Embracing the Cloud – Providing Simulation as a Service

Dr. Daniel Lacks
Cole Engineering Services, Inc. (CESI)
Orlando, FL
Daniel.Lacks@coleengineering.com

Lawrence A. Rieger, CMSP
U.S. Army Capabilities Integration Center
Fort Eustis, VA
Lawrence.a.rieger.civ@mail.mil

ABSTRACT

In August 2013, The Army Capabilities Integration Center (ARCIC) Director challenged that ARCIC could save big dollars if we could put OneSAF "in the cloud" in the Army Battle Lab Collaborative Simulation Environment (BLCSE). What seemed like a simple technical migration was actually a significant change in the way simulations would be structured and operated within a major distributed simulation environment (Battle Lab Collaborative Simulation Environment – BLCSE). While typical distributed simulation environments use either the High Level Architecture or Distributed Interactive Simulation protocols to exchange data between federates, a cloud environment seeks to remove as much of these data exchanges, and the resulting network infrastructure and latency, as possible. Software as a Service (SaaS) is the push of the DoD Cloud Computing strategy, and Simulation becomes the Software being provided as a service within cloud simulation. This paper details a simulation cloud testbed and several technical evaluations conducted to determine Simulation as a Service within the DoD Cloud Computing Strategy. It provides lessons learned and practical planning steps for the migration of a distributed simulation network into a community cloud environment with particular attention to the intricacies of establishing a robust Virtual Machine simulation environment. The authors also address the technical architecture problems associated with cloud computing, community issues of network redesign and the DoD Information Assurance Program (DIACAP) as well as the resource investment and cost benefit analysis for distributed workstations vice central blade servers. The more demanding configuration management and configuration control issues of simulations in the cloud, providing Modeling and Simulation as a service, are also addressed. The paper is based on an IRAD simulation cloud testbed and a series of distributed technical tests demonstrating SaaS over an Army secure simulation network.

ABOUT THE AUTHORS

Dr. Daniel Lacks is the Chief Scientist at Cole Engineering Services, Inc. (CESI) and Conceptual Modeler on the One Semi-Automated Forces (OneSAF) Integration and Interoperability Support (I2S) program. He received Computer Engineering BS (2001), MS (2002), and PhD (2007) degrees from the University of Central Florida specializing in distributed computing, software simulation, and software engineering. Dr. Lacks has spent the last 13 years working as a software and systems engineer in the DoD M&S industry on Live, Virtual, and Constructive programs to include Warfighter's Simulation (WARSIM), Joint Simulation System (JSIMS), One Tactical Engagement Simulation System (OneTESS), Joint Land Component Constructive Training Capability (JLCCTC), Live Virtual Constructive-Integration Cell (LVC-IC), Live Virtual Constructive-Integrating Architecture (LVC-IA), and OneSAF.

Lawrence A. Rieger is the Futures Technology Simulations Analyst for Joint & Army M&S Division (JAMSD), TRADOC ARCIC. He received a BA from Belmont Abbey College in 1976, an MS from Troy State University in 1982 and an MA from Kings College in 2013. He is also a graduate of the Army Command and General Staff College and the Army Management Staff College. Following active and reserve commissioned service with both light and mechanized forces, he has spent the last 29 years in the management and development of simulations for training and combat developments, working in live, virtual, and constructive environments. His prior assignments include Technical Configuration Manager for the Army Battle Lab Collaborative Simulation Environment (BLCSE) and Deputy TRADOC Project Manager – OneSAF. He received the CMSP in 2008

Embracing the Cloud – Providing Simulation as a Service

Dr. Daniel Lacks
Cole Engineering Services, Inc. (CESI)
Orlando, FL
Daniel.Lacks@coleengineering.com

Lawrence A. Rieger, CMSP
U.S. Army Capabilities Integration Center
Fort Eustis, VA
Lawrence.a.rieger.civ@mail.mil

BACKGROUND

Between the DoD push to cloud computing, and Army Data Center Consolidation Plan, and the steadily declining resources available to support distributed simulation events, the military simulation community is being pushed to move our simulation infrastructure and activities "into the cloud." The challenge comes in that most simulations are not designed or architected for cloud operations, and as a community we are unfamiliar with the technical and cultural issues involved in moving from a physical simulation facility and traditional distributed simulations into a cloud environment.

THE DOD MOVE TO CLOUD COMPUTING

DoD began its movement to the cloud in 2010 with the Secretary of Defense (SecDef) Gates' Efficiencies Initiative and the Federal Data Center Consolidation Initiative (FDCCI). Driving the FDCCI was the belief that "By shutting down and consolidating under-performing data centers and optimizing the data centers in our Federal inventory, we stand to save taxpayers billions of dollars and curb spending on underutilized infrastructure (CIO.gov, 2014)." As the FDCCI moved through DoD to the Army, we faced the need to move much of our simulation resources, particularly complex distributed constructive simulations and serious gaming, into a cloud environment. Recognizing that cloud computing is very attractive to our senior leadership, particularly the IT leadership, and that cloud computing simulations offer significant resource savings, there is good reason to embrace Simulations as a Service.

DoD Cloud Computing initiatives have been well underway, with the DoD Cloud Computing Strategy being published in July, 2012. However the primary thrust has been within typical enterprise Information Technology (IT), rather than in the simulations arena. The Army has pushed forward with the Army Data Center (ADC) Consolidation Plan (ADCCP), which is aimed at significantly reducing the number of enterprise and local data centers, but which has also included in its program various Army simulation centers and faculties, which physically meet the definition of an ADC although not in application practice. Within the ADCCP, simulation facilities are designated as Special Purpose Processing Nodes and generally remain exempt from consolidation (DOD CIO, 2013). However, the potential for savings, in real facilities and computing resources are causing senior leadership to push simulations into the cloud. Although there have been many movements in simulations in virtual machine processing, computer resource consolidation, and remote operations of simulations, to date none of the primary Army simulation systems have actually migrated into a cloud environment.

Defining the Cloud

The DoD Cloud Computing Strategy notes that while the traditional IT delivery model is focused on the development, maintenance and operation of computing hardware and software, the cloud computing model focuses on providing IT as a service. Within IT as a service, there are three services that are the core of cloud computing: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). It is important to note that while there is a major focus on virtualizing simulation hardware, or using High Performance Computers, this is not the basis of cloud simulations. It is the change from "us doing" the simulations to "them providing" simulation functionality as a service that marks the definition of simulations in the cloud. The National Institute of Standards and Technology (NIST) defines cloud computing as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources" and provides four cloud models: Private, Public, Community and Hybrid (Liu et al., 2011). Of these four cloud models, only Private and Community are likely to be of interest to the military simulations community. Although Public and Hybrid cloud models may be used for more open military gaming, including a hybrid model being likely for the Army Early Synthetic Prototyping effort (Darken, 2014), most military simulation work will be either Private or Community due to information security (including classified materials) requirements. Under the NIST definition, a Private Cloud Infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and may exist on or off premises. The only difference to a Community Cloud is that the infrastructure is used by a "specific community of consumers which have shared concerns" rather than "a single organization comprising multiple consumers." From a military viewpoint, it seems the difference is merely whether all the users belong to a single command, or whether we are part of a community of practice.

From our perspective, moving to simulations in the cloud means shifting from a series of distributed simulation facilities, each hosting their own various proponent and shared simulations and tools, interacting through the High Level Architecture (HLA) and a common Federation Object Model, and adopting Simulation as a Service, where a single "cloud facility" maintains the integrated software for the full simulation federation, and the simulation centers become an SaaS user, sharing some level of PaaS and IaaS. The simulation centers are GUI providers to the using soldiers, with the configuration management and software integration being a service to the various user facilities.

Defining Simulations in the Cloud

Simulation as a Service (SAAS) is the concept that the person using a simulation does not provide the hardware, software, and network infrastructure for a simulation event, but rather uses the provisioning of another servicing entity, the vendor, as the user requires. The intent is that a local simulation facility no longer maintains simulation software, updates simulation versions, provides computational hardware, and builds out a simulation event on various machines. Instead it operates solely at the user interface level, what could be called the Graphic User Interface (GUI) level, such as those available from thin client applications accessible from a web browser. When a unit shows up at a simulation center, neither they nor the facility staff load the simulation software and scenario files, wire computers in local networks, determine how many computers are need to handle the simulation load, or make sure all the peripherals are properly connecting and operating. They just sit down and start the event. In the same way, using SaaS, enables the simulation center, whether a battle simulation training center or an experimentation Battle Lab, to move out of the complex event setup and integration process, as well as the care and feeding of a large computer hardware and software inventory, and merely become the providers of desks and tables with simple web browsers pointed to the appropriate cloud provider.

THE ARCIC ENVIRONMENT

The Army Capabilities Integration Center (ARCIC) operates within a classical distributed simulation environment designated the Army Battle Lab Collaborative Simulation Environment (BLCSE). These various Battle Labs, not all TRADOC facilities, sustain a standing network which operates at the Secret level, and maintains the HLA (HLA 1.3NG, IEEE 1516) as its communications standard. In a typical simulation event, there are ten Battle labs generating entities using OneSAF, FireSIM, ATCOM and EADSIM as the primary simulations, and up to five additional sites on the network, usually subscribing to the Run Time Infrastructure simulation traffic. VMware is not a common use within the BLCSE at present, with just one Battle Lab operating virtual machines on any significant basis; using a mix of blade server, rack server and physical machines.

Because of the classified data, BLCSE is a closed network, using firewalls with deny by default as security filters, and functions as a private virtual network on the Defense Research and Engineering Network (DREN) backbone. As BLCSE looks toward cloud computing, the restrictive security requirements drive to the DoD CIO definition of a Closed ("community") Cloud (Takai, 2012). This does make a significant difference from what is classically viewed as moving "in the cloud", but probably aligns what will become normal practice within DoD simulation activities.

The Objective BLCSE

The ARCIC objective is to transition the BLCSE simulation architecture, currently a traditional distributed simulation federation, into a cloud environment where the users; military role-players and simulation inter-actors, interact with the simulation federation through SaaS. The users will remain at their current locations, within the simulation Battle Labs co-located with the proponent centers for the various Warfighting Functions. Since SaaS replaces locally hosted and prepared simulations, the enabler function of cloud computing will significantly change as local enablers will be simple Information Technology technicians and the local Information Assurance team. While proponent simulation developers will also remain around their current Battle Labs, setting up and integrating the simulations, including loading databases and instantiating the simulation event, will all be done by the providers at the cloud host site. The local IT technicians will be setting up simple web browsers that will point to closed cloud URL, which will then load the appropriate simulation interface and instantiate the user workstation interface. Establishing this objective environment demands significant changes in network architecture design, internal configuration management and configuration control processes, and federation integration processes. Although discussed in detail separately (Rieger, 2014), the larger complexities of federation integration in SaaS require close attention to the design of the cloud Local Area Network (LAN) and the interactions between various simulations and simulation tools that are provisioned there. This paper addresses the technical pieces of a single simulation being virtualized. The same process is necessary separately for each simulation and tool federate.

SOME ONESAF HISTORY IN DISTRIBUTED SIMULATIONS

OneSAF is an event driven distributed simulation capable of communicating over Local and Wide Area Networks (LANs and WANs) with other OneSAF nodes including the OneSAF InterOp and Object Database (ODB) Synchronization Services (OSS). The InterOp is capable of speaking common simulation protocols such as Distributed Interactive Simulation (DIS), HLA, and Test and Training Enable Architecture (TENA) so that OneSAF can federate with other simulations and devices. The InterOp also takes care of operations to map enumerations and time synchronization differences between protocols or federations. In the BLCSE environment, for example, OneSAF is grouped across distributed sites into clusters and then linked together via the InterOp as different HLA federates along with other simulations participating in the BLCSE federation. Other federations use techniques such as the RTI Distributor or mcDistributor nodes, DIS gateways, and TENA middleware to distribute data over WANs.

The OneSAF ODB protocol is optimized for LANs by using combinations of unicast, multicast, and broadcast network traffic. Unfortunately, typically without a major network overhaul, it is not possible to generically send multicast and broadcast traffic over a WAN. As of OneSAF v6.0, the OSS capability was introduced to distribute ODB data over WAN connections. The OSS provides many optimizations to consolidate multicast and broadcast traffic over a reliable unicast WAN connection, eliminate heart-beating activities, and selectively filter out bandwidth consuming data such as some articulated parts updates. This optimization, however, means that only a select few OneSAF system compositions can be run at remote sites including the Management Control Tool (MCT) workstations (non Battlemaster), Mission Command Adapter Web Service (MCA-WS), ARES, Unified Data Gateway (UDG). This limitation is mostly removed with the Version 8 enhanced Web Control Tool (WCT).

OneSAF does have previous experience in cloud operations, being used with remote operator stations in the CERDEC experiments which used Web Control Tools (WCT) at Aberdeen and Ft. Dix while the actual simulation was hosted in the Orlando RDECOM STTC facility. However, these experiments were conducted as a single simulation instantiation lacking the challenge of a federation over the RTI.

TECHNICAL DESCRIPTION AND LESSONS LEARNED, CESI CLOUD TESTBED

Beginning in February 2014, as an IRAD effort, CESI built a BLCSE cloud emulation environment using ESXi servers representing two virtual data centers separated by WAN emulation software (WANem) and vCloud Enterprise Suite running separate clouds in each virtual data center. The BLCSE "main site" consists of the OneSAF simulation running Simcore servers, the Unified Data Gateway (UDG web server), OneSAF Database (ODB) Synchronization Service (OSS) Server, and Battlemaster system compositions sending Distributed Interactive Simulation (DIS) packets to a Joint BUS (JBUS) representing a DIS simulation for the experiment. The BLCSE "remote site" consists of the OSS Client (linked to the OSS Server via the WANem link), a remote UDG web server, Management Control Tool (MCT) workstations, and Web Control Tool (WCT) thin client workstations. The cloud emulation environment is graphically depicted below in Figure 1.

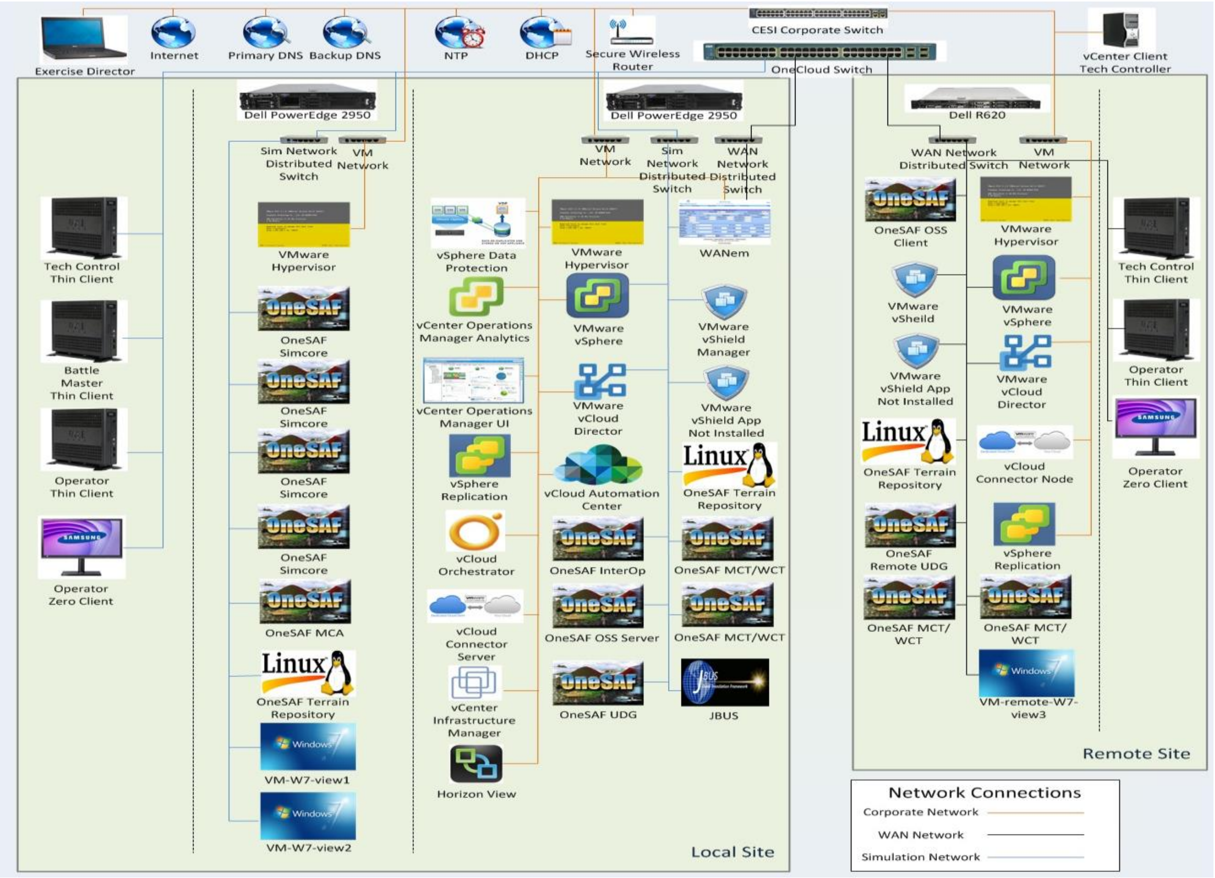

Figure 1. OneSAF Cloud Emulation Environment

While vCloud Enterprise Suite provides a multitude of capabilities, the capabilities this investigation focuses on are vCloud Director, vCenter Operations Management Suite Enterprise (vCOPS), vCenter Orchestrator (vCO), and vCenter Site Recovery Manager (SRM). vCloud Director allows the creation of a BLCSE Data Center and provides the ability to share leased OneSAF preconfigured VMs to client sites for a configurable time period. vCOPS provides the capability to examine detailed CPU, hard drive, memory, and network performance metrics as well as producing performance reports. SRM replicates the sites participating in the event to a remote site in the event of an outage; SRM provides the capability to rehearse a network outage to practice recovery and failover procedures. vCO provides the capability to automate functions using workflows which can be linked to vCenter Server (vCS) node menu options or run directly from vCO. Various vCO workflows were created to automate activities:

- Start OneSAF and optionally power on the VMs on demand.
- Stop OneSAF, collect logs, delete temporary files, and optionally power off the VMs on demand.
- Gather performance data on demand using metrics derived from the vSphere “Report Performance” feature.
- Respond to a custom CPU utilization alert to capture performance data using metrics derived from the vSphere “Report Performance” feature.

Three network subnets were used for this experiment: one for the VMware Management LAN contained in the CESI corporate network, one for the local site LAN, and one for the remote site LAN. WANem provides the Network Address Translation (NAT) capability to route the packets between the different subnets while applying bandwidth and latency constraints on the data when crossing subnets. VMware VXLAN3 technology provides the capabilities to create distributed virtual networks with the capability to transmit unicast, multicast, and broadcast packets over a WAN. The capability to transmit multicast and broadcast packets, however, was not needed because the DIS traffic is not sent to the remote site with this server site centric model which in-sources the server capabilities at the main site. Also, the OneSAF Simulation Object Relational Database (SORD) multicast and broadcast traffic was consolidated by the OSS Server/Client components and also replaced by the UDG REST web services protocol communicating with the WCTs.

In order to prove that the WCT is truly a thin client solution, WCTs were run on WYSE Z90SW thin client terminals. Thin client solutions prevent persistent data from being written to long-term storage. This assists with maintenance activities between training events since the WCT is a zero deployment thin client and the changes made by non-administrative users during an event are lost upon rebooting the thin client. While the WCT operates on the thin client hardware, scenarios consisting of 15,000 entities and greater require more than 2 GB of memory that the thin clients are configured with. The thin client machines were upgraded with an additional 1 GB of memory to accommodate this situation. The thin clients also provided a workstation for tech controllers running the vSphere Web Client. The vSphere Web Client requires a VMware browser plugin to run the console feature, used to manage the Battlemaster MCT via the browser, which requires approximately 200 MB of hard drive space. The thin clients did not have enough storage space to add this plugin, so an additional hard drive was added to provide the storage space for the plugin. Overall, the thin clients performed well, but it is recommend not buying thin clients which are “too thin.” Adding the memory and extra storage still kept the price of the thin client lower than a traditional workstation PC.

DEFINING THE VMWARE WITHIN THE EMULATION TESTBED

There are many hypervisors and software applications that allow users to create virtualized cloud environments. Kernel-based Virtual Machine (KVM) is a Type 1 Hypervisor built into the Linux kernel infrastructure and is available with Linux distributions such as Red Hat Enterprise Level (RHEL) 5.4 and above, Ubuntu 10.04 LTS and above, SUSE Linux Enterprise Server (SLES) 11 SP1, and other distributions. Xen Project is also a Type 1 Hypervisor maintained by the Linux Foundation which uses the Xen Management API (XAPI) to enable cloud services. OpenStack is an open source cloud solution which is compatible with KVM and Xen Project. Microsoft’s Hyper-V is available installable role in Windows Server 2012 or as a standalone product called Hyper-V Server. VMware ESXi developed its own vmkernel to load device drivers based on Linux device drivers to create an enterprise capable Type 1 Hypervisor solution. In technical definition, a Type 1 Hypervisor is directly resident to and boots with the machine hardware before an Operating System (OS) is applied. A Type 2 hypervisor interfaces to the machine operating system, and has an additional OS for the application resident to it. Currently the major actors within the Hypervisor community, including VMware, Microsoft and Citrix, operate as Type 1 hypervisors, as does both the CESI cloud emulator and the BLCSE evaluations. The BLCSE cloud, once implemented, will be a Type 1 Hypervisor.

VMware ESXi, part of the VMware vCloud Enterprise Suite v5.5, is the chosen hypervisor for this investigation for several reasons. DISA STIGs are available for VMware ESXi which provide guidelines for locking the hypervisor systems down in a secure environment. VMware agents running in the VMkernel are digitally signed by VMware thus creating a tightly locked-down system. VMware vCloud Enterprise Suite provides a multitude of hardened capabilities compared to lesser VMware packages (such as vSphere Enterprise Plus) and non-VMware virtualization and cloud solutions. The U.S. Army and VMware signed an Enterprise Level Agreement (ELA) for lowering the costs of the vCloud Enterprise Suite and other VMware licensing costs by 60% due to the volume of licenses that the Army will purchase; this also was a motivator for choosing VMware vCloud Enterprise Suite.

The emulation testbed VMware vCloud Enterprise provides the following high level capabilities:

- vSphere Enterprise Plus; includes the ESXi hypervisor required to host the VMs, vSphere Client and vSphere Web Client are used to manage the VMs through the vCS, and uses vSphere Data Protection (backup), vCenter Orchestrator (automation), and vCenter Single Sign On Service (SSO for common usernames and passwords across VMware applications)
- vCenter Operations Management Suite Enterprise (Monitor health of the cloud)
- vCenter Operations Management Suite Enterprise
- vCenter Infrastructure Navigator (Understand what is running on a VM)
- vCenter Site Recovery Manager Enterprise (Replicates sites across a WAN)
- vCloud Director to Create and manage Virtual Data Centers (VDCs)
- vCloud Networking and Security (vShield provides firewall and virus scanning optimizations)
- vCloud Connector Advanced Connects private clouds and public clouds

- vCloud Automation Center Enterprise Automate the VM deployment process
- VMware Horizon 6.0 which delivers virtualized Windows and remote desktop applications

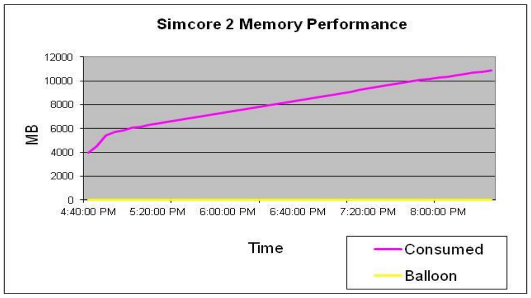

Figure 2. Simcore Memory Performance

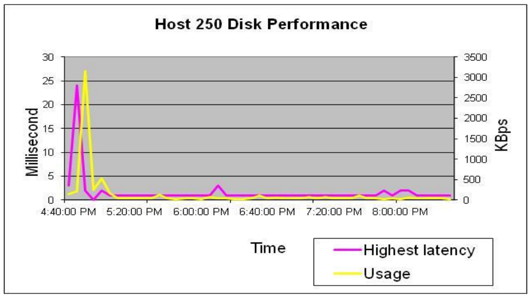

Figure 3. Host 250 Disk Performance

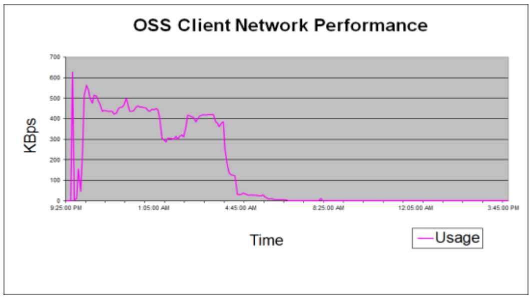

Figure 4. OSS Client Network Usage

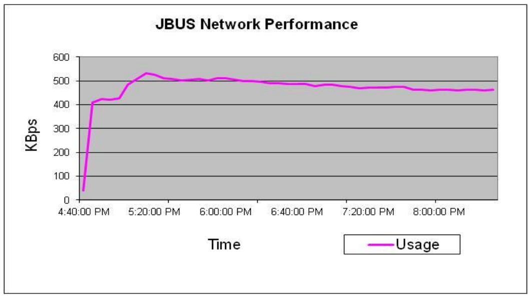

Figure 5. JBUS Network Usage

The configuration management of the virtual machines is essential. Ensure that you have a consistent naming convention for the VMs that account for the VM being a vanilla or template, the name and version of the simulation installed on the VM, etc. The vSphere Summary tab also provides a notes area which can be used to capture other configuration management notes. As the VMs are created, configured, and software is installed, snapshots and templates should be created to provide a fallback if an error occurred in the process or if an upgrade was needed to apply to all of the virtual machines. All VMs should be tested before making them a template. Depending on the virtual machine needed, any of the templates created in the three steps below can be cloned and migrated to a host to run on. The recipe for creating OneSAF virtual machines as recommended based on our findings is:

1. Create vanilla OS templates and take a snapshot. These templates have simply the OS installed with minimal configuration applied. All OS updates, patches, and service packs should be applied.
2. Clone that VM, configure the OS for running in your environment, and then create a template and a snapshot. This VM includes helper application installation, network, security, screensaver, desktop images, browser favorites, NFS configuration, etc. If your operation includes developers and testers, you may want to make additional clones, templates, and snapshots to have the tools needed on the VMs for those users.
3. Clone the previous VM(s) and install your simulation software, then create a template and a snapshot. In the case of OneSAF, a single image can be created to run any system composition (mode) of OneSAF. The onesaf.properties file should include any necessary option to run any system composition foreseen. The OS environment should be setup to run OneSAF, and any connections to NFS mounts should be mounted.

OneSAF v7.0 performance metrics were captured with vCOPS reporting and by using the vSphere “Report Performance” feature. vCOPS performance allows capturing very specific metrics at a granular level of detail. Graphs of the data are drawn or reports can be downloaded on any of the vCS nodes. The vSphere “Report Performance” feature reports host or VM CPU, disk, memory, and network performance based on parameters captured in the vCS database. Results are captured in Excel spreadsheets which include the numerical data and graphs as pictured below. Example findings for the particular scenario ran include Simcore escalating memory consumption which needs further investigation to determine if this is an issue or not (Figure 2). Disk performance spikes (Figure 3) when a scenario is loaded and when the simulation initializes; but during runtime activities, the disk usage is very low (with logging set to the default error level). The SORD traffic pushed over the WAN drops to almost nothing (Figure 4) when the scenario behaviors complete. This is contrasted to the InterOp which is still heart-beating DIS packets to JBUS (by DIS protocol design) even though the scenario activity is stagnant (Figure 5).

Another system configuration examined during the experimentation was regarding UDG deployment. The UDG can be deployed at the central site with browsers connecting to it to run the WCT at local or remote installations. The UDG connects to the SORD to transmit and receive OneSAF data and also hosts OneSAF terrain Tile Map Service (TMS) tiles. The UDG can also be deployed at forward operations to run in two different modes. The OSS Server and Client can be used to distribute SORD traffic to a remote UDG to host WCT data and terrain. This configuration reduces the amount of time that the WCT data needs to travel since web data is exchanged between the WCT and UDG on a LAN. However, the OSS distribution does not support exercise control functions, so the Battlemaster WCT (or MCT) cannot be run at a remote site. The other configuration supported has the remote WCT connecting to the UDG at the main site, but a terrain server UDG can be created at the remote site just to host terrain TMS tiles. This configuration does not require an OSS Server/Client connection for the remote UDG. A Battlemaster can be located at a remote site in exchange for the cost of WCT data being exchanged over the WAN.

As a side project, the client interaction experience was compared between the Wyse Z90SW thin client hardware and the Samsung NC220 zero client hardware. The WYSE thin clients provide traditional client-side resources (like web browsers and applications) which suited the OneSAF WCT thin application. Linux applications could be serviced using the VMware Web Client interface using a browser on the thin client as well. The thin clients have a feature that, by default, any user changed or saved settings are reverted when the thin client is rebooted. This creates a workstation environment that is easy for IT staff to maintain since they will not have to rebuild workstations between simulation exercises and experiments. While the primary focus of using the WYSE thin clients was to use the WCT or VMware vSphere Web Client interface, it is also possible to connect to VMs using RDP, PCoIP, and other protocols for vendors like Citrix or Microsoft.

The NC220 zero client connects the user directly to the VM desktop in a secure fashion using the VMware Horizon View connection server and transmits optimized Virtual Desktop Infrastructure (VDI) packets. We created shared desktop pools at the local and remote virtual sites and joined the thin and zero clients to the pools to operate the Windows VMs. Several tests were conducted to assess the experience, including using a dual extended monitor from the NC220. The NC220 manual has instructions to setup PC over IP (PCoIP) and Remote Desktop (RDP) protocols to transmit the virtual machine desktop to the zero client. However, the NC220 menu options only support PCoIP. The RDP features were not present in the menu options as pictured in the documentation. The RDP protocol is servable from Windows and Linux virtual machines; but due to the NC220 limitation and our testbed hardware configuration, PCoIP was used which required Windows VMs, VMware Horizon (with View), and additional configuration for Active Directory support. If the zero client supported RDP as advertised, the zero client would have supported OneSAF thick client Linux solutions.

Regarding the performance, the zero client connects to Windows VMs running the VMware Horizon View Agent and the performance responded well. Under normal desktop application use, the dual monitor feature of the NC220 worked well until two different videos were played on both monitors at full screen. In that case, the audio and video became choppy. Single monitor video playback on the NC220 was slightly choppy at full screen, but the audio was fine. Running the WCT on a Windows VM with the zero client performed well with no detectable issues other than consuming additional resources on the hypervisor to run the Windows VM compared to a UDG sending data to a WCT thin client. Overall, the WYSE thin clients provide a reliable platform to run the WCT or the VMware vSphere Web Client in a web browser, connecting to the UDG over LANs or WANs, so long as memory and hard disk space are considered into the thin client specs. Thin clients support a variety of VDI and other protocols whereas the zero clients evaluated were limited to one or two VDI protocols. The Samsung zero clients provide a secure way to connect to the hypervisors if loss prevention is an issue. However, our experimental lab limitation of PCoIP and Windows VMs did not suit this particular instantiation of OneSAF well for running on zero clients.

DIACAP IMPACTS ON CLOUD COMPUTING

BLCSE currently only has one location which routinely operates in a virtual machine environment, and their experience was extensively used when the ARCIC OneSAF cloud testbed was set up this summer. While detailed discussion of the DoD Information Assurance impacts on cloud simulations is are not part of this paper, DIACAP/RMF considerations for use of virtual machine environments and hypervisors require special attention.

For operation within stringent DoD IA security environments, virtual machines, the vCS, and ESXi hypervisors must be securely hardened in order to operate on secure DoD networks. For example, password policies declaring password length, character requirements, and reuse restriction parameters are applied to the host's BIOS, the ESXi hypervisor, the operating system present on each virtual machine, and the VMware Single Sign On (SSO) Service according to the Defense Information Systems Agency (DISA) Security Technical Implementation Guides (STIGs) (DISA for the DoD, 2014). The SSO allows administrators to manage users, roles, and permissions throughout most of the VMware login screens and simplifies password management for the users.

Another requirement present to ESXi systems running on DoD networks is the enabling of Lockdown Mode. Lockdown Mode prevents any vSphere API client, ESXi Shell, and SSH direct remote access to the hosts is blocked and enforces connectivity via the vCenter Server. This includes root access as only the special purpose *vpxuser* account is granted this permission. For example, the SSH daemon is not allowed to run on the ESXi host except during situations such as vendor support, initial configuration, troubleshooting, and break-fix situations. The ESXi SSH implementation is not a full-featured version and existing SSH sessions do not timeout when the SSH timeout value is changed. The STIGs also define policies for applying patches and updates to the hypervisors and virtual machines.

The vCS security posture requires routine patch application and STIG lockdowns. Software patches are routinely required on the vSphere systems for applying various patches, upgrading software, or removing applications. The operating systems, virtual machines, and virtual network can be manually patched by using vCS, VMware's Update Manager (vUM), and VMware's Update Manager Download Service/Server (vUMDS). Shavlik Protect (formerly VMware vCenter Protect) is also useful when patching guest operating systems and other VMware components (VMware vCenter Protect, 2014). Other vCS STIGs include disabling the web datastore browsing capability, preventing the local Windows administrator from accessing vCS as an administrator, and requiring a proxy gateway between the update managers and the Internet.

DISA VM STIGs require that ESXi controls access to host resources. Convenience features for copying and pasting, and drag and drop operations, between VM guest OSes and the remote console must be disabled. VMware Workstation and Fusion settings must be disabled, such as the Unity feature blending remote console access with the remote OS. Unauthorized floppy, IDE, USB, serial, and parallel devices must be disabled from accessing from the VMs. VM templates must be used to deploy VMs whenever possible.

THE OBJECTIVE NETWORK

One of the most challenging decisions when designing a network to run virtualized computers is to decide what type of physical storage solutions to use. While there are cost, power, performance, reliability, and other concerns surrounding storage hardware, the network topology of the solution will also impact the performance of the application. For example, a sophisticated Storage Area Network (SAN) may provide fast transactions for typical data lookups and storage, this solution will not perform well if 100 attached virtual machines decide to perform full-disk virus scanning activities at midnight. Note that VMware Network and Security prevents this type of situation, but it was used to illustrate the point of concurrent data access. In a situation where an application may perform like this, it may make more sense to install storage devices local to each server and not have a SAN solution. Here are some considerations when planning network storage requirements (VMware, 2011; VMware VROOM! Blog, 2014; ZDNet, 2012):

- Separate infrastructure traffic from virtual machine traffic.
- Enable jumbo frames for IP-based storage using iSCSI and NFS.
- Use NIC teaming for load balancing and passive failover capabilities.
- Create dedicated datastores to service database workloads.
- Use aligned VMFS partitions created with vCenter.
- ESXi, VMXNET protocols, and Para virtualized SCSI adapters are optimized for Oracle databases configured with Automatic Storage Management (ASM). Do not use Oracle failure groups due to CPU utilization.
- Wire speed is typically the limiting factor for I/O throughput.

- Guidelines to common physical storage configurations:
  - Use Solid State Devices (SSDs) when cost effective for performance reasons.
  - Local disks provide fast access to virtual machines on the same host but may prevent features like vMotion.
  - VMware Virtual SAN provides a hybrid approach to creating a SAN using local disks.
  - Network Attached Storage (NAS) provides a cost-effective approach for storing data that does not require fast access. It is also recommended to create a clone NAS for backups. For example, store VM templates on NAS hardware.
  - SANs provide high performance disk I/O at a premium price. The cost per megabyte (MB) of SANs may encourage engineers to consider pairing a SAN with a NAS to separate data performance and data size requirement differences.
  - Consolidate repetitive data across VMs into shared folders for VMs using the same datastore.

Approximately 79% of the OneSAF v7.0 storage footprint (does not include ARES models which would increase this value) can be moved to shared storage per datastore for static data such as terrain and documentation (Figure 6). Note that during exercise activities, terrain modifications are not written to the physical disk unless a checkpoint is taken and disk utilization is minimal (Figure 7). Although not investigated during this experiment, some installations run the OneSAF baseline in read-only mode with write-access only granted to the SOR, SYSTEM, and "ext" directories. This means that the entire OneSAF baseline could be installed on a VM as a shared baseline with other VMs creating links to the shared baseline thus reducing disk space consumption even more. Note that the "ext" directory is defined in the active configuration file to save data in a particular OneSAF extension.

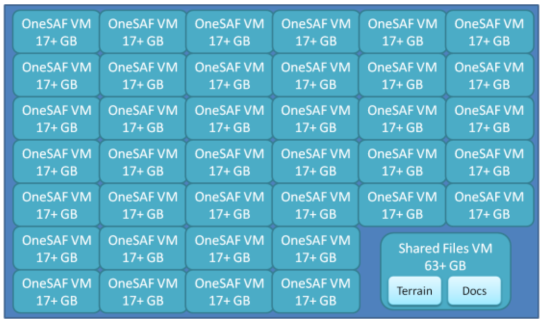
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB | OneSAF VM 17+ GB
Shared Files VM 63+ GB
Terrain
Docs

Figure 6. Datastore 709+ GB for 38 Copies of OneSAF

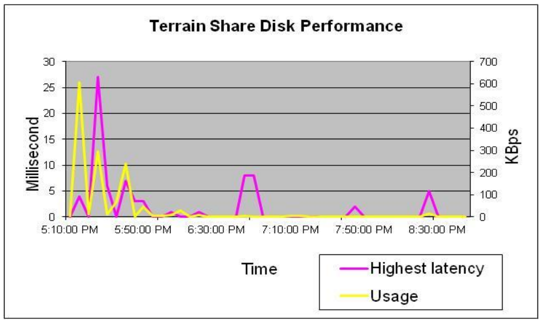

Figure 7. Terrain Share Disk Access

VMware vCloud Enterprise Suite provides the capabilities needed to run virtualized software in a cloud environment. Virtual machines are provisioned, monitored, cataloged, managed, backed-up, and replicated to provide robust simulation experimentation and training environments. IT activities are automated and simplified, and the robust networking and security features provide a secure foundation to conduct distributed simulation exercises. While many tech control procedures can be automated by using workflows, monitoring and controlling simulation tech control parameters are presently better achieved with simulation software. Overall, OneSAF is designed to be elastic and lends itself to operate in a virtualized environment. The recent addition of thin client tools in OneSAF, such as the WCT, and the implementation of the OSS server/client architecture provide cloud optimizations and flexibility to run the distributed simulation with a variety of choices on how best to distribute.

Cost savings can be realized using cloud technologies. The BLCSE strategy to in-source federates to a central site and only push out thin client or OSS traffic over the WAN will reduce bandwidth compared to the previous method of heart-beating data over an RTI. Datastore design can also provide significant cost to performance benefits when weighing differences between local datastores, datastores linked with VMware Virtual SAN, constructing a SAN or NAS, or using a hybrid approach of using local datastores, SAN, and NAS devices. Migrating from thick to thin or zero client hardware helps to reduce workstation costs and reduce power consumption. Distributing software via a cloud environment using VM images rather than printing DVDs for distribution not only saves cost for DVD creation and shipping, but also reduces calls for helpdesk support since the VMs are preconfigured correctly from the factory. Also, the vCloud capabilities to in-source control of hardware located at remote sites can reduce the personnel needed to operate the hardware and software located at the remote sites.

ACKNOWLEDGEMENTS

The authors thank Bryan Cole and John Stevens of CESI for open access to the CESI IRAD simulation cloud project, and gratefully acknowledge the extraordinary assistance of OneSAF Team Orlando, particularly John Logsdon and Amit Kapadia of Product Manager (PdM) Warrior Interoperable Simulation Environment (WISE), Dave Nash and Wayne Minton of LEIDOS, and Derrick Franceschini of STACKFRAME, who made possible the distributed network testing of Simulation as a Service. We also thank Anthony Pasquesi, Margaret Amori, and Jon Pawlowski from VMware for licenses and reach back technical support to the CESI testbed.

REFERENCES

CIO.GOV (2014). *Data Center Creation*. Retrieved April 16, 2014, from https://cio.gov/deliver/data-center-consolidation

Darken, Rudulph, PhD (2014). Early Synthetic Prototyping: Exploring New Designs and Concepts within Games. *Proceedings of the 2014 IITSEC, Paper 14133*.

DISA for the DoD (2014). VMware ESXi Version 5, Virtual Machine Security Technical Implementation Guide Overview. *DISA, Version 1, Release 2*.

Fang Liu, Jin Tong, Jian Mao, Robert Bohn, John Messina, Lee Badger and Dawn Leaf (2011). NIST Cloud Computing Reference Architecture. *NIST Special Publication, 500-292 September 2011*, p. 10.

Rieger, Lawrence (2014). Simulations in the Cloud – A Manager's Challenge. *Proceedings of the 2014 IITSEC, Paper 14104*, p. 5.

Takai, Teresa M. (2012), *Cloud Computing Strategy*, Washington D.C.: DoD CIO.

Takai, Teresa M. (2013), *Department of Defense Joint Information Environment: Continental United States Core Data Centers and Application and System Migration*, Washington D.C.: DoD CIO.

VMware (2011). *Oracle Databases on VMware Best Practices Guide*. Retrieved May 16, 2014, from http://www.vmware.com/files/pdf/partners/oracle/Oracle_Databases_on_VMware_-_Best_Practices_Guide.pdf

VMware vCenter Protect (2014). *vCenter Protect*. Retrieved May 16, 2014, from http://www.vmware.com/products/shavlik

VMware VROOM! Blog (2014). *Ten Reasons Why Oracle Databases Run Best on VMware*. Retrieved May 16, 2014, from http://blogs.vmware.com/performance/2007/11/ten-reasons-why.html

ZDNet (2012). *Another battle royale: SAN vs. local storage for VDI*. Retrieved May 16, 2014, from http://www.zdnet.com/blog/virtualization/another-battle-royale-san-vs-local-storage-for-vdi/4934