SAND2010-2964C

PERFORMING CYBER SECURITY ANALYSIS USING A LIVE, VIRTUAL, AND CONSTRUCTIVE (LVC) TESTBED

Brian Van Leeuwen, Vincent Urias, John Eldridge, Charles Villamarin, Ron Olsberg
Sandia National Laboratories**
Albuquerque, USA
{bpvanle, veuria, jmeldri, chvilla, rrolsbe}@sandia.gov

Abstract -- **Cyber security analysis tools are necessary to evaluate the security, reliability, and resilience of networked information systems against cyber attack. It is common practice in modern cyber security analysis to separately utilize real systems computers, routers, switches, firewalls, computer emulations (e.g., virtual machines) and simulation models to analyze the interplay between cyber threats and safeguards. In contrast, Sandia National Laboratories has developed new methods to combine these evaluation platforms into a cyber Live, Virtual, and Constructive (LVC) testbed. The combination of real, emulated, and simulated components enables the analysis of security features and components of a networked information system.**

When performing cyber security analysis on a target system, it is critical to represent realistically the subject security components in high fidelity. In some experiments, the security component may be the actual hardware and software with all the surrounding components represented in simulation or with surrogate devices. Sandia National Laboratories has developed a cyber LVC testbed that combines modeling and simulation capabilities with virtual machines and real devices to represent, in varying fidelity, secure networked information system architectures and devices. Using this capability, secure networked information system architectures can be represented in our testbed on a single computing platform. This provides an "experiment-in-a-box" capability. The result is rapidly produced, large scale, relatively low-cost, multi-fidelity representations of networked information systems. These representations enable analysts to quickly investigate cyber threats and test protection approaches and configurations.

I. INTRODUCTION

Securing our nation's critical information systems against cyber attack is an important and difficult task. Many of our nation's critical information systems are used by the DoD to conduct their operations and these information systems are often targeted for attack. The latest and most advanced security methods are used to protect these information systems from cyber attack. Also necessary are analysis methods and tools to measure the effectiveness of selected security approaches. Thus, tools are necessary for the DoD to analyze their information systems' security, reliability, and resilience against cyber attack.

The most widely-used security analysis technique used by computer information system (CIS) specialists is based on evaluation of hardware destined for placement in the information system. Here, specialists build and configure CISs from physical equipment that they have purchased. The CIS is instrumented using network diagnostic equipment and connecting computers to the networks to generate appropriate traffic. While very accurate, this approach is problematic for two reasons. First, the equipment can be very expensive to acquire, configure, and maintain. Second, instrumentation and experimentation can be very challenging. It is difficult to correlate traffic events that move across the CIS and, as a result, difficult to roll up studies and generate system-level information.

CIS specialists also use simulation extensively. There are numerous simulation tools in existence for studying CIS issues. Today's simulation tools have extensive capabilities and high accuracy. The simulation tools have extensive probing capabilities that make it possible to correlate events and generate system-level information. Simulation tools have been used primarily to analyze data capacity performance and help CIS users accomplish expansion studies. Currently, few simulation tools have the necessary network device fidelity that would enable specialists to effectively evaluate various security implementations and analyze threats and vulnerabilities at scale. Most simulation tools accurately represent the data link and network transport layers, but do not sufficiently model application programs.

To overcome the problems with security analysis using either an exclusive hardware CIS testbed or a simulation of a CIS, Sandia National Labs has developed a cyber security analysis capability using physical hardware, emulated machines, and simulation. This hybrid testbed approach is termed a Live, Virtual, and Constructive (LVC) approach to CIS analysis and evaluation. Key aspects of our LVC approach to cyber security analysis has been published [1][2][3].

Throughout this report the terms *simulated* nodes, *emulated* nodes, and *real* nodes are used. In this report, *simulated* refers to the nodes represented through simulation tools; in our case OPNET Modeler [4]. Simulated nodes generally use unique and abstracted implementations of the protocols and software running on virtualized hardware. *Emulated* nodes use real software, for instance an actual Windows OS, but run on emulated or virtualized machines. *Real* nodes are the real software running on real hardware.

II. LIVE, VIRTUAL, CONSTRUCTIVE (LVC) TESTBED DESCRIPTION

The LVC testbed Sandia National Labs used to perform cyber security analysis experiments is comprised of real nodes such as a number of Cisco routers and Cisco PIX firewalls, emulated nodes using the ESX Virtual Machine (VM) capabilities running

** Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

various OS and applications, and network simulation using OPNET Modeler. In some cases, the emulated nodes operated with surrogate applications; meaning, if the real application was not available, a similar application would be operated in its place. The following sections describe the various parts of the LVC testbed experiments and how they are combined to represent a CIS gateway of interest. Figure 1 illustrates a demonstration use case and identifies how components are represented in the experiment.

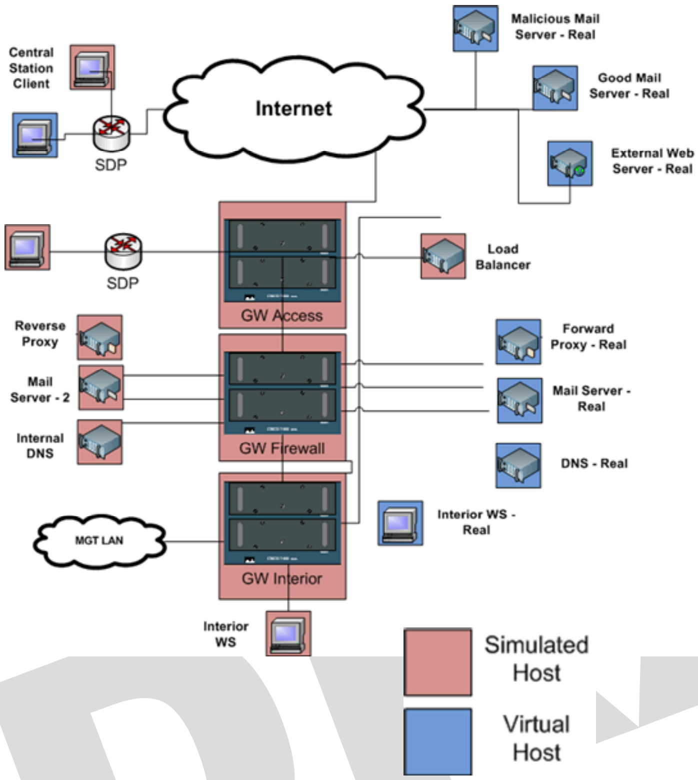

Figure 1: LVC Testbed Experiment with Simulated and Real Devices

III. Simulated Network Used in LVC Testbed Experiment

In many cases, having a standalone experiment network built with real devices on which to perform cyber security experiments is not possible due to reasons such as cost. Thus the capability to represent the network under study in the modeling and simulation domain is very attractive. A key aspect to our cyber security analysis capability is the availability of network device models. The OPNET Modeler network M&S tool meets this requirement. Network simulation tools such as OPNET Modeler are designed in part to allow analyst, engineers and researchers to understand how network algorithms perform under various traffic loads and device configurations. Analysts can implement and deploy these algorithms on networks of simulated devices, trace messages that the devices send between one another, and collect statistics on the resultant traffic including packet delays. Only recently has network M&S been identified as a tool to be used in cyber security analysis.

A key advancement that enabled using network M&S tools in cyber security analysis has been the capability to interface real network data traffic with simulated data traffic. The means of interfacing real network traffic with simulation traffic recently became available with OPNET's system-in-the-loop (SITL) capability. SITL uses the Winpcap library for Microsoft Windows machines and the libpcap library for UNIX-like machines to pass traffic packets from real or emulated nodes to or from simulated network devices.

The limitations of using M&S for cyber security analysis must be recognized. When using network M&S in a LVC testbed to perform cyber security analysis it must be understood that the modeled network components represent the behavior of real network devices in their configurations and capability to transport network traffic but accomplish this through different implementations of the network protocols. Device operating system (OS) and application vulnerabilities are ***not*** modeled with OPNET Modeler network modeling tools. Typically, vulnerabilities are implementation specific and vary with each version upgrade or patch installation. As a result, it is difficult to get accurate system-wide predictions from the models alone. Thus, a device model's behavior may not represent a real device's behavior when the vulnerability is exploited in the real device. In the case of vulnerability analysis, this difference limits the number of vulnerabilities that researchers might discover through the simulation models alone. As a result, the vulnerability researchers traditionally turn to the implementations for their analysis with the cost of limiting the size and diversity of the networks that they can analyze.

However, the model device can represent the real device in its configuration of security features such as filter rules and access control lists (ACLs). Most devices provide a variety of configuration options that users can set, based on their own security versus convenience tradeoffs. Because convenience is often valued more than security, many systems are, in practice, configured insecurely. If configurations in a real device permit or deny an attack, it is expected that the model with the same configuration will permit or deny the same attack vector.

A key part of our LVC testbed is the capability to interface real CIS devices and subsystems to simulated CIS devices and subsystems. The real part of the experiment could be a workstation connecting to a logically distant real server over an extensive simulated network or various traffic sources and sinks communicating over a network comprised of real and simulated parts. Combining real and simulated devices into a single experiment requires the SITL interface to translate data packets or datagrams between real and simulated domains. SITL employs translation functions to interface packets or datagrams between the two domains. Translation functions are necessary for cases where a datagram is created in one domain, either simulated or real, and interpreted in another domain. Packets created by specific protocol functions must have standard library translation functions available or translation functions must be developed. OPNET SITL currently supports a limited set of protocols [5]. In addition to standard SITL translation functions, Sandia National Labs has developed a Border Gateway Protocol (BGP) translation function and a Transport Control Protocol (TCP) translation function [1][2].

In cases, where the simulated network is transporting the data from one real device to another the translations are limited to the header portion of the data packets. The payload of the data packets can remain as a block of bits. Since the simulation may include filter rules in modeled routers and switches and ACLs in modeled firewalls, the data packet headers are read, interrupted, and acted upon as a real device with the same configuration would act upon the data packet.

IV. EMULATED DEVICES AND NETWORKS USED IN LVC TESTBED EXPERIMENT

In order to represent authentic network enterprise services, virtual machines (VMs) are utilized as surrogate systems functioning as hosts and servers. In the system under test, physical hardware solutions are utilized to provide services such as DNS, email and proxies. By utilizing VMs, several key advantages are encountered. First, given modern hardware, it is possible to virtualize a significant portion of the experiment, thus enabling numerous services and devices to be consolidated into a single, portable computing source, resulting in a cost efficient alternative to using proprietary hardware solutions. For example, similar functionality of a BlueCoat® Proxy [6] can be reached by implementing a Squid® proxy [7]. This approach provides ability to create authentic data traffic for several dozen systems without having to purchase several dozens of costly hardware platforms. However, there exist tradeoffs; primarily that the exact behavior and performance of using the actual hardware is not reached. We believe the benefits outweigh this limitation since this analysis approach leads itself to providing an "experiment-in-a-box" capability; meaning that through virtualization an entire experiment can be contained in a single computer (albeit a powerful machine). In addition, virtualization enables the developer to migrate and instantiate numerous instances of an experiment, which makes possible distributed activities such as training and testing/evaluation. Programmatic duplication of the virtual infrastructure enables the environment to be easily duplicated numerous times. It can be challenging for an analyst to build an entire infrastructure to test a particular component. Using virtualization the analyst can create experiments of the entire dataflow of the system.

Thus combining virtualization with simulation through system-in-the-loop enables analyst to create experiments with varying fidelity. The approach provides for placing fidelity, with hardware for example, in only the components or areas of interest without having to incur the cost of exactly duplicating the entire system.

V. SECURITY ASSESSMENT DEMONSTARTION EXPERIMENT - SETUP

This research activity included identifying and assessing a secure network gateway, essentially an interface between trusted and untrusted networks, which provides security for a large installation. The IT architecture, a system of security gateways are tasked to provide reliable and fault tolerant access to critical IT services in the event of single or multiple failures, including those resulting from cyber attacks. The gateway consists of network elements as well as Domain Name System (DNS) servers, Proxy servers, email relays, and an array of systems and services used to provide a complete and standalone IT capability.

The research team initially reviewed requirements documents of the security gateways to be assessed. This provided the research team with an understanding of the gateway's intended operation and the experiments to be performed. Additionally, the documents are a resource for the research team to create an experiment of real, emulated, and simulated devices. Experiments that assessed implementations of certain functions and devices are represented in the highest fidelity with real or emulated devices. Examples of real systems are actual operating systems implemented on VMs. Other components are represented as surrogates or simulated.

In the demonstration system, the gateway's network devices are represented in simulation. The simulated gateway network is comprised of OPNET Modeler high-fidelity models. High-fidelity models the models have similar behavior as the real devices they represent. The models have their own implementation of the same protocols and include similar variable parameters as the real devices. In many cases, including the demonstration system, the configuration parameters were extensive for each device. In general, the level of detail necessary to accurately create a model of a gateway device is the same level of detail required to build and configure a real gateway device. Thus, the optimal way to create models of the real devices is via direct import of actual configuration files. This is especially true in cases where there is extensive use of access control lists (ACLs) such as with firewalls. As it turns out, the gateway used in our demonstration assessment is comprised of Cisco devices with extensive configuration files.

Creation of the gateway network model was facilitated by an OPNET Modeler extension module called eXpress Data Import (XDI) [8]. XDI will translate a group of Cisco configuration files into an initial model of the network including device model configuration. However, in our experience, XDI is able only to create an estimate of the final model because either Layer-2 switch connectivity information is not available or specific device functionality may not be available in a model. The XDI import cannot be done blindly because there are cases where the real device implementation may not be available in the model or may be modeled in a different way. An astute developer must examine each resulting model for accuracy and completeness. This same astute developer must also be capable of recognizing the real device configuration objective and be certain that this objective is also configured in the model. In some cases, security mechanisms used in a real device must be represented differently in the model to result in similar behavior. However, the resulting XDI generated model is an incredible time saver since the vast majority of the tedious, mistake-prone human configuration is done automatically. The astute network engineer will, in almost all cases, start with an XDI import and then manually build out the network model.

Our initial target was to create a model representing as much of the gateway as possible. In other words, if a gateway device model was available in OPNET Modeler, it would be used in the experiment. All of the network devices, such as routers, switches, and firewall, had models available. Thus they were represented in simulation as shown in Figure 2. In addition, models of hosts are included on each network segment for debugging purposes. The gateway services, such as DNS, web proxies, and mail servers are represented with surrogate applications installed on VMs built with either Linux or Windows operating systems. Figure 3 illustrates the LVC demonstration experiment with both the simulated and real parts.

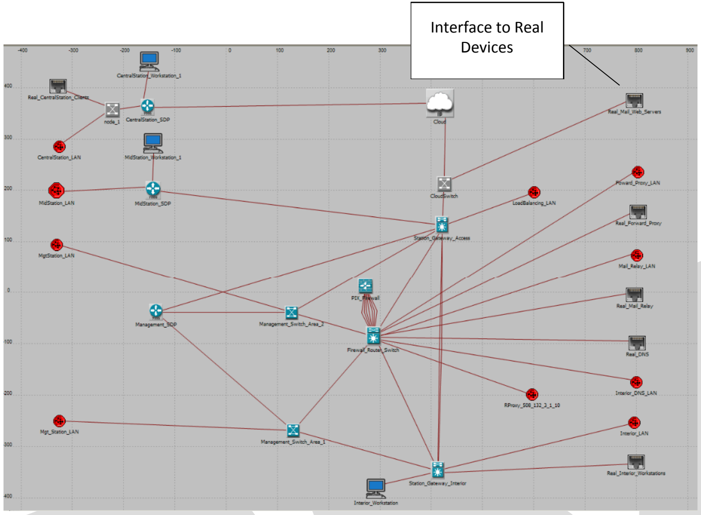

Figure 2: Simulated Part of LVC Testbed Experiment

In general, a model is built for a specific analysis purpose. The objective is to create a model that has precise representation of the specific areas of interest. Areas that are not of interest and do not have a significant impact on areas of interest can be abstracted to reduce model complexity. The goal is to obtain accurate results of interest while minimizing model and experiment development time. In addition, simulation computation resources may become an issue if models become too extensive.

VI. Security Assessment Demonstration Experiment – Security Mechanisms

Several security mechanisms used in the demonstration gateway are assessed in the experiment and warrant further discussion. The security mechanisms include firewall implementation, virtual private network (VPN) tunnels, and Cisco's VPN Service Port Adapter (VSPA) [9].

A. Cisco PIX Firewall

Our target demonstration system included a Cisco router/ switch with a Firewall Service Module (FWSM). Representing the FWSM in an experiment was a challenge since OPNET Modeler does not have a FWSM model that works with discrete event simulation (DES) nor works with the SITL interface. DES operation is necessary for experiments that interface simulation traffic with real devices. Lacking a model of the FWSM was overcome by recreating the FWSM behavior model with a model of the PIX firewall. Much of the FWSM functionality is similar to the PIX firewall functionality. Certain FWSM configurations can translate to PIX configurations. A single FWSM can be partitioned into multiple virtual devices, known as security contexts [10]. Each context has its own security policy, interfaces, and administrators such that each context is similar to a single standalone device. Since the FWSM used in our demonstration system used a single context, it was determined that the FWSM functionality could be reproduced with a PIX firewall. Importing the extensive FWSM firewall configuration files into a PIX model did require some modification to represent the real FWSM in a PIX device model. Ultimately, the combination of a switch model and a PIX firewall model was able to reproduce the functionality of the router/switch FWSM.

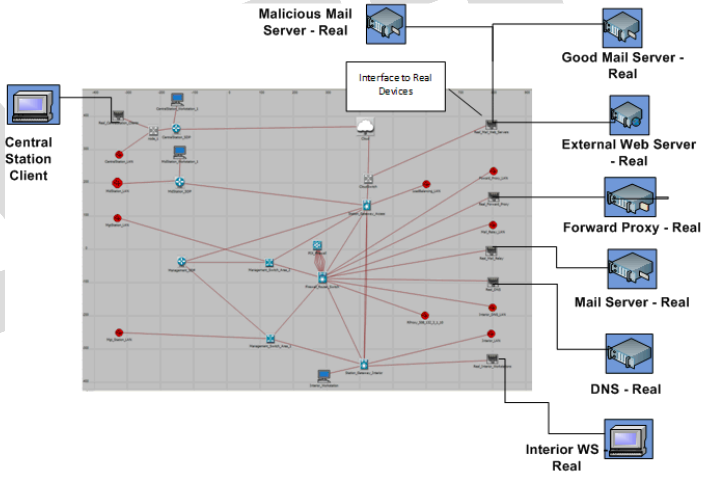

Figure 3: Combined Real and Simulated Parts of LVC Testbed Experiment

B. VPN Tunnels

Our target demonstration system, like many distributed enterprise networked systems, employs extensive use of encrypted virtual private network (VPN) tunnels to securely transport data over a public network. The demonstration system requirements included transporting both IP traffic and non-IP traffic (e.g., OSPF control data) between the remote location and the gateway interior over a public network infrastructure. To support both types of traffic the VPN is configured as a generic routing encapsulation (GRE) over IPSec tunnel. This configuration supports both traffic types by encapsulating all traffic destined for the VPN in GRE tunnel. IPSec can then be used to encapsulate the resulting GRE packet thus completing the GRE over IPSec VPN tunnel.

Since OPNET Modeler does not support full implementation of IPSec in DES an abstraction in our model is the lack of IPSec encryption. This is an acceptable abstraction since our example security analysis makes no attempt to hijack unencrypted packets transported on gateway connections. Further, the computational cost of each encryption and decryption is very expensive and the simulation would slow to a crawl. The GRE tunnel is precisely modeled and the additional packet size, resulting from additional headers, is accounted for in the model.

An additional challenge exists when modeled VPNs are combined with real VPNs if the source and termination are in different domains. For example, sourcing a VPN at a real router and attempting to terminate that VPN at a modeled router requires the insertion of additional real hardware in the experiment. Transition devices are required to terminate the real VPN and pass the resulting traffic into the modeled scenario. The traffic passed into the model is then re-encapsulated into a modeled VPN.

C. VPN Service Port Adapter (VSPA) Connectivity

Our demonstration gateway has connectivity to other distant gateways and remote locations. Connectivity is provided to the geographically dispersed locations via VPN over public infrastructure. The VPN implementations used in each gateway incorporate the Cisco VPN Service Port Adapter (VSPA) using the crypto-connection configuration approach. The Cisco module is implemented on the gateway interior router-switch. In this approach, VPNs are configured on the VSPA by attaching crypto maps to interface VLANs and then crypto-connecting a physical port to the interface VLAN [11]. This approach is considered a crypto-connect mode. Our demonstration system employed the VPN crypto-connect configuration approach with crypto maps attached to VLANs (using interface VLANs). Unfortunately OPNET Modeler does not support this VPN approach in discrete event simulation (DES). Modeler DES is not be able to associate a VLAN with a physical interface as required by crypto-connect. Our development team devised a workaround that produced the VPN behavior in the model. Our workaround is to manually set the physical interface to the IP address associated with the VLAN.

VII. SECURITY ASSESSMENT DEMONSTRATION EXPERIMENT - SIMULATION RUN-TIME

In an LVC experiment caution must be taken to be certain that the simulated part of the experiment can run at a real time rate. Since real or emulated devices operate at real time the simulation must also support that rate. As data packets progress through the modeled network the delays must be consistent with networks made of real devices or the interaction between the simulation and real devices no longer represents realism. As an example, real device TCP will interpret a slow simulation as a congested network and will throttle back its window size. This does not represent real TCP behavior and must not be permitted to occur. Thus caution should be exercised to be certain that the simulation can support the traffic loads under real-time operation.

Our team has developed estimation algorithms and test scenarios methods to effectively estimate whether or not a simulation scenario can run at real-time on the supporting compute platform [1]. Our current approach targets identifying simulation network characteristics such as number of SITL interfaces, SITL interface filter level, number of routers and other network devices, degree of connectivity, protocol usage, and expected traffic loads.

VIII. SECURITY ASSESSMENT DEMONSTRATION EXPERIMENT - RESULTS

A critical part of assessing the security of the system under test is to conduct a vulnerability assessment. Performing a vulnerability assessment on a representation of the system versus the actual system depends heavily on the composition of the representation of the system. Clearly the objective is to obtain identical or very similar behavior from the representative system when compared to the actual operational system.

In our research we assessed the cyber security behaviors of the representation or modeled system in comparison to an actual system. Several normal security assessment tools and techniques were used to evaluate the efficacy of the model. First, port and vulnerability scans were conducted against and through the model using traditional tools such as NMAP and Nessus®. The results are as expected. The simulated network devices enforced ACLs and polices of the system under test. This was demonstrated by finding that only certain types of traffic were allowed through particular components of the system, while others were complete dropped congruent with what would be expected with the real system. The port and vulnerability scans yielded the expected behavior by detecting the actual configurations and preplaced vulnerabilities in the experiment.

In our demonstration experiment, a simulated exploit of preplaced known vulnerabilities were conducted. This experiment used an open source vulnerability exploitation framework commonly used to assess the security posture of networked systems. In the experiment, common vulnerabilities were demonstratively exploited in the hybrid experiment. Malicious payloads created by the exploit tool were successfully passed through the hybrid representation of the system. The payloads passed through both, physical devices and modeled devices, and ultimately effected change on virtual hosts. The implication of this is significant. The experiments resulted in expected behaviors and thus lend itself to enabling distributed operation test and evaluation (OT&E). With this approach, a cyber security analyst or researcher can look at a particular component of the representation of the system or model, obtain physical devices of interest and test the components for the vulnerabilities, possibility of being exploited by know methods, and assess the effects on the entire system. The analyst can then deploy mitigation methods in the modeled system and assess their ability to prevent exploitation of the system. A key part of our cyber analysis approach is that experiments are standalone and are not connected to operational systems. After an analyst performs an experiment the modeled system can be quickly reconstituted back to its original state for further experiments.

A red team did an assessment of the demonstration hybrid representation or model and had positive results. An accurate logical representation of the network was able to be extracted through both active and passive techniques. Vulnerabilities were able to be exploited and mitigation strategies were tested. However, some noticeable differences in response time between modeled system and real system associated with network scans resulted.

IX. Conclusion and Further Study

In this research we have developed an important and capable cyber security analysis and experiment environment (i.e., testbed) to help perform analysis of communication networks and networked information systems. Our developments resulted in a LVC cyber analysis testbed comprised of simulated, emulated, and real components that leverages existing capabilities where possible. The LVC testbed enables higher fidelity representations of key computing applications or network devices while still leveraging the scalability and cost advantages of simulation tools. The result is rapidly produced large, yet relatively low-cost, multi-fidelity representations of networked information systems that enable analysts to quickly investigate threats then test different protection approaches and configurations.

In our research, we identified a secure information system use case that is comprised of LAN and WAN networks including routers, switches, and firewalls. Security mechanisms such as VPN tunnels, extensive access control lists (ACLs), network address translation (NAT), and virtual LAN (VLAN) separation are heavily utilized in the use case. Network device configuration files obtained from the use case system are used to create a high-fidelity model of the network that passes network traffic and performs like the real network. The use-case includes real computer systems that generate traffic for transport over the modeled network.

In our research, we examined the issues of real-time performance of the modeled components of the network and identified ways to increase its capability to transport higher traffic loads. Our approach supports replacing network devices that are represented in the constructive domain with real devices. Offloading the simulation by removing, for example, a simulated firewall and replacing with a real firewall enables the simulation to support higher traffic loads and run at real-time.

The cyber security LVC testbed provides high fidelity representations of key network nodes while still leveraging the scalability and cost advantages of simulation tools. Sandia National Laboratories applies the LVC testbed to its mission of enhancing computer security used in critical government and commercial applications

References

[1] B. Van Leeuwen, D. Burton, U. Onunkwo, M. McDonald, "Simulated, Emulated, and Physical Investigative Analysis (SEPIA) of networked systems," 2009 MILCOM Conference, IEEE, October 2009.

[2] B. Van Leeuwen, U. Onunkwo, M. McDonald, "BGP analysis using System-in-the-Loop (SITL) testbed," 2008 OPNETWORKS Conference, August 2008.

[3] E. Parker, N. Miner, B. Van Leeuwen, J. Rigdon, "Testing unmanned autonomous system communications in a Live/Virtual/Constructive environment," International Test and Evaluation Association Journal (ITEA), 2009; 30: 513–522.

[4] OPNET Technologies, Inc. www.opnet.com.

[5] OPNET Technologies, Inc., "OPNET Modeler documentation set - version: 15.0 - System-in-the-Loop (SITL)," May 2009.

[6] Blue Coat Systems, Inc., http://www.bluecoat.com/company/aboutbluecoat

[7] Squid Project, http://www.squid-cache.org

[8] OPNET Technologies, Inc., "OPNET Modeler documentation set - version: 15.0 - eXpress Data Import," May 2009.

[9] Cisco Systems Inc., "Cisco Security Modules for Routers and Switches - Cisco VPN Services Port Adapter," 2009.

[10] Cisco Systems Inc., "Cisco Services Modules - Firewall Services Module(FWSM) FAQ," 2009.

[11] Cisco Systems Inc., "Cisco VPN Services Port Adapter Configuration Guide - Overview of the IPsec Features," 2009.