Effective Deployment of LVC-TE on Wide Area Networks

Luis E. Velazquez
MARCORSYSCOM
Quantico, VA
luis.velazquez@usmc.mil

Lloyd Wihl, Ha Duong, Jeff Weaver, Jeff Hoyle
Scalable Network Technologies, Inc.
Culver City, CA
{lwihl,hduong,jweaver,jhoyle}@scalable-networks.com

ABSTRACT

The Marine Corps' Live, Virtual and Constructive Training Environment (LVC-TE) connects training systems at geographically separate bases to enable collective and battle staff training. The long-haul circuits that provide the connections are not dedicated to training exercises but are shared and simultaneously carry other network traffic for the Marine Corps. Excess latency and jitter injected into training exercises from these circuits can invalidate results and bias the results of the exercise for one side.

A major existing deterrent to the planning of large scale exercises is the inability to accurately estimate the load that will be placed by a local, regional, or country-wide training exercise on the underlying communication networks. This significantly prolongs the planning and approval processes.

In this paper, we present a new simulation-based framework to predict the impact of connecting training systems across different types of long-haul network circuits, validate key performance parameters, and streamline the planning of distributed training exercises. The framework profiles different training simulations/simulators and correlates captured traffic to scenario events. Traffic models can be scaled to represent higher numbers of entities, simulators, and time-varying, overlapping scenario events. Authoritative Marine Corps descriptions of the network on which the training exercise is run, in the form of Visio or similar formats, are converted into an executable, dynamic network simulation model. The traffic models are overlaid on the simulated network to predict how traffic generated during a training exercise, competing with non-training traffic, will be delivered, using metrics such as throughput, latency, packet loss and jitter. The framework enables reconfigurable, on-demand tradeoff analysis to derive optimal solutions.

Utilizing this framework, the authors present findings for the network performance impact of running a Virtual Battlespace 3 (VBS3) training exercise on the 29 Palms network.

ABOUT THE AUTHORS

Luis E. Velazquez is a federal employee (NH-IV/GS15) for the Marine Corps serving as the Branch Head for Future Capabilities and Innovation, Marine Corps Systems Command (MCSC), Systems Engineering, and Acquisition Logistics (SEAL). Over 30 years of professional experience between the Navy, Marine Corps, Industry, and Federal services. Subject Matter Expert (SME) with extensive engineering experience in employing LAN/WAN, software engineering, and systems integration performing advance services to include Live, Virtual, and Constructive (LVC) simulation systems integration and interoperability with tactical C4i devices.

Lloyd Wihl is Director of Application Engineering at Scalable Network Technologies. He has 30+ years' experience in the Modeling, Simulation and Training industry, developing system architectures and leading multi-million dollar projects in the areas of synthetic degraded digitized battlefields, distributed mission operations, network-centric systems, air combat, marine systems, space systems, visual systems, and flight simulation. He is a recipient of the NASA achievement award, has published and presented several papers on synthetic environments, and guided development of Scalable's live-virtual-constructive cyber training system that integrates cyber and kinetic warfare.

Dr. Ha Duong is Principal Engineer at Scalable Network Technologies where he has worked on modeling JTRS waveform for the Communication Effects Server (CES) project in the context of the Future Combat System (FCS) and Brigade Team Modernization (BCTM) programs. Over the past several years, Dr. Duong has focused on modeling vulnerabilities and cyber attacks in the StealthNet project at SCALABLE, and leveraging those models into the JTRS Network Emulator (JNE) product and the Cyber Test Analysis and Simulation Environment (CyberTASE) project. Dr. Duong has also led the Human-centric Training and Assessment System for Cyber Situational Awareness project. His current research interests include LVC-based cyber-attack representation, modeling and simulation techniques to represent complex operations in simulation environments, and analysis of cyber effects on DoD tactical networks.

Dr. Jeffrey Weaver is Vice-President of Engineering at SCALABLE. He obtained his Ph.D. degree in Electrical Engineering as an NSERC-PGS Scholar from Western University in Ontario, Canada. Dr. Weaver has held key technical and executive engineering roles during his career and has over twenty years of product development experience in hardware and software systems. His research interests include digital communication and propagation modeling using switched stochastic differential equations, signal processing and hybrid analytical-numerical modeling techniques. Dr. Weaver has seven patents in the areas of IP routing, VLAN, QoS, and high-performance hardware design.

Captain Jeff Hoyle (US Navy, Retired) leads Scalable Network Technologies communications and networking developments for the Department of Defense and Intelligence communities. Prior to joining Scalable, he served as Director of Advanced Maritime Programs for Northrop Grumman Aerospace Systems, a leading provider of military autonomous systems, and Director of Technology and Navy Programs for AtHoc, Inc., a leading provider of crisis communications capabilities to multiple Federal agencies. While on active duty, Captain Hoyle supervised all aspects of US Navy operations on five submarines and one aircraft carrier, including command of USS MAINE (SSBN 741) and ten deployments to forward operating regions on missions vital to national security. As a Defense Acquisition Program Manager, he led development of submarine exterior communications systems and joint tactical networking capabilities for the Army, Navy, Air Force and Marine Corps.

Effective Deployment of LVC-TE on Wide Area Networks

Luis E. Velazquez
MARCORSYSCOM
Quantico, VA
luis.velazquez@usmc.mil

Lloyd Wihl, Ha Duong, Jeff Weaver, Jeff Hoyle
Scalable Network Technologies, Inc.
Culver City, CA
{lwihl,hduong,jweaver,jhoyle}@scalable-networks.com

INTRODUCTION

As stated in the Marine Requirements Oversight Council (MROC) (2010). The United States Marine Corps (USMC) Live, Virtual, and Constructive Training Environment (LVC-TE) combines any of the three training domains (live, virtual, and constructive) to create a common battlefield or environment, by which units can seamlessly interact across live, virtual and constructive domains as though they are physically located together in the same battlespace. The LVC-TE will provide the means to conduct realistic, collaborative training and exercise of warfighting functions across the full range of military operations (ROMO). To enable the LVC-TE there are four major capability gaps that need to be resolved:

1. ***Integrating Architecture*** – provides the ability to allow for the easy, rapid and seamless integration of the live, virtual and constructive domain mission partners.
2. ***Integrated Dynamic Virtual and Constructive Synthetic Battlespace Representations*** – provides the ability to replicate entities across the full ROMO when executing fully integrated LVC operations.
3. ***Integration and Stimulation of Operational Systems*** – provides the ability for warfighters to train and execute mission rehearsal events utilizing their operational systems.
4. ***User Services*** – provide the ability to easily and rapidly conduct collaborative planning, preparation, execution, and assessment for LVC training, exercise, and mission rehearsal events.

Each of the above capability gaps rely on stable standardized network framework that facilitates data exchange across multiple security domains, geographic locations, and with information assurance.

Per the Training and Education Modeling and Simulation Master Plan of 2010, there will be a reliance of interconnecting simulations across distributed environments. The limitations of the current approach to predict an alternative solution for an LVC-TE enabling network requires access to existing network infrastructure to conduct ongoing experimentation for potential future "to-be" network analysis thus requiring time, resources, and valuable analytical rigor to evaluate potential tradeoffs. These tradeoff analyses have to consider every aspect of the LVC-TE "to-be" network design to include the impact of latency, scheduling of critical software upgrades, and accessibility to data repositories essential in the synchronization of the LVC-TE training environment. All of these analytical tradeoffs have to be conducted in a very limited resource-constrained and austere environment.

Live Virtual Constructive (LVC) systems will include training ranges connected to simulators, which will connect to each other and to constructive simulations. Aircraft, ships or vehicles and live Command and Control can participate in the exercise. Post-mission data will be captured, and analysis could reside in a data center. These connections result in a significant amount of data traversing the networks.

The number of potential connections among LVC components, both within a site and among geographically-distributed sites, along with the network traffic loads which are scenario-dependent, make current analyses labor intensive and time consuming. These are recurring engineering costs, as new analyses must be undertaken prior to each exercise. Our solution, which makes use of modeling and simulation of the network, will reduce these recurring costs and lead times and provide an easier way to perform more training reps to warfighters.

Network Performance and Training

All networks face common challenges like bandwidth limitations, bottlenecks, security attacks, session management, scalability, traffic congestion, and quality of service trade-offs. While network-induced delays may be a minor annoyance when reading e-mail or accessing web page, they can spell doom for a networked training exercise that

links trainees interacting in real-time in a common synthetic battlefield. Latency, jitter, and packet drops can all negatively affect an exercise to the point that it becomes an unfair fight, rendering the outcome and the trainee scoring invalid. High "gain" interactions among participants such as close formation flying, close combat maneuvering, ship deck landing, air-to-air refueling, and integration of maneuver with artillery and close air support all require minimal latency in the transfer of entity states, failing which instabilities, overcorrection and collisions among entities could arise. The variation in latency, or jitter, can be more of a problem than the latency itself. If jitter exceeds one iteration interval, and is left unchecked, then random position stepping can occur. Dropping packets instead of delivering them to their destination can have a significant effect on a fair fight. For example, a trainee may not be aware that he is being fired at, and so would not take cover, thus increasing his chances of being killed. This would be due to network performance rather than a trainee mistake, and thus be unfair for trainee scoring.

Current Approach

Current exercise planning approaches manually estimate bandwidth requirements for data transfer among simulators and seek to guarantee this bandwidth availability on the shared network. However, data transfers can vary widely as an exercise progresses. What effect will peak data transfers have on other network traffic? Will other network applications slow to a crawl? Will a multi-day training exercise with hundreds of participants be found to be invalid midstream or after the fact, due to peak traffic that exceeded allocated network resources? There are design decisions and tradeoffs to be made (some training traffic is more latency-tolerant than others), and accuracy is needed in the analysis. Attaining this accuracy with manual approaches is labor intensive and time-consuming, and must be redone every time the configuration of LVC components in the exercise changes. This can impact the start of a training exercise by months.

SOLUTION

Software Virtual Network

Software Virtual Networks (SVNs) make it possible to represent the communication network infrastructure at sufficiently high levels of fidelity to accurately determine the success or failure, and timing, of every packet delivery. The SVN provides an exact, high quality, emulation of network behavior that is indistinguishable from the real system.

Scalable Network Technologies' EXata is a commercial off-the-shelf (COTS) tool that uses an SVN to emulate the entire network, the various protocol layers, routers, switches, wireless access points, encryptors, simulators, and other devices. It can interoperate with real equipment to provide hardware-in-the-loop capabilities, and can also be connected to real applications, which run on the SVN just as they would run on real networks. EXata can model a variety of hybrid networks with thousands of emulated nodes exchanging different types of traffic.

A benefit of network emulation is detailed instrumentation. As network emulations execute, users can watch traffic flow through the network and view dynamic graphs of critical performance metrics. A statistical graphing tool displays hundreds of metrics collected during simulation of a network scenario. Multi-experiment comparison reports are also available to enable optimization of configurations. EXata also provides a high-performance interface that allows time-series and statistical data to be stored in a database during the simulation. The database can be configured to record statistics at different levels of granularity: from summary statistics at the system level to detailed statistics at the event level.

Exercise planners can use the SVN as a cost-effective method in which LVC connection decisions can be easily changed through a drag-and-drop user interface, and their impact evaluated, to predict how training traffic will perform with competing traffic when the planned exercise is deployed on the target networks.

Authoritative Marine Corps Descriptions of the Network

Per the Marine Corps Order (MCO) 5230.20, and MCSCO 5510.2 the Marine Corps invested resources in documenting all its networks in the form of Visio diagrams within the Marine Air-Ground Task Force (MAGTF) Collaborative Architecture Environment (MCAE). The MCAE serves as the Authoritative Source for Solution Architecture in the United States Marine Corps and promotes the collection, distribution, and reuse of authoritative architectures and primitives for architecture development. Specifically, the MCAE is an Authoritative Source for

Architectures in the Marine Corps. Marine Corps Systems Command is the technical resource responsible for development and sustainment of toolset. Its responsibilities include the system component of Architecture Development stored within MCAE. Marine Corps Combat Development and Integration (CD&I) is an operational stakeholder leveraging MCAE for management and distribution of authoritative operational data. Program Offices leverage MCAE daily to satisfy MCAE Architecture Compliance requirements. All accepted MCEN Baseline Architecture products are available on the MCAE Web Portal. The MCEN Baseline Architectures are developed with ubiquitous desktop tools provide stakeholders supports reusability and analysis without special tools and training. Conventions used are familiar to users across the spectrum of stakeholders. Customized stencil set provides users with the ability to create their own views from the authoritative architecture.

Network Modeling

To make efficient use of the architectures in the MCAE, it was necessary to extend EXata to import Visio™ diagrams to directly create executable models of the networks within EXata. This utility handles L2 and L3 switches and their VLAN configurations, hubs, gateways, bridges, routers, servers, firewalls, and many more Visio objects. User-specific information from the properties section of Visio shapes are parsed and used in creating the network model. Using this utility, creation of an executable model of the 29 Palms network from the MCAE architecture could be directly achieved.

In addition to Visio™ files, router configurations can be provided by the customer that simplify development of scenarios by providing auxiliary information that the Visio™ files do not provide. A second utility was developed directly import router configuration files from various manufacturers and use these to automatically configure the EXata router models.

These two utilities minimize the manual effort needed to create network models and have the added benefit that updates to the MCAE can be directly incorporated into the EXata models.

Once built, additional detail can be added to the network model. For the EXata model of the 29 Palms network, we further enhanced the model by setting OSPFv2 as the routing protocol for the routers, setting default routes for hosts, configuring multicast membership, and adding a remote server at the other end of a SIPRNet link.

New devices were added to the EXata palette, namely USMC training simulators/simulations including Virtual Battlespace 3 (VBS3), Supporting Arms Virtual Trainer (SAVT), and Combat Convoy Simulator (CCS). This enables instances of these simulators to be easily dragged and dropped onto the network laydown canvas and connected to the network simply by drawing a line from the simulator to the desired connection point (Figure 1).

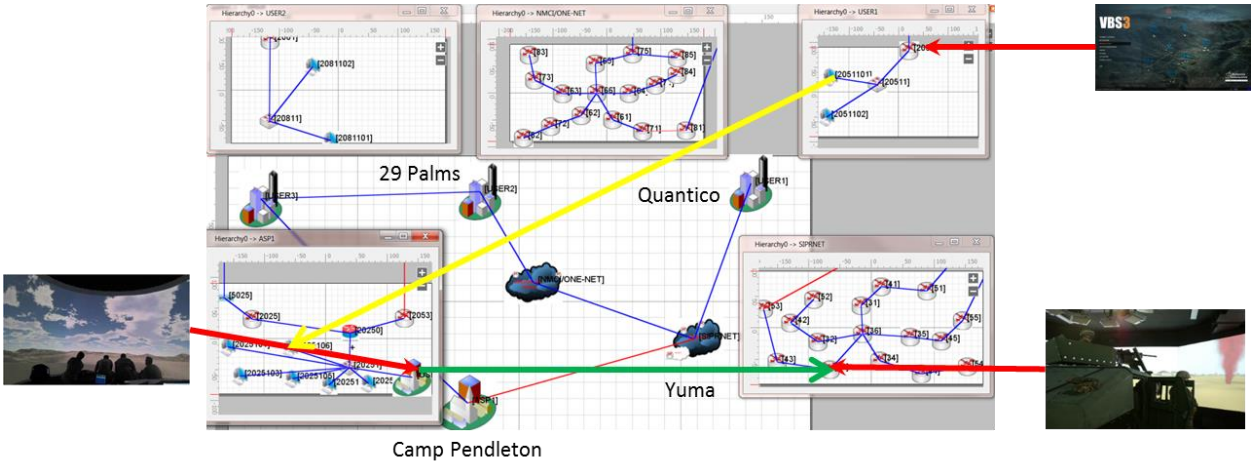

Figure 1: Example of Connecting Virtual Simulators to the Network Laydown

Instrumentation of Simulations

To achieve high fidelity in simulation of network performance during a training exercise, it is important to accurately model traffic loads. Instrumentation in the training simulation captures the network traffic generated during various scenarios and during specific events in the scenarios. EXata imports this captured traffic and infers a baseline application profile, which can be scaled up to model traffic loads from adding simulations and/or increasing number of entities, and/or temporally overlapping multiple traffic-generating events.

EXata Extractor is a tool that creates equivalent EXata models of the battlefield communications networks used in constructive simulations. It works by joining a DIS or HLA federation and listening for entities and radio transmitters. Using this information, it automatically creates a corresponding network configuration in EXata, so that EXata can act as a communication effects server for the federation. The operational capability of EXata Extractor was expanded to allow it to listen to all simulation traffic in addition to only entities and radio communications.

Deployable Virtual Training Environment (DVTE) is a suite of simulation applications which supports the training of Marines from the individual up to staff level. These simulations enable units to execute complex missions in advance of live exercises. Turn-key scenarios focus on training requirements such as Call for Fire, Joint Terminal Air Control, IED Defeat, Reporting Procedures, and Decision-Making Skills. VBS3 is a component of DVTE that trains Marines on everything from command and control to convoy standard operating procedures. The DVTE test network consisted of five DVTE computers and a laptop running the EXata Extractor tool connected to a mirrored switch port connecting the simulations. EXata Extractor captured all the traffic generated from Marine Corps-supplied VBS3 training scenarios and used it to create the application baseline.

Traffic Modeling

A key part of this project was to incorporate application analysis tools to infer traffic models from packet capture traffic, both cumulatively and by traffic type. Matching captured traffic peaks to the exercise event timeline showed correlation between traffic loads and specific scenario events. These tools created specific traffic models for exercise events such as Combat Net Radio (CNR) calls, firing, and bomb explosions. The parametric baseline traffic models were fit to existing data and could be scaled based on the number of characters, and the captured unicast, multicast and broadcast traffic flows among the DVTE computers.

In EXata simulation scenarios, traffic is modeled by application flows. Individual traffic characteristics for a scenario are stored in a JSON file. The JSON file describes the sequence of events in a training scenario. The number of VBS3 segments (or networks), number of entities in each segment and the events that occur in the scenario are configurable. This enables scaling of the traffic model for number of entities, additional simulators/simulations, LAN vs WAN traffic, and event timing. A modular set of software processes were used to transform the JSON file and generate the EXata application configuration files.

Non-simulation (competing) background traffic on the network can be replicated as synthetic traffic applications (e.g. CBR or FTP sessions) or be data-driven replayed as previously-captured PCAP files and modeled in the EXata simulation, e.g. to represent peaks loads at the start of day when Marines log in and check e-mail. Using our application characterization capabilities represents a third alternative, inferring captured traffic into application profiles.

Approach to Analysis

Each simulation run generates a statistics file, which reports summary statistics, and a statistics database, which records time-stamped statistics at various levels of detail. We used these to (a) analyze traffic loads as they relate to the selected training scenarios and determine the most critical scenario segments; (b) identify bottlenecks in the network performance and how to resolve them; (c) provide analysis to validate Key System Attributes (KSA's) and Key Performance Parameters (KPP's) for networks, training simulations, and traffic of LVC-TE.

FINDINGS

Analysis of Captured Traffic

The simulation experiment was built around a typical DVTE "Search and Destroy" training exercise. This exercise was executed several times with varying numbers of characters. All scenarios had 5 live Blue Force players (hosted on the 5 DVTE computers) and a varying number of additional Artificial Intelligence (AI) Blue Force characters and Opposing Force characters, ranging from 15 to 150. Recordings of all traffic among the computers were used to characterize the simulation traffic loads based on sources and destinations (example in **Table 1**).

The baseline traffic, when there are no significant events such as CNR calls, bombing or firing, is primarily composed of multicast traffic. The VBS3 server sent 60-90% of all multicast traffic and the clients sent the rest. As expected, the baseline traffic levels increased with the number of characters in the scenario. However, certain events in the scenario created bursts of unicast traffic between client and server, and the scaling of these bursts with the number of characters is clearly nonlinear, as shown by the peaks in the red and blue traces in **Figure 2**. Upon investigation, the traffic generation is governed by hidden (non-network) variables. For example, traffic is generated by a simulated bomb explosion, which can be correlated to the size of the bomb and the number of characters within the blast radius.

Table 1: Traffic by Destination

Destination	Percentage of All Traffic	Number of Packets	Protocol
Multicast	71.7%	306668	LAPD
Broadcast	9.8%	42094	UDP
VBS3 Server + Player 1	7.8%	33466	UDP
Player 2 client	2.5%	11181	UDP
Player 3 client	2.5%	10714	UDP
Player 4 client	2.4%	10061	UDP
Player 5 client	2.4%	10436	UDP

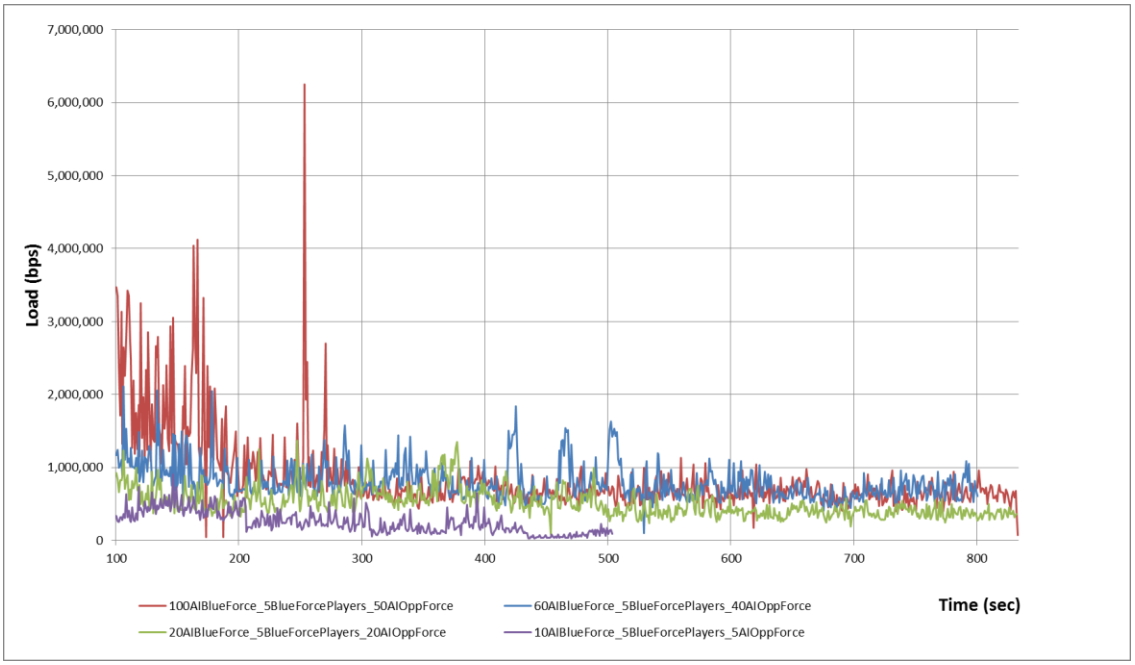

Figure 2: Traffic Captured from VBS3 Scenarios

Test Cases

Three test cases were used to compare network performance and packet delivery for different scales of training scenarios. Note the current instrumentation setup used for this paper did not include inter-site simulation traffic. Instead inter-site traffic was a simulation parameter provided by the user and expressed as a percentage of local traffic. In the future inter-site traffic levels will be substituted into the model when measured.

- Individual: 1 trainee, 50 AI-controlled characters, no distributed traffic
- Small Unit: 5 trainees, 100 AI-controlled characters, distributed traffic = 75% of local traffic
- Collective: 5 trainees, 150 AI-controlled characters, distributed traffic = 100 % of local traffic

For all test cases, the simulated traffic profile includes traffic loads per character and per trainee measured during instrumentation of DVTE and scaled to the size of the test case:

- Scenario initialization
- Characters moving and searching (baseline traffic)
- Voice communications over CNR (% of time active)
- Bomb drops (specific events)
- Firing (specific events)

There are multiple WAN gateways available at 29 Palms. For this study a mix of distributed training exercises were used.

Simulation Runs

Using these three traffic cases over the emulated 29 Palms network, the first step in the analysis was to examine key statistics about packet delivery to the receiving DVTE machines. For the individual case, the delays on the local area network did not exceed 2 ms, with the network fully capable of handling the simulation traffic. For both the small unit and collective cases, the inter-site simulation traffic due to moving and searching, and voice communications over CNR was adequately handled over the WAN; however, certain scenario events triggered a surprisingly large latency at a remote server of 20-30 seconds. These higher than expected latencies triggered further investigation using the EXata tools.

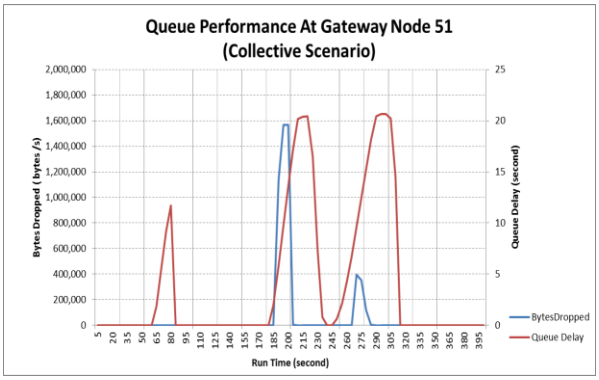

Figure 3: Bytes Dropped and Queue Delay – Collective

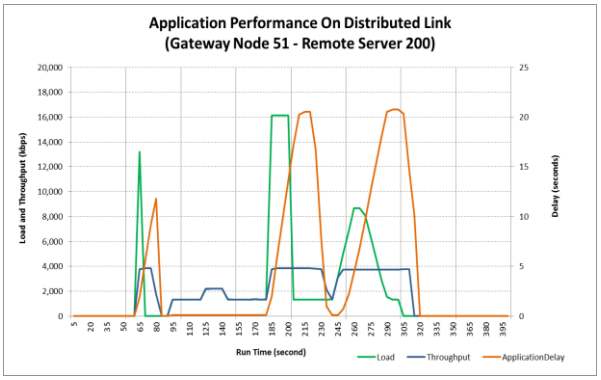

Figure 4: Load, Throughput and Delay - Collective

Correlating the timing of queue delays (red lines in **Figure 3**) to scenario events that generated simulation traffic, the initialization burst (start=60 sec), bombing (start=180 sec), and firing (start=240 sec) events cause significant queue delays. Note that bytes were dropped (blue lines in Figure 3) during the bombing and firing events, but not the initialization. To delve further into the reasons for this, the traffic load (green trace) is imposed as offered load on the distributed link (**Figure 4**) and achieves the indicated carried load (blue trace). From 90-180 seconds, the carried load matches the offered load exactly, and the green and blue lines are superimposed.

During the load spikes these lines diverse. The network queue is able to deal with the first load spike without packets dropping from the queues, as this load spike occurs for a short period of time. During this period, the packets are

queued up and result in a delay of 12 seconds. During the first load spike the throughput increases and approaches the link capacity of 4 Mbps.

There is similar behavior in the second load spike except that this spike is higher and lasted longer. The area under the green curve represents the total packets sent during any time interval (data rate multiplied by time). Clearly, the area under the second load is much greater than the area under the first. The greater number of packets waiting to be delivered over the low bandwidth WAN causes the queue size to increase more than under the initialization load. As a result, packets already in queue experience longer delay (20 seconds, compared to 12 seconds). As the queue is pushed to its capacity, it starts dropping packets. This packet drop does not occur at initialization due to the lower load. The throughput is pushed to the link capacity (4 Mbps) during the bombing and remains at this level much longer. This is due to the queue holding a large number of packets.

The firing load (third green peak in **Figure 4**) is less intensive than the bombing load (second green peak) and sends less total traffic (33.96 MB vs 40.32 MB). However, significantly more of the firing traffic is delivered compared to the bombing traffic both in absolute (29.66 MB vs 18.96 MB) and percentage of demand (87% vs 47%). This is because the lower demand on the link allowed more packets to get through before the queues overflowed. The link is held at the maximum throughput longer than during the bombing.

It is evident from the preceding analysis that Small Unit and Collective exercises that generate DVTE traffic loads comparable to those we measured, which connect to a remote site over a 4 Mbps WAN connection, would result in delays of 12-20 seconds and packet drops during bombing and firing. Both of these effects would be unacceptable during an actual training exercise. Dropping packets instead of delivering them to their destination can have a significant effect on a fair fight and bias the simulation exercise against one or more characters.

KPPs and KSAs

Net-Ready Key Performance Parameters (NR-KPP) for the exchange of information during a large-scale training exercise might specify threshold and objective latency between simulations of, say, 200 ms and 80 ms respectively for high gain interactions such as firing at nearby moving opposing players. The LVC-TE model can be used to validate such KPP's and Key System Attributes (KSA's). The benefit of this approach is that proposed changes to the network can be quickly made in the model and the effect of these network changes on the same training exercise can be quickly assessed.

As an example, we return to the collective exercise described previously. Our analysis showed that delays and packet drops originated at the gateway to the WAN. This was due to queuing resulting from the 4 Mbps available bandwidth. What effect would there be if we would increase the available bandwidth?

Referring again to **Figure 4**, the peak load generated over the WAN by the bombing in our Collective DVTE exercise was 16 Mbps. The end-to-end delay was reduced significantly, from 25 seconds to 880 ms. We observed that the packets dropped were reduced from 140 to 1. Note that these improved results might still fall short of threshold values for a KPP, and if so, further analysis could determine the next chokepoint, but this was beyond the scope of this study.

CONCLUSION

While the scenarios used for the proof of concept are simple enough to be analyzed by hand, they illustrate that performance of the 29 Palms network and its wide area connection could potentially lead to an unfair fight during a training exercise. Actual LVC-TE exercises could involve hundreds of participants at various bases with a plethora of potential connection points. Traffic loads among simulations publishing and subscribing to data will be scenario-dependent and dynamic. Further, simulations will send data across the MCEN or alternative long-haul networks and compete with non-training network traffic. As has been shown in the past, analyzing these large-scale training events manually becomes unwieldy and very time-consuming.

Our framework, consisting of simulator traffic recording, analysis and scaling, network topology importing, network emulation with modifiable connections, and detailed statistical reports provides significant improvements over the current manual methods. The analysis provided is much more than bandwidth: it predicts specific delays between sender and each recipient (some may be tolerable, others not, depending on relative entity positions and the “gain” of

the interaction), jitter, dropped packets, effects on non-training network traffic, and provides assistance to locate network chokepoints and resolve them. It offers further benefits in the ability to emulate wireless connections between live devices and training ranges including mobility, interference, terrain, and other factors that would not affect wired networks. An additional benefit is related to cybersecurity. Weapons performance, Tactics, Techniques and Procedures (TTPs), and Concepts of Operations (CONOPS) must be protected as they traverse the LVC networks. The network emulation's ability to respond exactly like a live network can play a key role in testing security and helping to defend the LVC environment against evolving cyber threats.

Our solution uses emulation to drastically reduce the effort and time needed to analyze and approve network configurations for training exercises and to conduct tradespace analysis that impacts long term acquisitions, and can be directly applied to exercises over other networks such as Distributed Mission Operation Network (DMON) or Navy Continuous Training Environment (NCTE). The end result is that warfighters can now get quicker access to LVC training and more reps, resulting in more preparedness for future conflicts.

ACKNOWLEDGEMENTS

This work was supported by SPAWAR Systems Center Pacific under contract No. N66001-17-D-5201-0002. Training and Education Command, Training and Education Capabilities Development.

REFERENCES

Marine Requirements Oversight Council (MROC). (2010). Live, Virtual, and Constructive—Training Environment (LVC-TE) Initial Capabilities Document (ICD) (Decision Memorandum 48–2010). Washington, DC: United States Marine Corps.

Training and Education Command, MAGTF Training Simulations Division (MTSD). (2010). Training and Education Modeling and Simulation Master Plan 2010. Quantico, VA: MAGTF Training Simulations Division.

Training and Education Command (MTSD). (2014). Training and Education Modeling and Simulation Master Plan 2014. Quantico, VA: MAGTF Training Simulations Division.

Velazquez, L.E. (2014, May 5). II MEF Battle Simulation Center (BSC) Integration of Aviation Distributed Virtual Training Environment (ADVTE) with the Virtual Battle Space (VBS) in Support of Exercise "Emerald Warrior." Quantico, VA: Marine Corps Systems Command.

Marine Corps Order (MCO) 5230.20

Marine Corps Systems Command Order (MCSCO) 5510.2