2012 Sixth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing

Energy-efficiency and Security Issues in the Cisco Nexus Virtual Distributed Switching

Benjamin Peterson
Department of Computer and Information Technology
Purdue University
47906 West Lafayette, Indiana, USA
e-mail: bdpeters@purdue.edu

Sergio Ricciardi, Jordi Nin
Department of Computer Architecture
Technical University of Catalonia (UPC)
08034 Barcelona, Spain
e-mail: {sergior, nin}@ac.upc.edu

***Abstract*—Virtualization technologies have brought with them the promise of increased security and energy saving. Such was the case with the Cisco Nexus virtual switching environment. However, possible security issues of this environment have not been evaluated and the achievable energy savings have not been quantified yet. In particular, it was necessary to investigate whether the security vulnerabilities existing in physical switches had persisted into the virtual environment. This paper provides an evaluation of the energy saving and an analysis of the security implications of the Cisco Nexus virtual distributed switching environment.**

Keywords-component; Nexus 1000V, enegy consumption, security in switching virtualization, virtual architectures.

I. INTRODUCTION

The growing of the Information and Communication Society (ICS) in the last years has increased the quantity of data transmitted, accessed and stored on the Internet. The size of datacenters has grown tremendously with the development of the ICS, and the energy requirement has become their main limiting factor. To put this into perspective, a medium-size datacenter such as the Barcelona Supercomputing Center consumes 1.2 MW (as much power as a town of 1,200 houses) and has an energy bill of € 1 million per year [1][2]. Apart from the HVAC (heating, ventilation and air conditioning) and UPS (uninterruptible power supply) systems, the power consumption in datacenters comes from the devices providing the computing and storage resources as well as the network interconnections, with the servers being the most energy-hungry devices. The PUE (power usage effectiveness [3]) measures the ratio between the total power required by the facility versus the power required only by the computing, storage and interconnection resources. A PUE of 2 is considered an average value [4], meaning that HVAC and UPS double the energy requirements of datacenters.

In order to lower the energy requirement of datacenters, virtualization technologies have been deployed in large scale. Aside from providing benefits, such as isolation, virtual management, and abstraction, virtualization offers economic benefits both in terms of capital and operational expenditures. These benefits stem from the reductions in hardware requirements and server consolidation (multiple virtual servers placed into a lower number of physical servers, increasing servers usage efficiency). Virtualization is possible not only at the machine level, e.g. by virtualizing a set of functionalities into a separate virtual machine, but also at the interconnection level, e.g. by virtualizing the network devices connecting the virtual machines. To this end, Cisco and VMware teamed up to create the Cisco Nexus 1000V: a purely software-based switch, replacing the classic VMware dvSwitch. It offers all the features of the VMware vSwitch while providing a number of advanced switching features commonly found on Cisco and other standards based switches. However, despite the superior functionality offered, there was uncertainty as to whether the vulnerabilities found in physical switches had persisted. Such subtle security issues, could allow an attacker to exploit the features of the virtual switching devices or the associated servers, adversely affecting datacenters functionality and their energy consumption [5].

In this paper, the energy savings and a highlight of potential security issues residing within the Cisco Nexus 1000V's distributed switch environment are analyzed. In particular, energy saving obtained by the virtualization of switches and servers is evaluated and the security ramifications of using the Nexus 1000V switches are tested to determine whether the vulnerabilities found in physical switches have persisted into the virtual environment.

II. ARCHITECTURE

In order to assess the energy requirements and the security implications of the Nexus 1000V, it was necessary to create an architecture that would provide a virtual infrastructure in which the Nexus 1000V could reside. While designing this architecture, emphasis was placed on ensuring that it would be functionally representative of the architectures commonly implemented in datacenters. To this end, a network architecture utilizing the Nexus 1000V switch was built. This architecture consisted of three servers running VMware's virtualization software ESXi and one server running VMware's management software, vCenter. Each of the servers running ESXi had virtual machines (VMs) installed on them as well as a virtual machine running the Nexus 1000V. All four of the servers were connected to each other to allow communication among them. A fifth server was used to run the traffic capturing software tcpdump (Fig. 1 depicts the physical servers architecture).

978-0-7695-4684-1/12 $26.00 © 2012 IEEE
DOI 10.1109/IMIS.2012.60

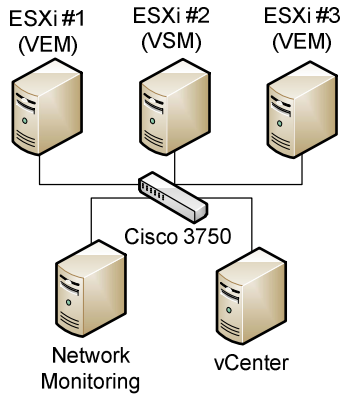

Figure 1. Physical servers achitecture.

A. Nexus 1000V Requirements

The Nexus 1000V was designed to provide network functionality amongst multiple virtual machines hosted on physical hosts. In order to integrate its network functionality into the virtual environment, it makes use of the VMware vNetwork Distributed Switch (DS) API [6]. Since this is the sole method for integrating the Nexus1000V into virtual environments, VMware's vSphere virtualization software must be used. VMware vSphere is a platform of management tools that are used to control the virtual environment. It has four component layers that handle various management tasks [7]. These layers consist of infrastructure, application services, VMware vCenter Server and Clients. The infrastructure layer's primary role is to manage and facilitate the sharing of resources, storage, and network capabilities amongst the virtual hosts. This is achieved by abstracting, aggregating, and allocating the physical resources residing on the virtual hosts. The application services layer's responsibility is to provide the virtual environment with high availability, security, and scalability. The VMware vCenter Server component layer's role is to provide the actual management of the virtual environment. The clients component layer provides the external facing point that is used to allow administrators to control and configure the virtual environment. Although vSphere manages the virtual environment, it is not the underlying hypervisor that resides on the virtual hosts. In fact, it is possible that vSphere resides on a virtual host that it manages. Two possible hypervisors can be placed on the virtual hosts to facilitate the creation and hosting of virtual machines: VMware's ESX and ESXi. Both of these hypervisors provide the same virtualization functionality [7]. The main differences between the two are that ESXi lacks the service console that is built into ESX and ESXi has the ability to be embedded into the firmware of a server. These hypervisors operate on "bare metal", running directly on the host's hardware and therefore not requiring an underlying operating system. vSphere uses its vCenter component layer to interact directly with the hosts running ESX or ESXi. It is also worth noting that not all versions of vSphere, ESX and ESXi support the Nexus 1000V. The Nexus 1000V requires that version 4.1 or later of vSphere Enterprise Plus be used [6]. VSphere can reside on its own server or, as stated earlier, it can reside on a virtual machine. If it is to reside on a virtual machine, version 3.5U2 or later of ESX or ESXi is required. Although vSphere requires version 3.5U2 or later of ESX or ESXi, each virtual host on which a Nexus 1000V VEM will reside must have version 4.0 or later of ESX or ESXi installed on it.

B. Design Considerations

In ideal circumstances, this research would have been carried out in a full-scale datacenter; however, due to the nature of security testing, doing so would put the integrity of the datacenter's functionality at risk. Because of this, it was impossible to utilize an existing datacenter. The purchasing of the equipment to create a full-fledged datacenter was also cost prohibitive. It was therefore necessary to design the test architecture so that it would be representative of the functionality found in datacenters. For this purpose, it was necessary to determine the minimum number of servers needed to replicate typical traffic handled by Nexus 1000Vs. The Nexus 1000V facilitates communication between the virtual machines and to other networked devices outside of the virtual environment. Communication can take place between the virtual machines on the same virtual host as well as between virtual machines residing on separate virtual hosts being serviced by the same Nexus 1000V virtual switching module (VSM). With this determination, it was ascertained that only two virtual hosts would be necessary to simulate the necessary network traffic.

It was also important to consider how the Nexus 1000V virtual Ethernet modules (VEMs) and VSMs communicated with one another. The VSMs communicate with the VEMs to pass configuration information [6]. One thing that was unclear was whether the VEMs communicated directly with each other when passing traffic between virtual machines on other virtual hosts or if the VEMs passed the traffic to the VSMs to have the VSM make the appropriate forwarding decisions. Despite this uncertainty, it was clear that there would need to be at least three virtual hosts in order to properly simulate the network traffic. With this setup, one virtual host would host the VSM and the other two virtual hosts would each host a VEM.

Another consideration was to evaluate the communication paths of which attackers take advantage. This determination was important because without it, it would be impossible to evaluate the different potential vulnerabilities residing in the switching functionality of the Nexus 1000V. In attacks, such as double tagging, attackers send direct communication to other machines [8]. Attackers often try to capture and analyze network traffic in an attempt to find other potential vulnerabilities or to simply gain more information about the compromised network. In such a way, an attacker could look to compromise a VEM so it miscommunicates with the VSM and has traffic meant for another VEM directed to the VEM the attacker resides on. Therefore, it was necessary to have two virtual hosts hosting their own VEM and another virtual host hosting the VSM.

Datacenters using the Nexus 1000V would likely have tens, hundreds, or possibly thousands of virtual machines; it is likely that they would be segmented in several networks. Therefore, two VLANs should be implemented to represent the segmentation typically found within datacenters.

The final consideration was the different operating systems that often reside in datacenters. Most datacenters use Microsoft, Linux, UNIX, or a combination of these operating systems. This was an important consideration because in most attacks, attackers would use virtual machines that had been compromised through other means to attack the Nexus 1000V in an attempt to escalate their intrusion. Because of this determination, it was decided that it would be necessary to have virtual machines running Microsoft, Linux, and UNIX operating systems.

As a result of these considerations, it was possible to determine that the minimum number of servers that would be necessary to replicate the functionality commonly found in datacenters was three virtual hosts running a combination of Microsoft, Linux, and UNIX operating systems. Two of these would need to host a VEM, and the third one would be necessary to host the VSM.

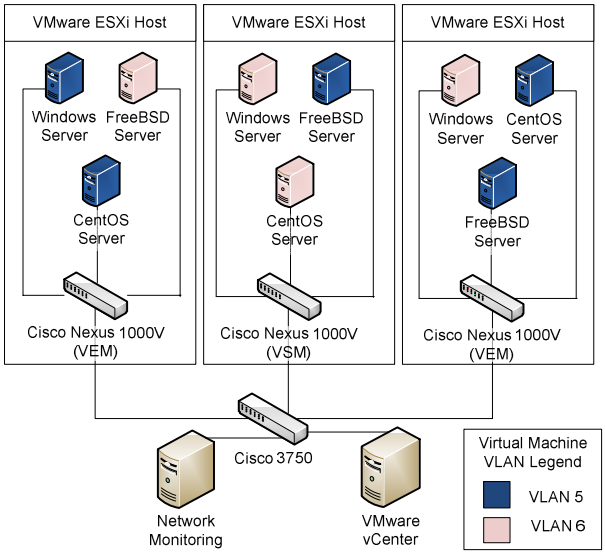

Figure 2. Logical server architecture.

C. Test Architecture

Based off the determinations from the Nexus 1000V's requirements and the design considerations for replicating the Nexus 1000V's functionality in datacenters, it was possible to begin creating the test architecture. First, a server running VMware vCenter was set up. Each of the three virtual hosts was populated with three virtual machines, having a server running Windows Server, FreeBSD, or CentOS. The servers were connected to a Cisco 3750 series switch. After vCenter had been properly configured to manage the virtual hosts, the Nexus 1000V was set up with two virtual hosts hosting VEMs and the other virtual host hosting a VSM. The virtual machines were then split into two VLANs to segment them from one another. Once proper network connectivity had been achieved, a separate server was added to the network, with the sole purpose of capturing network traffic that was passing through the physical switch. In Fig. 2 the logical architecture is depicted.

III. ENERGY CONSUMPTION CONSIDERATIONS

A. The Energy Model

The energy savings of virtualization are essentially due to the achievable consolidation, in which a number of low utilized virtual machines are put into physical machines, resulting in fewer powered servers with higher CPU utilization. Therefore, the energy consumption is related to the number of virtual machines that can be put into a single physical server. In order to assess a realistic case for the energy savings, two scenarios were considered, in which different assumptions are made. In the first scenario (A), no virtualization is allowed, whilst in the second scenario (B) a number of three virtual machines have been consolidated into each physical server. In the virtualization scenario (B), an instance of the distributed virtual switch Nexus 1000V is necessary on each physical machine in order to provide it with the required network connectivity and security features. In the virtualization scenario, apart from the virtual distributed switch, an additional host is required to run the VMware vCenter management software. It should be noted that the physical machines considered in the two scenarios are not required to be the same in terms of performance and energy consumption. Entry-level servers will work for hosting one single virtual machine each, whereas in the virtualization environment higher-level servers are required to host several virtual machines at a time without experiencing appreciable delays. The networking monitor host is not considered in the energy consumption evaluation since it is just needed for the security assessment and it is not required in the real world case. Table 1 summarizes the assumptions for the two scenarios.

Description	Scenario A (no virtualization)	Scenario B (virtualization)
Virtual machines per physical server	1	3
Nexus 1000V instances	0	1 for each physical server
vCenter instances	0	1
Physical servers category	Entry-level	Middle/high-level
Physical switch	1 Cisco Catalyst 3750	

Table 1. Assumptions for the two scenarios.

In the energy model (Fig. 3), the power consumption (P) of a physical server is made up of a fixed part (Φ), needed

for the machine to stay on, and a load-dependent variable part (δ), which varies proportionally with the CPU usage (L):

$$P = \Phi + \delta(L)\,. \tag{1}$$

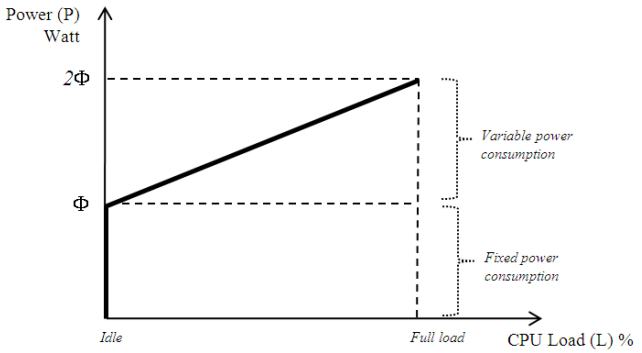

Figure 3. Server power consumption versus the CPU load.

The fixed part accounts for half of the total power consumption, and the variable part for the other half [9], varying linearly with the CPU load:

$$\delta(L) = \Phi \cdot L\,. \tag{2}$$

Therefore, from eq. (1) and (2), the power consumption P can be written as the line passing through the points (0, Φ) and (100, 2Φ), which is:

$$P(L) = \Phi + \Phi \cdot L = (L+1)\cdot\Phi\,. \tag{3}$$

Note that, from equation (3), if the load L = 100%, then:

$$\delta(L = 100\%) = 100\% \cdot \Phi = \Phi \tag{4}$$

i.e., at full load the variable power consumption is equal to the fixed power consumption, and thus:

$$P(L = 100\%) = 2\Phi\,, \tag{5}$$

the peak power consumption of a server is two times its idle power consumption, as wanted.

Therefore, given the server peak power and the current CPU load, it is possible to quantify its power consumption. The energy impact of both scenarios under typical power consumption and load values of the involved devices was evaluated. The power consumption of the two scenarios is the sum of the fixed and variable power consumptions of the physical devices (servers and switch). Therefore, the power consumption P^A of the scenario A is given by:

$$\begin{aligned} P^A &= N_{servers}\cdot\left(\Phi_{servers} + \delta_{servers}\right) + \left(\Phi_{switch} + \delta_{switch}\right) \\ &= N_{servers}\cdot\left(\Phi_{servers} + \Phi_{servers}\cdot L^A_{servers}\right) + \left(\Phi_{switch} + \Phi_{switch}\cdot L^A_{switch}\right) \\ &= N_{servers}\cdot\Phi_{servers}\left(1 + L^A_{servers}\right) + \Phi_{switch}\left(1 + L^A_{switch}\right) \end{aligned} \tag{6}$$

while the power consumption P^B of scenario B is given by:

$$P^B = N_{servers}\cdot\Phi_{servers}\left(1 + L^B_{servers}\right) + \Phi_{switch}\left(1 + L^B_{switch}\right) + P_{vCenter} \tag{7}$$

where $P_{vCenter}$ stands for the power consumption of the virtual machine hosting the VMware vCenter management software.

Since in scenario A each VM is assigned to a physical server, its average load is simply given by the load of the VM running on it:

$$L^A_{servers} = L_{VM}\,. \tag{8}$$

In scenario B, the average load of a physical server is given by the sum of the loads of the VMs running on it plus the load of the virtual distributed switch, i.e.:

$$L^B_{servers} = \sum_i L_{VM_i} + L_{Nexus}\,. \tag{9}$$

Description	Parameter	Scenario A	Scenario B
Number of physical servers	$N_{servers}$	9	3
Fixed power consumption of servers	$\Phi_{servers}$	225 W	250 W
Average virtual machine CPU load	L_{VM}	-	25%
Average Nexus 1000V load	L_{Nexus}	-	25%
Average physical machine CPU load	$L_{servers}$	25%	100%*
Physical servers peak power consumption [10]	$2\Phi_{servers}$	450 W	500 W
Cisco Catalyst 3750 power consumption [11]	Φ_{switch}, (130 W)	143 W @ L_{switch} = 33%	169 W @ L_{switch}= 100%

** 75% for the three VMs + 25% for the Nexus 1000V*

Table 2. Power load parameter values for the devices in both scenarios.

The parameter values used in the power consumption evaluation of both scenarios are detailed in Table 2. These values are for illustrative purposes and may vary according to the specific configurations of the datacenters.

B. Testbed Energy Consumption

In scenario A, the services offered by each virtual machine have to be deployed into individual physical servers. It is assumed that the average CPU load for each service is 25%. Each of the servers is connected to the Cisco Catalyst 3750 switch through a physical connection, and all the packets among the servers will always pass through the physical switch, which will result in a high utilization of the switch's CPU. The power consumption of such a configuration accounts for the fixed and variable power consumptions of the low-loaded physical servers in addition to the power consumption of the high-loaded switch [12], which sum to a total of 2.7 kW of power.

Keeping the same assumption for the VMs CPU load in the scenario B, three VMs are packed into each physical server, accounting for a 75% usage of its CPU, plus a 25% CPU load for running the Nexus 1000V distributed virtual switch, for a total a 100% CPU usage on physical servers. The $P_{vCenter}$ runs as an individual VM and its CPU load is assumed to be 25%. Three fully loaded high-level physical servers are needed to host the nine VMs, which are interconnected by means of the logical Nexus 1000V. Only the packets among different machines will actually pass

through the Catalyst 3750, which results in a lower switch CPU utilization with respect to scenario A. The vCenter management software is assumed to be hosted on an additional virtual machine. The resulting power consumption of scenario B sums to 1.7 kW, that is 1 kW of saved power with just three physical machines, corresponding to an energy saving of 37% with respect to scenario A.

The energy consumption calculations were scaled to reflect two types real world datacenters, a small one (100 servers) and a large one (5000 servers), obtaining similar results. In particular, with 100 servers, the energy savings settle at 40.3% and with 5000 servers the savings rise to 40.7%. It is worthwhile to note that with just a consolidation level of three VMs per physical servers, the energy savings are in the order of 40%, confirming that virtualization is an effective way to save a significant amount of energy.

IV. SECURITY ISSUES

The purpose of this section is to explore the security implications of using the Nexus 1000V. Thus, the focus is to determine whether the security issues with physical switches have persisted into Cisco's virtual switch.

A. Security Implications

The term security implication is often used loosely and, as a result, it has a somewhat ambiguous meaning. It is therefore necessary to describe its meaning within the context of this paper. With security implications, it is meant any action that causes the Nexus 1000V to deviate from its intended functionality. While this might seem drastic, it is important to keep in mind that such deviations are often used as "stepping stones" for achieving attacks that are for more sophisticated and nefarious.

B. Physical Switch Vulnerabilities

To determine whether security implications have persisted with the transition from physical switches to virtual switches, an array of vulnerabilities were tested. These vulnerabilities include ones that still affect physical switches and those that have been previously remediated. The following vulnerabilities will be covered in this paper:

- CAM overflows
- VLAN Hopping
- STP Manipulation
- ARP Poisoning

CAM overflows are a type of attack where an attacker attempts to flood a switch's CAM table with falsified information. Its goal is to provide the switch's CAM table with more information than it can hold and force the switch into broadcasting traffic it would otherwise sent directly to the intended recipient. Such vulnerability presents an attacker with the opportunity to receive traffic for which it was not intended. Testing for said vulnerability was possible with the macof tool [13]. It was found that, like physical switches, the Nexus 1000V also was susceptible to CAM overflows.

VLAN hopping is an attack that aims to ignore the logical segmentation imposed by VLANs. One potential way an attacker can ignore the segmentation imposed by VLANs is the technique known as double tagging. Double tagging is carried out by an attacker that crafts a packet containing two 802.1q tags [8], with one of the tags being the attacker's legitimate VLAN information and the second being the VLAN to which the attacker wishes to reach. When this packet is sent to the legitimate switch to which the attacker is connected, the switch will remove the correct VLAN tag while leaving the other VLAN tag intact. This traffic will then be passed to the next switch, which will discover the remaining VLAN tag and trust it since the traffic was passed from the initial switch. Because of this, the switch will pass the packet to the destination residing on the VLAN to which the attacker would have been otherwise unable to communicate. Through the use of Yersinia [14], it was found that it was possible to create packets with two 802.1q headers that were then processed and forwarded by the Cisco Nexus 1000V.

Spanning Tree Protocol (STP) is a protocol that has the purpose to identify and eliminate loops in switched environments [8]. In brief, this protocol works by having the switches communicate with one another to determine a "root" switch. This determination is based off two criteria, a potentially configured priority value and the switch's MAC address. Once the switches have agreed on a root switch, the other switches seek out paths to the root switch. After doing so, the switches are able to determine potential loops existing in the environment. If any loop is identified, the appropriate switch ports can be shutdown to prevent the loop from detrimentally influencing the network. It is possible for an attacker to take advantage of this protocol by tricking the switch to which the attacker is connected into believing the attacker's machine is a switch with a priority causing it to be elected as the root switch. While the other switches are determining the paths to the new switch, the network will be rendered unusable. To test whether this vulnerability existed within the Nexus 1000V, Yersinia was used [14]. It was found that, unlike typical physical switches, the Nexus 1000V does not run STP because it will deactivate all but one uplink to an upstream switch, preventing full utilization of uplink bandwidth. Instead, each VEM is designed to prevent loops in the network topology. Because of this, such attacks did not affect the switch.

Address resolution protocol (ARP) poisoning is yet another way that attackers can take advantage of physical switches. ARP is a protocol used by networked devices to map IP addresses to MAC addresses [15]. When hosts are unaware of the MAC address to which they should send traffic, they broadcast an ARP request. If a network device sees a request for its MAC address, it responds appropriately. The requestor will then associate the response's MAC address with the appropriate IP address. Another type of ARP message is referred to as a gratuitous ARP [8]. Many network devices broadcast gratuitous ARP

messages when they initially connect to a network to announce their address information. Attackers can take advantage of ARP transaction by using ARP poisoning [15]. With ARP poisoning, attackers send manipulated ARP messages to trick other network devices into believing their MAC address should be associated with a victim's IP address. In doing so, traffic destined for a host will be sent to the attacker's network device. To test for this vulnerability on the Nexus 1000V, the tool Ettercap [16] was used and falsified ARP messages were created. Ultimately, it was found that the Nexus 1000V was susceptible to ARP poisoning.

V. CONCLUSIONS

This study has helped to shed light on the effects of using the Nexus 1000V virtual switching. The information from this research was twofold: the energy consumption benefits of the Nexus 1000V were quantified and contrasted with the security implications of using the Nexus 1000V.

While the impact of using the Nexus 1000V on power consumption had been widely touted in marketing material, there existed little literature that quantified such improvements. After appropriately modeling the Nexus 1000V, it was possible to ascertain such information. The use of virtualization along with the use of the Nexus 1000V, does indeed bring with it energy consumption benefits. The modeling revealed that datacenters choosing to employ such technology should expect an energy reduction of 40% with just a consolidation level of three VMs per physical server. It is worth noting that this energy reduction does not include the reductions in cooling, which is quantifiable in another 40% assuming the average PUE. Virtualization and energy reduction not only help to reduce the capital and operational expenditures of datacenters, but they will also help to lessen the impact on the environment through a reduction in the emissions associated with energy production.

Prior to this research, there was little network security information in regards to the effect of using a Nexus 1000V. This research has helped to provide a glimpse into such security implications. In particular, it has shown the issues that have persisted from physical environment into the virtual environment. Conversely, in cases such as STP, security issues found in physical switches have been mitigated.

With the completion of this research, it is clear that the Nexus 1000V brings with it a clear benefit in terms of energy consumption. However, it does not eliminate all of the security vulnerabilities that have affected physical switches. Therefore, it is by no means a perfect solution, but rarely do such solutions exist. If the Nexus 1000V is properly configured, those employing it will be able to reap the energy consumption benefits while still maintaining a similar level of vulnerability to the level they would have had, had they implemented their datacenter in a strictly physical environment.

ACKNOWLEDGMENTS

This work was supported in part by the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems, the Spanish Ministry of Science and Innovation under the DOMINO project (TEC2010-18522) and ARES – CONSOLIDER INGENIO 2010 CSD2007-00004, the Catalan Government under the contract SGR 1140, the DIUE/ESF under the grant FI-201000740 and the U.S. Department of Education through the EU-U.S. Atlantis grant (P116J090064) from the Fund for the Improvement of Postsecondary Education (FIPSE).

REFERENCES

[1] Jordi Torres, "Green Computing: the next wave in computing", Ed. UPCommons, Technical University of Catalonia (UPC), February 2010, Ref. http://hdl.handle.net/2099.3/33669.

[2] Peter Kogge, "The tops in flops", pp. 49-54, IEEE Spectrum, Feb. 2011.

[3] The Green Grid, "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE", Technical Committee White Paper, 2008.

[4] W. Vereecken, W. Van Heddeghem, D. Colle, M. Pickavet, P. Demeester, "Overall ICT footprint and green communication technologies", in Proc. of ISCCSP 2010, Limassol, Cyprus, Mar. 2010.

[5] F. Palmieri, S. Ricciardi, and U. Fiore, "Evaluating Network-Based DoS Attacks Under the Energy Consumption Perspective", International Conference on. Broadband, Wireless Computing, Communication and Applications (BWCCA), 2011, pp. 374-379. doi: 10.1109/BWCCA.2011.66.

[6] Cisco Systems, "Cisco Nexus 1000V series switches", August 2011.

[7] VMware, "Introduction to VMware vSphere", November 2010.

[8] G. Bastien, S. Nasseh, and C. Degu, "CCSP self-study: CCSP SNRS exam certification guide", Indianapolis, IN: Cisco Press, 2006, pp. 279-302.

[9] X. Fan, W.-D. Weber, and L.A. Barroso, "Power Provisioning for a Warehouse-Sized Computer", In Proceedings of the ACM International Symposium on Computer Architecture, San Diego, CA, Jun. 2007.

[10] Dell PowerEdge R610 Rack Server data sheet, online. Available: http://www.dell.com/us/business/p/poweredge-r610/pd.

[11] Cisco Catalyst 3750 data sheet, online. Available: http://www.cisco.com/en/US/products/hw/switches/ps5023.

[12] S. Ricciardi, D. Careglio, G. Santos-Boada, J. Sole-Pareta, U. Fiore, and F. Palmieri, "Towards an energy-aware Internet: modeling a cross-layer optimization approach", Telecommunication Systems (2011), in press, doi: 10.1007/s11235-011-9645-7.

[13] D. Song, "Macof(8) – linux man page", Retrieved from: http://linux.die.net/man/8/macof.

[14] A. Omella, and D. Berrueta, "Yersinia", Retrieved from: http://www.yersinia.net/

[15] D. Bruschi, A. Ornaghi, and E. Rosti, "S-ARP: A secure address resolution protocol", Proc. 19th Annual Computer Security Applications Conference (ACSAC 2003), IEEE Computer Society, Dec. 2003, pp. 66-74, doi:10.1109/CSAC.2003.1254311.

[16] A. Ornaghi and M. Valleri, "Ettercap", May, 2005, Retrieved from http://ettercap.sourceforge.net/index.php.