Cloud Simulation Infrastructure – Delivering Simulation From the Cloud

Michael R. Macedonia, PhD, Christina Bouwens, James Shiflett
SAIC
Orlando, Fla., USA
michael.r.macedonia@saic.com, christina.bouwens@saic.com, james.e.shiflett@saic.com

ABSTRACT

This paper discusses the practical aspects of architecting a Semi-Automated Forces (SAF) system for a cloud computing environment and describes some of our recent experiments with SAF technology in the context of a cloud-enabled environment. The rapid transitioning of traditional computer applications such as email to cloud computing is beginning to extend to military simulation. The ubiquity of the global Internet and advances in mobile computing are allowing the military to reexamine its business model for constructive simulation. Distributing exercises has been common for over a decade; however, the model has been based on scheduled, dedicated and often temporary infrastructure. Cloud solutions offer the potential of "anytime, anywhere," on-demand simulation and training capabilities. The primary challenge has been in architecting simulations for virtualization and providing the requisite security for military operations. Solutions to these problems are being vigorously addressed. This paper explores some potential implementations of a Cloud Simulation Infrastructure (CSI) concept – how a simulations system could be hosted and accessed via the cloud. Although not the same as cloud computing, high performance computing (HPC) has some useful similarities to cloud computing and may offer an alternative delivery infrastructure for simulation services. We offer results from our work in HPC and SAF systems as a partial contribution to understanding and defining the CSI concept. In addition, we present results from our work with a web-based interface for managing and deploying SAF resources. Combining the results of these two bodies of work, the HPC and the web-based interface, we have developed prototypical model of SAF computing in the cloud. From this vantage point, we also examine the benefits of the CSI concept, such as ubiquitous access, common (across Services) content, technical and operational standards for training, and potential for tactical mission planning.

Key Words: SAF, Cloud computing, Virtualization, Simulation

ABOUT THE AUTHORS

CHRISTINA L. BOUWENS is a senior research scientist for Science Applications International Corporation (SAIC) in Orlando, Fla. She has over 20 years experience performing research and development with Forces Modeling and Simulation (FMS) and high performance computing technologies, in particular with distributed networked simulation (distributed interactive simulation, high-level architecture). She holds a B.S. in mathematics from Geneva College and an M.S. in mathematical science from the University of Central Florida, where she is currently pursuing her Ph.D. in operations research.

MICHAEL R. MACEDONIA, Ph.D., is a vice president and chief scientist for SAIC. Prior to joining SAIC, Macedonia was the general manager for Forterra Federal Systems in Orlando, Fla., which was acquired by SAIC. Macedonia also was a senior national intelligence executive and the director of the Disruptive Technology Office (DTO, now IARPA) for the Office of the Director of National Intelligence. Prior to DTO, he was the chief technology officer for the U.S. Army Program Executive Office for Simulation, Training and Instrumentation (PEO STRI) in Orlando, Fla. Macedonia was responsible for developing the technology strategy for the U.S. Army's lead training technology acquisition organization. Macedonia has a Ph.D. in computer science from Naval Postgraduate School. He also has an M.S.T. from the University of Pittsburgh and a B.S. from the United States Military Academy, West Point.

JAMES E. SHIFLETT is a vice president and chief strategist for SAIC. He has over 30 years of experience in the management, research and development of DoD software systems. He has focused on strategic policy, project advocacy, operational requirements, test and evaluation, fiscal management and successful program execution. He formerly served as the director for FCS Training Systems and as the program manager for the largest simulation effort in the United States Army, the Combined Arms Tactical Training (CATT) program.

Cloud Simulation Infrastructure – Delivering Simulation From the Cloud

Michael R. Macedonia, PhD, Christina Bouwens, James Shiflett
SAIC
Orlando, Fla., USA
michael.r.macedonia@saic.com, christina.bouwens@saic.com, james.e.shiflett@saic.com

IMPERATIVES FOR CLOUD COMPUTING

Cloud computing has received a significant amount of attention leading to large information technology (IT) initiatives at the corporate and government levels. The National Institute for Standards (NIST) defines cloud computing as:

a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.[1]

It is now part of national and DoD IT Strategy by virtue of the "Cloud First" guidance developed by the chief information officer (CIO) for the United States. The guidance states:

To harness the benefits of cloud computing, we have instituted a Cloud First policy. This policy is intended to accelerate the pace at which the government will realize the value of cloud computing by requiring agencies to evaluate safe, secure cloud computing options before making any new investments. [2]

These policies are coming about because the technical imperatives for cloud computing (both public and private) are very clear. Computing environments are becoming more asymmetric with the development of large, high-performance data centers with virtualized hardware connected via broadband networks such as 4G LTE to resource-limited, albeit ubiquitous, mobile devices such as the iPad® (Apple Inc.). These environments are creating new and exciting capabilities for both civilian and military purposes that we address in this paper.

These imperatives, and the emergence of military cloud environments discussed later in this paper, provide major new opportunities for modeling and simulation. In particular, we have demonstrated this potential with the OneSAF® (U.S. Department of the Army) entity-level resolution simulation.[3]

Current Practices for Distributed Simulation

For over 22 years, the modeling and simulation community has been exploring the concept of distributed simulation and the interoperability of defense simulations beginning in January 1989 with the first workshop on interoperability standards. At that time, the Internet was only used by a small community (mostly academic) and there was no such thing as the World Wide Web.

Over the years, the distributed simulation community has expanded modeling and simulation (M&S) capabilities to include the following:

- Interoperability standards have been defined and are in use serving a broad range of users, from high-fidelity virtual simulations (distributed interactive simulation, DIS; high-level architecture, HLA) to faster than real time analysis applications (HLA) to support for live test and evaluation activities (test and training enabling architecture, TENA).
- Networks such as the Defense Research and Engineering Network (DREN) and the Joint Training and Experimentation Network (JTEN) have made coordinated training and testing events possible – linking locations across the country and the world.
- Gaming technology development has led to advances in graphical rendering of simulation environments, highly interactive immersive worlds, and an introduction to new applications for interactive distance learning and highly engaging training environments.

Much of the distributed, remote accessibility that the cloud community offers has already been enjoyed by the M&S community. However, in contrast to the direction of cloud computing, the M&S community has been ***distributing*** its resources, contrary to the cloud model that seeks to consolidate them.

Legacy Challenges for Distributed Simulation

Distributed simulation comes at a technical and operational price that limits it utility in everyday training and experimentation:

- "Fair fight" is difficult to guarantee in training and experimentation in a long-distance environment. Different latencies and computing resources are a direct result of the distributed model. Fair fight also become problematic even when similar simulations are run on a local area network.

- Current DIS and HLA models of simulation do not support persistence because they do not have a central store or control. Therefore, it is currently very difficult to have simulations of long-duration in a distributed environment.
- Distributed simulations generally do not support the use of handheld mobile devices. Handhelds such as Android® (Google Inc.) tablets have limited computing, memory and battery resources that are quickly overwhelmed by simulation requirements. Therefore, tactical commanders generally cannot run intense simulations on their future command and control devices.

- Each participating M&S site has to maintain its own facilities and equipment in order to participate in exercises.

    - This requires facility space, cooling, power and computational hardware.
    - Operating systems and software need to be installed and maintained.
    - Resources are idle when not in use for an exercise – this can be a substantial amount of time.
    - People are required to maintain the environment including regular maintenance, hardware and software upgrades – and tracking how upgrading or not upgrading affects interoperability with various participating systems.
    - Significant time and expense is expended retooling / reconfiguring existing hardware for different exercise events.

- Set up for a particular distributed simulation exercise can take months for coordination and weeks on the ground at various sites for installing and integrating participating simulation systems.
    - Costs are incurred for engineers to travel to site for exercise support.
    - Integration and testing cannot take place until all are available on site.
    - A team of operators is required to support the execution of an exercise. They must be available ahead of and during the exercise and sit idle if the exercise goes down.

RATIONALE FOR MODELING AND SIMULATION IN THE CLOUD

The federal "Cloud First" strategy has led our research team to recognize the potential value of cloud computing concepts, and we are actively exploring the use of cloud computing. We are looking at overcoming the current challenges of distributed simulations by centralizing simulation resources and effectively delivering training and simulation services to a broad set of distributed users at both the enterprise and operational levels.

We also believe that a cloud simulation infrastructure would be more defensible in the context of a cyber challenge. Cyber Command has recommended DoD move to cloud-based architectures for its intelligence systems. General Alexander, Commander, U.S. Cyber Command, recently stated*: "How do we create the next set of architecture that is more defensible and can ensure the integrity of our data? I think it's in the cloud."*[4]

We believe this will lead DoD to utilize data centers at existing Army facilities to deliver secure, high-performance, cloud-based simulation and gaming over Army networks. This approach simplifies a number of different issues related to utilizing live, virtual, constructive and gaming (LVCG) for training.

One of the largest of these issues is ensuring that soldiers have hardware capable of running the training at an acceptable frame rate. For example, gaming technology has historically driven the development of video cards with recent years seeing a doubling of relative graphics performance each year. Centralizing the processing of these video games in a data center greatly simplifies testing and deploying new hardware that enables the top-flight features of the latest games.

The Army will only have to upgrade servers at a relatively small set of data centers, and the benefit will seamlessly extend to all of the computers connected to the network. Another common issue is ensuring that soldiers have the most up-to-date training available. Training applications that are installed on a dispersed set of computers are much more difficult to upgrade than training applications that are installed at a relatively small set of data centers. Updating the training at a data center makes the latest version immediately available to everyone on the network without having to touch each individual computer.

Implications for Modeling and Simulation

Unlike typical IT needs, M&S applications tend to use the underlying virtualized hardware more extensively for prolonged periods of time. The applications have higher memory requirements, intensive central processing unit (CPU) usage, minimum CPU counts per node, multiple distributed nodes, and a low latency / high bandwidth network. During execution, the demand on the virtualized hardware will be at a sustained high load for significant portions of the simulation exercise.

Not all M&S applications will reside in the cloud. The integration of live, virtual and constructive simulations along with command and control (C2) systems and other operational equipment requires M&S cloud implementations to allow for a mix of cloud and non-cloud resident applications.

With a fairly mature distributed simulation infrastructure, remote access with participation from individuals has been well established. M&S has solved the harder problem of how things actually connect, where cloud promises that things will work without ever defining how that happens.

THE BENEFITS OF CLOUD / VIRTUALIZATION SUPPORT

Cloud seeks to consolidate resources, providing the ability to co-locate many of the “back-end” components and all within a common hardware infrastructure. The virtualization component of cloud reduces the need for hardware (and space and power). For example, 50 standard PC configurations for running a very large OneSAF exercise would require a substantial lab space for all the machines and monitors (chairs, etc.). The same exercise using virtualization takes a fraction of the space with 50 virtual machines on 50 (or less) cores and running on a server. Zero or thin clients are needed only for the participants or pucksters of the exercise. This reduces the system footprint to the space necessary to support the operators and eliminates traditional PC workstations from the exercise. The overall hardware, power and space footprint is greatly reduced. Resources can be reallocated for other applications and exercise configurations without the need to wipe and reinstall the operating system and applications every time.

Cloud continues to support remote access, utilizing the same network infrastructure already in use with more traditional distributed simulation configurations. Zero, thin or thick client access options are available depending on display performance needs and location.

Pre-exercise scenario development, analysis and dry runs can take place on the cloud resident resources. Users can upload, run and share scenarios without having to download and run on locally managed systems. Software is centrally updated so users have access to the same, most recent software version as well as the old version.

The cloud capability for on-demand self service and multi-tenancy provides M&S users with availability of the simulation as needed and accessible even while other users might be using the same application.

Offering M&S as a Service provides many benefits relative to what we do today while allowing innovation in the way we utilize M&S.

- Centralized hosting of simulation resources decreases the cost of ownership by reducing licensing requirements, hardware and software maintenance / upgrades, and facility resources (such as power and space).
- With flexible, scalable environments, ramping up new users and exercise environments are performed more quickly and at relatively modest cost, resulting in faster implementation and “time to value”. In addition, the environment can scale according to need, increasing in size as expanded capability is needed, decreasing when needs are reduced – with overall ability to adapt in environments of sporadic use.
- This provides an environment that is device and location independent, expanding the accessibility of the resources.
- Together this provides for increased collaboration amongst the users of the environment – with common access to the same resources. Updates to the resources provide all the users with access to the same capabilities.

With these tools in place, we have the ability to change the entire process model for how we compose and use M&S capabilities, thus providing an opportunity for innovation. With freedom from lengthy implementation timelines, one can quickly and inexpensively "try out" new ideas. Users are also not held back from utilizing new tools and capabilities because of the need to support legacy systems.

This is a fundamental shift in how simulation will be delivered to the user community.

Our M&S-as-a-Service Exploration Approach

Our approach to addressing M&S as a Service was to begin by consulting with our local IT cloud experts. A number of hardware options were offered to us, but as we defined the virtual machines we needed for running our applications, it became clear that our resource needs were very different from what the standard IT offering could provide. An important lesson here was not to count on the IT department to provide the cloud capabilities needed for simulation.

Based on the NIST definition for cloud, we decided to focus on capabilities that were uniquely cloud (and not simply virtualization). Though virtualization was part of our solution, we wanted to see how the REAL cloud would support our needs.

SIMULATION IMPLEMENTATION WITH ONESAF

Our initial prototype used a simulation with a flexible, mature software architecture (OneSAF) and a virtualization infrastructure that is relatively mature and supports the functionality needed to support simulation's unique use of the environment (VMware®, Vmware, Inc.).

OneSAF is the Army's next-generation entity-level simulation that provides a composable, distributed and scalable simulation of real-world battlefield situations using validated physical models and doctrinally correct behavior models. It can support analysis, acquisition, planning, testing, training, and experimentation. OneSAF allows users to compose a wide range of complete simulation systems from a set of component-based tools, develop new or extend existing tools, as well as compose new single or multi-resolution entities, units, and associated behaviors from existing physical and behavioral software components.

VMware has a mature product line for virtualization services. Their cloud computing products provide a flexible, tailorable environment for automation and control of infrastructure resources. Access to their software tools has allowed the CSI project to quickly provision resources for simulation use.

OneSAF was a particularly good candidate for testing targeted cloud capabilities because of its composable architecture and flexible interface. OneSAF uses a gateway to communicate with the provisioning / service broker. This allows the broker to feed OneSAF with control commands from the user interface – allowing a user to configure, initialize and execute a scenario without needing access to the OneSAF Management and Control Tools (MCT). The gateway also allows the simulation to direct the provisioning and configuration of virtual machines (VMs) for use in a user designated configuration and exercise.

VMware's Vcenter functionality allowed our gateway to issue commands to provision and unprovision resources, to install the OS, software and exercises needed by the simulation. This ensures that the details of the configuration of the infrastructure for any exercise is "hidden" from the user – who focuses on the simulation and does not require any knowledge of how the simulation is hosted within the environment.

We developed an implementation that explored how simulations would execute in such an environment and how the virtualization and cloud tools provided by our selected hypervisor could be leveraged to create new M&S capabilities.

Based on our experience, we developed a framework for defining the environment that could help the community define standard approaches for cloud delivered M&S services.

We created a cloud-based service for OneSAF, called Cloud Simulation Infrastructure (CSI), that enables the Army to deploy simulation solutions directly to warfighter locations or to centralized simulation centers via enterprise networks. The result is a solution for providing training with lower operator overhead requirements, reduced exercise lead times and lower overall hardware capital costs associated with legacy simulation approaches.

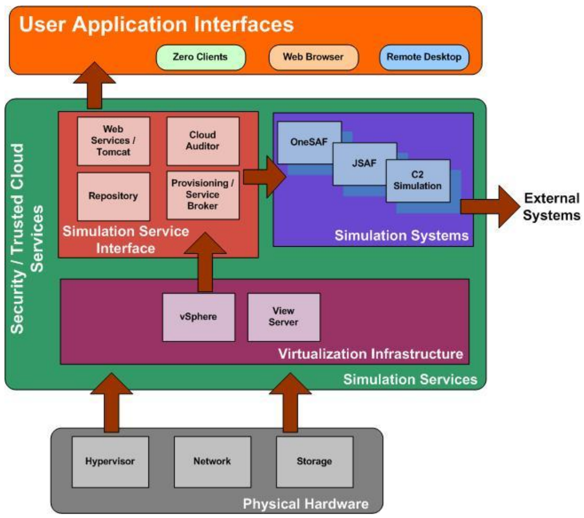

Figure 1. M&S- as- a- Service Framework

Modeling and Simulation as a Service Framework

The M&S as a Service Framework consists of three main components (Figure 1): User Application Interfaces, Simulation Services, and Physical Hardware.

Framework Overview

The User Application Interfaces include zero, thin or thick client interfaces that allow the end user to interact with the simulation services. The Physical Hardware provides the CPU, memory and network hardware infrastructure for hosting the simulation services component. The Physical Hardware implemented is a high-end network server utilizing virtualization for hosting simulation service components. Future work will explore how HPC hardware might be addressed in this framework or something similar.

The Simulation Services component contains the actual simulations used for creating the simulation services being delivered. These simulations can interact locally with other locally hosted simulations or may connect to external systems. In addition, the simulation services component contains the virtualization infrastructure. In the case of an HPC not utilizing virtualization, this subcomponent might instead provide a set of job scripts that will assign simulation component processes to specific nodes in the underlying HPC infrastructure.

Prototype: Physical Hardware

Our initial configuration was built using three-year-old, high-end workstations (six total), each with Intel® Core™ i7 (quad core) processors (Intel Corporation), 12Gb RAM, 256Gb storage, and an NVIDIA® graphics card (NVIDIA Corporation). With this six-machine configuration we were able to stand up a virtualized environment and generate upwards of 25 VMs running two to three separate OneSAF exercises running on separate virtual networks.

Our current configuration consists of two host systems, each with

- 40 physical CPU cores (80 hyperthreaded)
- 256Gb physical memory
- 8Gb fiber backbone between hosts and data store with a 1Gb network connection outside

In addition to the host systems, we have a 2TB data store. This configuration provides support for approximately 200 simulation-configured VMs.

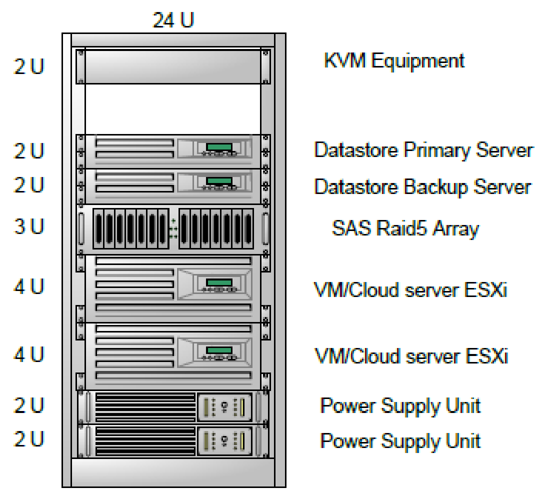

Figure 2. Cloud Server Configuration

Prototype: Virtualization Infrastructure

SAIC worked with VMware to assemble the right tools to support the capabilities required by our simulations.

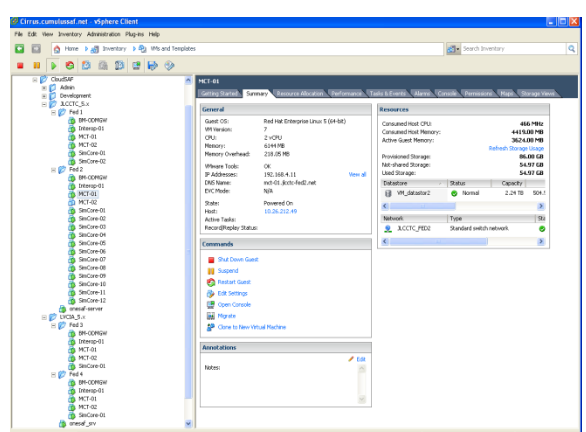

Figure 3. vSphere Access

In the end, we implemented:

- VMware vSphere®5 (Vmware, Inc.)
- VMware View™ 5 (Vmware, Inc.)

These products (which were themselves virtual appliances) provided a number of important capabilities:

- Linked cloning: Allowed our configuration to quickly clone VMs per the exercise configuration specified by the user
- Cloud / VM monitoring: Provided view into the performance of the individual VMs so we could monitor the state of the virtualized exercise as it responded to our user interface inputs
- Data redundancy and failover. We were able to isolate faults in individual VMs without interrupting the overall exercise.

In addition to the VMware products, we also used two products from Teradici Corporation: Zero Client and PCoIP®.

INITIAL CLOUD SERVICES IMPLEMENTED

Based on the NIST definition of the five essential characteristics for cloud computing we have implemented the following:

On-Demand Self Service: The configuration interface allows the user to select a OneSAF configuration and scenario – then causes the infrastructure to provision the needed resources using existing templates and linked clones. Metering is not yet provided.

Broad Network Access: Access to OneSAF is available over the network and has been demonstrated using thin client (web app, mobile app) as well as zero client and VM viewers. We have displayed our implementation (controls and OneSAF displays) using laptops, tablets and mobile phones.

Resource Pooling: OneSAF resources are dynamically assigned at the time of exercise configuration and then released when the resources are unprovisioned by the user in their application. The end user has no knowledge of where or what host the simulation services are running on. Currently, there is no control over the location of the resources since current delivery from a single location.

Rapid Elasticity: This is a key capability that has been implemented where resources are rapidly provisioned based on user configuration inputs. Initial configuration is performed automatically. We have also prototyped on the fly configuration of resources when the provisioned set proves to be inadequate to the simulation task at hand.

Measured Service: This has not yet been implemented but is part of our 2012 research effort.

The CSI driven OneSAF is more than just Software as a Service (SaaS) where users are provided application level support. Our implementation controls the underlying platform and infrastructure services – and they are modifiable based on *simulation management and simulation exercise events*. The implementation crosses cloud layers so that the delivered simulation services are supported by the underlying infrastructure as needed – yet this capability is hidden from the user. From a users perspective, they are configuring a OneSAF exercise – the infrastructure provisions the required virtual resources (VMs, virtual nets) on the fly (they don't need to be pre-configured and sitting in VMs ready for use). Our goal in this implementation was to demonstrate the cloud philosophy that hides all the details of what was going on underneath the hood in order to offer a pure simulation service. The user accesses a service – CSI seamlessly delivers the service. It is this capability that takes the implementation beyond virtualization and more toward the cloud.

RESULTS

Our solution was unique in that we have developed the ability for OneSAF to dynamically provision new processors and VMs on demand in order scale to very large scenarios (more than 20,000 entities).

We developed middleware for OneSAF and VMware to exploit mechanisms that let CSI know when the simulation needs more compute resources as entities and scenario complexity impact performance. This middleware can be used with other simulation systems. CSI automatically cloned VMs to distribute the processing of OneSAF. The internal network of 1Gb/s provides more than sufficient data transfer capacity for several hundred VMs.

CONVERGENCE OF CLOUD SIMULATION WITH OPERATIONAL ARCHITECTURES

The Army and the Air Force (USAF) have recognized this force for change at both the enterprise and tactical levels of the military. For example, the CIO for the USAF recently explained their rationale for pursing cloud computing in the enterprise:

"So I think the cloud may have a benefit there and may be a way for us to get at the mobile apps and mobile computing in a manner that now allows a greater mobility and at the same time doesn't increase the security or decrease our security posture."[5]

The Army is quickly moving cloud computing and virtualization into the tactical realm with its Distributed Common Ground System (DSCGS-A).

In some cases, cloud computing has shortened the time needed to analyze information from days and hours to literally seconds, said Clark Daugherty, Distributed Common Ground System-Army program manager for Lockheed Martin Global Training and Logistics.
"This program directly saves soldiers' lives when it comes to dealing with improvised explosive devices, intelligence about planned ambushes and attacks," Daugherty said."[6]

DSCGS-A is central to the Army's Operations-Intelligence (Ops-Intel) Convergence (OIC) concept identified by Program Executive Office for Command, Control, and Communications-Tactical (PEO C3T) for the Common Operating Environment represents an enormous step forward for the Army. The existing Army command and control environment at tactical level has a large level of complexity because of the existing pattern of deploying individual computing hardware with each software element of the Army's mission command systems. Tactical cloud environments have the potential to save both personnel and dollar resources while making the Army more tactically agile and powerful by virtualizing these applications on high-performance, multiprocessor computers.

The Future

Imagine the scenario where a tactical commander at the edge of the battlefield will soon be on their handheld device, using the power of the cloud to ask

- Computer – Plot three routes from my current location (point A) to my new assembly area/linkup (point B) where I will link up with local friendly forces. I want four routes, fastest, shortest, best concealment, and best coverage given current threat intel.
- Computer – In the point B assembly area, plot recommended location of my comms devices for the best coverage within the assembly area.
- Computer – I have three sensors with me, two EO/IR and one thermal. Plot recommended locations for these sensors to cover expected avenues of approach into the assembly area.
- Computer – I have two machine guns and six rifles, plot recommended locations and range cards to provide the best 360 protection for these weapons.
- Computer – I see a group of people on cell phones moving to my location. Who are these people and what do you know about the cell phones?
- As the group moves to link up and have provided the password, each person is asked to look into a handheld device and state his/her name. The computer takes a picture and checks his/her voice to verify that these are the expected people.

These are not science fiction scenarios. They are capabilities that are, for the most part, possible now with technologies that combine speech understanding, cloud and mobile computing. Add to this mix entity-level simulation with high-resolution geospatial databases and we have a revolution in training, mission planning and command systems.

We are now continuing this research in a working relationship with the Army Research Labs under a Cooperative Research and Development Agreement (CRADA) for Cloud-based Simulation. We also are now conducting experiments with Communications-Electronics Research Development and Engineering Center (CERDEC) in using the Cloud-Simulation Infrastructure to support long-distance experiments. Finally, we are examining the potential of employing CSI on tactical systems to be used for mission planning.

The convergence of OneSAF with CSI promises new capabilities undreamed of a decade ago and is being driven by advances in technology and the Army's experience in moving high performance computing forward to tactical units in combat areas. We see the possibilities of CSI in that environment and are now exploring ways to make this future real.

ACKNOWLEDGEMENTS

We would like to acknowledge the support of the Army Research Labs under CRADA #12-05. We also thank the VMware team to include Scott Drucker and Jay Bokulic and appreciate the technical support of the SAIC research team, Kurt Hawkes and Jay Duncan.

REFERENCES

[1] National Institute of Standards and Technology, The NIST Definiion of Cloud Computing, Special Publication 800-145, September 2011, http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

[2] Vivek Kundra, Federal Cloud Computing Strategy, US CIO, February 8, 2011, http://www.cio.gov/documents/federal-cloud-computing-strategy.pdf

[3] http://www.onesaf.net/community/

[4]http://www.nextgov.com/nextgov/ng_20111027_3085.php?oref=rss?zone=NGtoday

[5] "Air Force CIO pushes for operational mindset", Defense Systems, 9 August 2011, http://defensesystems.com/Articles/2011/08/08/Interview-Lt-Gen-William-Lord.aspx?admgarea=DS&Page=1

[6] http://www.jsonline.com/business/state-companies-helping-army-with-cloud-computing-7258roe-150105585.html