The Johns Hopkins University
APPLIED PHYSICS LABORATORY

Enclosure to NSAD-L-2010-099
JHS01

NSAD-R-2010-031

Live-Virtual-Constructive Architecture Roadmap Implementation

Common Gateways and Bridges Characterization Report

May 2010

The Johns Hopkins University
APPLIED PHYSICS LABORATORY

NSAD-R-2010-031

Live-Virtual-Constructive Architecture Roadmap Implementation

Common Gateways and Bridges Characterization Report

May 2010

Prepared by:

RR Lutz, JHU/APL
DL Drake, JHU/APL
D Bergin, Trideum
D Cutts, AEgis Technologies Group
K Lessman, Trideum
M O'Connor, ITT Corporation

This page intentionally left blank.

TABLE OF CONTENTS

EXECUTIVE SUMMARY ES-1

1. INTRODUCTION 1-1

2. REPORT FORMAT 2-1

3. PROJECT APPROACH 3-1
   3.1 DENOTATION OF "GATEWAYS" VICE "GATEWAYS AND BRIDGES" 3-2
   3.2 DATA GATHERING 3-2
      3.2.1 Developer Questionnaire Methodology 3-3
      3.2.2 Gateway Developers 3-3
      3.2.3 Developer Questionnaire Format 3-5
      3.2.4 Gateway User Questionnaire Methodology 3-5
      3.2.5 User Questionnaire Format 3-7
   3.3 EXECUTION PLAN 3-8

4. GATEWAY CHARACTERIZATION VIA QUESTIONNAIRE RESULTS AND ASSESSMENT 4-1
   4.1 FUNCTIONAL CAPABILITIES 4-1
      4.1.1 SDEM Translations 4-1
      4.1.2 SDEM Behaviors 4-3
      4.1.3 Architecture Translations 4-5
      4.1.4 Architecture Behaviors 4-6
      4.1.5 Exercise Management Behaviors 4-8
   4.2 OPERATIONAL CAPABILITIES 4-10
      4.2.1 User Interface 4-10
      4.2.2 Performance 4-12
      4.2.3 Operation Modes 4-14
      4.2.4 Extension Modes 4-16
      4.2.5 Platform 4-17
      4.2.6 Documentation and Support 4-17
      4.2.7 Maturity 4-19

5. SUMMARY 5-1

APPENDIX A: REFERENCES A-1
APPENDIX B: GATEWAY-RELATED GLOSSARY OF TERMS B-1
APPENDIX C: ABBREVIATIONS AND ACRONYMS C-1

LIST OF FIGURES

Figure 1-1: LVC Simulation Interactions 1-2
Figure 3-1: Initial Data Gathering and Synthesis Methodology 3-1
Figure 3-2: Organization, Consolidation, and Analysis Methodology 3-1
Figure 4-1: Developer Responses Regarding Performance (Number of Responses) 4-13
Figure 4-2: User Site Responses Regarding Performance (Number of Responses) 4-14

LIST OF TABLES

Table 1-1: Activities for Gateways and Bridges Task 1-4
Table 3-1: Gateway Developers Responding to the Questionnaire 3-4
Table 3-2: Participant Gateway Users 3-6
Table 3-3: Aggregate List of Gateways Employed by Users 3-7
Table 4-1: Developer Abbreviations 4-1
Table 4-2: Responses Regarding SDEM Translation 4-2
Table 4-3: Responses Regarding SDEM Behavior 4-4
Table 4-4: Responses Regarding Architecture Translations 4-6
Table 4-5: Responses Regarding Architecture Behaviors 4-7
Table 4-6: Responses Regarding Exercise Management Behaviors 4-9
Table 4-7: Responses Regarding User Interfaces 4-10
Table 4-8: High Correlation for Top 3 Performance Areas 4-13
Table 4-9: Responses Regarding Operational Modes 4-15
Table 4-10: Responses Regarding Extension Modes 4-16
Table 4-11: Responses Regarding Platform Characteristics 4-17
Table 4-12: Responses Regarding Gateway Documentation and Support 4-18
Table 4-13: Responses Regarding Gateway Maturity 4-19

EXECUTIVE SUMMARY

The Live-Virtual-Constructive (LVC) Architecture Roadmap (LVCAR) Study developed a vision for achieving significant interoperability improvements in LVC simulation environments. The study recommended several activities intended to reduce the time and cost required to integrate mixed-architecture events.

Three of the key LVCAR Phase I recommendations were to determine whether existing gateway and bridge applications were effective in meeting user requirements, whether improvements in gateway/bridge capabilities were necessary to address identified gaps, and how these improvements could be best implemented to maximize the Department of Defense (DoD) return on investment. The term "bridge" in this context refers to intelligent translators that link together enclaves of simulations that use the same underlying simulation architecture. A "gateway" is also an intelligent translator but is designed to link simulation enclaves that use dissimilar architectures. To address this recommendation, the LVCAR Phase II Common Gateways and Bridges task began with three major activities: performing gateway and bridge literature research, compiling the team's gateway and bridge usage and development experience, and developing formal gateway and bridge operation terminology. With this starting point, the team was able to create an initial delineation of gateway and bridge capabilities based on its members' background and experience.

At this point in the project, it became clear that the distinction between "gateway" and "bridge" was moot from a development and usage standpoint. For that reason, the term "gateway" is used for both types of applications in this document. Starting with the initial delineation of capabilities, the team compiled a Gateway Capabilities Matrix Template. This template allowed the team to create two structured questionnaires. The first was for commercial and government-funded gateway developers, for which an on-line web interface was provided. The second was for site visits to users of gateways within DoD, mainly selected from large-scale exercise coordinators for the United States (US) military services and joint operations. Although the questionnaires were written for different audiences, they were written in a parallel fashion that allowed a correlation of answers across the two audiences. The background literature research also aided the team in creating a candidate list of developers and users to fill out the questionnaires.

The characteristics investigated were in the following twelve areas: Simulation Data Exchange Model Translations, Simulation Data Exchange Model Behaviors, Architecture Translations, Architecture Behaviors, Exercise Management Behaviors, User Interface, Performance, Operation Modes, Extension Modes, Platform Support, Documentation and Support, and Maturity. Each of these characteristic types was codified in a set of terms of reference and explained to the developers and users so that their answers would strongly correlate.

In addition to the developer questionnaire and the gateway user site visits, the team also held a one-day workshop, the "LVCAR Common Gateways and Bridges Workshop," to present the findings of those questionnaires. The workshop was held at the Virginia Modeling, Analysis, and Simulation Center (VMASC) in Suffolk, VA, on March 4, 2010. This workshop not only allowed the team to verify its results with the community, but it also allowed additional feedback and data gathering. This additional information has been included in the data that is presented in this document.

At the completion of the data gathering and analysis aspects of this task, eleven gateway characteristic assessment points have been identified. Listing them in order of most agreement between developers and users to least agreement produces the following delineation.

1. A capability gap in Simulation Data Exchange Model (SDEM) Translations does not exist.
2. A capability gap in Architecture Translations does not exist.
3. A capability gap in Architecture Behaviors does not exist.
4. No capability gaps were identified for Operation Modes.
5. In regard to Support, the provided support generally meets user needs.
6. In terms of Gateway Performance, users and developers provided consistent responses. When asked to provide the critical performance characteristics, the top three characteristics for both developers and users matched perfectly.
7. Overall, regarding Exercise Management Behaviors, the match between developer-supplied capabilities and user-needed capabilities is strong.
8. In regard to Platform Support, there did not appear to be any gaps, as there seem to be two primary platforms in use by the community at this time.
9. There are no identified capability gaps in the User Interface capabilities; however, further research is needed to determine if there is a capability gap for remote control of multiple gateways.
10. In regard to Extension Modes, most users expressed a need to have the capability to extend the functionality of developer-provided gateways.
11. In regard to SDEM Behaviors, there is a disconnect between the developer-indicated capability and the user expressed needs, particularly in the areas of dynamic information translation where one SDEM requires it and the other does not, and supporting differing behaviors between SDEMs. As for the latter, about half of the users indicated a need for this capability, but no developers reported providing this capability. This is either an actual gap in the capability required by the user or a misunderstanding of the capability by either the users and/or developers.

Despite the apparent lack of gateway capability gaps, there was wide agreement at the 4 March workshop that there are several potential improvements that can be made that will lower the technical and cost risks generally associated with the use of gateways. Such improvements, along with team recommendations, will be forthcoming in a later deliverable. However, the information provided in this document provides the necessary foundation for these recommendations by articulating the current state of gateway usage within the DoD.

This page intentionally left blank.

1. INTRODUCTION

Modeling and simulation (M&S) has long been recognized as a critical technology for managing the complexity associated with modern systems. In the defense industry, M&S is a key enabler of many core systems engineering functions. For instance, early in the systems acquisition process, relatively coarse, aggregate-level constructive models are generally used to identify capability gaps, define systems requirements, and examine/compare potential system solutions. As preferred concepts are identified, higher fidelity models are used to evaluate alternative system designs and to support initial system development activities. As design and development continues, very high fidelity models are used to support component-level design and development, as well as developmental test. Finally, combinations of virtual and constructive M&S assets are frequently used to support operational test and training requirements. Note that other industries (e.g., entertainment, medical, and transportation) also make heavy use of M&S, although in somewhat different ways.

The advent of modern networking technology and the development of supporting protocols and architectures has led to widespread use of distributed simulation. The strategy behind distributed simulation is to use networks and support simulation services to link together existing M&S assets into a single unified simulation environment. This approach provides several advantages as compared to development and maintenance of large monolithic standalone simulation systems. First, each individual simulation application can be collocated with its resident subject matter expertise rather than having to develop and maintain a large standalone system in one location. It facilitates efficient use of past M&S investments, as new, very powerful simulation environments can be quickly configured from existing M&S assets. It provides flexible mechanisms to integrate hardware and/or live assets into a unified environment for test or training, and is much more scalable than standalone systems.

There are also some disadvantages of distributed simulation. Many of the issues related to distributed simulation are related to interoperability concerns. Interoperability refers to the ability of disparate simulation systems and supporting utilities (e.g., viewers and loggers) to interact at runtime in a coherent fashion. There are many technical issues that affect interoperability, such as consistency of time advancement mechanisms, compatibility of supported services, data format compatibility, and even semantic mismatches for runtime data elements. The capabilities provided by today's distributed simulation architectures are designed to address such issues, and allow for coordinated runtime interaction among participating simulations. Examples of such architectures include Distributed Interactive Simulation (DIS), the Test and Training Enabling Architecture (TENA), and the High Level Architecture (HLA).

In some situations, sponsor requirements may necessitate the selection of simulations whose external interfaces are aligned with more than one simulation architecture. This is what is known as a *multi-architecture simulation environment*. There are many examples of such

environments within the Department of Defense (DoD) (see references[1] for examples). When more than one simulation architecture must be used in the same environment, interoperability problems are compounded by the architectural differences. For instance, middleware incompatibilities, dissimilar meta-models for data exchange, and differences in the nature of the services that are provided by the architectures must all be reconciled for such environments to operate properly. This not only raises additional technical risk, but the additional resource consumption necessary to adjudicate these architectural differences affect cost and schedule risk as well.

Due to perceived increases in the number of multi-architecture simulation events anticipated in the future, along with the associated increase in costs, the DoD sponsored an initiative to examine the differences among the major simulation architectures from a technical, business, and standards perspective, and develop a time-phased set of actions to improve interoperability within multi-architecture simulation environments in the future, particularly Live-Virtual-Constructive (LVC) multi-architecture simulation environments (see Figure 1-1). This initiative is called the Live-Virtual-Constructive Architecture Roadmap (LVCAR). The first phase of this effort began in the spring of 2007 and continued for approximately 16 months. The result of this activity was a final report and supporting documentation that collectively totaled over 1000 pages. The second phase of this initiative is focusing on the implementation of the recommended actions from this report.

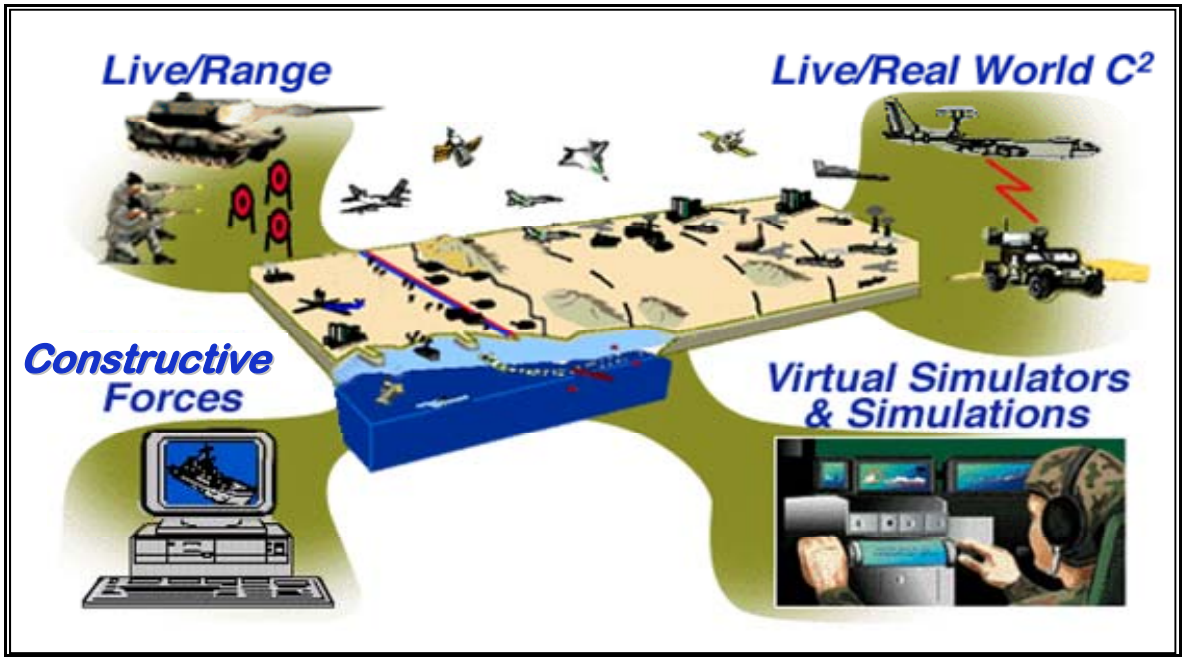

Figure 1-1: LVC Simulation Interactions

[1] References may be found in Appendix A.

One of the key actions from the LVCAR Phase I Report focused on the development of a common suite of gateway and bridge capabilities. The term "bridge" in this context refers to intelligent translators that link together enclaves of simulations that use the same underlying simulation architecture, such as a bridge between two existing HLA federations. A "gateway" is also an intelligent translator but is designed to link simulation enclaves that use dissimilar architectures, such as a gateway between simulations that use TENA as its external interface on one side of the translator and DIS on the other. Since LVCAR Phase I was primarily focused on multi-architecture interoperability, most of the emphasis in this area of this Phase I study was with respect to gateways. However, both gateways and bridges are ubiquitous in the LVC community today, and continue to represent one of the most widely used means of addressing interoperability concerns in multi-architecture LVC environments.

Although there are many success stories with respect to the use of gateways and bridges on multi-architecture developments, there have also been some reported problems. Since there is no such thing as a "common" gateway/bridge across (or sometimes, even within) user communities, managers of most LVC environments find themselves to be generally unaware of reuse opportunities for needed gateway and bridge capabilities. Thus, from a historical perspective, many programs have just built their own gateways and bridges from scratch based on their immediate needs, with little or no attention being paid to potential reuse. This has led to a large number of gateways and bridges in the LVC community today, with a significant portion having little documentation and with no visibility outside the projects for which the gateways and bridges were designed. This is, of course, very inefficient from a DoD corporate perspective, as many similar basic functional components keep getting developed over and over again, and maintenance costs are spread over a large set of redundant capabilities. Also, the continuous consumption of valuable project resources to design, develop, and test new gateways and bridges increases technical, schedule, and cost risk to user programs.

The main goal of the LVCAR Phase II Gateways and Bridges task is to identify the various technical and business issues associated with current gateway usage within the DoD and to identify recommended actions that improve the efficiency and effectiveness of gateway/bridge solutions in the future. The strategy that was implemented to accomplish this objective was composed of eight activities, listed in Table 1-1.

Table 1-1: Activities for Gateways and Bridges Task

Step	Activity
1	Develop a template of gateway/bridge characteristics based on existing documentation and the background and experiences of study team members. Include "soft" factors such as usability and availability.
2	Conduct a literature search to identify case studies of past multi-architecture LVC environment developments for the purpose of identifying a representative set of gateway/bridge developers and users.
3	Based on the literature search (as augmented with existing team member contacts), work with the identified gateway and bridge developers to properly characterize the existing gateways according to the template format.
4	Conduct a series of site visits with gateway and bridge user organizations to further characterize existing capabilities and to elicit future gateway requirements.
5	Assess the data from the developer questionnaires and site visits to identify capabilities that are either over-served (i.e., duplicative capabilities) or under-served (i.e., gaps) according to current and future user requirements.
6	Conduct a workshop with gateway and bridge developers and users to verify the findings of the gateway/bridge characterization effort, and to solicit input for the common gateways implementation strategy.
7	Deliver a report documenting the gateway and bridge characterizations based on the findings from the questionnaires and site visits and the feedback received at the workshop.
8	Based on perceived user requirements, develop and deliver a plan for providing improved gateway/bridge capabilities in the future, including distribution and maintenance strategies for new software.

This report addresses the final product produced through execution of Steps 1 through 7. A separate document will report on the results of Step 8. Section 3 of this document will provide additional information on the methodology employed on this study, including the participating organizations and details of the workshop identified in Step 6.

2. REPORT FORMAT

This report is constructed in four major parts: (a) the executive summary, introduction and report format; (b) the technical approach taken in this project; (c) the gateway/bridge characterization based on data gathered from both gateway developers and users; and (d) a summary of the conclusions.

This report has three appendices: Appendix A is a list of referenced documents; Appendix B provides a glossary of the gateway- and bridge-related terms that were codified for this project, and Appendix C is a list of abbreviations and acronyms.

This page intentionally left blank.

3. PROJECT APPROACH

The steps delineated in Table 1-1 describe the strategy that is being implemented to accomplish the project objectives. A detailed pictorial view is presented in Figure 3-1 and Figure 3-2.

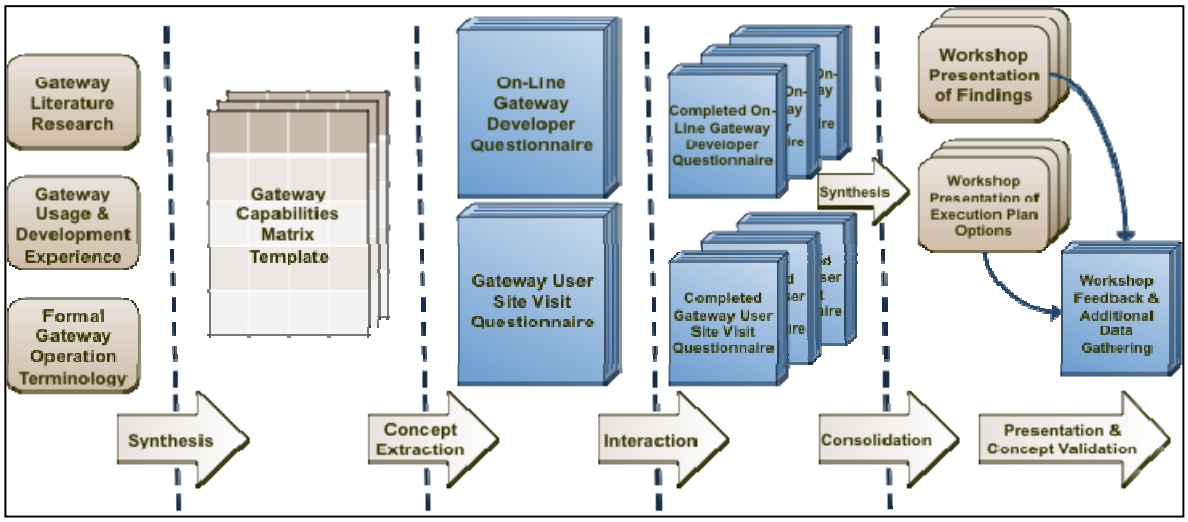

Figure 3-1: Initial Data Gathering and Synthesis Methodology

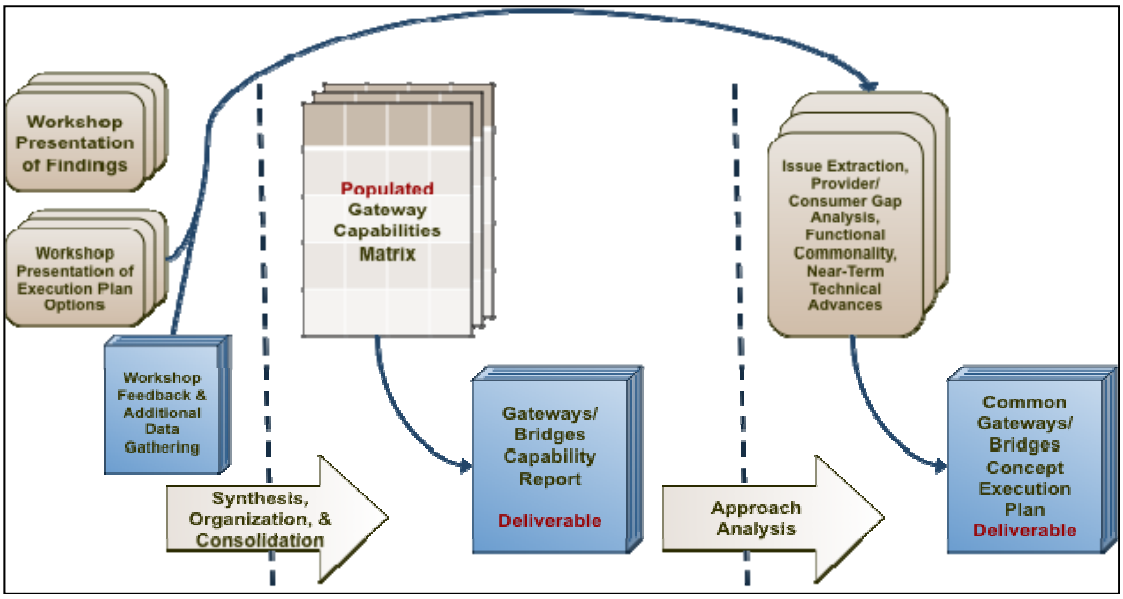

Figure 3-2: Organization, Consolidation, and Analysis Methodology

The initial work for this project was composed of three major activities: performing gateway/bridge literature research, compiling the team's gateway/bridge usage and development experience, and developing formal gateway/bridge operation terminology. With this starting point, the team was able to create a delineation of gateway/bridge capabilities based on its

members' background and experience. From this delineation, the team compiled a Gateway/Bridge Capabilities Matrix Template. This template allowed the team to create two structured questionnaires. The first was for gateway and bridge developers, for which an on-line web interface was provided. The second was for site visits to users of gateways and bridges within the DoD. Even though the questionnaires were written for different audiences, they were written in a parallel fashion that allowed a correlation of answers across the two audiences. The background literature research also aided the team in creating a candidate list of developers and users to fill out the questionnaires.

The capabilities that were of interest consisted of three general areas: functional capabilities, operational capabilities, and finally, a general question about additional capabilities or functionality. To ensure consistent answers from the community, the team formalized its terminology to the degree that it can be understood across different gateway and bridge architectures and usages.

3.1 DENOTATION OF "GATEWAYS" VICE "GATEWAYS AND BRIDGES"

It is at this point in the project that the team realized that the difference between "gateways" versus "bridges" became a subtle distinction depending upon which architectures were being translated. Although this may seem critical at a high-level of understanding, at the point of development and usage, the distinction becomes practically moot. The crux of this reasoning is that gateway systems are written to translate from one system (whatever that system may be) to another, so the distinction between translating between different architectures and the same architectures is really a matter of configuration. For this reason the team consistently referred to both gateways and bridges as "gateways." Thus, throughout the remainder of this document, the singular word "gateway" is used to refer to both gateways and bridges. This was fully explained to the questionnaire participants who agreed that this was a reasonable approach.

3.2 DATA GATHERING

The goal of the developer questionnaire was to identify and characterize gateway functionality across a range of gateways. It also provided a consistent mechanism for developers to document their gateway functionality, allowing the team to validate its initial list of capabilities, and to identify additional capabilities. In the development of the developer questionnaires the sensitivity of asking developers about what capabilities their gateways had or didn't have was taken into consideration. Much like a census, the team wanted to gather the information about the existing gateways, not determine a "best" gateway or to compare gateways. The team fully understood that the gateways were developed to meet certain needs, and, therefore, would not embody all of the characteristics that were being asked about within the questionnaire. Specifically, the team was asking about the capabilities of a gateway to identify what requirements were being requested by its intended user community.

In addition to the developer questionnaire and the gateway user site visits, the team also held a one-day workshop, the "LVCAR Common Gateways and Bridges Workshop," to present the findings of those questionnaires. The workshop was held at the Virginia Modeling, Analysis, and Simulation Center (VMASC) in Suffolk, VA, on March 4, 2010. This workshop not only allowed the team to verify its results with the community, but it also allowed additional feedback and data gathering. This additional information has been included in the data that is presented in this document.

3.2.1 Developer Questionnaire Methodology

Based on an analysis of the LVCAR Phase I recommendations and gateway literature research, a Gateway Capabilities Matrix was developed. This capabilities matrix provided a general-purpose template for describing gateway functionalities. The matrix was then populated with an initial assessment of known gateway capabilities based on the background and experiences of the Gateways Team.

From the Gateway Capabilities Matrix, the team developed a set of questions to be answered by gateway developers to describe the capabilities of their particular gateway. The gateway developer's responses to these questions were intended to allow developers to characterize gateway functionality across a range of gateways and to provide a consistent mechanism for documenting gateway functionality. The responses were not intended to rank or compare gateways but rather reflected the general understanding that individual gateways were designed for specific purposes and would not likely have capabilities beyond those necessary for their intended purposes.

The gateways literature research and gateway team experiences identified a total of 19 gateways. Points of contact for the gateways were found and a request was sent to each. The questionnaire was implemented as an online web form by Johns Hopkins University Applied Physics Laboratory (JHU/APL) personnel. A total of eleven developers completed the online questionnaire. The format of the questionnaire and the results are described in Section 3.2.3.

3.2.2 Gateway Developers

Table 3-1 identifies the gateway developers who responded to the questionnaire and also identifies the purpose of the gateway and the architectures and/or protocols between which the gateway communicates.

Table 3-1: Gateway Developers Responding to the Questionnaire[2]

Gateway	Purpose	Architectures / Protocols
Range Data Gateway (RDGW) - Developer	Provides support to range applications	TENA RDS, TENA KABIC (WSMR), TENA TECCS (Edwards AFB), and TENA NACTS (Nellis AFB)
GMU Gateway	Built to show how selectively reliable multicast can provide efficient network group communications in distributed simulation environment (no longer supported)	DIS and HLA
Network System Under Test Gateway (NSUT)	Provides translations between DIS and TENA 6, and TENA 6 and TENA 5.2 (it also provides translation between two slightly different SDEMs)	DIS and TENA (v5.2 and 6)
Gateway Builder (GWB)	Generate gateways between any protocol for which an adapter is available (adapters can be developed as necessary)	TENA, DIS 1278.1, DIS 1278.1A, HLA 1.3, HLA 1516, SIMDIS 1.5, Generic Packet Protocols, and ProtoCore
VR-Exchange	Provide gateway between various architectures / protocols	HLA, DIS, TENA, and others (via an API)
DIS Filter	Operates as a TSR bridge for DIS filter	DIS to DIS
Theater Air Command and Control Simulation Facility (TACCSF)	Operates as a TACCSF Software Router (TSR)	Operated at the network level, not at the architecture or protocol level
Gateway between HLA and SCORM	Operates as a HLA to SCORM gateway	HLA 1.3 and SCORM
Battle Lab Collaborative Simulation Environment (BLCSE)	Operates as a High Level Architecture adapter	DIS and HLA
Joint Simulation Bus (JBUS)	Provides a generic plug-in gateway	HLA, DIS, TENA, Link 16, and other C2 message formats
Pitch Gateway	Operates as a gateway between various architectures / protocols	HLA, DIS, and TENA

[2] See Appendix C: Abbreviations and Acronyms for the full names of architectures and project-related protocols that are abbreviated in this table.

3.2.3 Developer Questionnaire Format

Sixty questions were developed by the Gateways Team and sorted into the two categories of functional capabilities and operational capabilities. The functional capabilities include Simulation Data Exchange Model (SDEM) Translations, SDEM Behaviors, Architecture Translations, Architecture Behaviors, and Exercise Management Behaviors. The operational capabilities are related to User Interface, Performance, Operation Modes, Extension Modes, Documentation /Support and Maturity. Because it was an online questionnaire, the responses were constrained to the set developed by the Gateways Team, although respondents were allowed to provide explanatory text with each response. The response choices were numbered from 0 through 5, although the legends didn't give a description of every possible choice. This allowed intermediate values to be used for responses. As an example, if 0 indicated that the gateway had no ability to calculate dynamic information, and 3 indicated that there was a fixed set of dynamic information calculations, then the respondent could use 1 to indicate that this capability was partially addressed by allowing a limited variation of fixed dynamic information calculations, and the limitation was explained in the accompanying text response. Not all allowed responses to all questions were observed in the results.

3.2.4 Gateway User Questionnaire Methodology

The basis of the Gateway User Questionnaire is the same as that of the Gateway Developer Questionnaire. The Gateway Capability Matrix that described typical gateway functionalities and capabilities was also used by the team members to develop a set of questions to be used when interviewing gateway users about the application and configuration of the various gateways that they have used. The gateway user's responses to these questions were intended to provide a consistent mechanism to capture the community's use and configuration of gateways in support of multi-architecture LVC simulation gateway functionality. The intent was to capture organizational need rather than gateway capabilities, while allowing the feedback to be used to compare with gateway developer input to identify user needs, coverage and gaps. The responses were not intended to evaluate organization or facility capabilities nor were they intended to evaluate gateway usage or configuration.

Through literature research and JHU/APL Gateway Team experiences, there were a total of nine gateway users identified, with the intent to interview users that have used multiple gateways in the past. Gateway Users were selected independent of Gateway Developers, and included Service participation from the Test and Evaluation (T&E), Research and Development (R&D), and Training and Acquisition communities. Points of contact for the gateways were found and a request was sent to each. The Gateway User Questionnaire was completed through in-person interviews led by the Gateway Team. Of the nine users contacted, all nine participated in the in-person interviews. For the completion of the user questionnaire, the team felt that more complete responses would be achieved through in-person interviews rather than having users complete an online questionnaire.

The format and results of the interviews and questionnaire are presented below. Table 3-2: Participant Gateway Users identifies the gateway users who were interviewed. The user's organization and branch of service, as well as the distributed simulation architectures for which the gateway was used, are included. Through the gateway user interviews, numerous gateways were identified as well as the programs on which these gateways have been or are being used. This list of gateways is summarized in Table 3-3.

Table 3-2: Participant Gateway Users[3]

Service	Organization	Architectures
Army Training/Testing/ Acquisition	PEO-STRI Orlando, FL	CTIA, TENA, and DIS
Army R&D	RDECOM MATREX Orlando, FL	HLA and DIS (future plans for TENA and Live 32)
Army Training	TRADOC(HQ) Fort Monroe, VA	DIS, HLA, and Tactical protocols
Army T&E	Redstone Test Center (RTC) Distributed Test Control Center Redstone Arsenal, AL	HLA, DIS, and TENA
Army T&E	Inter Range Control Center White Sands Missile Range, New Mexico	TENA, DIS, HLA, and KABIC to TENA (Intra-range protocol)
Navy T&E	PAX River (Navy) Patuxent River Calvert, MD	HLA, DIS, and TENA
Air Force Training, T&E and Experimentation	Distributed Mission Operations Center (DMOC), Kirtland AFB, NM	TENA, DIS, and HLA
Air Force Acquisition and Test	Simulation and Analysis Facility (SIMAF), Wright-Patterson AFB, Ohio	HLA and TENA LVC to DIS
OSD/Joint Test and Evaluation	TENA/JMETC Technical Team	TENA, DIS, HLA and several range-specific formats
Joint Training	JFCOM Suffolk, VA	HLA, DIS, TENA, and Tactical protocols

[3] See Appendix C: Abbreviations and Acronyms for the full names of services, organizations, and architectures that are abbreviated in this table.

Table 3-3: Aggregate List of Gateways Employed by Users[4]

Referenced by Multiple Users [number of users]	Referenced by a Single User
JMETC Gateways developed using Gateway Builder (GWB) [6]	HLA Adaptor in OneSAF
MaK Gateway Products [4]	JSAF Gateway
InterTEC NSUT Gateway [3]	SIMPLE Gateway
SSEGW [2]	DIS Interoperability Manager for OneSAF
BLCSE Middle Ware [2]	RTC developed Shared Memory Multi-Architecture Interface (SMMAI)
Command & Control (C2) Adaptor Gateway [2]	Gateway for interconnecting Trideum Corporation HLA to TENA to GERTICO
JBUS [2]	EXCIS (message translation) Gateway
	TSIU MATREX platform to TAIS (air track) Gateway
	Global Magic TENA to DIS Gateway
	OTB-based Gateway
	JIMM to Shared Memory plug-in architecture

3.2.5 User Questionnaire Format

The questionnaire used for the user interviews was also developed by the Gateways Team and followed the same format as that of the developer questionnaire. The only difference was that the questions were from the perspective of a user vice a developer. The categories also included functional capabilities: SDEM translations, SDEM behaviors, architecture translations, architecture behaviors, and exercise management behaviors; and operational capabilities: user interface, performance, operation modes, extension modes, platform, documentation/support, and maturity.

[4] See Appendix C: Abbreviations and Acronyms for the full names of services, organizations, and architectures that are abbreviated in this table.

3.3 EXECUTION PLAN

As shown in Figure 3-2, the remaining activity on this project is to deliver an Execution Plan based on the team's findings. The Execution Plan will detail the temporally-sequenced set of recommended actions to improve the efficiency and effectiveness of gateway usage in the future. The Execution Plan will be documented in the project's second deliverable.

4. GATEWAY CHARACTERIZATION VIA QUESTIONNAIRE RESULTS AND ASSESSMENT

This section discusses the insights into the developer and user responses. Most of the subsections include a table with the combined user and developer responses. Within these tables, the names of the users' sites are not presented next to their responses to preserve the anonymity that was afforded them to help ensure their candid responses. For the responses for the developers, the tables do list their affiliations, since it allows one to correlate which gateway has certain characteristics, and because anonymity was not requested. Table 4-1 lists the abbreviations that were used for each developer's affiliation.

Table 4-1: Developer Abbreviations

Abbreviation	Developer Affiliation
RDGW	InterTEC Range Data Gateway - Developer
GMU	George Mason University (GMU) Gateway
NSUT	InterTEC Network System Under Test Gateway
GWB	JMETC Gateway Builder
VR-EX	MaK VR-Exchange
DISF	Distributed Interactive Simulation (DIS) Filter
TSR	Theater Air Command and Control Simulation Facility (TACCSF) Software Router
SCORM	Gateway between HLA and SCORM
BLCSE	Battle Lab Collaborative Simulation Environment
JBUS	JFCOM Joint Simulation Bus (JBUS) Generic Plug-in gateway
PITCH	The gateway developed by Pitch Technologies AB for various architectures/protocols

4.1 FUNCTIONAL CAPABILITIES

As previously defined, functional capabilities represent actions required by the gateway to meet SDEM, architecture, or exercise management needs. These are considered the core capabilities of a gateway. The functional capabilities usually determine if the gateway meets the user functional needs when interfacing between multi-architectural LVC resources.

4.1.1 SDEM Translations

SDEM Translations are the set of capabilities to convert data from one SDEM to another SDEM. This includes unit conversion, coordinate conversion, and enumeration mapping and may require translating data stored in a single object in one SDEM to multiple objects in another SDEM. Examples of SDEM Translations are unit conversion on a single attribute and single element enumeration to a multi-element enumeration. The questions presented to the developers and users regarding SDEM Translations, and their associated answers, are shown in Table 4-2.

Table 4-2: Responses Regarding SDEM Translation

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Perform unit conversion on a single attribute (SDEM element)? **0:** no unit conversion **3:** unit conversion within a defined set **5:** conversion between arbitrary units	5	3	0	5	5	0	0	0	3	5	5	Y	Y	Y	Y	Y	Y	?	Y	Y
Perform complex data type conversions using multiple attributes? **0:** no multiple attribute conversion **1:** multiple attribute conversion for 5 or fewer fixed types **3:** multiple attribute conversion for 15 or fewer fixed types **5:** arbitrary multiple attribute conversion	5	0	0	5	5	0	0	1	1	5	5	Y	Y	N	Y	Y	Y	Y	Y	Y
Translate single attribute enumerations between SDEMs? **0:** no single enumeration mapping **1:** one to one enumeration mapping **3:** many to one enumeration mapping **5:** arbitrary single enumeration mapping	1	1	2[5]	5	5	0	0	1	3	5	5	Y	Y	Y	Y	Y	Y	Y	Y	Y
Translate multi-component enumerations between SDEMs? **0:** no multi component enumeration mapping **1:** 1 type of multi-component enumeration mapping **3:** 5 or fewer types of multi-component enumeration mapping **5:** arbitrary multi-component enumeration mapping	5	0	0	5	5	1	0	0	3	5	5	Y	Y	N	Y	Y	Y	Y	Y	Y
Translate identifiers between SDEMs? **0:** no translation of SDEM identifiers **1:** 1 specific translation **3:** 5 or fewer specific translations **5:** arbitrary translations	5	0	?	5	5	1	0	5	3	5	5	Y	Y	Y	Y	Y	Y	Y	Y	Y
Provide information that is required by one SDEM but not present in the other? **0:** no ability to add data **1:** ability to add fewer than 5 fixed types of data **3:** ability to add fewer than 15 fixed types of data **5:** ability to add arbitrary data	0	1	1	5	5	1	0	0	3	5	5	Y	Y	Y	Y	Y	Y	?	Y	Y

[5] The developer indicated that the gateway can perform some many-to-one mappings, but it cannot perform many-to-one mappings in general.

SDEM Translations are one of the most fundamental activities of a general-purpose gateway. The set of capabilities defined in this category (see Table 4-3) are essential to translating between SDEMs. Most of the general-purpose gateways provide the majority of the defined capabilities at some level. Two of the gateways targeted at multiple users provide extensive support. Nearly all of the users indicated a need for the capabilities as listed. Although all of the gateways do not support all of the capabilities, it was determined that the general-purpose gateways provided sufficient functionality. Therefore, it was determined that a capability gap in SDEM Translations does not exist.

4.1.2 SDEM Behaviors

SDEM Behaviors refer to the capability to correctly represent behaviors required when using a particular SDEM. This may include representing a behavior in one SDEM that is present in the other SDEM. SDEM behaviors may require the gateway to maintain state. An example of a SDEM Behavior is dead reckoning of positional attributes. The questions presented to the developers and users regarding SDEM Behaviors, and their associated answers, are shown in Table 4-3.

The need for a gateway to translate SDEM behaviors is dependent on the SDEMs being used. Different SDEMs define different sets of behaviors. The user response shows less need for some of the SDEM behaviors than for SDEM translations. This is most likely the result of the specific SDEMs in use by the users.

Table 4-3: Responses Regarding SDEM Behavior

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Provide dead-reckoning between SDEMs that use different dead-reckoning mechanisms? **0:** No dead-reckoning support **1:** 1 DIS dead-reckoning algorithm **3:** 4 DIS dead-reckoning algorithms **4:** all DIS dead-reckoning algorithms **5:** arbitrary dead-reckoning algorithms	1	0	?	0	5	4	0	0	0	4	5	Y	Y	Y	Y	Y	Y	?	N	Y
Publish attributes based on updates exceeding thresholds? **0:** No ability to publish attributes based on thresholds **3:** set of fixed thresholds to publish attributes **5:** ability to publish attributes based on arbitrary thresholds	5	0	5	5	5	5	0	0	0	5	5	Y	Y	Y	Y	Y	Y	Y	Y	Y
Calculate dynamic information that is required by one SDEM but not present in the other? **0:** no ability to add dynamic information **3:** fixed set of dynamic information **5:** arbitrary ability to add dynamic information	5	0	0	5	1[6]	3	0	5	0	3	5	N	N	N	Y	Y	Y	?	Y	N
Support behaviors defined in one SDEM but not in the other? **0:** no capability to support required SDEM behaviors **1:** capability to support 1 fixed behavior **2:** capability to support 2-3 fixed behaviors **3:** capability to support 4-5 fixed behaviors **5:** ability to support arbitrary behaviors	0	0	0	0	0	0	0	0	0	2	5	N	Y	Y	Y	?	Y	?	Y	N
Support differing SDEM publication rates? **0:** no support for SDEM publication rates **3:** support for 5 or fewer fixed publication rates **5:** arbitrary SDEM publication rates	5	0	0	0	3	5	0	0	3	5	5	Y	N	Y	Y	Y	Y	?	Y	Y

Dead-reckoning is a fundamental concept within many SDEMs. Almost all of the users indicated a need for gateways to support different dead-reckoning mechanisms. Only one of the developers provides full support for this capability. One other gateway provides nearly full

[6] The developer indicated that their gateway does not have the built-in capability, but a plug-in could add the capability.

support. Although almost all users indicated the need for this capability, none indicated the need was not being met. Therefore, there is a disconnect between the developer indicated capability and the user expressed needs.

All of the users and most of the general-purpose gateways provide the capability to publish attributes based on thresholds. Dead-reckoning is a specialized subset of this capability. Because of the importance of dead-reckoning to the community, it was listed as a separate capability. The developer responses to this capability may explain the apparent disconnect between the user and developer with regard to dead-reckoning in that the gateways are providing the needed capability.

It is not unusual for one SDEM to require dynamic data that another SDEM does not. This can be a challenge for gateways. Sometimes the required dynamic data can be calculated by the gateway. An example of this functionality is velocity, which can be calculated from the position of the entity at two different times. Half of the users reported a need for this capability. Most of the general-purpose gateways provide some level of this capability.

As shown in Table 4-3, one user organization was not sure about the available SDEM behavior provide by the gateway that was used. On further questioning of the interviewee, it was clear that the user was very successful in using the gateway. This highlights a subtle but important point, that a well functioning gateway is able to provide the user required functionality without the user having in-depth knowledge of how the gateway operates.

The most significant gap in user needs and developer capabilities is the support for behaviors that are uniquely defined in one SDEM. About half of the users indicated a need for this capability. No developers reported providing this capability. This is either an actual gap in the capability required by the user or a misunderstanding of the capability by either the users and/or developers. Support for differing SDEM publication rates is reported as a need by most of the users. However, only four of the developers provide this capability. It is possible the user needs are being met by the set of gateways that provide this capability. There is an indication of gaps in both SDEM behaviors that are uniquely defined in one SDEM and SDEM publication rates. More research in both areas is required to determine if there are user needs that are currently not being meet by available gateway functionality.

4.1.3 Architecture Translations

Architecture Translations are the set of capabilities to convert architecture-defined data between different architecture executions. This covers translation of data that is not defined in the SDEM, but is present in all architecture executions. Examples of Architecture Translations are translation of object identifiers and timestamps. The questions presented to the developers and users regarding Architecture Translations, and their associated answers, are shown in Table 4-4.

Table 4-4: Responses Regarding Architecture Translations

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Translate identifiers between Architectures? **0:** no ability to translate architecture identifiers **1:** ability for scheme based on 2 architectures/ implementations **3:** ability for scheme based on 5 architectures/ implementations **5:** arbitrary identifier translation	?	0	2[7]	2[7]	5	0	0	0	1	5	5	Y	N	Y	Y	Y	Y	Y	Y	Y

Almost all users expressed a need for the capability to translate architecture-defined identifiers between architectures. Most of the general-purpose gateways provide this capability, Therefore, it was determined that a capability gap in Architecture Translations does not exist.

4.1.4 Architecture Behaviors

Architecture Behaviors are the set of capabilities to perform actions required by the architecture. These behaviors are required by the architecture and apply to all SDEMs using the architecture. An example of an Architecture Behavior is a request to publish attributes of a specified object. The questions presented to the developers and users regarding Architecture Behaviors, and their associated answers, are shown in Table 4-5.

[7] The developer indicated that their gateway is able to translate architectures identifiers for three architectures schemes.

Table 4-5: Responses Regarding Architecture Behaviors

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Support architecture publication rates? **0:** Cannot support architecture publication rates **3:** can support fixed set of architecture publication rates **5:** can support an arbitrary set of publication rates	5	0	1[8]	3	5	0	0	0	3	5	5	Y	N	Y	Y	Y	Y	?	Y	Y
Publish all the attributes of an object in an Architecture that does not support partial updates when translating from an Architecture that permits partial updates? **0:** cannot support publish all attributes **5:** can publish all attributes based on partial attribute update	5	5	5	5	5	0	0	5	5	5	5	Y	N	Y	Y	Y	Y	?	Y	Y
Provide services available in one architecture but not in another? **0:** no additional services **5:** arbitrary additional services	5	?	5	0	5	0	0	0	5	5	5	Y	Y	Y	Y	Y	N	?	Y	N
Publish only changed attributes of an object in a protocol that permits partial updates when translating from a protocol that only supports publication of all attributes for each update? **0:** cannot publish only changed attributes **5:** can publish only changed attributes	0	5	0	5	5	0	0	0	5	0	5	Y	N	Y	Y	Y	Y	?	Y	Y
Translate Remote Procedure Calls (RPC's) between architectures that support RPCs? **0:** no support for RPCs when both architectures do not support **3:** some ability to support fixed set of RPCs **5:** arbitrary ability to support RPCs	3	0	0	0	1[9]	0	0	3	0	0	3	Y	N	N	N	N	Y	N	Y	N
Translate RPCs from architectures that support RPCs to an architecture that does not support RPCs? **0:** no support for RPCs when both architectures do not support **3:** some ability to support fixed set of RPCs **5:** arbitrary ability to support RPCs	0	0	0	0	1[9]	0	0	3	0	0	5	N	N	Y	N	N	Y	N	Y	N

[8] The developer indicated that their gateway is able to support a limited set of architecture publication rates.

[9] The developer indicated that their gateway does not have the capability built-in, but a plug-in could add the capability.

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Translate RPCs from an architecture that does not support RPCs to one that does? **0:** no support for RPCs when both architectures do not support **3:** some ability to support fixed set of RPCs **5:** arbitrary ability to support RPCs	0	0	0	0	1[9]	0	0	3	0	0	3	N	N	?	N	N	N	N	Y	N
Remove translated objects based on the rules of the original publisher SDEM or protocol? **0:** no ability to remove objects based on creator's rules **3:** ability to remove based on fixed set of rules **5:** arbitrary ability to support remove rules	3	0	5	5	5	0	0	0	3	5	5	Y	N	?	Y	Y	Y	?	Y	Y

For the most part, the developer and user responses for Architecture Behaviors are in alignment. The general-purpose gateways provide strong support for most of these capabilities.

The primary exception is in the area of Remote Procedure Calls (RPCs). Currently, only one of the architectures provides native support for remote procedure calls - TENA. Although users do use the TENA RPC mechanism, they do not express a strong need for gateway support to translate RPCs to other architectures that do not support them via emulation. It is possible to emulate the TENA RPC in DIS or HLA; however, there does not appear to be a strong need for it at this time. Therefore, it was determined that a capability gap in Architecture Behaviors does not exist.

4.1.5 Exercise Management Behaviors

Exercise Management Behaviors are capabilities that are not directly related to the SDEM or architecture. These capabilities are used to meet objectives of the overall exercise independent of the SDEM or architecture. These capabilities may not involve the translation of data or behaviors. An example of an Exercise Management Behaviors is filtering for bandwidth. The questions presented to the developers and users regarding Exercise Management Behaviors, and their associated answers, are shown in Table 4-6.

Table 4-6: Responses Regarding Exercise Management Behaviors

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Support Exercise Management publication rates? **0:** no support for exercise management rates **1:** support for one type of object class **3:** support for a fixed set of object classes **5:** support for arbitrary set of object classes	3	0	3	0	0	5	0	0	0	0	0	Y	N	N	Y	N	Y	N	Y	N
Filter (not translate) specified object classes between SDEMs? **0:** no ability to filter **1:** ability to filter based on 1 class/attribute **3:** ability to filter based on 5 or fewer attributes **5:** ability to filter on any attribute	5	0	3	5	3	4[10]	0	3	5	5	5	Y	Y	N	Y	Y	Y	N	Y	Y
Filter updates from specified senders? **0:** no ability to filter based on sender **5:** can filter based on sender	0	0	2[11]	0	1[11]	5	0	0	5	5	5	N	Y	Y	Y	N	Y	N	N	N
Filter updates based on conditions or events? **0:** no ability to filter based on conditions or events **3:** ability to filter based on 5 or fewer fixed events or conditions **5:** ability to filter on arbitrary conditions or events	0	0	3	5	4[12]	4[12]	0	3	3	5	5	N	Y	N	Y	N	Y	N	Y	N

An important objective of some gateways is to support constraints imposed by the infrastructure in place to support an exercise. These constraints include long-haul networks, computer processing, and differences in data interests. Gateways may be used to create enclaves in the exercise. An enclave is a segmented part of the overall network for which data traffic may be limited. Gateways are then used to control the data flow between the enclaves. The users that require all of the Exercise Management Behavior capabilities generally use this enclave approach. Overall, developer-supplied capabilities and user-needed capabilities are well-matched.

[10] The developer indicated that their gateway has the ability to filter on more than 5 attributes, but not on any attribute.

[11] The developer indicated that their gateway has a limited ability to filter updates.

[12] The developer indicated that their gateway has the ability to filter on more than 5 fixed events, but not an arbitrary number.

The need most often stated by the user is the capability to filter (i.e., not translate) specified object classes. Most of the developers provide this capability. This capability is often used to limit the data being sent to the enclaves.

Approximately half of the users stated the need for the capability to filter data from specific users. Less than half of the developers provide this capability. Not surprisingly, one of the gateways developed specifically to support enclave-based exercises provides the most support for this capability.

Approximately half of the users also stated the need for condition-based filtering between enclaves. The gateway developers generally provide strong support for this capability. It was determined that generally there is not a gap in Exercise Management capabilities; however, the user need and gateway support for filtering user-specific data should continue to be monitored to ensure that it is being met at an acceptable level.

4.2 OPERATIONAL CAPABILITIES

As previously defined, operational capabilities describe features of the gateway not directly related to translation. These capabilities generally relate to the usability of the gateway.

4.2.1 User Interface

The User Interface capabilities allow the user to interact with the gateway during pre-exercise, exercise, and post-exercise activities. A robust user interface is not always required for a gateway, but is seen as a usability enhancement as it displays real-time gateway-specific information to the user. An example of a User Interface capability is the ability to view the execution time of translated objects. The questions presented to the developers and users regarding Exercise Management Behaviors, and their associated answers, are shown in Table 4-7.

Table 4-7: Responses Regarding User Interfaces

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Provide translation statistics during runtime? **0:** no translations statistics provided during runtime **3:** 3 or fewer translations statistics provided during runtime **5:** 5 or more translations statistics provided during runtime	3	0	5	5	5	3	0	0	0	5	5	N	Y	N	Y	Y	Y	Y	Y	N
Display gateway status during runtime? **0:** no status provided during runtime **3:** 5 or fewer status indicators during runtime **5:** 10 or more status indicator during runtime	3	3	1[13]	3	5	3	3	0	0	4[14]	3	Y	N	N	Y	Y	Y	Y	Y	N
Generate debug logs and specify the level of logging? **0:** no debug log generated **3:** debug log generated **5:** debug log generated with specified level of detail	5	3	5	5	5	0	3	0	5	5	3	Y	Y	Y	Y	Y	Y	N	Y	N
Allow the user to modify translation rules during runtime? **0:** no ability to modify translation rules at runtime **1:** ability to modify one translation run at runtime **3:** ability to modify 5 or fewer translation rules at runtime **5:** ability to modify at least 10 translation rules at run time	5	0	5	0	5	5	0	0	0	5	0	?	N	N	Y	N	N	N	N	N

Most users stated a need to view translation statistics during runtime. This capability includes displaying the number of objects being translated, number of updates processed, and update rates. Most of the gateways provide strong capability to display this type of data to the user.

The users also indicated a significant need to monitor the status of the gateway during runtime. This includes displaying the state of the gateway, including its use of resources. This functionality can be important in determining if the gateway is overloaded. The developers provide a sufficient level of support for this capability.

[13] The developer indicated that their gateway has the ability to display a status in a debug window.

[14] The developer indicated that their gateway has the ability to display more than 5, but fewer than 10, status indicators during runtime.

The need most often stated by the user is the capability to debug log files. From the comments, users indicated this capability was key to determining how to fix problems encountered during the gateway operation. Most of the developers provide significant capability in this area.

Only one user stated a need to change the translation rules during runtime. This capability is sometimes used to prevent objects with inappropriate data from contaminating an exercise, or to stop updates from a particular sender that are causing interoperability/performance problems because of their participation. Although only one user stated a need for this capability, four of the developers provide a strong capability. This is potentially an example of a capability that is available from developers, but is not known by the users.

One new capability was identified during the data collection. Feedback was received from both developers and users regarding remote management of multiple gateways. This is a particularly important capability for users with a large number of enclaves. Because this capability was identified during the data collection, there is no data on the number of users expressing this need or the number of developers who provide it. There are no identified capability gaps in the User Interface capabilities; however, further research is needed to determine if there is a capability gap for remote management of multiple gateways.

4.2.2 Performance

The category of Performance contains execution performance capabilities. This is the set of performance metrics that are relevant for gateway users. There are no "correct" or "better" values for these capabilities. The values are based on the needs of the federation using the gateway. An example of a performance characteristic is the number of translated attributes per second.

Defining meaningful performance capabilities proved to be a significant challenge for the team. There are many complex and unique use cases that are needed to ensure all performance aspects are properly covered. Therefore, the team decided to ask the developers and users what performance characteristics were important to them. The difficultly in defining measurable performance capabilities lead to the inclusion of a gateway benchmark standard as a future option.

Some developers and users provided more than one answer to the question of performance characteristics. Table 4-8 lists the frequency of each answer. As the table shows, the top three performance characteristics for users and developers are the same. Although this is a small set of responses, it is still significant that the responses line up so well. Thus, users and developers have consistent perspectives.

Figure 4-1 and Figure 4-2: represent developer and user site responses to a question about the most important gateway performance characteristics. The question was stated as "What performance factors were considered most important for your gateway?" Even though prepared

multiple choice responses were not provided in either the on-line or the interview form, there was a high level of correlation of responses. As can be seen in Figure 4-1 and Figure 4-2:, there is a strong correlation between the user needs and the gateway support with regard to the priority and functionality for 1) latency, 2) throughput (number of object translations) and 3) robustness and stability, as shown in Table 4-8.

Table 4-8: High Correlation for Top 3 Performance Areas

Top 3 Developer Responses Regarding Performance	Top 3 User Site Responses Regarding Performance
1. Speed at which translations can be made (low latency)	1. Latency/Speed at which translations can be made
2. Number of object translations that can be supported	2. Number of entities / Number of object translations that can be supported
3. Robustness, stability, long running	3. Robustness, stability, long running, Mean Time Between Failures (MTBF)

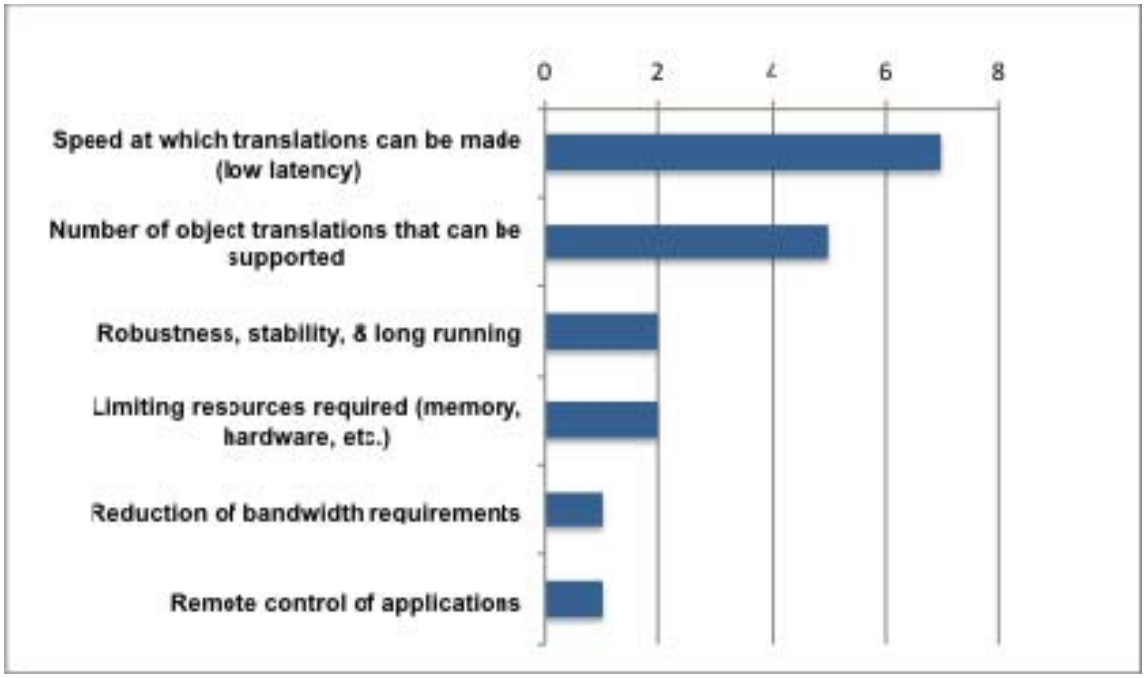

Figure 4-1: Developer Responses Regarding Performance (Number of Responses)

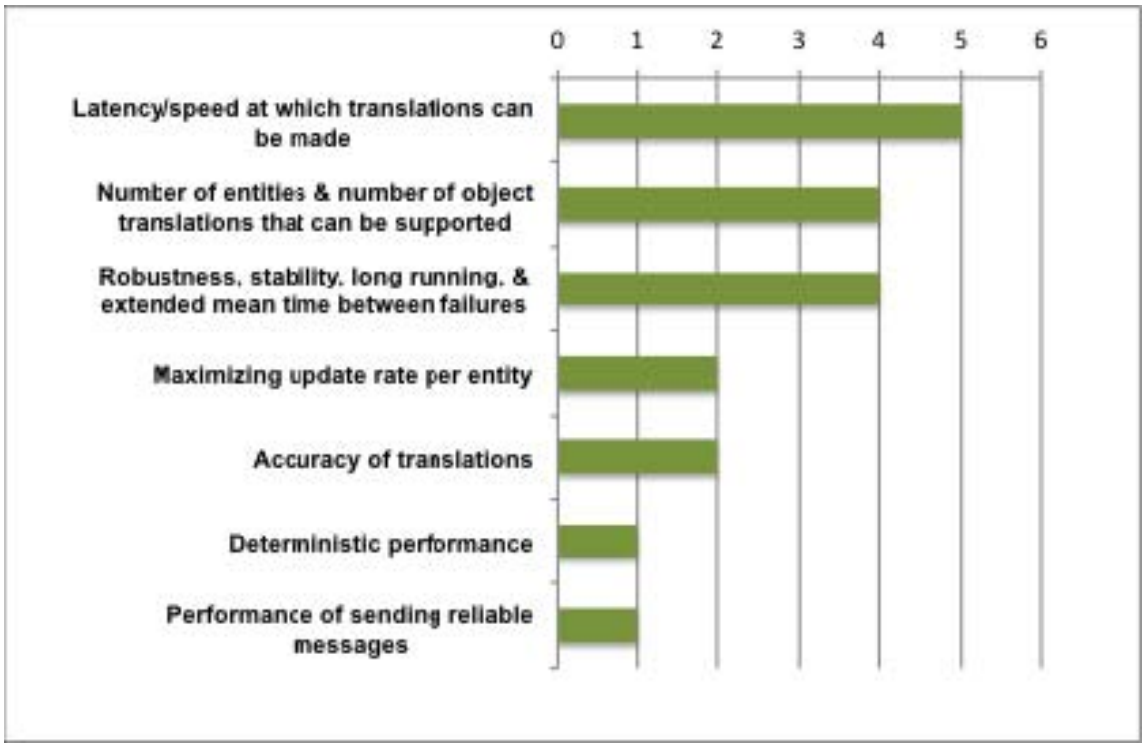

Figure 4-2: User Site Responses Regarding Performance (Number of Responses)

4.2.3 Operation Modes

The Operation Modes category is the set of possible operation modes for a gateway. A gateway may be able to operate in one or more modes. The modes are based on the ability of the gateway to perform translations between SDEMs and architectures. An example of an Operation Mode is only performing translations between SDEMs on a single architecture. The questions presented to the developers and users regarding Operation Modes, and their associated answers, are shown in Table 4-9.

Table 4-9: Responses Regarding Operational Modes

Question: Does the gateway have the ability to:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses **Y:** Yes **N:** No **?:** Did not know								
Translate between two different SDEMs that are using the same architecture (two-sided gateway)? **0:** does not support **5:** does support	0	0	5	5	5	5	0	0	0	5	5	Y	Y	N	Y	Y	Y	N	Y	N
Translate between two different SDEMs that are using two different architectures (two-sided gateway)? **0:** does not support **5:** does support	5	5	3[15]	5	5	0	?	0	5	5	5	Y	Y	Y	Y	Y	Y	Y	Y	Y
Translate between more than two different SDEMs that are using more than two different architectures (n-sided gateway)? **0:** does not support **3:** 3-sided **5:** n-sided	0	0	2[16]	0	5	0	0	0	0	5	5	N	N	Y	Y	Y	N	N	N	N
Translate between different versions of an architecture using the same SDEM? **0:** does not support **5:** does support	0	0	0	5	5	5	0	0	0	5	5	Y	N	N	Y	Y	N	N	Y	N
Translate between two different architectures using identical SDEMs? **0:** does not support **5:** does support	0	5	0	5	5	0	0	0	0	5	5	N	N	N	N	N	N	N	N	Y

Gateways can be thought of as having "sides," where each side conforms to a single SDEM/architecture. By far the most common user need is for a two-sided gateway where each side has a unique SDEM and architecture. This operating mode is supported by all of the general-purpose gateways.

Although two-sided gateways are the most common, multi- or n-sided gateways do exist. Two developers provide this capability. Only three users indicated a need.

There is strong user need for the capability of two-sided gateways to support different SDEMs within a single architecture. This capability is also well supported by the developers.

[15] The developer indicated that their gateway is able to translate between two variants of a SDEMs running on two variants of TENA.

[16] The developer indicated that their gateway is able to translate between DIS to TENA 6 or TENA 6 to TENA 5.2.

There is a user need for gateways which can support different versions of the same architecture. This situation occurs when a new version of an architecture is released, but all applications have not migrated to it yet. Three of the developers provide strong support for this operating mode. No capability gaps were identified for Operation Modes.

4.2.4 Extension Modes

The Extension Modes category includes the set of possible modes in which a gateway can be extended to support SDEMs and Architectures requirements for a particular end user. A gateway may support more than one extension mode. Some gateways may not be extendable, and may only support a specified set of SDEM/Architecture combinations. An example of an Extension Mode is requiring all translations to be hand-coded by the end user. The questions presented to the developers and users regarding Extension Modes, and their associated answers, are shown in Table 4-10.

Table 4-10: Responses Regarding Extension Modes

Question: How do you support users in extending the capability of your gateway:	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses								
By modifying its base code. (requires access to source code)	Y	Y	Y	N	N	Y	Y	Y	Y	Y	N	Y	N	?	Y	Y	Y	N	Y	Y
Via an API/SDK/plug-in	N	N	N	Y	Y	N	N	Y	N	Y	Y	Y	N	N	Y	N	Y	?	Y	Y
Extend core translations via a developer-provided code generator	N	N	?	Y	N	N	N	N	N	Y	Y	N	Y	N	N	Y	Y	Y	Y	Y
All translations via a developer-provided code generator	N	Y	N	Y	N	N	N	N	N	N	N	N	Y	N	N	Y	N	Y	Y	Y
Runtime translation performed using interpreted mapping files which user may write or extend?	N	N	N	N	Y	Y	N	N	Y	Y	Y	?	N	N	Y	N	Y	Y	Y	N
End users are not allowed to extend the gateway	Not asked											N	N	N	N	N	N	N	N	N
General Legend	**Y:** Yes **N:** No **?:** Did not know																			

The data on extension modes was collected to provide insight into how gateway developers intend for users to extend their gateways. The developers provided several different

approaches to this capability. Most users expressed a need to have the capability to extend the functionality of developer-provided gateways.

4.2.5 Platform

The Platform category is the set of hardware and software configurations to support the gateway. Some gateways may have multiple variable configurations. An example of a Platform characteristic is the types of supported operating systems. The questions presented to the developers and users regarding Platform characteristics, and their associated answers, are shown in Table 4-11.

Table 4-11: Responses Regarding Platform Characteristics

Question: What platform / software / OS do gateways that you have used run on or require?	Platforms: The set of hardware and software configurations to support the gateway (some gateways may have multiple variable configurations) Y: Yes N: No ?: Did not know								
For this gateway, what Hardware Platforms have you used? I – Intel, S – SPARC, M – MIPS, OT - Other, A - Any	I	A	I	I	I	ISM	I	I	I
For this gateway, what Operating Systems have you used? W - Windows, L – Linux, J – JDK, X - OS-X, R – RTOS, I – IRIX, A – Any, S – SPARC	LW	A	L	LWJ	LWR	LWSI	W	W	LW
For this gateway, what is the supporting software that is required that the user has to provide/install? C- C++/C, J - Java, O - Oracle, E – Eclipse, M – MySQL, R – RTI Lib, V – Visual Studio, T-TENA m/w, OT - Other	Y	Y	N	Y	Y	N	N	N	Y
	JR	VJM	n/a	CJ	J	n/a	n/a	n/a	JERT
What supporting hardware is required to use the gateway? M - Multiple Network Interface Cards, E - Excess Memory, OT – Other	N	Y	N	Y	Y	N	N	N	Y
	n/a	M	n/a	M	M	n/a	n/a	n/a	M

Data on the platforms supported by developers and used by the users was collected for informational purposes. Therefore, there did not appear to be any gaps, as there seems to be two primary platforms in use by the community at this time.

4.2.6 Documentation and Support

The Documentation and Support category is defined as the different types of support available from the gateway developer, configuration manger, or sponsor. An example of a Documentation and Support characteristic is the level of user documentation. The questions presented to the developers and users regarding Documentation and Support characteristics, and their associated answers, are shown in Table 4-12.

Most gateway developers provide a variety of support materials. The provided support generally meets user needs.

Table 4-12: Responses Regarding Gateway Documentation and Support

Question: What types of support are available for your gateway?	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses									
Does the gateway have associated documentation available?	5	2[17]	2[17]	5	5	5	5	2	5	4	3	Does the gateway have associated documentation available? If so, have you used it?	Y/Y	Y/Y	Y/Y	Y/Y	Y/Y	Y/Y	Y/N	Y/Y	Y/Y
	0: no documentation **1:** informal only **3:** 1 formal document **5:** multiple formal documents																				
Does the gateway have the Help Desk / Support available?	5	0	3	3	5	4	4	0	2[18]	1	3	Does the gateway have the Help Desk / Support available? If so, have you used it?	Y/Y	Y/Y	N	N	Y/Y	Y/Y	Y/N	Y/Y	Y/Y
	0: no help desk **1:** email w/o funded support **3:** email with funded support **4:** email & telephone support, manned during business hours **5:** funded on-site support																				
Does the gateway have an associated Feature Request form or process?	3	0	0	3	5	5	5	0	0	5	5	Does the gateway have an associated Feature Request form or process? If so, have you used it?	Y/Y	Y/Y	N	N	Y/N	Y/N	Y/N	Y/Y	Y/Y
	0: no feature request process **3:** informal feature request process **5:** formal feature request process																				
Does the gateway have an associated Problem Report form or process?	5	0	3	3	5	5	5	0	5	5	5	Does the gateway have an associated Problem Report form or process? If so, have you used it?	Y/Y	?	N	Y/Y	Y/Y	Y/Y	Y/N	Y/Y	Y/?
	0: no problem report process **3:** informal problem reporting **5:** formal problem report process																				
What is the business model for your gateway (**G:** GOTS, **C:** COTS, **M:** Mixed GOTS/COTS, **A:** Alternative)?	G	A	A	G	M	G	G	A	G	G	C	How did you obtain the gateway (**G:** GOTS, **C:** COTS, **D:** Developed)?	G	G	G	G	GCD	G	G	G	G
Is a license [& payment] required?	N	N	N	N	Y	N	N	N	N	N	Y	Was a license [& payment] required?	N	N	Y	Y	N	N	N	N	Y
General Legend	**Y**: Yes **N**: No											**Y/Y**: Yes, Yes **Y/N**: Yes, No									

[17] The developer indicated that the documentation supplied includes a User Manual and additional information; however, the documentation that is available is not a formal set of documents.

[18] The developer indicates that email and telephone support is available. This support is funded for ARCIC-led experiments only.

4.2.7 Maturity

The Maturity category is defined as the different levels of software maturity and use of the gateway. An example of a Maturity characteristic is the number of organizations using the gateway. The questions presented to the developers and users regarding Maturity characteristics, and their associated answers, are shown in Table 4-13.

Data on the maturity of gateways was collected for informational purposes.

Table 4-13: Responses Regarding Gateway Maturity

Question: How extensively is your gateway used?	RDGW	GMU	NSUT	GWB	VR-EX	DISF	TSR	SCORM	BLCSE	JBUS	PITCH	User Site Responses								
How many programs / federations has your gateway been used on?	1	0	1	1	3	1	1	0	1	3	4	+	?	+	+	?	+	+	+	+
How many users of your gateway are there (number of deployments)?	2	0	2	2	4	5	3	0	2	4	4	Many	?	3	?	?	?	?	?	Many
How many exercises has your gateway been used in?	1	1	1	2	3	5	5	1	1	3	4	Many	?	8	?	?	12	3	6/yr	4/yr
Does the gateway have a formal sustainment or maintenance process? Frequency?	Y	N	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	N	Y 3/yr	N	Y	Y	Y	Y	?
What was the last version update of the gateway that you are using (is version known)?												V	?	?	V	V	V	?	?	V
Has any Verification & Validation (V&V) been performed on your gateway?	Y	?	Y	Y	Y	N	N	N	?	Y	Y	N	N	Y	N	N	N	?	N	N
If Yes for V&V performed, what type? Formal, Informal (e.g., testing), or by Customers	I	?	I	I	C	n/a	n/a	n/a	?	C	C									
Legend	**0**: None **1**: 1 to 5 **2**: 6 to 10 **3**: 11 to 30					**4**: 31 to 50 **5**: more than 50 **Y**: Yes **N**: No						**+:** Other programs are known **Y**: Yes **N**: No **V:** Version known **?**: Did Not Know								

This page intentionally left blank.

5. SUMMARY

In modern multi-architecture distributed simulation environments, gateways play an important role in reconciling the various differences among the different simulation architectures. Despite the success many programs have achieved with respect to their gateway implementations, the results of LVCAR Phase I have helped to identify several problems and inefficiencies with respect to how gateways are applied in today's LVC simulation applications, which result in undesirable increases in cost and technical risk. The objective of the LVCAR Phase II Gateways and Bridges task is to examine these perceived problems and offer solutions that will allow multi-architecture distributed simulation environments to be built "better, faster, and cheaper" in the future.

An important element of defining the "to be" state of where the DoD would like to get to with respect to gateway technology is to fully define the "as is" state. Thus, the main purpose of this report has been to characterize the gateways in wide use across the DoD today. Since this characterization indirectly defines current user requirements for gateways, this report will provide an important input into the Execution Plan that will dictate the "way ahead."

This characterization effort has resulted in a total of eleven key findings. Listing them in order of most agreement between developers and users to least agreement produces the following delineation.

1. A capability gap in SDEM Translations does not exist.
2. A capability gap in Architecture Translations does not exist.
3. A capability gap in Architecture Behaviors does not exist.
4. No capability gaps were identified for Operation Modes.
5. In regard to Support, the provided support generally meets user needs.
6. In terms of Performance, users and developers have a consistent perspective. When asked to provide the critical performance characteristics, the top three characteristics for both developers and users matched.
7. Overall, regarding Exercise Management Behaviors, the match between developer-supplied capabilities and user-needed capabilities is strong.
8. In regard to Platform Support, there did not appear to be any gaps, as there seem to be two primary platforms in use by the community at this time.
9. There are no identified capability gaps in the User Interface capabilities; however, further research is needed to determine if there is a capability gap for remote control of multiple gateways.
10. In regard to Extension Modes, most users expressed a need to have the capability to extend the functionality of developer-provided gateways.

11. In regard to SDEM Behaviors, there is a disconnect between the developer indicated capability and the user-expressed needs, particular in the areas of dynamic information translation where one SDEM requires it and the other does not, and supporting differing behaviors between SDEMs. As for the latter, about half of the users indicated a need for this capability, but no developers reported providing this capability. This is either an actual gap in the capability required by the user or a misunderstanding of the capability by either the users and/or developers.

Collectively, these conclusions seem to suggest that there are few real gaps in capability, and that existing gateways are satisfying most user needs. However, these results also highlight how key capabilities are spread out among many different gateways, and that few gateways are both general-purpose and fully-featured. Also, potential users of gateways have no convenient way to compare the features and capabilities offered by different gateways (other than this report), or even to know what gateways are available for reuse. These additional concerns will be assessed in conjunction with the data from this report to produce a complete picture of the “as is” and “to be” states, which in turn will provide the foundation for the roadmap between these two states as defined in the Gateways and Bridges Execution Plan.

APPENDIX A: REFERENCES

The following documents were the sources for the technical details in this report.

(a) Kuhl, F., Weatherly, R., & Dahmann, J. (1999, October 31). Creating Computer Simulation Systems: An Introduction to the High Level Architecture, Upper Saddle River, NJ: Prentice Hall.

(b) Richbourg, R., and Lutz, R. R. (2008, September). "Live Virtual Constructive Architecture Roadmap (LVCAR) Comparative Analysis of the Architectures," Alexandria, VA: Institute for Defense Analyses.

(c) Simulation Interoperability Standards Committee. (2001, March 9). "Standard for Modeling and Simulation High Level Architecture - Federate Interface Specification," *IEEE Std IEEE1516.1-2000.* HLA Working Group.

(d) Simulation Interoperability Standards Committee. (2000, December 11). "Standard for Modeling and Simulation High Level Architecture - Framework and Rules," *Institute of Electrical & Electronics Engineers, Inc. (IEEE) Std IEEE1516-2000.* HLA Working Group.

(e) Simulation Interoperability Standards Committee. (2001, March 9). "Standard for Modeling and Simulation High Level Architecture - Object Model Template Specification," *IEEE Std IEEE1516.2-2000.* HLA Working Group.

(f) TENA Software Development Activity. (2005, February). "TENA The Test and Training Enabling Architecture - Architecture Reference Document," Version 2005.

This page intentionally left blank.

APPENDIX B: GATEWAY-RELATED GLOSSARY OF TERMS

The following terms were clarified according to the specific characteristics needed for a clear understanding of the analysis. Additionally, this gave the questionnaire takers a clearer understanding of the terms used in the questions.

Term	Definition	Example
Simulation Data Exchange Model (SDEM)	The set of data exchanged during the execution of a federation.	RPR FOM, MATREX FOM, and TENA Standard OM
Architecture	The mechanism to transfer data between simulations.	DIS, HLA, TENA, CTIA
SDEM Translations	The set of capabilities to convert data from one SDEM to another SDEM. This includes unit conversion, coordinate conversion, and enumeration mapping. This may require translating data stored in a single object in one SDEM to multiple objects in another SDEM.	Unit conversion on a single attribute, single element enumeration to a multi-element enumeration.
SDEM Behaviors	The capability to correctly represent behaviors required by the SDEM. This may include representing a behavior in one SDEM that is present in the other SDEM. SDEM behaviors may require the gateway to maintain state.	Dead reckoning of positional attributes.
Functional Capabilities Categories	Functional capabilities represent actions required by the gateway to meet SDEM, Architecture, or exercise management needs.	SDEM Translations, SDEM Behaviors
Architecture Translations	The set of capabilities to convert architecture-defined data between different architecture executions. This covers translation of data that is not defined in the SDEM, but is present in all architecture executions.	Object identifiers, and timestamps
Architecture Behaviors	The set of capabilities to perform actions required by the architecture. These behaviors are required by the architecture and apply to all SDEMs using the architecture.	Requests to publish attributes of a specified object

Term	Definition	Example
Exercise Management Behaviors	The set of capabilities that are not directly related to the SDEM or architecture. These capabilities are used to meet objectives of the overall exercise independent of the SDEM or architecture. These capabilities may not involve the translation of data or behaviors.	Filtering for bandwidth
User Interface	The set of capabilities that allow the user to interact with the operation of the gateway during pre-exercise, exercise, and post-exercise activities. This category includes display of real-time information to the user.	Execution-time view of translated objects
Performance	The set of performance metrics that is relevant for gateway users. There are no "correct" or "better" values for these capabilities. The values are based on the needs of the federation using the gateway.	Translated attributes per second
Operation Modes	The set of possible operation modes for a gateway. A gateway may be able to operate in one or more modes. The modes are based on the ability of the gateway to perform translations between SDEMs and Architectures.	Translation between SDEMs on a single architecture
Extension Modes	The set of possible modes in which a gateway can be extended to support SDEMs and Architectures. A gateway may support more than one extension mode. Some gateways may not be extendable, and may only support a specified set of SDEM/Architecture combinations.	All hand-coded translations
Operational Capabilities Categories	Operational capabilities describe features of the gateway not directly related to translation.	Operation Modes, Extension Modes
Platform	The set of hardware and software configurations to support the gateway. Some gateways may have multiple variable configurations.	Supported operating systems
Documentation / Support	The different types of support available from the gateway developer, configuration manager, or sponsor.	Level of user documentation
Maturity	The different levels of software maturity and use of the gateway.	Number of user organizations

APPENDIX C: ABBREVIATIONS AND ACRONYMS

AB	Aktiebolag (Swedish: Limited Company on shares)
AFB	Air Force Base
API	Application Programmers Interface
ARCIC	Army Capabilities Integration Center
BLCSE	Battle Lab Collaborative Simulation Environment
C2	Command and Control
COTS	Commercial off the Shelf
CTIA	Common Training Instrumentation Architecture
DIS	Distributed Interactive Simulation
DISF	Distributed Interactive Simulation (DIS) Filter
DMOC	Distributed Mission Operations Center
DoD	Department of Defense
EXCIS	Extensible C4I Instrumentation Suite (US Army)
GERTICO	German Run-Time Infrastructure based on CORBA
GMU	George Mason University
GOTS	Government off the Shelf
GWB	Gateway Builder
HLA	High Level Architecture
HQ	Headquarters
IEEE	Institute of Electrical & Electronics Engineers, Inc.
InterTEC	Interoperability Test and Evaluation Capability
JBUS	Joint Simulation Bus
JDK	Java Development Kit
JFCOM	Joint Forces Command
JHU/APL	The Johns Hopkins University Applied Physics Laboratory
JIMM	Joint Integrated Mission Model
JMETC	Joint Mission Environment Test Capability
JSAF	Joint Semi-Automated Forces
KABIC	A component of the TENA White Sands Missile Range, the format for which the InterTEC's Range Data gateway supports
LVC	Live-Virtual-Constructive
LVCAR	Live-Virtual-Constructive Architecture Roadmap

MaK	VT MÄK, a company of VT Systems Inc.
MATREX	Modeling Architecture for Technology, Research, and Experimentation
MIPS	Microprocessor without Interlocked Pipeline Stages
M&S	Modeling and Simulation
MTBF	Mean Time Between Failures
MySQL	A relational database management system acquired by Oracle and Sun corporations
m/w	Middleware
n/a	Not Applicable
NACTS	Nellis Air Combat Training System
NSUT	Network System Under Test
OneSAF	One Semi-Automated Forces
OTB	OneSAF Test Bed
PAX	Patuxent
PEO-STRI	Program Executive Office Simulation, Training and Instrumentation
PITCH	Pitch Technologies AB
R&D	Research and Development
RDECOM	US Army Research, Development, and Engineering Command
RDGW	Range Data Gateway
RDS	Range Data Systems
RPC	Remote Procedure Call
RTC	Redstone Test Center
RTI	RunTime Infrastructure
SCORM	Sharable Content Object Reference Model
SDEM	Simulation Data Exchange Model
SDK	Software Development Kit
SIMAF	Simulation and Analysis Facility
SIMDIS	A trademarked name for an analysis and display tool developed by the Naval Research Laboratory for simulation visualization
SIMPLE	Standard Implementation for Multiple Platform Link Evaluation
SMMAI	Shared Memory Multi-Architecture Interface
SPARC	Scalable Processor Architecture
SQL	Structured Query Language for database access
SSEGW	Sensor Signal Emulator Gateway
TACCSF	Theater Air Command and Control Simulation Facility
TAIS	Tactical Airspace Integration System
TECCS	TENA Expeditionary Command and Control Suite
TENA	Test and Training Enabling Architecture
T&E	Test and Evaluation

TRADOC	Army Training and Doctrine Command
TSIU	Tactical Systems Interface Unit
TSR	TACCSF Software Router
US	United States
VMASC	Virginia Modeling, Analysis, and Simulation Center
VR	Virtual Reality
VR-EX	MaK Virtual Reality (VR)-Exchange
V&V	Verification & Validation
WSMR	White Sands Missile Range

This page intentionally left blank.

NATIONAL SECURITY ANALYSIS DEPARTMENT
THE JOHNS HOPKINS UNIVERSITY • APPLIED PHYSICS LABORATORY
Johns Hopkins Road, Laurel, Maryland 20723-6099