Enabling Effective Training with Mission Partners using Resilient Multilevel Architectures

Jennifer Lewis, Diana Pineda, Iain Ferguson
SAIC
jennifer.e.lewis@saic.com, diana.m.pineda@saic.com, iain.d.ferguson@saic.com

ABSTRACT

To create an integrated and ready warfighting force, we must be able to train with highly classified assets in live virtual constructive (LVC) simulation environments alongside our mission partners, where each US and partner participant sees the appropriate level of fidelity for their classification or caveat. Multilevel security (MLS) solutions that enable an LVC environment to simultaneously process multiple classification levels have been a challenging problem from both a technical and a policy perspective. This paper describes the implementation of successful MLS training environments, up to and including Special Access Program (SAP) data, with US and mission partner participants, allowing ground truth to be appropriately obfuscated in near real-time based on classification caveat and user need-to-know, resulting in real world-like situations where some participants do not have visibility into the use of classified assets. Intended for a technical audience, this paper also describes a Resilient Multilevel Architecture (RMA) approach to develop and continuously evolve a multilevel security training system to meet its specific training and security needs based on the overall system maturity. RMA addresses the need to share complex information in diverse ways beyond a single security domain or need-to-know (access) caveat through the application Zero Trust principles, data centric frameworks, and active risk management. This paper also provides data and metrics from executed training environments and discusses lessons learned from the implementation team.
.

ABOUT THE AUTHORS

Jennifer Lewis is a Senior Director for Application Modernization who has developed innovative solutions for distributed LVC environments for the US Army and US Air Force for more than 20 years. Her solutions have focused on holistic integrated systems with capability to provide actionable data and insights to users and leadership. She holds a Master of Science in Computer Science with an emphasis in Telecommunications and Networking from the University of Texas at Dallas and is a Master Certified Modeling and Simulation Professional.

Diana Pineda is a Chief Simulation and Software Engineer with more than 20 years of experience in software systems engineering, technical program management, and project leadership for distributed training and analytic LVC simulation environments for the US Army, Air Force, and Space Force. She holds a Bachelor of Science Degree in Engineering Physics from Embry Riddle Aeronautical University and a Master of Science in Engineering Management form University of Central Florida, and she is a Certified Modeling and Simulation Professional.

Iain Ferguson is the Chief Technology Officer (CTO) for the Air Force and Combatant Command (AFCC) business group at SAIC with more than 25 years of experience in flight operations, developmental test, acquisitions, and digital transformation. He holds a Bachelor of Science Degree in Engineering Mechanics from the US Air Force Academy, a Master of Science Degree in Aeronautical Engineering from the Air Force Institute of Technology, and a Master of Science Degree in Flight Test Engineering from the US Air Force Test Pilot School.

Enabling Effective Training with Mission Partners using Resilient Multilevel Architectures

Jennifer Lewis, Diana Pineda, Iain Ferguson

SAIC

jennifer.e.lewis@saic.com, diana.m.pineda@saic.com, iain.d.ferguson@saic.com

INTRODUCTION

To create an integrated and ready warfighting force, we have long held that we need to train the way we fight. In some cases, though, this can induce its own complications. In real-world operations, mission partners, competitors and even enemy combatants may witness US weapon systems performing classified actions, thereby gaining information about US capabilities through simple observation. In a Modeling and Simulation (M&S) environment, not only can we control what is seen by outside observers, but we can also control the way the system shares specific data elements among Live Virtual Constructive (LVC) users. This allows us to train alongside mission partners in a way that enables each US and mission partner to see the appropriate level of mission fidelity for their classification or caveat. The Committee in National Security Systems (CNSS) defines multilevel security (MLS) as "the capability of an information system that is trusted to contain, and maintain separation between, resources (particularly stored data) of different security domains" (Committee on National Security Systems, 136). Traditionally, the Department of Defense (DoD) uses physical, network-centric approaches to ensure the security of sensitive data, but with advent of new data-centric technologies and modern zero trust security principles, new opportunities exist to gain efficiencies and improve interoperability by allowing data of varying sensitivities to be physically co-located while logically separated. MLS is a complex topic that must address a variety of policy and technology challenges in both operational and training domains.

This paper is not intended to provide a complete discussion of MLS. Rather, it focuses on the use of MLS in LVC training environments. Because of its data centric nature, a robust MLS implementation requires a holistic approach to enterprise Information Technology (IT) that is not typically found in M&S environments. Rather than relying on one system or concept to secure M&S data, a Resilient Multilevel Architecture (RMA) implements multiple features, controls and layers of security concepts that provide resiliency, even if one or many security features are compromised. This paper first introduces MLS concepts in terms of LVC environments. It then describes a successful implementation of MLS for multinational joint task force training. Finally, it recommends next steps to achieve secure M&S data sharing with mission partners.

MLS CONCEPTS

To access information pertinent to national security, a user must meet the standards for access to classified information and related control and dissemination caveats, and they must have a need-to-know. This paper uses the term "user" to represent any system interacting with the system, whether it is a person or a non-person entity (NPE), such as a hardware device or a software application. The process of verifying the user is who he says he is, i.e. authentication, is a separate process from determining whether the user can access specific data, i.e. authorization. MLS primarily concerns itself with authorization, although as we discuss in later sections, both processes are required to ensure secure data sharing.

The DoD defines three classification levels, i.e. Top Secret, Secret, and Confidential, also known as security domains. Security caveats specify additional requirements for information access. Caveats use a defined taxonomy to describe whether information is releasable to citizens of particular countries or coalitions, e.g. UK, FVEY, or NOFORN. Caveats can also describe that information is releasable only to specific Special Access Programs (SAP) or Sensitive Compartmented Information (SCI). Original Classifying Authorities (OCAs) include both classification level and caveats in markings on all documents related to national security. Need-to-know is more individualized. It describes whether a user requires access to information to perform his or her job function, regardless of their security clearance level or other approvals. Need to know can change quickly, as a user's job function changes or as the role that job

function plays changes within the program. Need to know does not have a defined taxonomy to describe job functions across the DoD, so it cannot be readily described in document markings like classification and caveat can, although as we discuss in later sections, need-to-know can be designed into data-centric MLS approaches.

Security Domains

The DoD typically separates data from different security domains and caveats using a network-centric approach where data is stored on different physical networks that are authorized up to the highest level of classification for information on that system. One network such as Secret Internet Protocol Router Network (SIPRNet) can hold data up to Secret while another network such as Joint Worldwide Intelligence Communication System (JWICS), can hold data up to Top Secret. The challenge, however, is in sharing data from one network (or system) to another, especially when the data must pass from a higher level of classification to a lower one (e.g. Top Secret to Secret). A common way to share data across these networks is the SneakerNet, a term used to describe transferring data using air-gapped, portable storage devices such as universal serial bus (USB) drives or even manually reinputting data into the other system. The term comes from the idea that the method of transport of the USB drive is a person walking from one machine to another, wearing sneakers.

Real-time data sharing cannot happen through the SneakerNet. Instead, the primary way to share data across security domains without human intervention, also known as machine-to-machine (M2M), is a Cross Domain Solution (CDS). A CDS can be used to transfer data files between physical networks or to allow authorized users access to the file system on other physical networks. While multiple implementation approaches exist for CDSs, a hardware-based solution is a common choice for high security environments like the DoD. Hardware-based solutions typically have strict implementations, such as data diodes that ensure data can flow only in one direction, e.g. low to high, but also ensure there can be no malicious access of the higher system by an actor on the lower system, thereby maintaining a trust boundary between the two networks. This approach maximizes security but limits functionality and can impact mission success. A CDS can also implement software-based approaches that are more flexible, including the ability to share data bi-directionally if needed. The rules that CDSs use to share data are included in the Authority to Operate (ATO) that DoD organizations require to allow the CDS to execute on the network. This minimizes the benefit gained from a software-based approach, making it difficult to adapt sharing policies as the mission need evolves and may not always be sufficient based on the Authorizing Official's (AO's) requirements. Early efforts in MLS also implemented a proof-of-concept CDS with a Protection Level 4 (PL4)-interface, which permits access to users that lack sufficient clearance for some of the information on an information system, if all users have a Secret clearance (Danner and Valle, 2005). More recently, cloud-based CDSs, such as Amazon Web Services (AWS) Diode or Microsoft Azure Data Transfer, enable the use of cross domain solutions while leveraging the advantages of cloud service providers (CSPs), although the policies for their implementation are still evolving.

Security Caveats

The DoD allows storage of data with different security caveats on the same physical network if the data's security domain is the same, provided the necessary security controls are implemented in that system (commonly called Protection Level 3). A common approach to protecting security caveats is the use of network segmentation, which protects data using concepts such as subnets and firewalls. This network-centric approach limits the ability to conduct training exercises with mission partners because the network devices that implement these approaches cannot evaluate the content of the data packets, only the data associated with the routing of the data packet. For a Distributed Interactive Simulation (DIS)-based environment, this means that an Entity State Protocol Data Unit (PDU) can either pass through the firewall or it cannot. If the user is not eligible to see one element in the PDU, he will not see any data related to that entity.

Data Distribution Service

Authorizing officials can allow data centric approaches to provide a more granular view of the LVC environment. Data Distribution Service (DDS) is a software connectivity standard allowing real-time information exchange in distributed systems. Managed by the Object Management Group (OMG), DDS uses a data-centric publish/subscribe communication model based on data topics, the data elements the user is allowed to access or is interested in accessing. While the DDS standard does not address security, DDS has been applied to multiple secure applications within the DoD (RTI, 2024). To achieve MLS via DDS, the publishing system must map specific data elements to security

caveats and publish those elements to appropriate data topics. Subscribing systems must ensure they only subscribe to appropriate data topics based on their security caveats and need-to-know. This model relies on the DDS architecture to ensure appropriate delivery of data. System designers must make policy decisions related to authorization manually during data topic configuration. Data topics offer a more granular option for determining which data elements a user may access. For example, DDS could deliver the location and orientation from an Entity State PDU on one data topic while delivering its fuel level on a different data topic. This approach allows a mission partner to interact with the entity in the training environment without knowing all ground truth about that entity. At I/ITSEC 2018, the US Air Force Agency for Modeling and Simulation (AFAMS) Operational Training Infrastructure (OTI) Ecosystem used DDS to demonstrate MLS in an LVC environment, showing the ability to deliver appropriate LVC flight data to visualization systems supporting specific security domains (Tingey, 2019).

Attribute Based Access Control

Another data centric approach to sharing between caveats is Attribute-Based Access Control (ABAC). ABAC is a paradigm for controlling data and application access based on descriptive attributes. The defined attributes can vary widely depending on the design and context of the dataset. However, in its most basic form, ABAC relies on the evaluation of attributes for the subject, attributes of the object, environment conditions, and the policies defining allowable operations for subject-object attribute combinations, as shown in Figure 1 (Hu et al, 2014). Subject attributes describe the user making the request while Object attributes describe the data being requested. Environment conditions describe additional context regarding the data, such as date, time, or geographic location. While ABAC is relatively straightforward to maintain, its implementation requires significant up-front design. System administrators and stakeholders must define, document, and implement the full scope of attributes and combinations of attributes to ensure principles of least privilege and compliance with classification markings for MLS use cases. In a DIS environment intended to support MLS, each element of a PDU has its own Object attributes assigned to it indicating, for example, what security domain or caveat to which it is releasable. We will discuss this process in more detail in the next section.

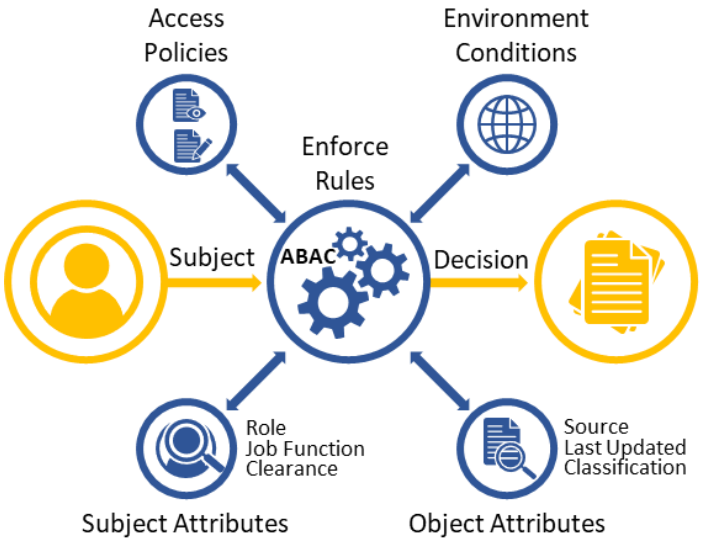

Figure 1: ABAC Enables Granular Authorization Decisions

MULTINATIONAL JOINT TASK FORCE TRAINING

In demonstrations of multinational joint task force training in 2022 and 2023, SAIC used Next Generation Threat System (NGTS) for the constructive environment, DIS as the primary LVC communication method, and an ABAC-based data platform called Koverse to obfuscate sensitive sensor and weapon systems capabilities in real time, as shown as Figure 2. NGTS is a synthetic environment generator used to support training, testing, analysis, and research and development. NGTS models threat and friendly aircraft, ground and surface platforms, and corresponding weapons and subsystems. System designers configured the environment with a blue network NGTS and a green network NGTS. Within the blue network, core simulations and applications showcased Blue Forces and Red Forces (Threats), alongside the Koverse data platform and a routing service. Within the green network, NGTS #2 features core simulations and applications tailored specifically for Green Forces.

Founded by former National Security Administration (NSA) data architects who were part of the original development of the now-open-sourced Apache Accumulo, Koverse is an Accumulo-based data platform that enables ABAC. Using an approach called cell-level security, Koverse allows for high-speed, high concurrency queries on sensitive data, resulting in near real-time MLS processing of DIS PDUs in the training environment with less than 10 ms average latency. This demonstration also identified a number of requirements for ABAC in an M&S system such as this one.

To configure the Koverse's ABAC, system designers began by considering the simulation information available for a specific entity, which we will refer to as Entity W. First, the system must recognize Entity W's existence as an attribute. Second, the system must both recognize the existence of and associate Entity W with a set of one or more capabilities, which we will call Capability Y_{W1}…Capability Y_{WN}. Next, the system must both recognize the existence of and associate Entity W with any effects prompted by Capability Y_{WX}. Please note, we generalize the term capabilities to include communications, sensing, and weapons systems. Therefore, the associated effects could include a wide range of outcomes from successful communications to entity or terrain battle damage.

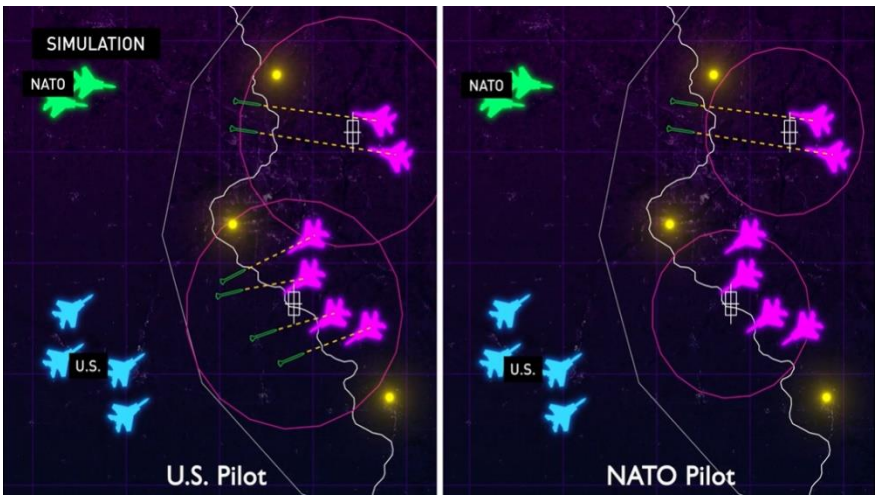

Figure 2: Koverse Enables Obfuscation of NGTS-Modeled Sensor and Weapon Systems Capabilities in Real Time

Object Attribute Considerations

In a DIS environment, an entity's existence and other critical information is documented in the Entity State PDU. The Entity State PDU describes an entity's capabilities in its Capabilities or Articulation Parameters records, although in many cases, the simulation application infers these capabilities through the Entity Type field or associated Transmitter, Electromagnetic Emission, Weapon Fire, and Detonation PDUs. DIS communicates the effect of a capability both in the Detonation PDU and the Entity State PDU's Entity Appearance field. Because of the interaction between DIS PDUs, system designers created inter-related ABAC object attributes for each DIS PDU field, e.g. which other data elements impact this data element, lowest allowable security domain, and enumeration of dissemination caveats. This shows the need for context aware and stateful attribute assignment for ABAC to be effective in an M&S system.

Subject Attribute Considerations

Whether an ABAC subject can view a particular entity, capability or effect requires access policies that model the implications of Capability Y_{WX}'s deployment or activation. In some cases, a subject may not even know that an event has occurred, e.g. when the effects are too subtle or when they impact an entity or transmitter the subject does not know exists. In many cases, the subject may know that an event happened but not why the event happened. For example, the subject may see that an entity is smoking but cannot tell what caused the smoke. These considerations introduce further system complexity, as shown in Figure 3. A well-designed MLS training environment that effectively shares across multinational caveats and need-to-know creates inconsistencies between their respective simulation applications, creating the potential for two different views of the scenario events and varying behavior based on the user's perspective of "ground truth". In this environment, only two users exist, i.e. US and NATO. However, system designers included subject attributes to allow differentiation between job function, need-to-know, and clearance level to allow secure dissemination of information within the blue and green networks. To ensure correct ABAC behavior with multiple types of users per network, the system must implement an authorization mechanism that verifies applicable user metadata. In addition, the data platform should use both ABAC and Role Based Access Control (RBAC), when appropriate, to support complex security scenarios.

DIS-Based Implementation

Because a DIS environment is typically populated in a random order based on the arrival of DIS PDUs sent via a best effort protocol, system designers ensured the efficacy of the ABAC approach at system startup using a predefined enumeration list to assign object attributes. This list assigns distinct Entity IDs to an array of entities. Due to security sensitivities in the multinational force, some entities are identified by low-fidelity classifications such as national identifications while others are identified by specific aircraft models. Next, the system uses a state-based ruleset to manage unidentified Entity IDs. If the system receives an Entity State PDU with an unknown Entity ID, it blocks any interaction until the entity has been characterized and its ABAC object attributes assigned. Upon receiving the initial PDU, the system examines and records essential fields such as Entity Type, Entity ID, aircraft model, and country of origin. Once the entity is authenticated and categorized, its Entity ID is added to the enumeration list, triggering tailored rulesets based on the entity's specific ABAC object attributes. Finally, the system uses rulesets to handle inter-related PDUs. The system blocks Entity State or Weapon Fire PDUs representing weapon flyouts that correspond to the armament of an entity with sensitive weapon systems, in this case a F-15C. The result of this policy is that the green network never sees the weapon flyout, even if they have the real-world technology to detect it. For Detonation PDUs, the system masks the associated Entity ID for detonations originating from an F-15C. The result of this policy is that the green network sees the detonation but cannot associate it to a weapon. For Electromagnetic Emission PDUs, the system blocks track and jam emissions originating from an F-15C. The result of this policy is that the green network can detect some emissions from an F-15C but will not have ground truth about what entities the F-15C is tracking or jamming. Note that sharing of emission data via the Electromagnetic Emission PDU is especially complex in the DIS environment, involving many embedded records. The system must parse, label, and assign attributes with minimal latency to maintain real-time.

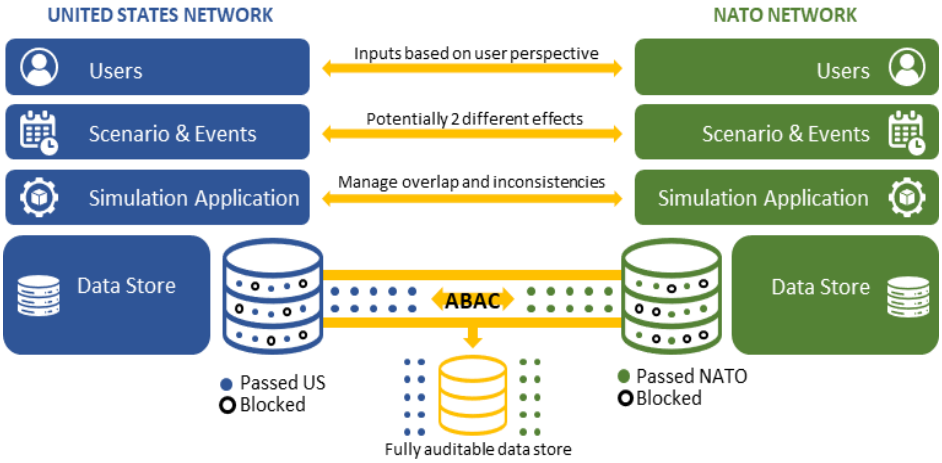

Figure 3: Effective MLS Creates Challenges Through the M&S Tech Stack

RESILIENT MULTILEVEL ARCHITECTURE

While many significant strides have been made in the implementation of MLS, there is still a significant amount of confusion in the variety of different circumstances that apply and the necessary solutions for each case. Many stakeholders assume that there is a one-size-fits-all for MLS, creating confusion in how complex problems can be solved across the variety of M&S use cases. We recommend a new wholistic approach to MLS called Resilient Multilevel Architecture (RMA) that functions as a maturity model framework for incrementally implementing MLS based on the overall enterprise information technology (IT) modernization and the evolving mission needs for data sharing. The term resilient is used to incorporate the principles of cyber resilience, defined by the NIST 800-164 as "the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that include cyber resources (Ross et al, 73)." This term emphasizes the need to move beyond simple cybersecurity compliance and think wholistically about the need to "develop survivable, trustworthy secure systems" that accomplish their intended mission (ibid, ii). RMA also follows the lead of the DoD's Cyber Reference Architecture Principle 2, which mandates that systems increase mission assurance through resilience (DoD, 2023) and builds on the Zero Trust Architecture outlined in NIST 800-207 (Rose et al, ii). RMA is not a single solution but rather a security strategy to enable the sharing of information across traditional security boundaries through the application of Zero Trust principles, data centric frameworks, and active management of the risk of sharing versus the impact mission.

One implementation of RMA is the Mission Partner Environment (MPE) Information Domain Data Centric Security Framework, as shown in Figure 4. First, mission partners need to establish mutual trust through a trust framework that identifies shared policy decisions and agreements. The trust framework defines "what" the partners intend to share while subsequent layers implement "how" the data is shared. Second, the network must know who the users are and what their role is through foundational identity, credential, and access management (ICAM). ICAM provides enterprise authentication and authorization decisions in conjunction with other underlying layers. Next, the Zero Trust implementation ensures ICAM verification occurs with each data request. ABAC acts as a further Policy Enforcement Point (PEP), inspecting each request for access to data and returning the decision on whether to permit or deny the access. Furthermore, Trusted Data Format (TDF) provides a protective wrapper around data elements as they are pulled from their associated data sources. TDF encrypts the data and wraps the data into a TDF file. The receiving TDF server re-verifies that the receiver is eligible to access the data before decrypting for the end user. Therefore, this framework provides MLS through an RMA with five layers of assurance.

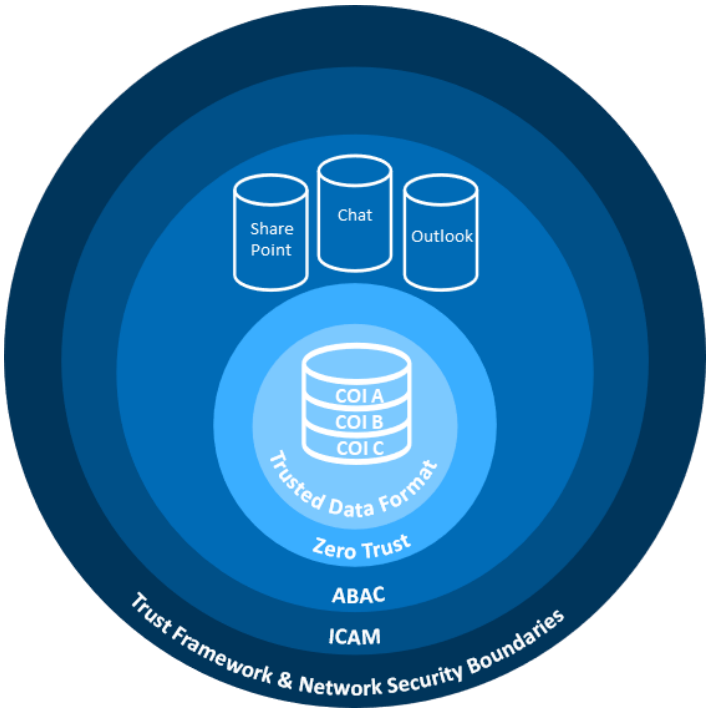

Figure 4: Layering Security Approaches Builds Architectural Resiliency

Outside of MLS, applying additional best practices for enterprise IT will allow M&S environments to deploy more quickly. We recommend a full enterprise campaign to modernize the M&S environment, including virtualization using orchestrated containerization and Infrastructure as Code (IaC). Containerization and orchestration when used in conjunction with IaC allow the entire training environment to be configured and deployed within minutes, minimizing the lengthy setup some M&S environments perform. Furthermore, M&S as a Service (MSaaS), standardized by the NATO M&S Group builds on this concept to provide rapid deployment of interoperable and credible simulation environments (Siegfried and Berg, 2015).

CONCLUSION

MLS is a complex topic that must address a variety of policy and technology challenges, especially in operational environments. For training environments, robust implementation of key enterprise IT concepts along with a more granular data-centric approach to data sharing, can dramatically improve our ability to train in multinational environments.

REFERENCES

Center for Development of Security Excellence. (2024) Marking National Security Information. Accessed online June 30, 2024. https://www.cdse.edu/Portals/124/Documents/jobaids/information/Marking_National_Security_Information.pdf

Committee on National Security Systems (CNSS). (2022) Committee on National Security Systems (CNSS) Glossary *Committee on National Security Systems Instruction (CNSSI) No 4009, 2 March 2022.*

Danner, B. and Valle, T. (2005). Multilevel Security Assessment for the Distributed Mission Operations Network (DMON). *Proceedings of the 2005 Interservice/Industry Training, Simulation, and Education Conference.*

DoD Chief Information Officer. (2023). DoD Cybersecurity Reference Architecture. Accessed online June 10, 2024. https://dodcio.defense.gov/Portals/0/Documents/Library/CS-Ref-Architecture.pdf

Hu, V., Ferraiolo, D., Kuhn R., Schnitzer, A., Sandlin, K., Miller, R., and Scarfone, K. (2014). Guide to Attribute Based Access Control (ABAC) Definition and Considerations. *National Institute of Standards and Technology (NIST) Special Publication 800-162*.

Rose, S., Borchert, O., Mitchell, S., Connelly, S., (2020) Zero Trust Architecture *National Institute for Standards and Technology (NIST) Publication 800-207.*

Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., McQuaid, R., (2021) "Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, *National Institute of Standards and Technology (NIST) Special Publication 8001-160v2 r1.*

RTI. (2024). Special Operations Forces. Accessed online June 24, 2024. https://www.rti.com/hubfs/_Collateral/capability-briefs/rti-capability-brief-special-operations-forces.pdf

Siegfried, R. and Berg, T. (2015). M&S as a Service: Paradigm for Future Simulation Environments. *Proceedings of the 2015 Interservice/Industry Training, Simulation, and Education Conference.*

Tingey, P. (2019). From Simulation to Deployment: DDS the Connectivity Solution. *Proceedings of the 2019 North Atlantic Treaty Organization (NATO) Computer Assisted Analysis, Exercise, Experimentation (CAX) Forum.*