

Security Control Variations Between In-house and Cloud-based Virtualized Infrastructures

Ramaswamy Chandramouli

*Computer Security Division, Information Technology Laboratory
National Institute of Standards & Technology
Gaithersburg, MD, USA
mouli@nist.gov*

Abstract-Virtualization-related components (such as Hypervisor, Virtual Network and Virtual Machines (VMs)) in a virtualized data center infrastructure need effective security controls. However, the differences in scope of control (among stakeholders) over this component set between in-house and cloud-based virtualized infrastructures introduce variations in security control measures that can be deployed by the respective stakeholders. In this paper, we analyze those variations and their efficiency and security impacts. We also suggest technology enablers that can minimize those impacts and improve the overall security robustness of the basic computing units of a virtualized infrastructure, (i.e., VMs).

Keywords-Virtual Machine; Virtual Network; Hypervisor; Virtualized Host; Cloud Service Model

I. INTRODUCTION

Server Virtualization, in some instances augmented with storage virtualization, is becoming the trend for data center infrastructures in both enterprises and cloud provider environments. In fact, Virtualized Servers and Virtualized storage have become the platform for many enterprise applications such as Enterprise Resource Planning (ERP) [1] because of the following:

(a) Efficiency in the utilization of processor and memory resources in a virtualized host because of the ability to run multiple Virtual Machines (VMs) as opposed to a non-virtualized host where only a single O/S stack can be run. Similar efficiencies can be achieved in the case of virtualized storage because of the presence of an abstraction layer above the physical storage layer, (e.g., disk arrays).

(b) Scalability and Elasticity that are enabled in Virtual Servers and Virtual Storage by the very nature of virtualization. An example is the capability to add VMs at will to a physical host with underutilized capacity and ability to add disk arrays transparent to the programs that gather, store, retrieve and process data.

There is general agreement in the security community that the security control measures used for protecting servers that run a single O/S stack (referred to as non-virtualized hosts) alone are not sufficient for protecting servers that run multiple O/S stacks (referred to as virtualized hosts). The reason for the agreement is the presence of a single trusted layer, (i.e., the hypervisor) below multiple VMs in virtualized hosts and the risk of compromise to this layer posing the risk of compromising the integrity of all VMs running in that host [2]. The detailed differences between virtualized hosts and non-virtualized hosts are given below:

(a) In a non-virtualized host, the interface to the hardware is through a regular O/S, whereas in a virtualized host, the interface to the hardware is through a software module, called a hypervisor, which contains just the kernel of an O/S with some necessary additions such as device drivers, etc.

(b) A virtualized host has resident in it multiple Virtual Machines (VMs), each with its own stack of O/S and Applications. All of the VMs share the same physical resources provided by the virtualized host – such as the processor, memory and directly attached storage. The hypervisor mediates access to shared resources by the various VMs, and provides isolation between the VMs.

(c) To enable VMs to communicate to the physical network and to provide isolation among them, a Virtual network is defined within each virtualized host. A Virtual network can be looked upon as a set of logical (sub)networks within a shared physical network. A virtual network can be configured using a combination of software-defined communication interfaces called virtual network interfaces (or vNICs) inside a VM as well as software-based switches, called Virtual Switches, that can be defined within the hypervisor.

In a virtualized enterprise data center, catering to internal information technology (IT) processing needs of an

enterprise (henceforth referred to as in-house virtualized infrastructure), all the components of a virtualized host are owned and controlled by the single entity, i.e., the enterprise. However, in the virtualized data center owned and operated by cloud service providers (henceforth referred to as cloud-based virtual infrastructure), while the virtualized host (the physical machine) and the software that provides the virtualization, (i.e., the hypervisor) are owned by the cloud service provider, the VMs in it are created and operated by the cloud service consumer. Hence, the internal configuration of VMs in a cloud-based virtual infrastructure belonging to an Infrastructure as a Service (IaaS) cloud provider is under the control of the cloud service consumer although the capabilities to configure a virtual network linking these VMs will still rest with the cloud service provider. Thus, we see that there are differences in the scope of control over the components of a virtualized infrastructure between in-house and cloud-based virtualized environments.

The main objective of this paper is to illustrate the variations in security control measures between the in-house and cloud-based virtualized infrastructures that these scope of control differences introduce. A second objective, or rather a by-product of the illustration process is to show the impact of these variations (in security control measures) on the effectiveness and efficiency of the total set of security controls and how they can be addressed to improve the overall security robustness of the basic computing units of a virtualized infrastructure, (i.e., the VMs).

The organization of this paper is as follows. In Section II, we identify the differences in scope of control among stakeholders over the components of a virtualized infrastructure by looking at the layers of a cloud service architectural stack. Section III identifies threat scenarios relating to virtualization-related components. In that section we also look at the broad class of security control measures and identify whether these control measures may be affected by differences in scope of control between in-house and cloud-based virtualized infrastructures. Section IV describes in detail the security control variations for VM protection at the Virtual Network layer and VM end-point. Section V provides the summary and conclusions.

II. SCOPE OF CONTROL IN CLOUD-BASED INFRASTRUCTURES

In the case of in-house virtualized infrastructures, since all components are owned and operated by a single entity, i.e., the enterprise, differences in scope of control over the

overall set of components do not arise. Hence we limit the scope of control analysis only to cloud-based virtualized infrastructures. In order to do that we digress a bit and look at the broad picture of cloud service models. The three widely accepted cloud service models are [3]: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS), all of them consisting of two major players – the cloud provider and the cloud consumer. In the case of SaaS, the cloud provider makes available a software application while PaaS offers a set of development tools (or an application development environment such as J2EE [4] or .NET [4]) to develop and possibly host applications. In these two service offerings, (i.e., SaaS and PaaS) the underlying data center infrastructure used (or leased from another provider) by the corresponding category of cloud service providers need not be based on virtualized servers. In the case of IaaS, what is made available to the cloud consumer are computing units in the form of VMs. Hence, the data center infrastructure in the case of IaaS cloud service providers should consist of only virtualized hosts. However, in this paper, we assume that the data center infrastructure of a cloud service provider irrespective of the cloud service model (because of efficiency, scalability and elasticity considerations) is a virtualized one consisting of virtualized hosts, virtual network and VMs. Based on this assumption, we can start taking a look at the various layers in a cloud service architectural stack. One such architectural stack based on slight variations from the model given by the Cloud Security Alliance [5] is given in Figure 1 below. In this stack, we notice that the facility, networking infrastructure and the physical host layers are common to all IT infrastructures - whether virtualized or not- and hence these layers are not relevant for our scope of control analysis. Going up one more layer in the stack, we find that it is in the resource abstraction layer that the main engine providing the virtualization, (i.e., the hypervisor) and virtual network are defined. Virtual Machines - the main computing units of a virtualized infrastructure reside in the VM layer. Our focus of attention for identifying the differences in scope of control between in-house and cloud-based environments is limited to these two layers. This is due to the fact that these are the two layers whose composition differs between virtualized and non-virtualized infrastructures. We give below our observations on the scope of control among stakeholders over components in cloud-based virtualized infrastructures used in all three cloud service-models.

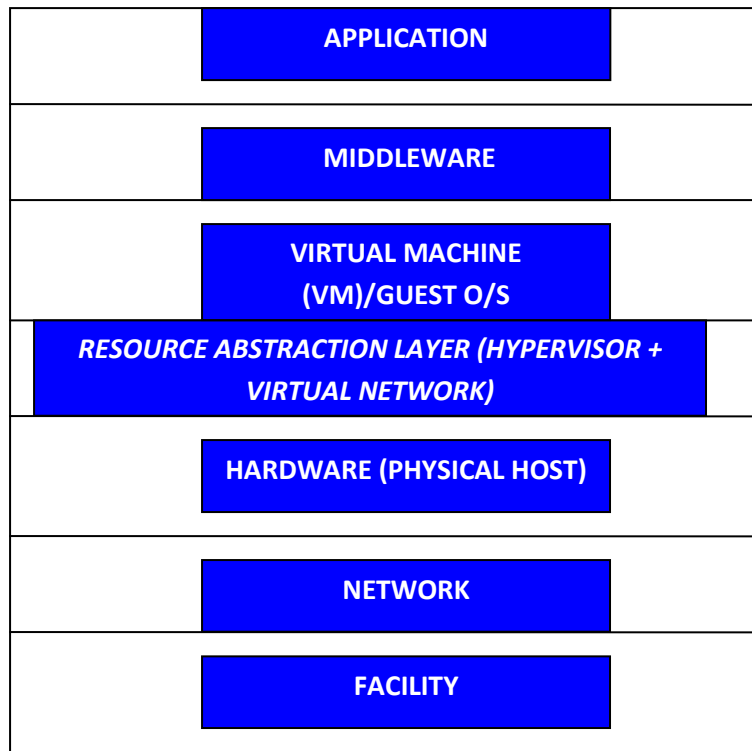


Figure 1. Cloud Service Layers.

(a) In the resource abstraction layer (virtualization layer), all components - the hypervisor and the virtual network (in all three cloud service models) - are totally under the control of only one entity, (i.e., the cloud provider).

(b) In the case of VM layer, the single component it holds - the VM instance with its embedded Guest Operating System - is under the control of cloud service provider in the case of SaaS and PaaS service models but controlled by the cloud service consumer in the IaaS model (mainly due to the fact that the SaaS provider does not want to assume any administrative responsibility for it after a VM instance is configured and instantiated by an IaaS consumer).

Although not shown in the architecture diagram, conceptually, one can think of a data layer which contains the components for the management software and the physical artifacts for storing the data that applications generate and use. This layer is entirely under the control of a single entity, i.e., the cloud provider - in all three cloud service models and hence is not relevant for our scope of control analysis. It is worth mentioning that although all storage-related technologies (both virtual and physical) are under the control of cloud service providers, the

responsibility for appropriate security control measures for data protection (through encryption of data in transit and data at rest - for the portion of data generated and used) still rests with the cloud consumer [6].

III. THREAT SCENARIOS & SECURITY CONTROL MEASURES FOR ENTIRE VIRTUALIZED INFRASTRUCTURES

Our scope of control analysis narrowed our focus to just two layers that contain all the artifacts relating to virtualization, i.e., the Resource Abstraction layer and the VM layer. Hence, our threat analysis is also limited to scenarios involving the components contained in these layers. These components are re-listed here for facilitating further discussion:

- (a) Hypervisor and Virtual Network - from the Resource Abstraction Layer
- (b) Virtual Machines (VMs) with their Guest O/S - from the VM Layer

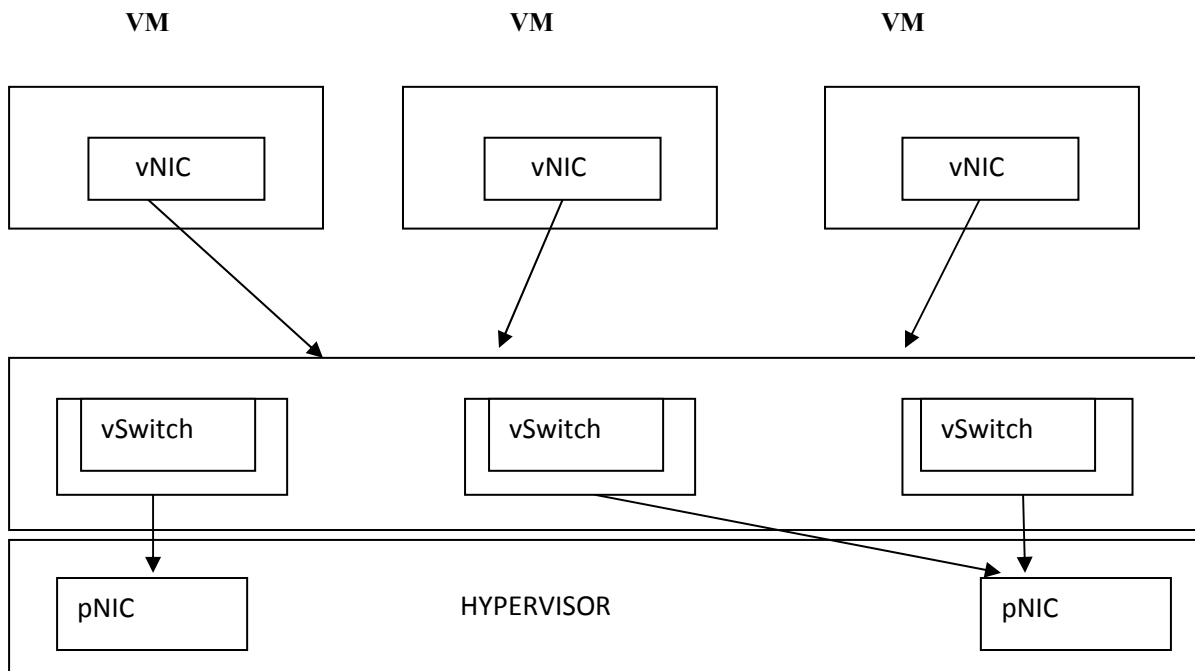


Figure 2. A Virtual Network Configuration.

The virtualization environment involving the above components can be described as follows. The VMs provide the complete, encapsulated computing stacks built up of Guest O/S with middleware and applications riding on top of the chosen O/S. A network linking these VMs, thus enabling communication among them is called a Virtual Network. An example of a virtual network configured using software-defined virtual network interface card (vNIC) within each VM and software-defined virtual switches (vSwitch) defined within the hypervisor is shown in Figure 2. A virtual network provides communication not only among VMs residing on a single virtualized host but also connectivity with the outside world (the physical network), if any of the virtual switches is also connected to a physical network interface card (pNIC) of the virtualized host. In this virtualization environment, the most common (though not exhaustive) threat scenarios are identified below:

TS-1: The compromise of the hypervisor (by exploiting the vulnerabilities in the kernel - which is rare) can potentially compromise the security of multiple VMs (in fact potentially all) resident on that virtualized host [7].

TS-2: A single VM that has been compromised can be used as a launching pad for attacking other VMs (especially if they share some common resources such as memory or

there exists a communication channel between them due to the fact that these VMs constitute the different tiers (webserver, application server, database server etc) of a multi-tier application) [8].

TS-3: Normally, the effect of all operations performed within a VM must be restricted to that VM. However, under some circumstances, a malicious program running in a VM, by exploiting vulnerabilities in the hypervisor software, can alter the state of other VMs, the hypervisor itself or even the hardware. Such a VM is called a rogue VM and it poses the threat of subverting the isolation property required to be provided by the hypervisor.

TS-4: Attacks on running guest VMs by a malicious hypervisor. The impact of this threat is the same as that of TS-3 but the threat source here is the hypervisor as opposed to a rogue VM in the case of TS-3.

For each of the threat scenarios, let us see the broad class of security control measures that are required and identify whether these control measures may be affected by differences in scope of control among stakeholders between in-house and cloud-based virtualized infrastructures.

TS-1 has to do with the compromise of the hypervisor. The usual security control measure adopted is to keep the

patches from the hypervisor vendor up to date. This is a hypervisor-based security control measure and since the hypervisor is always under the control of the single entity in both in-house and cloud-based virtualization infrastructure, the potential for variation in security control measures due to differences of who has control does not arise.

TS-2, TS-3 & TS-4 have to do with compromise of the VM. Security control measures can be provided for VMs through a combination of: (a) Virtual Network-resident security controls and (b) VM-resident security controls [9]. In the case of in-house and SaaS/PaaS cloud-based virtualized infrastructures, both the Virtual Network and the VMs are under the control of a single entity. However, in the case of IaaS cloud-based virtualized infrastructure, the VMs to be protected are under the control of IaaS cloud consumer. This is where there is potential for variations in control measures for protection of VMs between in-house and cloud-based virtualized infrastructures. These variations are analyzed and discussed in the next section.

IV. SECURITY CONTROL VARIATIONS FOR VM PROTECTION

Security controls for VM protection can be deployed both at the Virtual Network layer as well as in the VMs themselves. As we have already seen, both these layers are accessible to a single entity only in the case of in-house and SaaS/PaaS cloud-based virtualized infrastructures. However, only the Virtual Network layer is accessible for IaaS cloud provider and the VMs are accessible for IaaS cloud consumer. Hence variations in security controls for protection of VMs are introduced due to these differences in scope of control among stakeholders as discussed below:

A. Security Control Variations at the Virtual Network Layer

In virtualized infrastructures for in-house IT as well as for providing cloud services (all three service models), the configuration of a Virtual Network (the network linking all virtual machines within a single virtualized host) is entirely under the control of the data center owner/operator. Leveraging the virtual network, there are two approaches to providing security for VMs. They are [10]:

(a) Extending the concept of Virtual LAN (VLANs) into the virtual network and

(b) Virtual Network Configuration-based solutions for VM protection

In the VLAN-based approach, the virtual switches defined in a hypervisor are made to recognize the VLAN tags and thus the concept of network isolation in the physical network is extended to the virtual network inside a virtualized host. As far as Virtual Network Configuration-based solutions go, there are two types: In the first approach, VMs hosting sensitive resources such as fileshare VM or hosting sensitive applications such as data warehousing or payroll processing are connected to isolated virtual network segments which are not behind any firewall or Network Address Translation devices. In the second approach, a special-purpose VM called a Virtual Security Appliance [11] that contains a hardened Guest O/S and one or more security applications, is installed and configured to provide the necessary protection for VMs. The type of protection depends upon the security application(s) that is (are) packaged as part of the Virtual Security Appliance. Examples of popular security applications are Firewall and Intrusion Prevention [12]. In a firewall solution, the data center operator can provide protection to VMs by defining VM-specific rules that can control traffic to and from virtual machines. There are tools in the market place [11], which, using a combination of a management server and a set of security appliances, can control traffic in and out of an arbitrary set of VMs irrespective of the VLANs to which they belong. By placing the entire set of virtual machines to be protected on an internal-only virtual switch of the virtual network, all traffic is made to flow through the security appliance and thus the security appliance can act as a Layer 2 bridge that controls all traffic flowing to and from the protected VMs without reconfiguring them in any way. The consequence of this is that firewall rules encompassing layers 2, 3 & 4, (i.e., including IP addresses and specific TCP or UDP port) can all be defined using this virtual network-based security appliance. However, in the case of an IaaS Cloud consumer, who owns and operates a set of VMs, the virtual network layer is not under his/her control. Hence a virtual network-based firewall cannot be deployed in this situation and this class of user can only provide the necessary traffic control by having a VM-based firewall. Deploying a VM-based firewall solution imposes a great deal of performance overhead compared to a virtual network-based firewall as this security application competes for the same resources, (i.e., CPU, Memory, etc.) as functional (business) applications do on each of the VMs. For example, if there are 10 VMs in a virtualized host, 10 firewall applications will be competing for its

resources. On the other hand, providing the same firewall functionality using a Virtual Security Appliance will involve running a single instance of a firewall application instead of ten (in the case of VM-based solution), thus providing a significant performance improvement while accomplishing the same security objective of blocking unwanted/malicious traffic in and out of the VMs that the cloud service consumer has rented and is operating.

B. Security Control Variations at the VM end-point

The VMs (often called the endpoints) in a virtualized infrastructure need to have the same protection measures as a physical server in a non-virtualized environment. These protection measures include but not limited to [8]:

- Anti Virus/Anti Malware Software
- Intrusion Prevention

In the case of a virtualized infrastructure providing IaaS cloud service, the VMs are owned and operated by cloud consumers and hence the security of these VMs rest with them. Because of the fact that the hypervisor controlling these VMs is owned and under the control of the IaaS cloud provider, the only security control measure available to the cloud consumer is to run individual instances of above classes of software (anti-virus, etc.) in each of the VMs rented and operated by them.

In the case of virtualized infrastructures providing in-house IT needs or providing SaaS or PaaS cloud services, both the hypervisor and the VMs are owned and operated by a single entity - data center owner or operator. Using the published interfaces (called the hypervisor introspection APIs) provided by the hypervisor vendor, the data center owner/operator can either develop (or procure from a third-party), a security application that can be installed as a security appliance on a special hardened VM. This security appliance thus runs as a single instance of security software and this instance would be running separate from all other instances of Guest O/Ss running in the various VMs that it would protect. For example, a security appliance for an antivirus solution can perform the functions of - memory scanning, monitoring of processes and investigation of network traffic - for all VMs and Guest O/Ss running on that virtualized host. The variations in security control measures that can be deployed by the stakeholder between the two virtualized infrastructure environments, (i.e., IaaS cloud Versus SaaS/PaaS cloud) for VM protection has the following efficiency and security implications:

(a) Multiple instances of say (an anti-virus solution) in a single physical (virtualized) host makes of the order of magnitude huge demands on the processor and memory cycles of that host as opposed to a single instance of security software. Related to this is the management issue of keeping these security solutions in synch in all of the VMs in a virtualized host.

(b) A rogue process within a Guest O/S can potentially shut down the anti-virus solution running in that same Guest O/S. However, the single instance of such a solution running in a security appliance in a hardened VM that runs in the same virtualized host is not only able to thwart such attacks on itself, but is also able to provide protection to all other VMs running in that virtualized host.

V. CONCLUSION

In this paper, we saw that protection of VMs can be obtained through efficient and effective means in the case of in-house and SaaS/PaaS cloud-based virtualized infrastructures through the following:

- Implementing VLANs in the virtual network
- Creating isolated network segments for VMs running sensitive applications, and
- Virtual Security Appliances utilizing hypervisor introspection API

However, the protections provided by the above measures have to be obtained through a less efficient way by means of individual VM-based solutions in the case of IaaS cloud-based virtualized infrastructures. This is due to the fact that the IaaS cloud consumer who needs to protect VMs does not have access to components of Resource Abstraction layer such as the Hypervisor and the Virtual Network. Analysis of such variations in security control measures and their relative effectiveness and efficiency impacts can lead to exploration of ways to minimize those impacts.

Specifically, to address the effectiveness and efficiency gap between security controls available for stakeholders in the in-house and cloud-based infrastructures, we propose the following:

- (a) Provide selective visibility to cloud customers to VMs rented by them through hypervisor introspection API
- (b) Define pre-defined VLAN segments for cloud customers to place the VMs created/rented by them..

The justification for the above measures to improve the overall security robustness of VMs can be made based on the following observations:

(a) Any security/monitoring solution based on virtual network can significantly reduce the demand on the CPU cycles of the virtualized host compared to a host-based (i.e., VM-based) solution

(b) The vulnerability of the Guest O/S itself is not a factor with respect to the integrity of the security solution

as these types of solutions that are based on visibility into the virtual network are run on dedicated VMs with hardened Guest O/Ss.

REFERENCES

- [1] M.Hoyer, K.Schröder, Daniel Schlitt, and D. Wolfgang Nebel, "Proactive Dynamic Resource Management in Virtualized Data Centers", ACM Conference on e-Energy, New York, NY, USA, May 2011.
- [2] T. Garfinkel and M. Rosenblum, "When Virtual is harder than Real: Security Challenges in Virtual Machine Based Computing Environments", Stanford University Department of Computer Science. <http://www.stanford.edu/~talg/papers/HOTOS05/virtual-harder-hotos05.pdf> [Retrieved: July, 2012]
- [3] P. Mell and T. Grance, "A NIST Definition of Cloud Computing," NIST SP 800-145, <http://csrc.nist.gov/publications/#SP-800-145>, [Retrieved: May, 2012]
- [4] D.M. Whittinghill, and K.D. Lutes, "Teaching Enterprise Application Development: strategies and challenges", ACM SIGITE' 11 Conference, West Point, NY, Oct 2011.
- [5] Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing, v2.1," www.cloudsecurityalliance.org/csaguide.pdf, pp. 61-64.
- [6] L.M. Kaufman, "Data Security in the world of cloud computing." IEEE Security and Privacy, Vol. 7, No. 4, 2009
- [7] S. Jin, J.Ahn, S.Cha, and J.Huh, "Architectural Support for Secure Virtualization under a Vulnerable Hypervisor," ACM MICRO'11 Conference, Porto Alegre, Brazil, Dec 2011, pp. 272-283.
- [8] J. Sahoo, S.Mohapatra, and R.Lath, "Virtualization: A Survey On Concepts, Taxonomy And Associated Security Issues," IEEE 2nd International Conference on Computer and Network Technology, Bangkok, Thailand, Apr 2010, pp. 222-226.
- [9] J. N. Matthews, 'et al.' "Running Xen – A Hands-On Guide to the Art of Virtualization," Prentice Hall, 2008
- [10] Five exciting VMware networking features in vSphere 5 <http://searchvmware.techtarget.com/tip/Five-exciting-VMware-networking-features-in-vSphere-5> [Retrieved: March, 2012]
- [11] S. Lowe "Mastering VMware vSphere 4," Wiley Publishing, 2009.
- [12] L. Garber, "The Challenges of Securing the Virtualized Environment", IEEE Computer, Volume 45 Issue 1, Jan 2012.