A Capabilities-Based Architecture for Simulating WMD Emergency Response

Mary Ann Pigora, Theresa Tamash
SAIC
Orlando, FL
pigoram@saic.com, tamasht@saic.com

ABSTRACT

The Automated Exercise and Assessment System (AEAS) is a simulation that enhances training and coordination of emergency response to incidents involving Weapons of Mass Destruction (WMD). The system simulates WMD incidents, and allows emergency responders to utilize their own Incident Command System and available resources to exercise decision-making and other cognitive skills. The WMD scenarios cover a variety of incidents, including chemical, biological, radiological, nuclear, and high explosive (CBRNE) attacks.

In order to optimize the training experience, the exercising jurisdiction must be able to respond to the simulated incident with the resources and capabilities that they would normally have, and using their own command structure. There is no standard equipment set or nomenclature for emergency responder agencies, and resources available to response agencies vary greatly. Agency structure is also highly variable, with HazMat resources reporting to the Fire Department in some jurisdictions, and to Law Enforcement or even Emergency Medical Services in others. It is impractical to ask jurisdictions to specify their equipment from an exhaustive and perpetually outdated list, so a capabilities-based approach to building resources was created. This approach allows unlimited resource composition and a flexible command structure, allowing user specification of specialized units such as Urban Search and Rescue or National Guard WMD Civilian Support Teams. This paper will discuss the AEAS Agency Survey component used for specifying a jurisdiction's resources and the capabilities-based implementation for simulating those resources in a WMD response.

ABOUT THE AUTHORS

Mary Ann Pigora is a software engineer at SAIC where her work has included architecting AEAS, creating CGF and visual databases for human simulation, and work on the OneSAF Objective System. She has over ten years of experience in 3D graphics including as a Technical Director for Walt Disney Feature Animation where she did computer animation for the movies *Mulan* and *Fantasia 2000*. She received her MS in Computer Science from the Georgia Institute of Technology.

Theresa Tamash is a software engineer at SAIC where she has worked primarily on military training Semi-Automated Forces (SAF) systems. She has 3 years experience working on Computer Generated Forces (CGF) in military applications. She is currently working towards her MS in Computer Science at the University of Central Florida.

A Capabilities-Based Architecture for Simulating WMD Emergency Response

Mary Ann Pigora, Theresa Tamash
SAIC
Orlando, FL
pigoram@saic.com, tamasht@saic.com

INTRODUCTION

The Automated Exercise and Assessment System (AEAS) is a simulation that enhances training and coordination of emergency response to incidents involving Weapons of Mass Destruction (WMD). The system simulates WMD incidents, and allows emergency responders to utilize their own Incident Command System and available resources to exercise decision-making and other cognitive skills. The WMD scenarios cover a variety of incidents, including chemical, biological, radiological, nuclear, and high explosive (CBRNE) attacks.

The primary purpose of AEAS is to evaluate a community's ability to respond to a WMD event. AEAS can be used for community-wide readiness assessment, for evaluation of different resource compositions to analyze need or justify equipment expenditure, and for planning a community's response to a possible incident type. The training audience consists of command-level personnel on the scene and in the Emergency Operations Center, rather than the personnel on the ground performing physical tasks. AEAS allows decision-makers to command their simulated resources and to coordinate with other decision-makers, and to see the consequences of their decisions.

In order to optimize the training experience, the exercising jurisdiction must be able to respond to the simulated incident with the resources and capabilities that they would normally have, and using their own command structure. There is no standard equipment set or nomenclature for emergency responder agencies, and resources available to response agencies vary greatly. Agency structure is also highly variable, with HazMat resources reporting to the Fire Department in some jurisdictions, and to Law Enforcement or even Emergency Medical Services in others. It is impractical to ask jurisdictions to specify their equipment from an exhaustive and perpetually outdated list, so a capabilities-based approach to building resources was created. This approach allows unlimited resource composition and a flexible command structure, allowing user specification of specialized units such as Urban Search and Rescue or National Guard WMD Civilian Support Teams. This paper will discuss the AEAS Agency Survey component used for specifying a jurisdiction's resources and the capabilities-based implementation for simulating those resources in an AEAS exercise.

The AEAS product consists of four components: Survey, Player Station, Controller Station, and After Action Review (AAR) (see Figure 1). In the Survey, the user defines the jurisdiction's emergency response resources, as discussed in the following section. These resources can be used for readiness assessment and imported into the exercise so that the players can deal with the simulated scenario using their own jurisdiction's capabilities.

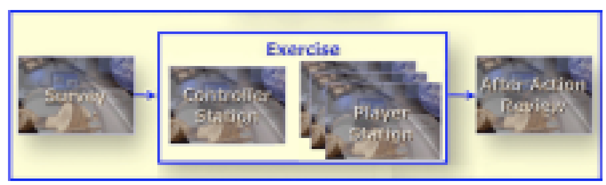

Figure 1. AEAS Components

The Controller Station is used by the exercise facilitator, or *controller*, to set up the exercise parameters and monitor the progress of the exercise. There is a Player Station for each player, on a network with the Controller Station. The players can give commands to the resources they control, build the command hierarchy, and communicate with each other and with simulated outside agencies (see Figure 2). The AAR provides fodder for discussing the results of the exercise, because not all questions have "right" answers. For further details on the AEAS system, please refer to (Pigora, et al., 2002).

Figure 2. AEAS Player Station Interface

THE AGENCY SURVEY

Traditional military simulations rely on predefined unit and entity compositions. Some architectures provide for domain-independent behaviors (Van Veldhuizen and Hutson 1996) or user-defined entity composition based on physical components (Courtemanche, 2000). AEAS takes the tenet of separating physical and cognitive processes and from it derives a capabilities-based architecture that allows resource specification by the user and provides flexibility in domain, abstraction level, and command hierarchy.

AEAS provides geotypical resource configurations that can be used in an exercise, but to impart the optimal training experience, the exact resource mix available to the exercising jurisdiction should be used to initialize the simulation. A jurisdiction's resource definition is accomplished in the Agency Survey, typically in advance of the actual exercise date. A requirement of the Survey was that it could be filled out by anyone with basic computer skills. A representative of each exercising agency or department fills out the Survey to reflect that agency's resource composition, including any mutual aid pacts, which are agreements for the loan of supplies and personnel during times of need with agencies in surrounding jurisdictions. The results are brought together in the Survey's Community Profiler component. The Profiler contains a database of cities and populations that can be used to find the geotypical location that best matches the local jurisdiction. Three geotypical areas are available in AEAS: rural county, midsize city, and large city. Agency Surveys are read in and correlated, and the resulting Community Profile contains all the information about a jurisdiction's available resources, facilities, and staffing. This Profile is used for resource assessment and to initialize the simulation.

Command Hierarchy

One of the challenges in creating a simulation for emergency responders is in defining the command hierarchy. Jurisdictions vary widely in how they assign responsibility to different agencies. HazMat resources, for example, report to the Fire Department in some jurisdictions, and to Law Enforcement or even EMS in others. To address this requirement, a domain-independent functional view of the command hierarchy was created. An agency chooses its areas of responsibility from a list of Functional Areas (FAs). FAs are data driven and can represent any functional decomposition of command, such as the staffing structure of a hospital. For AEAS, the FAs are derived from the Emergency Support Functions used for delineating responsibility in emergency management. AEAS has defined 36 FAs, ranging from Law Enforcement to Debris Management. One agency may have several FAs. The Fire Department, for example, may control Fire, Hazmat, and Urban Search and Rescue teams. The controller uses the Functional Areas to create the command hierarchy for the simulation, allowing the local jurisdiction to control what resources report to which commander (see Figure 3). In the exercise, a player is assigned responsibility for one or more nodes in the command hierarchy. The hierarchy can change over the course of the simulation as resources are loaned out or assigned to report to different commanders. An example of loaning resources would be when the player who is performing the Law Enforcement function requests city busses from the player who is performing the Transportation function to assist in evacuating an area.

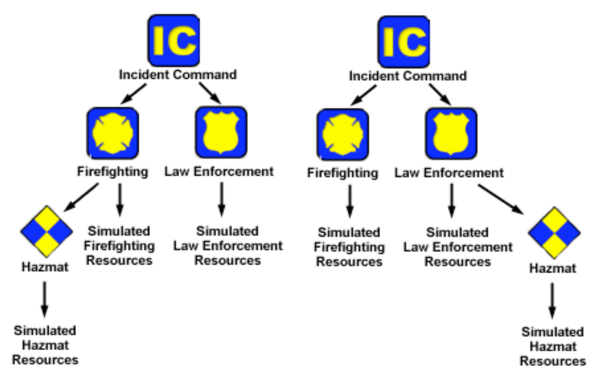

Figure 3. Sample Command Hierarchies

Resources

Once the Agency selects its areas of responsibility, simulated resource types must be defined for those FAs. Firefighting, for example, would probably have several types of resources such as engines, rescue vehicles, and ladder trucks. Providing a list of common resource types to choose from is not sufficient, as there is no consistent nomenclature between Firefighting departments across the country. What is called a ladder truck in one jurisdiction might be called an aerial in another, and they will most likely have differences in the number of crew they carry. Furthermore, the specific equipment each carries will vary, and providing an exhaustive list of equipment types, brands, and variants for the user to choose from would be impractical. The simulation, however, is only concerned with the ways in which the resource can be applied to the scenario being exercised. Recognizing this, resource types can be built by specifying the *capabilities* each has, such as a fire engine with fire suppression, rescue, and basic life support capabilities. Decomposing common resource functionality into its most basic capabilities allows for maximum reuse across domains. A land transport capability, which has an associated numeric supply of capacity, can be used to define a city bus that transports civilians, paired with an advanced life support capability to define an ambulance that transports patients, or used for a military personnel carrier (see Figure 4).

Ladder	**City Bus**
Fire Suppression	Land Transport
Rescue	
Pump Water	**MedEvac Helicopter**
Aerial	Advanced Life Support
High Angle Rescue	Air Transport
HazMat Truck	**EMS Transport**
HazMat	Advanced Life Support
Rescue	Basic Life Support
Advanced Life Support	Land Transport

Figure 4. Resource Types Defined by Capabilities

Capabilities have specific commands associated with them, therefore, if a resource type is assigned a capability, all instances of that resource can be issued those commands. A resource with a HazMat capability, for example, can be commanded to produce a plume model or make a recon entry into the hot zone, among other things. There may also be default capabilities assigned to a class of resources, such as a capability all responding units have that brings common commands like "go to rehab", or "assist another unit".

Capabilities also have associated data-driven supplies (see Figure 5), whose amount is specified by the user and which may be required and/or consumed by commands. Making a recon entry would require and consume both Level A suits and air cylinders, while successfully producing a plume model would require that the resource have some sort of plume modeling software. A resource with a resupply capability can transfer supplies to other resources, such as when a cascade truck refills air tanks. This gives the simulation the ability to track logistical issues, and charts of on-scene supplies and resources can be overlaid in the AAR with the commands that were or were not issued to get them to the scene in time to mitigate the incident. In addition, the exercising jurisdiction can be made aware of capabilities or supplies that none of their resources possess that may have aided in their response to the WMD situation.

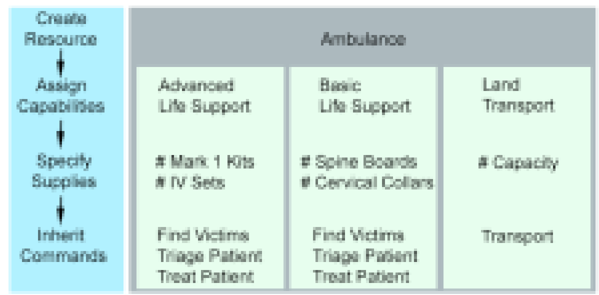

Figure 5. Building a Resource Type

This capabilities-based approach allows maximum flexibility in defining resources, allowing all jurisdictions from the largest to the smallest to specify exactly what they have in a reasonable length of time. Large jurisdictions will define a general type, such as heavy rescue, and specify a large number of them as being available in the jurisdiction. A small jurisdiction, however, may have very few heavy rescues, each with a special function such as mass casualty or HazMat. The smaller jurisdiction can define each heavy rescue as its own type, and specify which should be dispatched for certain types of alarms. These approaches can be mixed and matched to quickly define a heterogeneous pool of resources that may contain unique or specialized pieces of equipment.

Facilities

In the simulation view, facilities are treated as stationary resources. They have facility-specific capabilities, such as being designated as an emergency shelter, or being evacuatable. Each Functional Area can define the facilities they utilize. For ease of use, many FAs have predefined facilities available to them, such as fire stations for the Firefighting FA or dumps for the Debris Management FA. Facilities can be named to reflect the ones in the local jurisdiction, and resources can be stationed at them. They have pre-defined locations in each of the three geotypical locations and will appear on the Player Station map.

Facilities can also have supplies specified for them, specific to their particular facility capabilities. For example, medical facilities require supply information such as the number of patient beds, the normal occupancy rates, and the number of additional beds that could be made available in the event of a large emergency. When a resource transports a victim to a medical facility, the medical facility capability admits the patient, using one available bed. The medical facility may also be given commands, such as to discharge all elective surgery patients, which would increase the number of available beds. If the simulated WMD event produces enough casualties, the medical facility may eventually run out of beds. Players are forced to deal with the consequences of having a limited number of hospital beds during large epidemics such as a smallpox outbreak.

Some facilities are defined to appear in all exercises without being specified in the survey. Schools, apartment buildings, office buildings, nursing homes, and others appear in every exercise, all with populations appropriate to the geotypical location. These facilities may need to be evacuated during a simulation if the facility is in danger. Many of the default facilities have capabilities that are crucial to success in exercises. For example, some can be opened as emergency shelters, again with capacities appropriate for the population of the geotypical location.

Currently, facilities are not as configurable as resources. They may not be relocated on the map, and have default population or capacity supplies appropriate to that geotypical location. It would be a straightforward addition to the system to allow the user to specify locations and populations, but would require much more set-up time for the simulation, as well as a more detailed knowledge of the exact resources of a jurisdiction. AEAS endeavored to strike a balance between fidelity and usability in the simulation.

ASSESSMENT

A valuable aspect of AEAS is its ability to provide unbiased assessments to the jurisdiction based not only on their performance during the simulation, but also on the capabilities and supplies they have. After a jurisdiction has used the Survey application to describe their resources and facilities, an assessment is performed (see Figure 6). This assessment is also available from the AAR application for discussion after the exercise.

Capabilities and supplies are assessed separately against baseline values for each type of WMD event: chemical, biological, radiological, nuclear, and high explosive. Some capabilities or supplies are not applicable to certain types of WMD events and this is noted in the assessment with an "NA". For example, the supply type "chemical agent detector" is only needed for scenarios where a chemical agent is involved. By assessing the jurisdiction's resources separately for the different kinds of attacks, the jurisdiction can more easily gauge their preparedness.

Sets of baseline values were developed with input from Subject Matter Experts (SMEs) for each geotypical location depicting the minimum numbers of each capability and supply necessary to effectively handle each kind of WMD event. Each geotypical location has different baseline values that properly reflect the varying population densities and the probable effects of each kind of WMD attack on that location. Since many jurisdictions' populations will fall between those of the geotypical locations available, the assessment allows the user to choose which geotypical location to assess against.

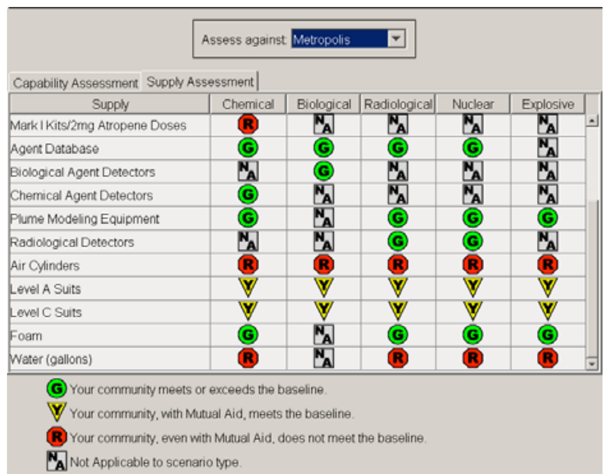
Assess against: Metropolis

Capability Assessment | Supply Assessment

Supply	Chemical	Biological	Radiological	Nuclear	Explosive
Mark I Kits/2mg Atropene Doses	R	NA	NA	NA	NA
Agent Database	G	G	G	G	NA
Biological Agent Detectors	NA	G	NA	NA	NA
Chemical Agent Detectors	G	NA	NA	NA	NA
Plume Modeling Equipment	G	NA	G	G	G
Radiological Detectors	NA	NA	G	G	NA
Air Cylinders	R	R	R	R	R
Level A Suits	Y	Y	Y	Y	Y
Level C Suits	Y	Y	Y	Y	Y
Foam	G	NA	G	G	G
Water (gallons)	R	NA	R	R	R

G Your community meets or exceeds the baseline.
Y Your community, with Mutual Aid, meets the baseline.
R Your community, even with Mutual Aid, does not meet the baseline.
NA Not Applicable to scenario type.

Figure 6. Capability and Supply Assessment

In the Assessment, a green "G" symbol is shown if the jurisdiction's resources meets or exceeds the baseline values. If the jurisdiction's resources along with the resources available through their mutual aid pacts with neighboring communities meet or exceed the baseline values, a yellow "Y" is displayed. If the jurisdiction, including their mutual aid pacts, cannot meet the minimum capability and supply needs, a red "R" is assigned. Distinction between the jurisdiction's capabilities alone and with mutual aid is necessary because in extreme situations, mutual aid may not be available. For instance, in the case of a large-scale smallpox outbreak, neighboring jurisdictions will need their resources to handle patients in their own community.

The process of defining a jurisdiction's resources, facilities, capabilities and supplies can be a valuable asset even without running the simulation. The assessment can allow a community to see how prepared they are for WMD situations. A community could also include additional resources in their Survey to see if, during the simulation, they can manage the emergency better.

RESOURCES IN AN EXERCISE

The Agency Survey saves the resource specification in an Extensible Markup Language (XML) format, which is used to initialize an exercise. Other resources may also be created by the exercise, such as crowds, bombs, and other agents used in the selected scenario. A resource may represent an individual (a person in a crowd), several individuals acting together (a rescue unit), lots of individuals (a population center in a neighborhood, indicating the number of people at risk there), or a scenario object (a bomb or facility). A resource that represents multiple individuals may be treated as an abstract aggregate until something of interest at the individual level happens, then be split into more than one resource. This would occur when a member of a response unit is injured. The injured responder would be split off from his unit and be treated as an individual victim on the scene. This variable level of fidelity in resources allows computation to remain tractable for the thousands of individuals required to simulate a large city, while providing the detail required by the scenario.

In the exercise, the simulated resources are controlled by commands. The commands a resource can carry out are determined by its capabilities, since a resource has only those commands that it inherited by being assigned one or more capabilities. A command given to a resource may have data associated with it, such as "report to location," which would have a map location. Commands can be generated by the user, by other resources as they execute commands they have received, or by the executing scenario. All stimulation of the simulation is accomplished by processing commands, and the simulation can be rewound to a critical point or rerun by processing the logged commands.

When a command is issued, it is scheduled in an execution queue and delivered to the recipient at the correct simulation time step. As it executes, it may generate more internal commands as a result, which may be executed in the current time step or at a future time. For example, if a resource is given a command to find victims, it will locate an injured entity, then generate new commands for itself to triage, treat, and possibly transport the victim. Triage, treat, and transport may all be stand-alone commands given by a user, or may be used as building blocks to create more complex commands. This allows complex activities to be broken down into reusable component commands, which can be executed individually, or even divided between multiple resources.

Each capability is implemented as a Java class, and each command that can be given to a resource is a method in a capability class. Thus, new capabilities can be added to and removed from the system without affecting anything else. When a resource is created, each capability that the resource type was given in the Survey is instantiated and added to the resource, along with label-value pairs indicting supply type and number.

When a resource receives a command, its capabilities are searched to find a method whose name matches the command, and whose parameters match the data object passed with the command. If no such method is found, the command is ignored, as the resource does not have the capability to handle it. A command may also have a cleanup method, named the same as its main method but with the suffix "end". This is executed when a behavior finishes or is superceded by another behavior, to clean up any work in progress.

There are two types of commands: normal and immediate. For a normal command to be carried out, it requires that the resource be available, that is, the resource is not busy doing something else. If the resource is busy, the command will go into a queue for the resource to execute when it is done with its current task. Typically, the resource will give some feedback indicating this, such as a radio message saying that it is busy, but it will get to the new task as soon as possible. An immediate command does not have this restriction. An immediate command is carried out regardless of what other activity the resource is engaged in. The immediate command may execute concurrently with whatever else the resource is doing. An example of this would be giving a status report, which the unit does without stopping its current activity. Or the immediate command may supercede and cancel the current activity, such as when a resource is told to evacuate an area because of an imminent danger (see Figure 7).

Figure 7. Command Execution

CONCLUSIONS AND FUTURE WORK

A capabilities-based architecture allows a user to specify exactly the applicable resource composition available to address a situation, even if they utilize unusual or unconventional resource types. Supplies are data-driven and restricted to what will have an effect in the current task at hand, avoiding the problems associated with an equipment-driven approach such as large databases of equipment types, and outdated models. The capabilities-based approach to resources used in AEAS allows a jurisdiction to exercise scenarios using their own nomenclature, force composition and command hierarchy. New capabilities and supplies are easily added, allowing system extensibility. The approach can easily be applied to any sort of heterogeneous force composition, such as unconventional military units.

The National Guard is providing a copy of AEAS to several thousand local jurisdictions in the United States. Follow-on work could include creation of new WMD scenarios, and support for new scenario types, such as cyberterrorism. Enhancements may also include geospecific locations based on GIS data so that responders can train on a scenario that takes place in their own jurisdiction, and expansion into other response communities such as the Hospital Incident Command System (HICS). Because of the performance assessment component of AEAS, it could be especially effective in team training for first responder or military operations (Pigora, et al., 2002).

ACKNOWLEDGEMENTS

SAIC would like to thank the National Guard Bureau and Mr. Dan Donohue for their insight and support of this project.

REFERENCES

Courtemanche, A. J. (2000). The Incorporation of Validated Combat Models Into a Discrete Event Simulation (DES) CGF. *Proceedings of the 9th Conference on Computer Generated Forces and Behavior Representation*.

Pigora, M. A, Barshatzky, K., Kerrigan, M. and Murphy, H. (2002). A Task-Based Approach to Training, Evalulating, and Simulating WMD Response Behaviors. *Proceedings of I/ITSEC 2002*.

Van Veldhuizen, D. A. and L. J Hutson (1996). A Design Methodology for Domain Independent Computer Generated Forces. *Proceedings of the Eighth Midwest Artificial Intelligence and Cognitive Science Conference*. 86--90.