Analysis of Virtual Networking Options for Securing Virtual Machines

Ramaswamy Chandramouli
Computer Security Division, Information Technology Laboratory
National Institute of Standards & Technology
100 Bureau Drive, Gaithersburg, MD, USA
E-Mail: mouli@nist.gov

***Abstract* – Virtual Machines (VMs) constitute the primary category of resources to be protected in virtualized infrastructures. Out of the two types of protection for VMs – Host-level and Network-level – it is the approaches for the Network-level protection that are different in virtualized infrastructures as compared to those for non-virtualized environments. This is due to the fact that the VMs are end nodes of a virtual network as opposed to being end nodes of a physical network. In this paper, we provide a detailed analysis (in terms of advantages and disadvantages) of some of the key approaches for two Network-level protection measures for virtualized infrastructures – Network Segmentation and Traffic Control using Firewalls. The choice of these two Network-level protection measures is due to the fact that they form the foundation for the network configuration of the entire virtualized infrastructure. We also provide the overall conclusions from the analysis in the form of recommended deployment choices based on approaches for these two network-level protection measures for securing VMs.**

Keywords - Virtual Machine; VLAN; Hypervisor; VXLAN; Virtual Firewall.

I. INTRODUCTION

Virtualized hosts (also called hypervisor hosts) are increasingly deployed in data centers because of efficiency, scalability and cost considerations. The virtualized infrastructure resulting from the deployment of virtualized hosts has three main categories of components - Hypervisor Software, Virtual Machines (VMs) and Virtual Networking components such as Virtual Network Interface Cards (vNICs), Virtual Switches and Virtual Firewalls.

Out of the three categories of components above, the VMs constitute the fundamental resource to be protected in a virtualized infrastructure, since they are the compute engines on which business/mission critical applications of the enterprise run. These VMs are virtual counterparts of physical servers and hence just like their physical counterparts, security for these VMs has to be provided through host-level and network-level measures. These measures may also vary depending upon whether the overall virtualized infrastructure (in which the VM resides) is used for in-house enterprise applications or for offering cloud services to external entities (e.g., Infrastructure as a Service Public Cloud). We provide a brief overview of the two types of protection mentioned above. (a) Host-level protections required for VMs include features for robust authentication, access using secure access protocols and secure session establishment. The mechanisms required for providing these features are no different for VMs compared to their physical counterparts (i.e., physical servers). (b) Network-level protections required for VMs are feature-wise similar to those that are required by their physical counterparts. However, the mechanisms or approaches required for providing these protections are different due to the fact that the VMs are end nodes of a virtual network as opposed to being end nodes of a physical network.

For any type of datacenter infrastructure (virtualized or non-virtualized), there is a general consensus that the following are some of the key Network-level protection measures [1]. They are: (a) Network segmentation or isolation, (b) Traffic control using firewalls, (c) Creating Redundant communication pathways, and (d) Traffic Monitoring and Prevention using IDS/IPS.

Out of the above four network-level protection measures, the first two - Network Segmentation and Traffic Control using Firewalls - form the foundation for the network configuration of the entire virtualized infrastructure. Hence, in this paper, we have chosen to focus on different approaches or mechanisms used for these two network-level protection measures, by performing a detailed analysis of the advantages and disadvantages of each of the approaches.

 ISBN: 978-1-61208-460-2

Before we describe the organization of the rest of the paper, a few observations regarding our chosen network-level protection measures in the context of virtualized infrastructure are in order. In a virtualized infrastructure, the distinguishing networking environment is the virtual network. Hence the network segmentation approaches discussed in this paper have to involve some virtual network components such as virtual switches. Similarly, a viable approach for traffic control using firewalls has to use a virtual firewall instead of a physical firewall. In Section II, we focus on two network segmentation approaches and discuss the advantages and disadvantages of each. Control of virtual network traffic using two different types of virtual firewalls and the advantages and disadvantages of each are analyzed in Section III. In Section IV, we provide the overall conclusions from the analysis, in the form of deployment choices based on approaches for the two network-level protection measures for securing VMs.

II. NETWORK SEGMENTATION IN VIRTUAL NETWORKS

The approaches to network segmentation in the context of virtualized infrastructures are the same as those used in physical network with some variations (these variations are underlined in our description below): (a) Using a combination of firewalls – the firewalls used are virtual firewalls (as opposed to a physical firewall) and are implemented as Virtual Security Appliance (VSA) and hosted on security hardened VMs with multiple Virtual Network Interface cards (vNICs). Each vNIC may be connected to a different network segment [2]. (b) Isolated network segments created as logical segments on top of a physical network segment – one is the VLAN approach that is based on tagging packets and switch ports with a unique identifier called VLAN ID, and the other is overlay-based virtual networking technology that creates an overlay network by encapsulating packets with IDs depending upon the type of overlay technology. Both approaches (VLAN and Overlay) rely on the capabilities in virtual switches of the virtualized host.

The three approaches for network segmentation in virtualized infrastructures outlined above are discussed in Sections A, B and C below. After a brief description of each approach, an analysis of each approach is provided with the relative advantages and disadvantages. Where ever applicable, the distinct advantage of a particular approach is also brought out.

A. Network Segmentation Using a Combination of Firewalls

Let us now consider a virtualized host with 4 VMs – VM1, VM2, VM3 & VM4. We can form a network segment (say a DMZ) using two virtual firewalls, one each on any two VMs - say VM1 & VM4. These firewalls are VM-based Virtual Security Appliances residing on VMs defined with multiple vNICs – each one connected to a different network segment. The firewall in VM1 then will have one vNIC connected to an external network (say the public Internet) of the enterprise and the other vNIC connected to the DMZ segment in the virtualized network within a virtualized host. Correspondingly the firewall in VM4 has to have one vNIC connected to the internal network of the enterprise and the other vNIC connected to the DMZ. The vNIC connection to the DMZ (from both firewall VMs - VM1 & VM4) is established by their pathway to an internal-only virtual switch. This internal-only virtual switch has no uplink connection to any physical NIC of the virtualized host and hence traffic from any VM connected to it cannot travel directly outside the virtual network segment (not to speak of outside the virtualized host). The internal-only switch can only forward traffic directly to VMs connected to it - say VM2 &, VM3. All incoming and outgoing traffic into the VMs connected to the internal-only virtual switch whose source/target is an internal/external network, has to go through the firewall in VM1 or in VM4. The firewalls in VM1 & VM4 thus form the gatekeepers for the virtual network segment (i.e., DMZ).

A.1 Analysis

The advantages of network segmentation within the virtualized network of a virtualized host using a combination of virtual firewalls are: (a) Simplicity of Configuration: It can be configured with commodity firewall VSAs hosted on multi-vNIC VMs. (b) Flexibility within a Virtualized host: More than one isolated network segment can be created within the virtual network of the virtualized host by simply adding another firewall VM.

The limitations of this approach to network segmentation in a virtualized network are the following: (a) The logical network segment created inside a virtualized host can neither be extended to the physical network of the data center nor to the virtual network in another virtualized host (since segmentation is obtained by virtual firewalls inside the virtualized host). This makes the approach to network segmentation non-scalable. (b) A consequence of creating a non-scalable network segment is that the migration of a VM in the network segment to any other virtualized host (due to

 ISBN: 978-1-61208-460-2

performance or availability or load balancing reasons) is out of the question, unless a network segment (with identical configuration) exists on the target virtualized host.

B. Network Segmentation Using Virtual LAN (VLAN) Technology

The second approach to network segmentation in virtualized infrastructures is broadcast-containment networking technologies, such as VLAN. The requirement for this is that the hypervisor should have capabilities to define virtual switches that are VLAN-aware [3][4]. The segmentation is obtained by assignment of an identifier called the VLAN ID to one or more ports of a switch and connecting the VMs designated for that VLAN segment to those ports. VMs on one VLAN can only communicate directly with VMs on the same VLAN and a router is needed for communication between VMs on different VLANs [5]. Assignment of a VM to a particular VLAN can be based on the application tier it is hosting (e.g., Web Server, DBMS server, etc.) or the client to which the VM belongs (in cloud data centers). These VLAN-capable virtual switches (VS) can perform tagging of all packets going out of a VM with a VLAN tag (depending upon which port it has received the packet from) and can route an incoming packet with a specific VLAN tag and MAC address to the appropriate VM by sending it through a port whose VLAN ID assignment equals the VLAN tag of the packet. An example of a VLAN-based virtual network segmentation inside a virtualized host is given in Figure 1.

B.1 Analysis

The advantages of a VLAN-based network segmentation approach are: (a) Network segments can extend beyond a single virtualized host (unlike the segment defined using virtual firewalls) since the same VLAN ID can be assigned to ports of virtual switches in different virtualized hosts. (b) The number of network segments that can be configured is reasonably large since a single virtual switch can typically support 64 ports and the 12-bit VLAN ID address space enables creation of 4000 VLAN segments.

The disadvantages of VLAN-based network segmentation approach are: (a) The configuration of the ports in the physical switch attached to a virtualized host must exactly match the VLANs defined on the virtual switches inside that virtualized hosts. This results in tight coupling between virtual network and physical network of the data center. The consequence of this tight coupling is that the port configuration of the physical switches has to be frequently updated since the VLAN configuration of the virtual network of the attached virtualized host may frequently change due to migration of VMs among VLANs as well as among virtualized hosts. (b) Another consequence of frequent migration of VMs among VLANs, as well as among virtualized hosts is that the VLAN configuration of ports in the physical switch may not match with that of the connected virtualized host. This may result in a situation where a particular hypervisor (or a virtualized host) may end up processing messages for every VLAN on the network, even when it is not hosting any active VM belonging to that VLAN [6] and (c) Segments created using broadcast-containment technologies cannot be allowed to have a large span since they will result in greater traffic in the overall data center due to multicast and broadcast traffic. But greater VM mobility (due to load balancing and availability reasons) may require VLANs with a large span resulting in an undesirable phenomena called VLAN sprawl [6].

C. Network Segmentation Using Overlay-based Virtual Networking Technology

In the VLAN-based approach, the logical network segments were created on a physical LAN using portgroups of virtual switches inside virtualized hosts. These logical network segments did span multiple virtualized hosts. The total number of these segments possible is limited to around 4000 due to the 12 bit address space of VLAN ID. Another limitation is the lack of independence between the physical and virtual networking infrastructure, since the port configuration of the physical switches attached to the virtualized hosts have to be consistent with the VLANs defined on the port groups of virtual switches inside those virtualized hosts. The overlay-based virtual networking approach to network segmentation overcomes these two limitations in the following two ways [7].

 ISBN: 978-1-61208-460-2

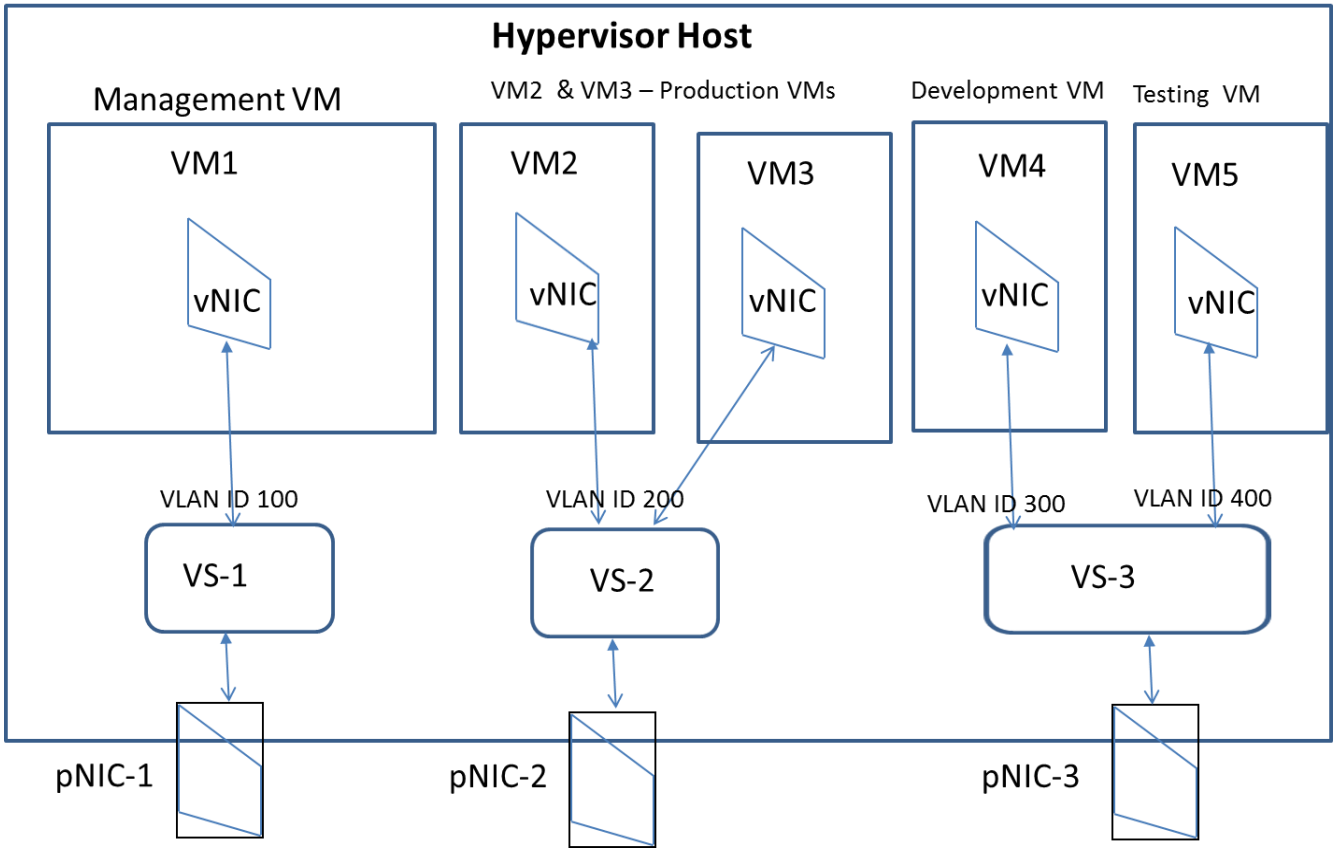

Figure 1 - VLAN-based Network Segmentation

(a) Overlay-based virtual networking technologies have a larger address space enabling larger number of virtual network segments. An example is the 24 bit VXLAN ID of the VXLAN overlay scheme that can enable 16 million virtual network segments to be defined, and (b) The overlay schemes by their definition, create a logical Layer 2 network (called the overlay network) over the physical Layer 3 backbone of the data center (called the underlay network). Since this scheme does not warrant any modifications to the physical network, it provides physical-logical network independence. As already alluded to, overlay-based virtual networking schemes achieve segmentation by creating a logical Layer 2 network over the physical Layer 3 network. The overlay network is created by encapsulating a native Layer 2 packet with another Layer 2 identifier. There are three common encapsulation schemes – VXLAN, GRE and STT [8].

Let us now look at the encapsulation process in VXLAN [9] through components shown in Figure 2. The Ethernet frame originating from a VM, that just holds the MAC address of the destination VM is encapsulated in two stages: (a) First with the 24 bit VXLAN ID (virtual Layer 2 (L2) segment) to which the sending/receiving VM belongs and (b) Second, with the source/destination IP address of the VXLAN tunnel endpoints called VTEPs. [10]. VTEPs are logical network endpoints (nodes) for the encapsulated VXLAN packets and they reside in the kernel of a hypervisor. A VXLAN-encapsulated packet originates at the VTEP in the kernel of the hypervisor where the source VM resides (carrying the VTEP's address as the source IP address) and terminates at the VTEP in the kernel of the hypervisor where the destination VM resides (carrying this VTEP's address as the destination IP address). Thus, we see that VXLAN encapsulation enables creation of a virtual Layer 2 segment that can span not only different hypervisor hosts but also different IP subnets within the data center.

C.1 Analysis

The advantages of a network segmentation approach based on Overlay-based networking technology are: (a)

 ISBN: 978-1-61208-460-2

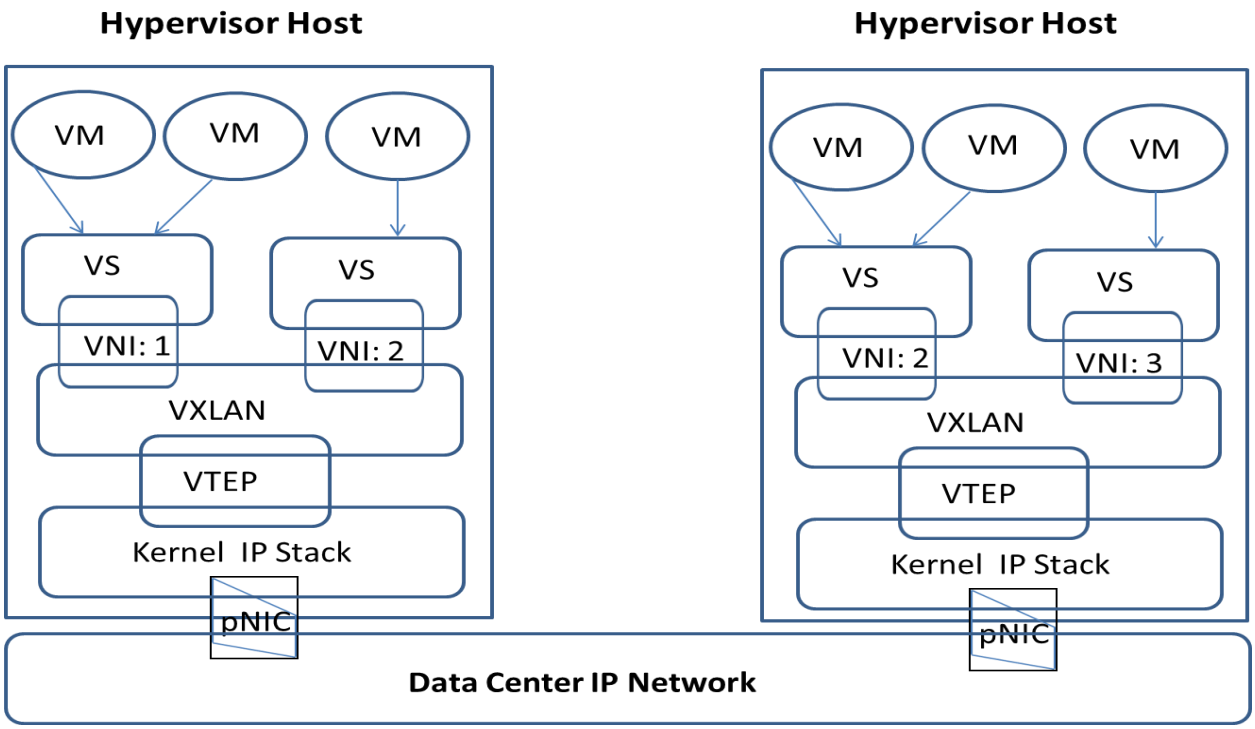

Figure 2 – Overlay-based Virtual Network Segmentation

Because of independence between the virtual network and the physical network, there is greater VM mobility compared to a VLAN- based virtual network environment, (b) The physical-logical network independence, together with the bigger overlay segment ID address space (e.g., a VXLAN ID is 24 bits as opposed to 12 bit VLAN ID allowing for 16 million segments compared to 4096 for VLAN), makes the overlay based network segmentation infinitely scalable. Another factor contributing to scalability of the overlay scheme is that the encapsulating frame is an IP/UDP packet. Hence, the number of virtual network segments is limited only by the size of IP subnets that can be defined within the data center and not by the number of ports in virtual switches as in the case of VLAN-based network segmentation. Further, by using internal, non-routable IP addresses for VMs (using DHCP and NAT capabilities) running within virtualized hosts, the number of virtual networks that can be realized is even higher and (c) The VLAN scheme uses the Spanning Tree Routing Protocol to forward packets, VXLANs can use the ECMP protocol of Layer 3 [11], thus efficiently utilizing all available network links in the network fabric of the data center.

The disadvantage of a network segmentation approach based on Overlay-based networking technology is that it requires large mapping tables in each virtual switch level in order to generate overlay packets – since the MAC address of the destination VM could be located in any IP subnet and any host in the data center. Building these mapping tables using just a flooding technique is inefficient. Hence, a control plane needs to be deployed in the virtualized infrastructure to populate the mapping tables for use by the overlay packet generation module in the hypervisor. This creates an additional layer of control and adds to the complexity of network management [11].

III. TRAFFIC CONTROL IN A VIRTUAL NETWORK

Traffic control in a virtual network can be performed using either a virtual firewall or a physical firewall. However, in a virtualized infrastructure, the computing nodes (whose incoming/outgoing traffic needs to be controlled) are VMs and are end nodes of a virtual network. Hence, the deployment of a physical firewall will require the traffic from the virtual network to be diverted into the physical network (where the physical firewall resides) and

 ISBN: 978-1-61208-460-2

then back into the virtual network. This extra route travelled by communication packets will result in latency and consequently reduced performance of applications hosted on VMs. Hence, in this paper, we consider only virtual firewall-based solutions for traffic control in virtual networks for securing VMs.

Earlier in this paper (Section 2), we saw that two or more virtual firewalls can be used to create network segments in a virtual network. Since the focus of this Section is on traffic control, we are going to analyze the use of virtual firewalls only for controlling inter-VM traffic.

Inter-VM traffic can be of two kinds: the traffic between two VMs residing on the same virtualized host (either connected to the same or different virtual switches) and traffic between two VMs hosted on different virtualized hosts. Traffic between two VMs residing on the same virtualized host can only be enabled if each VM has at least one vNIC (a VM can have multiple vNICs just like a physical server can have multiple physical network interface cards or network adapters) connected to a common virtual switch. This is due to the fact that two virtual switches in a virtualized host cannot be connected to each other. Although, theoretically, a pathway between two VMs on the same virtualized host can be establishing by routing the traffic from one VM (say VM1) to the physical network (through one physical NIC) and back into the same virtualized host (through another physical NIC) to the target VM (say VM2), this is not a viable option in most situations, due to the latency issue referred to above. Of course, for enabling traffic between two VMs residing on two different virtualized hosts, the traffic has to travel from the virtual network (in the originating virtualized host) where the originating VM resides, through the physical network of the data center and back again into the virtual network of the target VM (in the target virtualized host).

Virtual firewalls come in two flavors: (a) VM-based – this class of virtual firewalls, comes packaged as a virtual security appliance on a specially-configured VM and (b) Hypervisor Kernel-based – this class of firewalls operates as a kernel loadable module in the kernel of the hypervisor.

A. Traffic Control using VM-based Firewalls

A VM-based firewall, as the name indicates, is a firewall software that runs in a VM. It can be installed as a software module on a guest VM already running in a virtualized host or it can be packaged as a virtual security appliance on a specially prepared VM instance. Its location within the virtual network of a virtualized host is critical as its function is to monitor, drop or forward packets between sets of VMs belonging to different security zones. This is the reason that this class of firewalls is called bridge-mode firewalls as it also acts as a bridge between zones (since the only link between VMs connected to two different virtual switches is through a VM with vNICs connected to both virtual switches).

A.1 Analysis

The advantage of VM-based firewall for traffic control is that since it is available as a Virtual Security Appliance, it is easy to deploy and configure in a virtualized host. Its initial location within the virtual network of the virtualized host is relatively easy as it is dictated by the layout of the security zones based on the various virtual switches and this type of firewall only monitors and filters packet flows between one virtual switch and another.

The disadvantages of VM-based firewalls are: (a) It cannot monitor and filter traffic flowing between two VMs connected to the same virtual switch, (b) Its performance is limited by the number of virtual CPU cores allocated to the VM it is residing or packaged in, and (c) All traffic flowing into and out of all portgroups and virtual switches associated with zones pertaining to this firewall, has to be redirected to this firewall causing unnecessary traffic (a phenomena called Traffic Trambones [12]).

B. Traffic Control using Hypervisor Kernel-based Firewalls

Hypervisor kernel-based firewalls are also called hypervisor-mode firewalls and VM NIC firewalls. These firewalls install as a hypervisor module along with a VSA, the latter used purely for initial configuration (and re-configuration) for the hypervisor module. Hence, the main firewall functions of monitoring and packet filtering are done in the hypervisor kernel-module with the VM hosting the VSA portion playing the role of a Control or Service VM. Logically residing between a VM vNIC and the hypervisor virtual switch, this type of firewall provides a vNIC-level firewall enforcement point for traffic to and from VMs. Thus they can be used for selectively protecting a given subset or all the VMs in a host or a cluster. Because of the visibility at the vNIC level, these firewalls can protect traffic flowing between two VMs connected to the same virtual switch, unlike the bridge-mode firewalls. Another distinguishing feature of this type of firewalls is that the firewall does not require any changes

 ISBN: 978-1-61208-460-2

to the virtual network configuration inside a virtualized host (such as additional network pathways for redirecting traffic to the VM hosting firewall) or modification to IP addresses of VMs.

B.1 Analysis

The advantages of hypervisor kernel-based firewalls are: (a) Their performance is of an order of magnitude much higher than bridge-mode firewalls since they perform packet inspection from within the kernel at native hardware speeds rather in a VM where the performance is limited by the capacity of virtual CPU allocated to it [13], (b) These perform monitoring at the VM NIC (vNIC) level and hence all firewall rules (or ACLs) and state are logically attached to the VM interface. Hence these rules and state move with the VM when it migrates from one virtualized host to another, thus providing continuity of security protection for a VM irrespective of its location [12], and (c) Implementations that support firewall rules at a higher level of abstraction than IP addresses or ports, can be used to filter packets at data center, host cluster and port group levels.

The disadvantages of hypervisor kernel-based firewalls are: (a) This class of firewalls works as a managed kernel process and is therefore neither a VM resident program nor is part of the virtual network of the hypervisor. Hence conventional management tools having access only to VMs or virtual networks cannot be used to monitor this class of firewalls and (b) Some of the implementations of this class of firewall support only 5-tuple based rules (Source and Destination IP Address, Source and Destination Ports and Protocol). They do not support higher level abstractions such as Security Groups, Zones or Containers. However, some of the latest offerings do support firewall rules based On higher level abstractions and flow statistics as well.

IV. SUMMARY & CONCLUSIONS

In this paper, we performed a detailed analysis of two network segmentation approaches and two virtual firewall types for the protection of VMs in virtualized infrastructures. Comparing the features of the two network segmentation techniques, we find that the only distinct advantage that overlay-based network segmentation (such as VXLAN) holds over the VLAN-based approach is its infinite scalability. Hence, unless the number of VMs in the data center is in the order of thousands, the VLAN-based approach provides a satisfactory outcome in terms of performance and for meeting the goal of securing VMs. Because of this, the VLAN approach is economically justifiable in many environments. Further justification comes from the fact that overlay-based network segmentation schemes require sophisticated virtual switches, large mapping tables and the overhead of creating a control plane using SDN controllers.

Analyzing the relative advantages and disadvantages of the two firewall types – VM-based and hypervisor kernel-based – we find that the hypervisor kernel-based firewall is superior to the VM-based one in terms of three features. They are: (a) Performance (executes in the hypervisor kernel instead of in a VM), (b) Reduced network traffic (no diversion of traffic needed from various switches to the VM hosting the firewall) and, (c) Firewall rules are associated with the VM interface (since it is placed between a VM NIC and the virtual switch) and move with VM. The third feature is very critical from the point of view of the security of the VM, since it provides continued protection to it even when it migrates to different virtualized hosts or host clusters within the data center, without any additional re-configuration. It is this overwhelming security assurance feature that makes the hypervisor kernel-based firewall, the security software of choice in many virtualized infrastructures.

REFERENCES

[1] D. Shackleford, Virtualization Security – Protecting Virtualized Environments, Wiley Publishing Inc, Indianapolis, IA, USA, 2013

[2] "DMZ Virtualization with VMware Infrastructure". [Online]. http://www.vmware.com/files/pdf/dmz_virtualization_vmware_infra_wp.pdf [retrieved: Jan, 2016].

[3] "MAC Bridges and Virtual Bridged LANs". [Online]. https://www.ietf.org/meeting/86/tutorials/86-IEEE-8021-Thaler.pdf [retrieved: Dec, 2015].

[4] "IEEE 802.1Q Virtual LANs (VLANs)". [Online]. http://www.ieee802.org/1/pages/802.1Q.html [retrieved: Dec, 2015].

[5] A. Hameed, and A. N. Mian, "Finding Efficient VLAN Topology for better broadcast containment," Third International Conference on the Network of the Future (NOF), Gammarth, Nov 2012, pp.1-6.

[6] Introduction to Virtualized Networking. [Online]. http://www.ipspace.net/Introduction_to_Virtualized_Networking [retrieved: Dec, 2015].

[7] Overlay Virtual Networking. [Online]. http://www.ipspace.net/Overlay_Virtual_Networking [retrieved: Dec, 2015].

 ISBN: 978-1-61208-460-2

[8] “Overlay Virtual Networking and SDDC”. [Online]. http://my.ipspace.net/bin/list?id=xSDNOverlay [retrieved: Jan, 2016].
[9] “Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks”. [Online]. https://tools.ietf.org/html/rfc7348 [retrieved: Jan, 2016].
[10] “VXLAN Overview: Cisco Nexus 9000 Series Switches”. [Online]. http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/white-paper-c11-729383.pdf [retrieved: Dec, 2015].
[11] “Scaling Overlay Virtual Networks”. [Online]. http://content.ipspace.net/get/Scaling%20Overlay%20Virtual%20Networks.pdf [retrieved: Jan, 2016].
[12] Virtual Firewalls. [Online]. http://www.ipspace.net/Virtual_Firewalls [retrieved: Dec, 2015].
[13] Virtual Firewall. [Online]. https://en.wikipedia.org/wiki/Virtual_firewall [retrieved: Dec, 2015].

 ISBN: 978-1-61208-460-2