The Live Virtual Constructive (LVC) Architecture Roadmap: Foundations from the Past and Windows to the Future

Amy Henninger, Ph.D.
IDA
Alexandria, VA
ahenning@ida.org

Bob Richbourg
IDA
Alexandria, VA
rrichbou@ida.org

Robert Lutz
JHU APL
Laurel, MD
Robert.lutz@jhuapl.org

Randy Saunders
JHU APL
Laurel, MD
Randy.Saunders@jhuapl.org

Margaret Loper
GTRI
Atlanta, GA
Margaret.Loper@gtri.gatech.edu

ABSTRACT

Combining numbers of simulations in a single large exercise requires that the inter-communications architectures (e.g., HLA - High-level Architecture, TENA - Test & Training Enabling Architecture, etc.) that those systems use must also be present in the exercise. When systems that use different inter-communications architectures interact, the interaction must be transmitted between the architectures as a lossless communication. To many, this will seem a trivial consideration, given the commonplace and broad experience with communication between systems using the internet. However unexpected it may sound, HLA interactions will not be understood by simulation systems using TENA unless additional effort (typically, adding a gateway) is made supporting appropriate translation.

This small example symbolizes a broader but closely-related problem set that has impacted large simulation exercises for years. The Live-Virtual-Constructive Architecture Roadmap (LVCAR) Study was chartered to develop a future vision and supporting strategy for achieving significant interoperability improvements in LVC simulation environments, reducing the problem set to the trivial challenge that many internet-experienced consumers expect. The study addressed three main areas of concern; the desired future integrating architecture(s), the desired business model(s), and the manner in which standards should be evolved and compliance evaluated. For each area, the study provided near-, mid-, and long-term recommendations that together constitute a roadmap to guide the evolution of LVC architecture development to achieve a more seamless environment.

This paper reviews the study's assumptions, fundamental precepts, and conclusions and presents them as integral parts of a plan now being carried out. The paper also provides a view into the reasoning behind the study's recommendations and concludes with a description of the future for simulation architectures.

ABOUT THE AUTHORS

Dr. Amy E. Henninger is a member of the Research Staff in the Joint Advanced Warfighting Division at the Institute for Defense Analyses. She's worked as an analyst, developer, researcher and manager on a variety of DoD M&S programs for over 20 years. Henninger is formerly an Adjunct Faculty Member for the M&S Graduate Program at the University of Central Florida and has written over 45 papers on a variety of topics (e.g., neural networks, autonomous agents, multi-agent systems, embedded training systems, analytic methods for uncertainty, games for training, and performance metrics for intelligent systems). She holds six university degrees culminating in a Ph.D. in Computer Engineering.

Dr. Robert Richbourg is a member of the Research Staff in the Joint Advanced Warfighting Division at the Institute for Defense Analyses. He is a retired Army officer who earned his Ph.D. in computer science in 1987. In his last active duty assignment, he was an Academy Professor and Director of the Artificial Intelligence Center at the United States Military Academy, West Point. He has been working in the area of simulation environments for 15 years under sponsorship of DARPA, M&SCO, and STRICOM, JFCOM, and NGA.

Mr. Robert Lutz is a Principal Staff Ccientist at The Johns Hopkins University Applied Physics Laboratory in Laurel, MD. He has over 29 years of experience in the design, implementation, and evaluation of modeling and simulation (M&S) systems for military customers. Currently, he is leading the LVCAR "Systems Engineering Process" and "Common Gateways and Bridges" Phase II tasks, serves as the Navy's Broad Area Maritime Surveillance (BAMS) Program M&S lead in the Airspace Integration (AI) area, and supports several M&S standards initiatives within the Simulation Interoperability Standards Organization (SISO). He also serves as a regular guest lecturer in The Johns Hopkins University Whiting School of Engineering.

Dr. Margaret L. Loper is the Chief Scientist for Georgia Tech Research Institute's Information Technology & Telecommunications Laboratory. Margaret has more than twenty years of experience in M&S. Her technical focus has been on parallel and distributed simulation and she has contributed to the areas of temporal synchronization, simulation testing, and simulation communication protocols. Margaret's past projects include leading the standards study team for the Live Virtual Constructive Architecture Roadmap study and simulation cloning research for Deep Green. She earned a Ph.D. in Computer Science from the Georgia Institute of Technology, an M.S. in Computer Engineering from the University of Central Florida, and a B.S. in Electrical Engineering from Clemson University.

Mr. Randy Saunders is a Principal Staff Engineer at the Johns Hopkins University Applied Physics Laboratory. He has over 20 years of experience in the design, implementation, and integration of high-fidelity simulations for military and business customers. He received M.S. degrees in Engineering from Harvey Mudd College in 1980 and in Computer Science from the University of Southern California in 1985. Mr. Saunders has been involved in distributed simulation standardization since the first DIS Workshop, both DIS and HLA standards committees, and as a member of the SISO Standards Activities Committee. He is presently a vice-chair of the HLA Evolved Product Development Group. He joined APL in 2001.

The Live Virtual Constructive (LVC) Architecture Roadmap: Foundations from the Past and Windows to the Future

Amy Henninger, Ph.D.
IDA
Alexandria, VA
ahenning@ida.org

Bob Richbourg
IDA
Alexandria, VA
rrichbou@ida.org

Robert Lutz
JHU APL
Laurel, MD
Robert.lutz@jhuapl.org

Randy Saunders
JHU APL
Laurel, MD
Randy.Saunders@jhuapl.org

Margaret Loper
GTRI
Atlanta, GA
Margaret.Loper@gtri.gatech.edu

1.0 INTRODUCTION

In April, 2007, the DoD Modeling and Simulation Steering Committee (M&S SC) began sponsorship of the Live, Virtual, Constructive Architecture Roadmap (LVCAR) study. The M&S SC recognized that the modeling and simulation community had achieved great successes since the beginning of the SIMNET (Miller and Thorpe, 1995) program that allowed linkage of critical resources through distributed architectures. In part, the success was predicated on an iterative and evolutionary development of the intercommunication architectures, with progressive capabilities enhancements that allowed more varied and larger application of the technologies across multiple user domains. While the architectures displayed impressive capability to meet needs as designed, they were not implemented with a focus on ensuring architectural compatibility with existing capability. Thus, each time systems using different architectures required interconnection as parts of a larger simulation event, substantial design and engineering effort was required to achieve cross-architecture interoperability. Given this environment, the LVCAR study was chartered to:

"…*methodically and objectively develop a recommended roadmap (way forward) regarding LVC interoperability across three broad areas of concern: notional definition of the desired future architecture standard, the desired business model(s), and the manner in which standards should be evolved and compliance evaluated.*"

During the study, emphasis was placed on analysis of the technical options that would achieve or make transparent architecture interoperability. A comprehensive analysis of these technical requirements across all of the architectures illustrated that there was a high degree of commonality between them, particularly HLA (Kuhl, et al, 1995) and TENA (Powell, 2005). While a few key differences have been indicated in the specifications of requirement for these architectures, a considerable amount of capability overlap (considering only major characteristics) is evident (see Figure 1). At the implementation level, however, there are substantive differences among the architectures. Such differences are characterized as "wedge issues", potentially becoming barriers to achieving cross-architecture interoperability and are discussed in Richbourg and Lutz (2008).

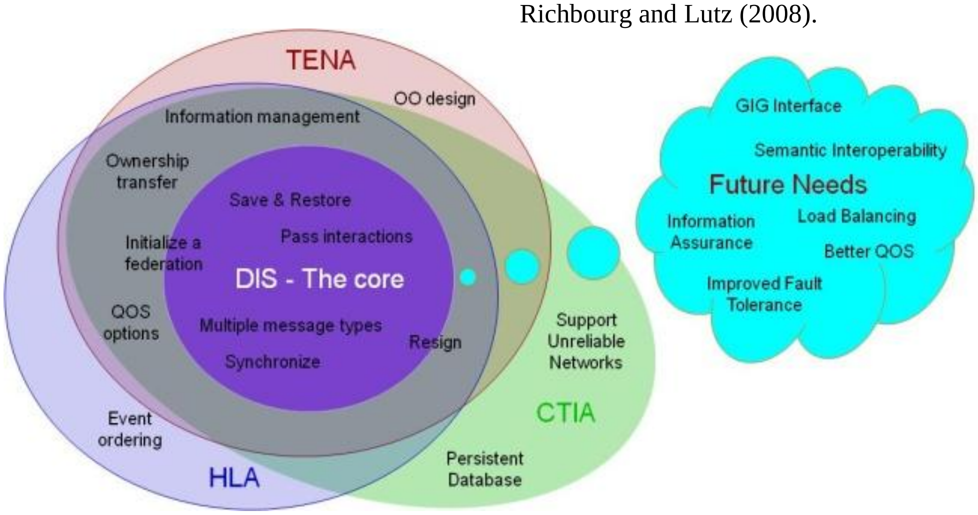

Figure 1. High-level Capability Overview

2.0 UNDERSTANDING THE PROBLEM

There is a perception by many in the LVC community that interoperability would be much easier (and less costly) if only there was a single architecture available for use. Included in this perception is the concept that the Department would benefit by eliminating the costs associated with maintaining overlapping capabilities. The desire to achieve a single-architecture state is based on a number of difficulties in the current situation that can be directly attributed to the existence of multiple architectures.

First, in many cases, mixed-architecture events can only use the set of capabilities common across all of the architectures to be included in the event (although this is as much a simulation system limitation as it is an mixed architecture constraint). Further, the costs required to integrate architectures rarely contribute directly to achieving simulation event goals. Instead, the associated costs usually provide point solutions, versions of which have likely been created in the past and probably will be paid for again in the future. Mixed architectures also impede "plug-and-play"; it is slower, more expensive, and sometimes impractical to compose simulation events using any of the wide range of assets (e.g., simulations, simulators, labs, ranges, C4ISR systems) available in the DoD inventory. In this view, one cannot simply choose an asset based on functional merit alone in all cases; frequently, the asset also is constrained to be compatible with a specific architecture. Typically, the associated cost cannot be ignored, so event designers will not even consider incompatible assets and will thus design events without having a complete picture of the resources that might have been used.

However, while each of these disadvantages can be attributed to the existence and use of multiple architectures, their existence does not necessarily justify an assumption that ridding the DoD of all but one architecture would result in an optimal state of affairs. There are at least five main factors suggesting that such an assumption is fallacious. First, legacy systems will continue to be used and it is unlikely that these systems will upgrade to using a new or different architecture. Thus, use of legacy systems is most likely to preclude the possibility of ever achieving a truly "single-architecture" state. Second, use of a single architecture may still require the use of supporting bridges, much as use of different RTIs can require bridges today. Third, gateways will be required for connecting any single simulation architecture to C4I systems, to the GIG, or, in general, to any type of system that has a primary purpose outside the simulation arena. Fourth, the alignment of a family of simulations on a single architecture represents a single point solution. Even if such standardization were attained, history points to the likelihood that the diverse group of simulation users would quickly diverge into specializations, leading to the need for gateways to bridge their differences. Fifth, the selection or creation of a single architecture assumes that the rapid advances of the commercial software industry will not lead to a better implementation in the future. When this does occur, the existing standard architecture would be abandoned by users who have needs for the superior architecture delivered by the commercial sources.

The concurrent existence of multiple architectures may allow benefits that are less likely to be achieved in a single architecture state. These include: 1) the ability to support multiple business and standards-use communities simultaneously; and 2) fostering the capability to "use the right tool for the job", avoiding the "one size fits all" problem.

In summary, there are advantages and disadvantages associated with having a number of architectures available for use. There is no paramount advantage or disadvantage that allows one to immediately recognize the best possible solution. A significant problem for the LVCAR roadmap effort is to navigate this trade space to arrive at an achievable solution that maximizes the benefit for all concerned while not exceeding the resources that will be necessary to realize that solution.

3.0 CANDIDATE STRATEGIES

The LVCAR Strategy decision space considered in the analysis of alternatives (AOA) (see Figure 2) was based on requirements and capability analyses. These five strategies represent an expert team consensus-based opinion of the possible high-level approaches to addressing LVC interoperability issues from which a roadmap of lower-level actions and activities could be derived. These strategies are elaborated on below. As indicated in the figure, several of these strategies were rejected early in the analysis. The rationale for strategy retention and elimination is provided in Section 6.0.

3.1 Strategy 1: Maintain the Status Quo

In this strategy, no specific actions are taken to unify the current distributed simulation architectures. This can be thought of as the "natural selection" or "distributed, uncoordinated management" strategy, which recognizes that the various architectures will evolve as needed to meet the future needs of each user base, and that when mixed architecture environments are required, the current (but admittedly ad hoc, inefficient, and decentralized) approach of using gateways and bridges will eventually become good enough to meet future needs.

3.2 Strategy 2: Enhance Interoperability of Mixed Architecture Events

In this strategy, the focus is to create solutions to improve the interoperation of existing architectures in a mixed-architecture environment. Examples of such solutions include establishing standard agreements (e.g., processes, terminology, object models) that cut across the various architectures and improving the performance, reliability and (re)usability of future gateways and bridges. The individual architectures would evolve to support their native user base.

3.3 Strategy 3: Encourage and Facilitate Architecture Convergence

This strategy is very similar to the preceding strategy, with the exception that policy actions and investment incentives would be added to cause the architectures to converge into either a single architecture or a reduced set of compatible and interoperable architectures. Thus, while the same roadmap actions would be taken with regard to improving both model and runtime interoperability in the near-to-mid term, this strategy would include additional actions as necessary to achieve some level of architecture convergence (including the potential for physical convergence) at a specified future date.

3.4 Strategy 4: Select One of the Existing Architectures

In this strategy, an evaluation of how well existing individual architectures satisfy all identified requirements (including projected future requirements) would be conducted, and the architecture that represents the "closest fit" to future needs would be chosen as the foundation of a single future architecture for LVC. Actions at that point would focus on adding (hopefully reusing from other architectures) features and capabilities needed by users that are not currently included in the chosen architecture, and on instituting policy and financial incentives to convince affected users to transition.

3.5 Strategy 5: Develop a New Architecture

In this strategy, an entirely new architecture would be developed based on current and future requirements for LVC environments. While reuse of the best ideas and implementations from existing architectures would be encouraged where appropriate, this strategy is intended to be a new start to incorporate emerging technologies and modern design paradigms into the baseline architecture structure rather than than trying retrofit such ideas into existing architecture(s). Policy and financial incentives would be used to spur adoption of the new architecture.

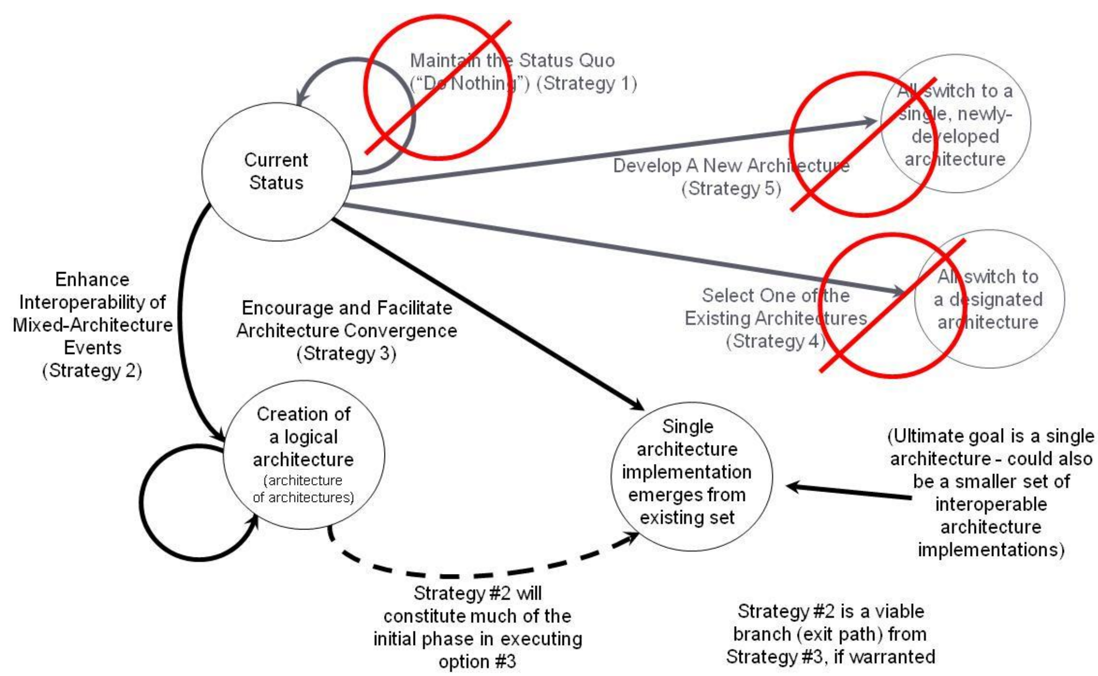

Figure 2. Candidate Strategies for the Analysis of Alternatives

4.0 OBSERVATIONS

The current state includes a wide range of user communities, and different architectures and protocols are used across those communities, with no single architecture dominant (see Figure 3). There is a range

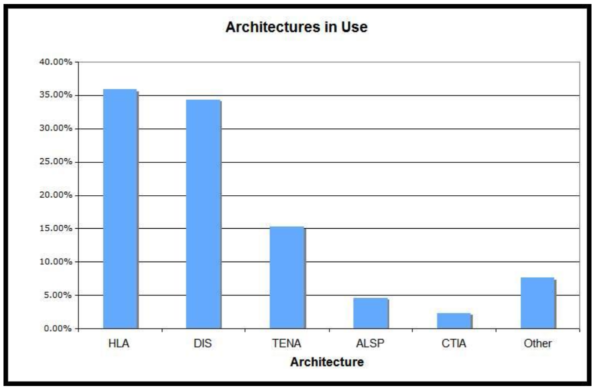

Figure 3. Sample of Architectures Being Used[1] (August, 2007)

of qualitative factors that must be considered to understand the current state; these also have implications for producing an informed decision for the best way forward. The following factors (or assertions) represent practical considerations regarding the application of distributed simulation architectures within the LVC community today, and are considered factual by the communities represented on the LVCAR effort.

- *Much can be accomplished with the architectures that are available today and nearly all of the existing architectures are being improved to better serve their communities of use.*
- *The Department of Defense has not always taken and is not currently using a consistent, coherent approach to managing LVC environments.*
- *The number of available architectures has increased since the early 1990's, at least partially, as a result of inadequate management.*
- *Mixed-architecture environments occur as dictated by needs of the using applications, not because of any inherent benefit in mixing architectures.*
- *When mixing architectures is necessary, point solutions to bridging the architectures work in most cases where syntactic interoperability is the main concern, although these kinds of solutions may introduce additional latency and information loss for some applications.*
- *Mixed-architecture approaches may introduce certain limitations on the range of services available to participants within the full simulation environment.*
- *Many legacy, and even some new, simulations will not transition to using a different architecture, unless there are compelling incentives to do so.*
- *GOTS-based and COTS-based business approaches are difficult to reconcile within the scope of a single product.*
- *Cultural and resource issues will be persistent barriers to convincing existing architecture users to switch to a different architecture.*
- *Architectural choices of how to transfer data between applications (syntactic issues, the concern of the LVCAR study) and application-level choices of how to interpret received and encode transmitted data (semantic issues, beyond the scope of the study) both have impacts on interoperability.*
- *Significant improvements in LVC interoperability can also be achieved via supporting data, tool, and process standards.*

In short, the currently available architectures were deemed to be generally meeting the primary needs of their constituent communities today and are evolving to meet future needs as well. History shows that the number of available architectures tends to increase over time and that once a community of use develops around an architecture, that architecture is very likely to continue to be used. By definition, the inter-architecture communication problem only occurs during mixed-architecture events. And, while there are advantages and disadvantages associated with the number of architectures available for use, there is no paramount advantage or disadvantage that allows one to immediately recognize the optimal number, given the current user-split across the architectures.

Figure 4 characterizes the high-level trade-space in two axes (Control and Marketplace) from the perspective of the DoD enterprise. In this model, Control represents the degree of influence the DoD corporate level has over an architecture and its related business and standards practices and Marketplace represents the degree to which the architecture, including its corresponding business and standards processes, promotes cross-stakeholder and cross-user participation.

[1] ALSP (Wilson and Weatherly, 1994) is no longer being used.

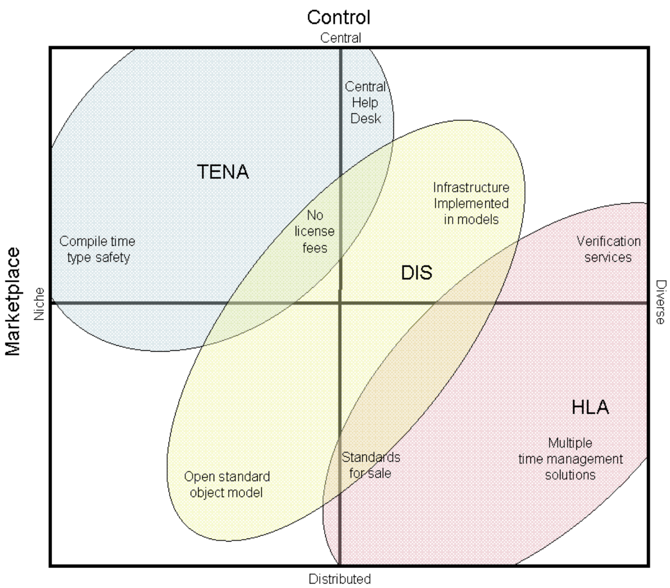

Figure 4. Conceptual Similarities, Contrasts and Progressions of LVC

By plotting the three major competing architectures on these axes, similarities, contrasts and progressions are visible. For example, characteristics such as "compile time type safety" and "(Government) Central Help Desk" are clearly present in the TENA model, but nowhere else. Characteristics such as "infrastructure implemented in models" and "open standard object model" are clearly present in the DIS (Hofer and Loper, 1995) model, but nowhere else. And, characteristics such as "multiple solutions for time management" and "verification services" are present in the HLA model, but nowhere else. In contrast, some characteristics are common. For example, the characteristic "no license fee" exists in both the TENA model and the DIS model. And, characteristics like "standards for sale" are shared by both the DIS and the HLA model.

This diagram also communicates trends in historical progressions. Namely, whereas DIS seemed to provide a good middle ground, HLA adjusted to improve diversity but ultimately at the expense of architecture ownership, and then TENA adjusted in the opposite direction to improve ownership but in so doing limited diversity. A significant problem for the LVCAR roadmap effort was to navigate this trade space to arrive at an achievable solution that maximizes the benefit for all concerned while not exceeding the resources that will be necessary to realize that solution.

4.1 Levels of Interoperability. Many of the known problems that impact LVC integration stem from technical incompatibilities among the various distributed simulation architectures. However, achieving the goal of a truly interoperable LVC operating environment requires that developers consider a wide range of issues beyond the basic question of how to pass runtime data along the simulation network. As examples, issues related to data modeling, coordinate systems, synthetic natural environment representation, and algorithmic consistency are frequently outside the scope of the problems that simulation architectures were ever designed to address. The consideration of such issues as part of a structured systems engineering methodology is critical if executions of the LVC environment are to produce valid results.

Tolk and Muguira (2003) identify five basic levels of simulation interoperability, with each level building on preceding levels. LVC interoperability is affected across all five levels. The distributed simulation architectures in use today all provide services for achieving technical and syntactic interoperability (e.g., levels 1 and 2); however, problems with how these services interact at runtime can adversely affect interoperability in mixed architecture environments. While solutions can be found to such problems, the LVCAR stressed the idea that most practical distributed simulation applications require interoperability at levels above the syntactic level. Addressing interoperability issues at the semantic level (and above) frequently transcends the architectures, and generally involves the establishment of cross-community agreements and standards on such supporting resources as data, processes, and tools. Thus, although the primary focus of the LVCAR was on the syntactic-level issues of mixed architecture integration, the general desire to reduce the technical, cost, and schedule risks associated with developing and operating future LVC environments also requires the consideration of higher-level interoperability issues.

5.0 FUNDAMENTAL PRECEPTS

During development of the desired state properties, implementation strategies, and subsequent recommendations for the LVCAR Study, the Expert Team converged on a core set of beliefs, axiomatic 'meta-recommendations' that underscore the approach and provide guiding principles for implementation and execution of the roadmap. These principles, presented below, represent the four fundamental precepts to the LVCAR Study Final Report (Henninger et al, 2008).

5.1 Fundamental Precept #1: Do No Harm. The DoD should not take any immediate action to discontinue any of the existing simulation architectures. There is a considerable degree of consensus within the

LVC user community that a long-term strategy based on architecture convergence would benefit the DoD. However, it is also understood that there are many cost and design issues that must be resolved prior to implementing such a strategy, and that the actual implementation needs to be a well-planned, deliberate, evolutionary process to avoid adversely impacting participating user communities. Because of these considerations, it would be unwise to eliminate support for any of the existing simulation architectures in the near-term. Rather, as the differences among the architectures are gradually reduced, it should be the users themselves that decide if and when it is appropriate to merge their architectures into some smaller set based on both technical and business concerns. Any attempt by the DoD to mandate a convergence solution on an unwilling user base is certain to meet strong resistance and likely to fail.

5.2 Fundamental Precept #2: Interoperability is not Free. The DoD must make the necessary investments to enable implementation of the activities described in the LVC Roadmap. LVC interoperability is not free. It is not reasonable to expect that LVC interoperability goals can be met with little or no investment. Since the return on LVC investments is nearly impossible to accurately quantify in the near-term, it is understood that major new up-front investments are difficult to justify. In recognition of this fact, the Roadmap has taken a long-term approach which requires only limited investment early in its implementation, with subsequent investments dependent on demonstrable progress. Without the necessary investments, the LVC Roadmap is nothing more than a blueprint of what is possible to accomplish, with no mechanism to realize the associated benefits.

5.3 Fundamental Precept #3: Start with Small Steps. The DoD should take immediate action to improve interoperability among existing simulation architectures. The vast range of technical problems currently associated with the development and execution of mixed-architecture LVC environments is well recognized. Such problems increase the technical risk associated with the use of these mixed-architecture environments, and require considerable resources to address. While architecture convergence would lessen (and even eliminate) several of these problems, it is not practical to expect any significant degree of convergence to occur for many years. Instead, LVC users need near-term solutions that reduce both cost and technical risk until such time as architecture convergence can occur. These solutions include actions such as improved gateways/bridges, common object models, and common development/execution processes. Many of these solutions can be implemented at a moderate cost, and provide significant near- and mid-term value to the LVC community.

5.4 Fundamental Precept #4: Provide Central Management. The DoD must establish a centralized management structure that can perform Department-wide oversight of M&S resources and activities across developer and user organizations. A strong centralized management team is necessary to prevent further divergence and to effectively enable the architecture convergence strategy. This team needs to have considerable influence on the organizations that evolve the existing architectures, and must also have influence on funding decisions related to future LVC architecture development activities. Without centralized DoD management, existing architecture communities will continue to operate in line with their own self-interests, and the broader corporate needs of the DoD will be treated as secondary issues that are likely to continue to be ignored as concerns that are not germane to the local problems.

6.0 STRATEGY ASSESSMENT

The initial strategy assessment involved the evaluation of the various qualitative factors and other more practical considerations discussed earlier in this paper. The first, most basic architecture strategy is Strategy 1, "Maintain the Status Quo.” The major concern with respect to this strategy is that the current inefficiencies and excessive integration and resource requirements inherent in today’s mixed-architecture environments are intolerable even in the present, and the natural divergence of existing architectures will continue to degrade the situation. Clearly, this strategy does nothing to improve the current situation, by definition.

Strategy 5, "Develop a New Architecture", suffers from the impracticality of its implementation, at least in the near-term. Most who work in the distributed M&S community have experienced the clear failure of policy mandates and, based on that experience, are convinced that the similar tactics are unlikely to work in the future. Further, incentives to users of existing architectures to transition will be ineffective if the users believe that their existing architecture solution is already fully meeting their needs (which most do). Developing a new architecture to replace the existing ones will more likely result in yet another architecture being added to the existing set; history indicates that few, if any, architectures will be retired as the new one comes on-line. Thus, the combination of potentially high up-front costs, long development time, and (currently) weak justification for an entirely new LVC architecture effectively eliminates this strategy as a viable near-term solution.

Strategy 4, "Select One of the Existing Architectures”, suffers from many of the same issues as the "Develop a New Architecture" Strategy. That is, users of existing architectures seem to be generally satisfied with them, and significant investments have already been made in supporting infrastructure. The users of architectures that were not the “one selected” would be asked to switch to an architecture that they may have already rejected or otherwise deemed to be cost ineffective. Thus, there would likely be significant resistance to migrating to an externally-designated architecture when that migration would require new investments (e.g., software, personnel training, supporting infrastructure) and may not work as well as their current architecture within their domain. Short of a policy mandate (which has historically been shown to be ineffective), an orderly transition is unlikely to happen. Thus, while this strategy could become a natural outgrowth of the "Encourage and Facilitate Architecture Convergence" Strategy, simply choosing an existing architecture and compelling other architecture users to migrate was considered to be impractical, ineffective, and unworkable as a near-term solution. As such, this strategy was also eliminated from further consideration.

Having eliminated these three strategies from further consideration, only Strategy 2 (“Enhance Interoperability of Mixed-Architecture Events”) and Strategy 3 (“Encourage and Facilitate Architecture Convergence”) remained as viable candidates for further evaluation. Based on analyses and expert opinion, these two strategies appeared to be the most promising for several reasons. Principally, they lack the impracticality of the other strategies while still providing needed improvements in LVC interoperability. Both are designed to prevent continued architecture divergence, and both provide sufficient flexibility to allow mid-term course corrections if evolving user requirements suggest a deviation. Both Strategy 2 and Strategy 3 could readily branch to either Strategy 4 or Strategy 5, or to one another, as the most desirable long-term course. In addition, both provide for the existence of an oversight body that will have vision over all of the separate groups that currently manage and evolve the existing architectures.

Strategy 2, “Enhance Interoperability of Mixed-Architecture Events”, is founded on the idea that having multiple architectures available for use is desirable and that the best way forward is to take actions that can reduce or eliminate the barriers to interoperability between the existing architectures and protocols. More specifically, this strategy acknowledges that the existing architectures have been created, have evolved, and are being maintained to meet the specific needs of their constituent communities. Elimination of any architecture should only occur as a natural result of disuse. Modification and management of the existing architectures is left to the owning communities as the best option to ensure meeting the needs of the various user communities, both throughout the DoD and among the Department’s coalition partners. To resolve the interoperability problems, efforts should be directed towards creating and providing standard resources, such as common gateways, common componentized object models, and common federation agreements, which can render integration of the multiple architectures an efficient and nearly transparent process. In effect, these actions will create the perception of a single architecture that supports all the diverse simulation systems, even though the systems will actually be serviced by an “architecture of architectures”, composed of as many different architectures and protocols as are required to interconnect the participating simulation systems.

Strategy 3 differs from Strategy 2 in the method of addressing the multi-architecture problem, the role of the oversight body, and in the primacy of effort. First, Strategy 3 takes the view that the many problems inherent in allowing redundant architectural capability to exist clearly outweigh the associated benefits. The existence of multiple architectures is a problem that must and shall be resolved. Gradual convergence of the architectures is a viable strategy to resolve the problem. As an example, if two architectures are so common in their capabilities that there is little, if any, significant technical difference between them, then those two architectures should be gradually converged into a single architecture. Similarly, if the complete set of technical capabilities offered by one architecture is a subset of the capabilities provided by another, the “smaller” architecture should be gradually converged into the larger one. Thus, Strategy 3 seeks to manage gradual convergence of the entire set of architectures where appropriate. Eventually, the convergence process could result in either a single architecture or a smaller set of compatible, interoperable architectures. Managing convergence in this way requires an oversight body that can influence the evolution of the architectures, using a combination of policy, incentives, and disincentives to shape the actions taken by the user communities that control architecture evolution. Finally, while Strategy 3 includes the same actions as Strategy 2 that will immediately reduce the costs and problems that arise when integrating multiple architectures, the primacy of effort is given to achieving convergence of the existing architectures.

As part of the final AoA, each of candidate strategies was decomposed and evaluated with a cost-benefits analysis by an expert team. The measure of "utility" applied was "Return on Investment" (ROI), which was defined as the ratio of the relative benefit of each activity versus the cost to perform each activity. "Benefit" in this case was defined to be the amount of savings that is achievable in future LVC environment development effort as a result of having performed that activity. So, for instance, if performing Activity Z results in a set of reusable products that will drive the cost of building an LVC environment down from "X" dollars to "Y" dollars in the future, then "X – Y" (multiplied by the number of expected LVC events) provides a reasonable measure of benefit for Activity Z. "Cost" in this context was the man-years associated with implementing the activity. The "breakeven point" is that point in time at which benefit equals cost. From a decision-maker's perspective, those activities that have early breakeven points and whose benefits increase rapidly with respect to cost in future years will be considered the most desirable.

To provide context to the ROI evaluation activity, the HLA Federation Development and Execution Process (FEDEP, IEEE [2003]) was employed in this analysis as a framework for identifying where certain costs and benefits are incurred while constructing LVC environments. While other process models could have been used, the FEDEP was chosen due its broad applicability to all simulation architectures and status as a recognized IEEE standard.

The input data to this analysis was primarily based on surveys from LVCAR Expert Team and Working Group members. Although real financial data from actual LVC user programs would have been a preferred source of information for this analysis, such data is not generally extractable from the Work Breakdown Structures (WBS) used by most programs. While the use of SME opinion rather than real cost data can be considered a possible criticism of this analysis, it should be emphasized that the surveys were specifically targeted to people with substantial experience in developing LVC environments, and thus the collective opinions of the targeted group was determined to provide the best possible estimation of both activity costs and the savings achievable via implementation of the various activities.

The results of this ROI evaluation exercise led to a "best-of-breed" that included some elements of both Strategy 2 and Strategy 3, in effect producing a blended strategy that includes the most desirable components of both. In addition to optimizing ROI, this combination also ensured that the strategy provided the most robust posture and the agility necessary to adapt to changing situations (e.g., C4I systems, impact of Global Information Grid, embedded training in operational systems, Multi-Level Security). The following section showcases some of the high-level recommendations resulting from the AoA, and provides detail on some of the implementation-level activities.

7.0 ROADMAP RECOMMMENDATIONS

This section provides the summary focus area recommendations that drive the investment roadmap activities along with general related recommendations. The LVCAR provides the complete development for each recommendation, but that rationale is not summarized in this paper given the length of the discussions.

7.1 Strategic Level Recommendations

These recommendations are further elaborated with supporting rationale and data throughout the final report. The DoD should:

<u>Technical Architecture</u>

- *Take actions that can reduce or eliminate the barriers to interoperability*
- *Direct efforts towards creating and providing standard resources, such as common gateways, common componentized object models, and common federation agreements*
- *Provide a free highly-customizable and well-documented set of gateway products to the LVC user community*
- *Move beyond the debate of technical interoperability and start focusing on the semantics of these systems*

<u>Standards</u>

- *Develop adequate spheres of influence in relevant standards organizations (e.g., SISO) and related communities (C4I, DISA, etc.)*
- *Develop standards evolution processes that can provide required stability, yet be flexible and responsive to users*

Business Model

- *Identify an LVC Keystone*[2] *to gather and disseminate information across the M&S community, representing a unified consensus of opinion*
- *Balance the marketplace across architecture approaches so that investments are made in terms of their overall benefit to the DoD enterprise*
- *Evaluate the potential impact of ongoing open source RTI efforts on the interoperability of M&S systems across DoD and consider the suitability of open source as a mechanism for balancing the marketplace*
- *Identify influential Federation Proponents (JNTC, NCTE, JMETC, large PEOs, etc) to integrate emerging developments in support of future architectural solution(s)*

General

- *Provide resources to address LVC issues that are not directly architecture-related (e.g., semantic interoperability, conceptual modeling, etc.)*
- *Lead efforts to standardize or automate translations of data/scenario inputs to simulations and data capture formats*
- *Provide technical positions in support of M&S enterprise decisions*
- *Develop and implement processes that support solid, performance-based decision-making to evaluate the efficacy of the roadmap, make mid-course corrections, and develop the next-generation of goals*

7.2 Tactical Level Activities

Table 1 below focuses exclusively on DoD-level investments, the Roadmap Activities. These are seen as common goods particularly worthy of DoD-level attention. Table 1 recommendations fall largely into three categories: architecture, business model, and standards. In tune with Fundamental Precept #4, the central management must direct technical efforts to perform the roadmap activities.

The Architecture activities are designed to enhance the interoperability of mixed-architecture events, while preserving options and positioning the community for some degree of architecture convergence in the future. As stated earlier, the activities are founded on the idea that having multiple architectures available for use is desirable and that the best way forward is to take actions that can reduce or eliminate the barriers to interoperability between the existing architectures and protocols. Effectively, the activities are designed to transparently create an "architecture of architectures" when necessary, while leaving the individual architectures relatively unchanged so that they continue to provide uninterrupted service to their constituent communities.

Table 1. Implementation Level Activities

Type	Activity
Architecture	Common components of architecture-independent object models
	Describe and document a common, architecture-independent systems engineering process
	Create common, reusable federation agreement template
	Analyze, plan and implement improvements to the processes and infrastructure supporting M&S asset reuse
	Produce and/or enable reusable development tools
	Investigate – Convergence feasibility determination and design
	Convergence plan
	Convergence implementation
	Produce common gateways and bridges
	Specify a resource or capability to facilitate pre-integration systems readiness
Standards	Make IEEE standards more accessible to LVC community.
	Engage SISO and the broader LVC community
	Coordinate activities and fund participation in commercial standards development groups
	Investigate - Increase sphere of influence in Standards Development Organizations
	Develop evolutionary growth path for LVC standards
Business	Identify LVC Keystone
	Investigate – Balance the marketplace
	Balance the marketplace
Mgmt	Decision Support Data

The Architecture work also places great emphasis on the need to expand the Department's vision for M&S interoperability by moving beyond the debate of

[2] See Swenson (2008) for a full explanation of the term. The central player in a healthy ecosystem is the keystone organism, which serves as the leader of the ecosystem.

technical interoperability and encouraging focus on the semantics of these systems (Richbourg et al, 2008). This more elegant focus will direct us to a path towards improving the effectiveness[3] of LVC applications, as well as the costs of LVC applications. Technical interoperability has been a problem, but it is clearly tractable; solutions to the technical interoperability problems exist and they should no longer consume all of our attention. From this point forward, the technical vision for the next phase of LVC in DoD must raise the bar.

Getting to the point where the bar can be raised, however, would seemingly be well served by a shift in business practices. Currently, M&S development and use is spread across a large number of program elements and authority for executing those funds is spread across an equally large number of organizations. There is no single organization that controls both policy and funding under a single mission umbrella. The differences in institutional investment and cost of entry for the users have resulted in a marketplace including an array of somewhat redundant key products that cannot compete on technical merit alone. The Business Model activities are designed to move the costs and control of the architectures and related tools to a common environment where access and risk are spread across a greater constituency. This also improves the potential for innovation and reduces barriers to entry. Thus, the Business Model work makes a case for harnessing the power of M&S intellectual capital and focusing diverse fiscal resources through the instantiation of a common workspace to share architecture and tool advancements and to serve as a unifying place for change to happen.

For change to propagate, however, adequate spheres of influence in relevant standards organizations and related communities (e.g., C4I, DISA) must be developed. This will better ensure that DoD interests are well served. Also, standards processes must be coordinated to provide the required stability, while preserving flexibility and responsiveness to users. The Standards activities are designed to develop this organizational influence, promote flexible standards evolution processes, and build a sense of community.

Finally, to measure the effects of these changes and plan for the future, the M&S SC requires improved decision-making data. This includes data from the technical domain, business domain, and standards domain. While the LVCAR did not focus on management or leadership issues, it did recognize and address the need to provide improved decision support data for management use.

[3] It will improve the validity of analyses and reduce the possibility of negative training.

8.0 CLOSING REMARKS

Ultimately, the goal must be an environment in which the M&S SC can leverage its millions to influence the billions spent on distributed M&S and LVC across the Department. This is possible. Microsoft, for example, has profound influence over the information technology (IT) marketplace; yet, "both its revenue and number of employees represent about 0.05% of the total" resources in that marketplace. (Iansiti and Levein, 2004) This example suggests that it is possible for a central M&S agency with a budget of merely $35M to have a substantial influence on the estimated $10B (Cuda and Frieders, 2007) spent annually on M&S in the DoD.

On 16 July 2007, Congress passed House Resolution 487, "recognizing the contribution of modeling and simulation technology to the security and prosperity of the United States, recognizing modeling and simulation as a National Critical Technology" and commending members of the modeling and simulation community in government, industry, and academia who have contributed. We believe that Congress has a vision for M&S in the United States and we believe that the DoD, as a corporate entity, can either be a driving force in shaping that vision or can go along for the ride. The vision for this Roadmap is for the DoD, as a corporate entity, to be a driving force in the way forward for distributed M&S and LVC as a technology supporting the security and prosperity of the United States.

ACKNOWLEDGEMENTS

This work was performed in support of the Modeling and Simulation Standards Process task sponsored by the Department of Defense Modeling and Simulation Steering Committee (M&S SC), resourced by the Modeling and Simulation Coordination Office, and managed by JFCOM J7. The support of all those organizations is gratefully acknowledged.

REFERENCES

Cuda, D. and Frieders, M. (2005). Estimating DoD Budgets for Modeling, Simulation, and Wargaming. Institute for Defense Analyses Technical Report to the Defense Modeling and Simulation Office.

Henninger, A., Richbourg, R., Lutz, R., Saunders, R., Loper, M., Swenson, S. (2008). Live Virtual Constructive Architecture Roadmap (LVCAR) Final Report. Report developed for OSD P&R.

Hofer, R.C. and Loper, M.L. (1995). DIS today [Distributed interactive simulation] Proceedings of the IEEE Volume: 83 Issue: 8 Aug 1995 Page(s): 1124-1137.

Iansiti, M., and Levien, R. (2004). Strategy as Ecology. Harvard Business Review. (March, 2004).

IEEE (2003). 1516.3 - IEEE Recommended Practice for High Level Architecture (HLA) Federation Development and Execution Process (FEDEP)

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

Loper, M., and Cutts, D. (2008). Live Virtual Constructive Architecture Roadmap (LVCAR) Comparative Analysis of the Standards Management Processes. Report developed for OSD P&R.

Miller, D., and Thorpe, J.A. (1995). SIMNET: the advent of simulator networking Proceedings of the IEEE Volume: 83 Issue: 8 Aug 1995 Page(s): 1114-1123.

Powell, E. (2005). Range System Interoperability. In the Proceedings of Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) '05. Orlando, FL.

Richbourg, R., and Lutz, R. (2008). Live Virtual Constructive Architecture Roadmap (LVCAR) Comparative Analysis of the Architectures. Report developed for OSD P&R.

Richbourg, R., Lutz, R., and Ceranowicz, A. (2008). My Simulation is from Mars; Yours' is from Venus. To be presented in the Interservice/Industry Training, Simulation and Education Conference. Orlando, FL.

Swenson, S. (2008). Live Virtual Constructive Architecture Roadmap (LVCAR) Comparative Analysis of the Business Models. Report developed for OSD P&R.

Tolk, A. and Muguira, J.A. (2003). The Levels of Conceptual Interoperability Model (LCIM) summarized online at: http://en.wikipedia.org/wiki/Conceptual_interoperability

Wilson, A.L. and Weatherly, R.M. (1994) The aggregate level simulation protocol: an evolving system. Proceedings of the 26th Winter Simulation Conference WSC '94. Orlando, FL.