LVC Common Gateways and Bridges

David L. Drake, Robert R. Lutz
JHU/APL
Laurel, MD
david.drake@jhuapl.edu, robert.lutz@jhuapl.edu

Kurt Lessmann
Trideum Corporation
Huntsville, AL
klessmann@trideum.com

Dannie E. Cutts
The AEgis Technologies Group, Inc.
Huntsville, AL
dcutts@aegistg.com

Michael J. O'Connor
ITT Corporation
Huntsville, AL
michael.oconnor@it.com

ABSTRACT

The Live Virtual Constructive (LVC) Architecture Roadmap (LVCAR) report sponsored by the Modeling and Simulation (M&S) Steering Committee recommended several actions to promote the sharing of tools, data, and information across the Enterprise and to foster common formats and policy goals to promote interoperability and the use of common M&S capabilities. One of these recommended actions pertained to the development of a common suite of gateways and bridges capabilities.

This paper describes the Common Gateways and Bridges aspects of LVCAR implementation. The paper initially presents the outcome of a series of gateway and bridges community interviews, questionnaires, and workshops that gathered information on existing gateway and bridge capabilities and limitations. Following this is an analysis of gateway and bridge implementation paradigms. Lastly, an examination of emerging software technologies is presented that may lead to improving gateway and bridge robustness, performance, and reusability. In conclusion, an implementation plan for the incremental development, testing, and distribution of common gateway/bridge capabilities is presented that addresses the issues that were highlighted in our research.

The anticipated execution of this plan will reduce the costs associated with large multi-architecture LVC events and will result in more timely and robust test and training environments for our warfighters in the future.

ABOUT THE AUTHORS

DAVID L. DRAKE is a Senior Technical Staff in the Global Engagement Department at Johns Hopkins University Applied Physics Laboratory performing research and development in the area of modeling and simulation supporting the U.S. government and military. Mr. Drake is currently the lead on the Live Virtual Constructive Service-Oriented Architecture Task, and supports the LVC Architecture Roadmap Implementation tasks for Common Gateways and Bridges and Architecture Convergence. Mr. Drake has 8 years experience in modeling and simulation design, development, and standards, and 23 years as a computer security professional in computer security design, implementation and evaluation. Mr. Drake received his Bachelors in Mathematics from SUNY at Buffalo. He is published in the areas of simulation, service-oriented architecture, grid computing, security, and risk assessment and has a patent on the process for enterprise-wide intrusion detection.

ROBERT R. LUTZ is a Principal Staff Scientist at The Johns Hopkins University Applied Physics Laboratory (JHU/APL). He has over 29 years of experience in the design, implementation, and evaluation of computer modeling and simulation (M&S) systems for military customers. He received his M.S. degree in Operations Research from the State University of New York at Stony Brook in 1980. Since joining JHU/APL in 1992, Mr. Lutz has assumed leadership roles on several M&S programs, including the Naval Simulation System (NSS), Joint Warfare System (JWARS), and the Simulation Based Acquisition (SBA) initiative. Mr. Lutz also served as the technical editor for IEEE 1516.2 (HLA Object Model Template) and as working group chair for IEEE 1516.3 (HLA FEDEP). Currently, he is supporting a variety of M&S initiatives in support of the U.S. Modeling and Simulation Coordination Office (MSCO), the U.S. Joint Forces Command (USJFCOM), and the Navy's Broad Area Maritime Surveillance (BAMS) Program. He also serves as a guest lecturer in The Johns Hopkins University Whiting School of Engineering.

DANNIE E. CUTTS is a Senior Computer Scientist with AEgis Technologies Group Inc. supporting the USJFCOM Technical Development & Innovation (TD&I) branch in Suffolk, VA. He has over 20 years experience in Modeling and Simulation for NASA and the DoD and has been involved with the High Level Architecture (HLA) since 1995, serving on the Interface Specification and Time Management Working Groups. He has provided HLA Training and Cadre support for DMSO, and currently serves on the IEEE Drafting Group for the HLA IEEE 1516 standard. Mr. Cutts is a Certified Modeling and Simulation Professional (CMSP). He has supported numerous federation development efforts, as well as projects bringing legacy and new simulations to HLA Compliance. At USJFCOM he is involved in efforts to improve interoperability between Live, Virtual and Constructive assets for Joint Training.

KURT LESSMANN is a Vice President at Trideum Corporation, specializing in providing solutions and services in the area of test and evaluation, modeling and simulation, systems analysis, and information technology. Mr. Lessmann has over 15 years experience in supporting Live, Virtual, and Constructive (LVC) integration activities for large scale Joint Distributed Test activities. Mr. Lessmann also supports the test community by providing Object Model design, software and range system design and integration, and application of M&S techniques. He currently supports the Test and Training Enabling Architecture (TENA) project and the Joint Mission Environment Test Capability (JMETC) Program Office as a senior Event Lead. Mr. Lessmann has a Bachelor of Aerospace Engineering from Auburn University.

MICHAEL J. O'CONNOR is a Senior Principal Engineer for ITT Corporation's Information Systems division. Mr. O'Connor has more than 20 years experience in modeling and simulation. He has been a key participant in the development of distributed modeling and simulation standards including IEEE 1278 and IEEE 1516. Mr. O'Connor was the editor for the SISO Real-time Platform Reference Federation Object Model (RPR FOM). He has held many positions in the community including Chairman of the SISO Standards Activities Committee and Vice-Chairman of the SISO Executive Committee. Mr. O'Connor oversees the development of ITT's Chemical, Biological, Radiological, and Nuclear Simulation Suite, which supports DIS, HLA, and TENA. He has a Bachelors of Computer Engineering from Auburn University and a Master's of Science in Computer Science from the University of Alabama in Huntsville.

LVC Common Gateways and Bridges

David L. Drake, Robert R. Lutz
JHU/APL
Laurel, MD
david.drake@jhuapl.edu, robert.lutz@jhuapl.edu

Kurt Lessmann
Trideum Corporation
Huntsville, AL
klessmann@trideum.com

Dannie E. Cutts
The AEgis Technologies Group, Inc.
Huntsville, AL
dcutts@aegistg.com

Michael J. O'Connor
ITT Corporation
Huntsville, AL
michael.oconnor@it.com

1 BACKGROUND

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 Study 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, in the LVCAR Implementation Project, the Common Gateways and Bridges task began with three major activities: performing gateway and bridge literature research, investigating gateway and bridge usage and development practices, and developing formal gateway and bridge operation terminology. To create a technical foundation, an initial delineation of gateway and bridge capabilities was created. After completing this task, it became clear that the distinction between "gateway" and "bridge" was moot from a development and usage standpoint. The difference between a gateway and bridge was primarily a matter of how the system was configured rather than how the system was coded. For that reason, the term "gateway" is used for both types of applications in our analysis, and also in this paper.

Starting with the initial delineation of capabilities, a Gateway Capabilities Matrix Template was compiled. This template allowed the creation of 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 in the creation of a candidate list of developers and users to fill out the questionnaires.

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 validation of the initial list of capabilities, and to identify additional capabilities. Much like a census, the goal was to gather the information about the existing gateways, not determine a "best" gateway or to compare gateways. It was 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, developers and users would be asked 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, a one-day workshop, referred to as the "LVCAR Common Gateways and Bridges Workshop," allowed the presentation of the questionnaires' findings. The workshop was held at the Virginia Modeling, Analysis, and Simulation Center

(VMASC) in Suffolk, VA, on March 4, 2010. The workshop not only allowed verification of its results with the community, but it also allowed additional feedback and data gathering.

The eleven gateway developers that responded to the questionnaire are listed in Table 3-1. The architectures addressed by the gateways included the Test and Training Enabling Architecture (TENA), Distributed Interactive Simulation (DIS), and the High Level Architecture (HLA), as well as a collection of other simulation- or communication-specific protocols.

The ten gateway users that responded to the questionnaire are listed in Table 3-2, and the range of the gateways that they employed are listed in Table 3-3. Seven of the gateways were employed by more than one of the users, as shown in the table.

The questionnaires covered approximately 80 questions regarding gateway characteristics, use, and goodness of fit to the task. From the responses, a good deal of information was gained, particularly if the gateway developers were designing gateways that were in sync with the gateway users needs.

3 GATEWAY CHARACTERIZATION

3.1 SDEM Translations

Simulation Data Exchange Model (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.

SDEM Translations are one of the most fundamental activities of a general-purpose gateway. The set of capabilities defined in this category 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 between what the developers are providing and what users need.

Table 3-1: Gateway Developers Responding to the Questionnaire

Gateway Developed	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)
George Mason University Gateway	Shows how selectively reliable multicast can provide efficient network group communications	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)	Allows the generation of 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	Provides a 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

Table 3-2: Participant Gateway Users

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

3.2 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 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.

Table 3-3: List of Gateways Employed by Multiple Users

Gateways Referenced by Users [number of users referencing gateway]
JMETC Gateways developed using Gateway Builder (GWB) [6]
MaK Gateway Products [4]
InterTEC NSUT Gateway [3]
SSEGW [2]
BLCSE Middle Ware [2]
Command & Control (C2) Adaptor Gateway [2]
JBUS [2]
HLA Adaptor in OneSAF [1]
JSAF Gateway [1]
SIMPLE Gateway [1]
DIS Interoperability Manager for OneSAF [1]
RTC developed Shared Memory Multi-Architecture Interface (SMMAI) [1]
Gateway for interconnecting Trideum Corporation HLA to TENA to GERTICO [1]
EXCIS (message translation) Gateway [1]
TSIU MATREX platform to TAIS (air track) Gateway [1]
Global Magic TENA to DIS Gateway [1]
OTB-based Gateway [1]
JIMM to Shared Memory plug-in architecture [1]

This is most likely the result of the specific SDEMs in use by the users.

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 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.

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.

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.

3.3 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.

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.

3.4 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.

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.

3.5 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.

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.

3.6 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.

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.

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.

3.7 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/exercise 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. There are many complex and unique use cases that are needed to ensure all performance aspects are properly covered. Therefore, the decision was made 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. 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. In the final comparison across the desired performance characteristics, the top three developer responses and the top 3 user responses had a high correlation, as shown in Table 3-4.

Table 3-4: 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)

3.8 Operational 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.

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 for this capability.

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.

There is a user need for gateways that 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.

3.9 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 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.

3.10 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.

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

4 ASSESSMENT OF GATEWAYS

Despite the success many simulation programs have achieved with respect to their gateway implementations, the LVCAR Implementation project has helped to identify several problems and inefficiencies with respect to how gateways are applied in today's LVC simulation applications, which may result in undesirable increases in cost and technical risk. The on-going endeavor to examine these perceived problems and offer solutions that will allow multi-architecture distributed simulation environments to be built “better, faster, and cheaper” is still ahead.

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 analysis has been to characterize the gateways in wide use across the DoD today, determining where the capability gaps exist, which 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, or even to know what gateways are available for reuse. These additional concerns will be assessed 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 an execution plan for supporting the gateway community.

5 IMPLEMENTATION PLAN FOR INCREMENTAL IMPROVEMENTS

At the completion of the data gathering and analysis aspects of this task, there was wide agreement that there are several potential improvements that can be made to lower the technical and cost risks generally associated with the use of gateways. Such improvements, along with the authors' recommendations, are three strategies for execution: *informing* the community as to gateway existence and capabilities; *enhancing* the efficiency and effectiveness by which existing and future gateway products are applied; or *creating* new gateways components or systems. For completeness, we also include a *status quo* strategy, which describes the impact on the DoD modeling and simulation communities if no action is taken. Below is a description of these strategies.

The first strategy presented involves taking a *laissez-faire* approach to the growth and maturation of the development and use of gateways in the modeling and simulation community, aptly named the "Status Quo" strategy. This strategy acts as a base case for the other strategies. It describes the outcome of not taking any action and allowing the current market forces to continue shaping the industry. The immediate benefit of using the Status Quo strategy is the lack of any new requirements for DoD investment and the fact that it minimizes any potential disruption to the existing LVC community.

The second strategy ("Inform") focuses on educating the user community about existing gateway availability and capabilities. This knowledge would assist potential gateway users in making better-informed decisions when considering use of a gateway. Understanding what gateways are available and what capabilities they have could reduce proliferation of gateways by promoting reuse of existing products vice building additional one-off gateway solutions. The execution of this strategy involves a number of education options, such as workshops, tutorials, and training courses. An illustration of its tasks and their dependencies are shown in Figure 5-1.

The third strategy ("Enhance") incorporates several of the fundamental elements defined in the previous strategy but extends the purely educational focus with several products intended to make more effective use of the gateway capabilities that exist today. Examples of products identified in this strategy include a common Gateways Description Language (GDL) to allow gateway capabilities to be described in a machine-readable form, a set of Gateway Performance Benchmarks (GPB) that would provide a common way of assessing the relative ability of competing gateways to provide needed capabilities, and a common Gateway Configuration Model (GCM) that would provide a standard means of initializing, tailoring, and configuring gateways. Widespread adoption of these products in the LVC community will result in users making much better choices as to the gateway products they use in their applications, and will also assist users with how to best employ these products to minimize technical and cost risks to their projects. An illustration of its tasks and their dependencies are shown in Figure 5-2.

The fourth strategy ("Create") is focused on creating new capabilities to meet the gateway needs of users. These new capabilities can range from relatively minor extensions or enhancements to existing gateways to whole new gateways (or gateway capability libraries) to address capability gaps and provide users with a common interface and common configuration processes/tools. While representing the most expensive option, it also provides the highest potential return. However, achieving this return is dependent on the degree of market penetration that is achieved with new gateway products. An illustration of its tasks and their dependencies are shown in Figure 5-3.

Of the four strategies, it was recommended that the "Enhance" strategy be executed, given that it has the highest perceived ROI. A project plan is presented that details the duration and dependencies of the various tasks in this strategy and estimates the level of effort for execution of this strategy.

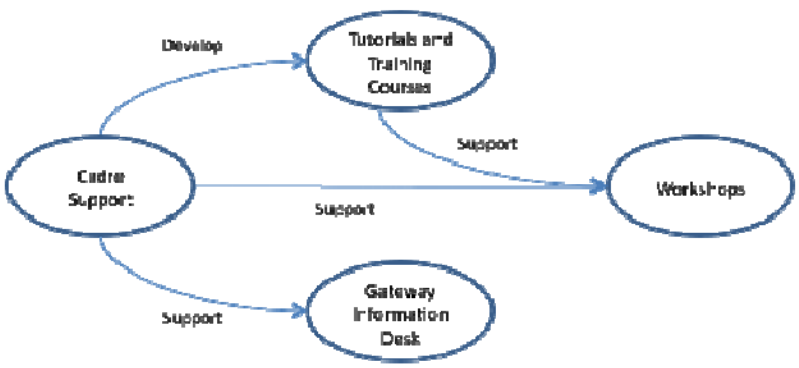

Figure 5-1. Dependency Relationships for the Inform Tasks

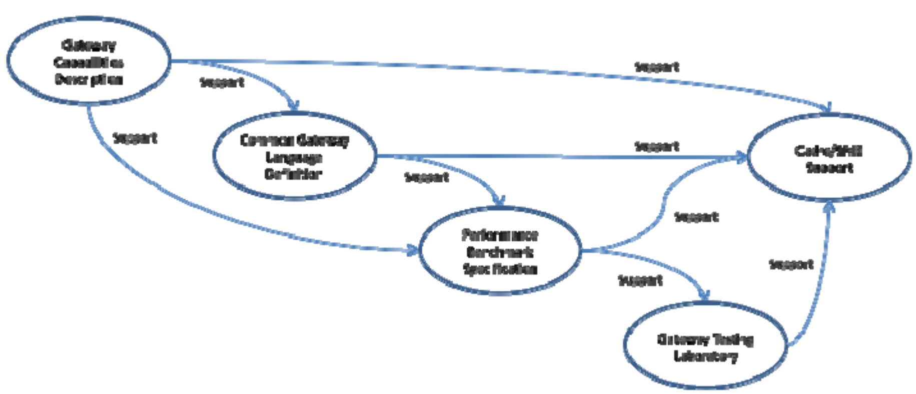

Figure 5-2. Dependency Relationships for the Enhance Tasks

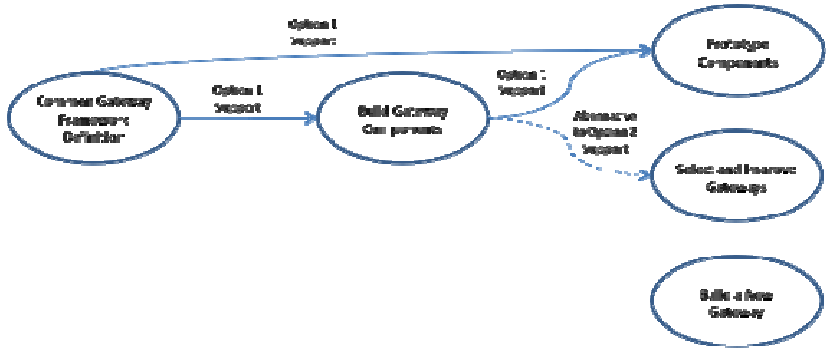

Figure 5-3. Dependency Relationships for the Create Tasks

6 SUMMARY

This paper provided an analysis of the current characteristics of gateways as they are developed, the characteristics needed by gateway users, and the existing capability gaps. From this analysis options were presented as to a way forward, with the recommendation of an "Enhance" strategy as the best technical- and cost-effective approach.

ACKNOWLEDGEMENTS

This work was performed in support of the LVCAR Implementation task sponsored by the Department of Defense Modeling and Simulation Steering Committee (M&S SC) under High Level Task S-C-1 and managed by Dr. Gary Allen of the Joint Training Integration and Evaluation Center.

REFERENCES

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.

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.

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.

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.

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.

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