this supposition in order to determine whether to mitigate this issue. Note: during a simulation a single MACS SimMgr would not be required to update the position of more than 50 aircraft. Large scenarios are distributed among several pseudo pilots. However, this may still be an issue, the high variability and dropped or duplicate tracks were also observed in the 50 aircraft traffic samples.

Figure 17 shows the average latency times of publishing the MACS SimMgr state data to the HLA (blue). Plus or minus two standard deviations are also represented in red and green, showing the large increase in variability of publishing the MACS SimMgr state data based on the aircraft traffic generation load.




Figure 17. State Data Publishing Latency for the MACS SimMgr with Standard Deviations

3.2.4 Documentation of Latencies Between Tested Facilities

The latencies of publishing constructive state data across the HLA to external facilities can best be examined by comparing the data from Test Configurations 2, 3, 5, 6, and 7. This provides a wide range of network and LVC configurations that may impact the HLA transit timing. Figure 18 shows each of these data plotted on the same graph. The data behave as we would expect. The first (lowest) line represents the latencies observed when running two instances of the ADRS toolbox connected to the HLA all on the same network. All other things being equal, this should be as fast as the system can run, and it is corroborated by the data. The next two data lines both represent connections the remote facility located at Ames. This should be faster than connecting to a network outside of Ames as the data shows. The difference between the instances was the process running at CVSRF and connected to HLA (see Figure 5 and Figure 7). The better performing instance (Configuration 3) ran an ADRS Toolbox at CVSRF connected directly to the HLA at DSRL. The slightly slower performing Configuration 5 ran an LVC Gateway at CVSRF connected to an LVC Gateway Toolbox connected to HLA running at DSRL. The reason the different configurations are of interest is because a Gateway allows for connection of multiple components at the remote facility, which provides the greater flexibility in system design. The observed time improvements when using the ADRS Toolbox may be a factor of one less process to send data through, an internal LVC Gateway functionality, or a combination of both. Since both Configurations performed similarly, investigating the difference is not a high priority. The final two data lines represent the average HLA transit latencies associated with send data to NASA Dryden via NISN and NASA Glenn via a VPN over the internet. While the NISN connection performed slightly better, at low traffic loads, both provided reasonable latency values. As traffic loads rose to 200 and 400 aircraft, the latencies continued to climb, but not beyond a value that would jeopardize the requirements..




Figure 18. HLA Transit Latency, CVSRF vs Dryden vs Glenn

3.2.5 Combined Analyses

Each of the Test Configurations provides latency and timing information for a piece of a possible LVC system. However, no one Configuration provides the complete end-to-end latency data. The primary reason for this was because during preparation for these analyses, the LVC requirements were still under consideration. In addition, the MACS DSR software does not yet have the necessary functionality to record when an aircraft state update is displayed. Just as an LVC system can be developed by combining the appropriate system components, the data from the eight test configurations can be combined to provide information on different LVC system architectures that are feasible and those that may not support a relevant simulation.

Figure 19 shows the average cumulative latencies calculated from each state data source to the HLA. It also plots the Terminal and Host radar detection to display thresholds (2.2 and 3.0 seconds respectively), since they are the most restrictive. For each bar, the latencies include the time to publish the state data to the local Toolbox, the time to transit from the remote facility to the DSRL, and the time to transit the HLA (fixed based on Configuration 2). Missing from the calculation is the time from HLA to the MACS DSR for display (since this is not yet available). Instead the time left over between the top of the bar and the corresponding operation threshold represents the amount of time available to complete the display of the traffic. Based on the MACS SimMgr performance, we are estimating no more than 0.5 seconds for display. The first bar represents the time latency for live aircraft. As expected the publishing time accounts for the majority of the latency and goes beyond the available latency to be usable for Terminal radar emulation. In addition it is only 0.18 seconds below the En-route threshold. Without more accurate publishing time data, use of the S-3B given the tested method of data transmission would be a risk that would have to be mitigated. The second and third bars represent the Ikhana and B747 latencies. Both are well below the Terminal and Host latency thresholds. Assuming the time to process and display the data on the MACS DSR is reasonable, these data source will be below the latency requirement. The last bar represents the MACS SimMgr constructive state data latencies. As expected the MACS publishing time accounts for the majority of the total latency. Even with these issues and accounting for the MACS DSR processing and display time, these data should be well below that required latency.




Figure 19. End-to-End Latency Comparison based on Average Latencies

While Figure 19 showed the average cumulative latencies: in order to have confidence the system can perform properly, Figure 20 plots the same data including a two standard deviation buffer for each of the latency contributors. While the virtual and constructive data sources are still well below the latency thresholds, with the standard deviation buffer, the S-3B live data is not. Given the current latency data, in order to effectively coordinate live, virtual and constructive traffic, it would be ideal for some delay to be added to the virtual data sources. Referring to Figure 20, in order to emulate all of the data coming from a common data source (as it would operationally), the goal would be for all of the data sources to ultimately have the same amount of latency as it arrived at its destination. In this case the source with the largest latency would set the threshold, if possible the other data sources would need to have their messages delayed by the appropriate amount to provide the “common” latency.




Figure 20. End-to-End Latency Comparison based on Average Latencies plus 2 Standard Deviations

4 Conclusion

The purpose of the LVC is to provide a simulation infrastructure bringing live and simulated components together to provide a relevant environment for the upcoming UAS in the NAS integrated events. The LVC is designed to be distributed so that assets are not required to be co-located for a simulation or flight test. The distributed nature of the proposed LVC environment is the purpose for conducting the LVC Characterization: Timing of input data and message passing between the distributed facilities needs to be well understood in order to develop an instantiation of the LVC that meets the requirements of each event. This includes emulating the operational requirements of surveillance data, and ensuring data from disparate sources are properly synchronized.

The LVC Characterization timing tests were designed to determine whether the existing aircraft state data generation capabilities were able to meet the known operational air traffic control system requirements, stress the LVC system to determine the limits of its capabilities and inform the design of event scenarios, and determine whether data from any of the types of data sources has unique features that may need to be mitigated during an event.

One of the first results from the testing demonstrates that the existing instantiation had significant problems handling very high traffic loads. This was primarily seen in the 400 aircraft traffic load tests, but was also an issue during several 200 aircraft traffic runs. The LVC system had several instances of dropped and duplicated data messages that were traced to the LVC processes that send data between different facilities and the MACS SimMgr, which produces the constructive state data messages. Scenarios with more than 100 active aircraft are not anticipated in this project, so this should not be an issue during the integrated events. The LVC system capacity will be investigated further, however, to understand whether throughput can be increased in the future.

Configuration 1 provided a baseline test for the latency of the virtual traffic provided by the MACS SimMgr state data generator. The 50 and 100 traffic load scenarios for this test yielded higher than expected latencies between the time the state data was generated and when it was available for display in the MACS DSR display (near 0.7 seconds, Table 10). Further analysis of the data (see Figure 13) showed a “sawtooth” cyclic nature to the latency times. Further investigation of the queuing and buffering mechanism in the MACS SimMgr and ADRS are required to fully understand this behavior. MACS and ADRS must first be understood then fixed or mitigated. Even with these observed behaviors, the average (and maximum) latencies were still below the operational required values of 2.2 seconds for Terminal airspace and 3.0 seconds for en route airspace and should not pose an issue for the use of the MACS SimMgr.

The addition of HLA and its associated Toolboxes in Configuration 2 did not appreciably affect the observed latencies. This supports the continued usage of HLA as a middleware solution, which is required for the more complex LVC instantiations where multiple LVC components are connected at a single facility.

Special care must be taken to properly synchronize the data streams when integrating the traffic generated by the MACS SimMgr with data from flight simulators similar to the B747 cab at NASA Ames and the Ikhana Simulator at NASA Dryden. The data shown in Figure 19 demonstrate a significant difference between the observed latency for the MACS SimMgr data and the candidate flight simulators. If the latency of the MACS SimMgr stays consistently greater than the flight simulators (and is not addressed through software modifications), steps to add latency to the virtual simulator data streams may be necessary to integrate the data with MACS sources. The importance of properly synchronizing these data streams will be a function of the required fidelity of the simulation and desired usage of the data. As such, latency requirements should be discussed with the researchers during the development of the LVC instantiation for each event.

While integrating data from live aircraft (whether UAS, surrogate, or manned intruders), the aircraft state data must also be synchronized with the virtual and constructive data sources in order to provide a consistent data stream. Data were collected from the live flight of the NASA Glenn S-3B Viking aircraft during the characterization testing. Though the data collection only supported whole second timing resolution, it still provided a valuable range of latency times indicating that the use of 3G cellular technologies to transmit data from the ground station to the LVC does not meet the operational requirements we are attempting to emulate. Further testing with other live aircraft, in particular with the candidate aircraft for the upcoming flight tests, prior to Flight Test 3 will be critical in order to determine the probable latency of that equipment and to develop mitigation strategies to properly synchronize the data with the constructive scenario traffic data. It should be noted that live flight of the aircraft is not necessary to test the latency, the position of the aircraft is not the primary concern, only that the location (even if on the ground) is transmitted via the same mechanism as it would if it were flying.

Due to the complexity of the LVC and the complexity of the unmanned air traffic environment it is designed to emulate, until a specific instance of the LVC is built to meet a set of simulation requirements including the implementation of the architecture and input scenarios, the degree of relevancy will not be completely understood. However, the intention of these analyses was never to complete the characterization, but inform the LVC developers of potential areas of risks as an instantiation of the LVC is developed to meet a simulation or flight test set of requirements. It is only this instance of the LVC that can be completely tested. For this reason, one of the test components of the LVC build-up for a simulation is to conduct a characterization of the proposed LVC system.

Appendix A: Acronyms

ADRS	Aeronautical Data link and Radar Simulator
ADS-B	Automatic Dependent Surveillance-Broadcast
ARC	Ames Research Center
ARINC	Aeronautical Radio, Incorporated
ATC	Air Traffic Control
CDTI	Cockpit Display of Traffic Information
COTS	Commercial Off-The-Shelf
CSD	Cockpit Situation Display (A specific type of CDTI)
CVSRF	Crew Vehicle Simulation Research Facility
DFRC	Dryden Flight Research Center
DIS	Distributed Interactive Simulation
DSR	Display System Replacement
DSRL	Distributed System Research Laboratory
FAA	Federal Aviation Administration
FOM	Federation Object Model
GCS	Ground Control Station
GDL	Gateway Data Logger
GOTS	Government Off-The-Shelf
GPS	Global Positioning System
HLA	High Level Architecture
HITL	Human-In-The-Loop
ICD	Interface Control Document
IRIG	Inter-Range Instrumentation Group
IT&E	Integrated Test and Evaluation
LVC	Live Virtual Constructive
MACS	Multi-Aircraft Control System
MIL-STD	Military-Standard
NAS	National Airspace System

NASA	National Aeronautics and Space Administration
NISN	NASA Integrated Services Network
NTP	Network Time Protocol
RAIF	Research Aircraft Integration Facility
RTI	Real Time Infrastructure
TENA	Training Enabling Architecture
UAS	Unmanned Aircraft System
UTC	Coordinated Universal Time
VPN	Virtual Private Network

References

-
- ¹ DoD: "Modeling and Simulation Master Plan", DoD 5000.59P, Oct 1995
 - ² Henninger, Amy E., Cutts, Dannie, Loper, Margaret, et al, "Live Virtual Constructive Architecture Roadmap (LVCAR) Final Report", Institute for Defense Analysis, Sept, 2008
 - ³ FAA: "NAS System Specification, Functional and Performance Requirement for the National Airspace System, General, Volume 1," FAA, Department of Transportation, NAS-SS-1000, 15 April 1995.
 - ⁴ FAA: NAS-SR-1000, National Airspace System (NAS) System Requirements, Department of Transportation Federal Aviation Administration, January 2005
 - ⁵ RTCA: "Minimum Operational Performance Standards (MOPS) for Aircraft Surveillance Applications (ASA) System," RTCA DO-338, Dec 13, 2011
 - ⁶ FAA: "Automatic Dependent Surveillance – Broadcast (ADS-B) Out Performance Requirements to Support Air Traffic Control (ATC) Service"; Final Rule, Date: 5/28/2010, Federal Register Volume 75, Number 103, Document ID: FAA-2007-29305-0289, Washington, DC.
 - ⁷ Prevot, T., and Mercer, J., "MACS: A Simulation Platform for Today's and Tomorrow's Air Traffic Operations," AIAA-2007-6556, AIAA Modeling and Simulation Technologies Conference, August 2007.
 - ⁸ Battiste, V., and Johnson, W., "Development of a cockpit situation display for free-flight," AIAA and SAE, 1998 World Aviation Conference. September 1998.
 - ⁹ Software website: <http://www.pitch.se/prti>, August 2013
 - ¹⁰ Lutz, R., LeSueur, K., Fast, P., Graeff, R., Simolte, A., Rutledge, J., and Motti, R., "A Persistent LVC Simulation Environment for UAS Airspace Integration," The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC), December 2010.
 - ¹¹ AviationSimNet Standards Working Group, "AviationSimNet Specification", Version 2.2, MITRE Corporation, June 2010
 - ¹² "IEEE Standard for Distributed Interactive Simulation – Application Protocols," IEEE 1278.1-2012, Dec 19 2012
 - ¹³ Edward T. Powell, Jason Lucas, Kurt Lessmann, and George J. Rumford, "The Test and Training Enabling Architecture (TENA) 2002 Overview and Meta-Model," published in the Proceedings of the Summer 2003 European Simulation Interoperability Workshop, 03E-SIW-041, found at <http://www.sisostds.org/>