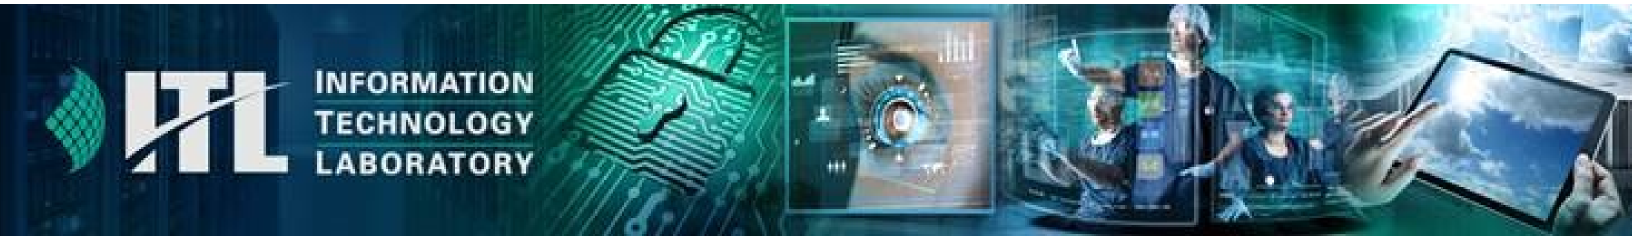

ITL BULLETIN FOR JUNE 2016

EXTENDING NETWORK SECURITY INTO VIRTUALIZED INFRASTRUCTURE

Ramaswamy Chandramouli, Larry Feldman,[1] and Greg Witte,[1] Editors
Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
U.S. Department of Commerce

Background

With the proliferation of the cloud computing model, and as the number of virtualized hosts in modern data centers continues to increase, organizations have a need to provide these hosts with the same level of security that has historically been provided to physical devices. A virtualized host (i.e., a physical host running a server virtualization product) is capable of supporting multiple computing stacks (called Virtual Machines or VMs), each with a different platform configuration (e.g., operating system [OS], middleware) and each with unique security needs.

Application programs loaded into a VM are often valuable server programs (e.g., webserver, database management system) that support important business processes and generally need more security protection than do other virtual hosts, such as workstations.

Introduction

To provide an analysis of various virtual network configuration options for protection of VMs and to present recommendations based on the analysis, NIST has released Special Publication (SP) 800-125B, *Secure Virtual Network Configuration for Virtual Machine (VM) Protection.* The publication discusses several relevant configuration areas including network segmentation, network path redundancy, traffic control through firewalls, and VM traffic monitoring. Each configuration option in each of these areas has different advantages and disadvantages, which are identified in SP 800-125B. Analysis of these advantages and disadvantages has led to the development of one or more security recommendations for each configuration area. A brief explanation of these configuration areas is provided below.

Network Segmentation Configurations

The main motivation for network segmentation is to achieve logical separation for applications with different sensitivity levels or belonging to different departments. The network segmentation approaches

[1] Larry Feldman and Greg Witte are Guest Researchers from G2, Inc.

discussed in this publication are organized by their increasing order of scalability. The initial approach is to host all applications of a given sensitivity level in one VM and host all VMs of the same sensitivity level (based on hosted applications) in a given virtualized host. This is not strictly a network segmentation approach, since it does not involve configuring a network parameter, but it is still included in the guidance because it meets the objective of providing VM protection. The publication also discusses approaches for creating virtual network segments inside a virtualized host using virtual switches and virtual firewalls.

Two truly scalable (data center-wide) approaches for creating virtual network segments that span multiple virtualized hosts are discussed in this guidance. These approaches are based on VLAN (virtual local area network) and overlay-based virtual networking technologies, respectively.

SP 800-125B provides advantages and disadvantages for each discussed network segmentation configuration approach, with associated security recommendations. Each recommendation has a unique identifier of format VM-NS-Rx, where VM stands for virtual machine, NS for network segmentation, and Rx for the recommendation sequence.

Network Path Redundancy Configurations

Because configuring multiple VM communication paths may be essential for ensuring required availability, network configuration for achieving this goal can be looked upon as an integral part of network-based protection for VMs. SP 800-125B discusses the establishment of this redundancy.

The physical network configuration in a data center is largely unaffected by the presence of virtualized hosts, except for some configuration tasks (e.g., VLAN configuration of ports in the physical switches connecting to the virtualized hosts, configuration of associated links as trunk links). The network path redundancy configuration options discussed in this publication are confined to those related to the virtual network inside the virtualized hosts, including their physical network interface cards (pNICs). Failover policy options, which are provided as a virtual network configuration feature in many hypervisor offerings, support network path redundancy.

Hypervisor offerings also provide a configuration feature called network interface card (NIC) teaming. NIC teaming allows administrators to combine multiple pNICs into a NIC team for NIC failover capabilities in a virtualized host. The members of the NIC team are connected to the different uplink ports of the same virtual switch. Failover capability requires at least two pNICs in the NIC team. One of them can be configured as “active” and the other as “standby.” If an active pNIC fails or traffic fails to flow through it, the traffic will start flowing (or be routed) through the standby pNIC, thus maintaining continuity of network traffic flow from all VMs connected to that virtual switch.

Different hypervisors offer different NIC teaming policy configuration options that may have an impact on failover. The publication analyzes different options pertaining to different ways in which the NIC team detects NIC/link failure and performs failover.

SP 800-125B provides recommendations to improve the fault tolerance (redundancy) already furnished by NIC teaming. Each recommendation has a unique identifier of format VM-NPR-Rx, where VM stands for virtual machine, NPR for network path redundancy, and Rx for the recommendation sequence.

VM Protection through Traffic Control using Firewalls

SP 800-125B describes the following two scenarios where traffic control for VM protection is to be exercised:

- Traffic flowing between any two virtual network segments (or subnets); and
- All traffic flowing into and out of a VM.

It comes as no surprise that control of both types of traffic (enumerated above) is enforced through firewalls. SP 800-125B provides the analysis of several use cases in relation to these scenarios and a brief overview of the three classes of firewalls, including physical firewalls, subnet-level virtual firewalls, and kernel-based virtual firewalls. As a result of this analysis, the publication summarizes advantages and disadvantages for each class of the firewalls, and security recommendations for firewall deployment architectures.

VM Traffic Monitoring

Firewalls only ensure that inter-VM traffic conforms to organizational information flow and security rules. However, to identify any malicious or harmful traffic coming into or flowing out of VMs and to generate alerts or take preventive action, it is necessary to set up traffic monitoring capabilities to monitor all incoming and outgoing traffic of a VM. This requires functionality—port mirroring—to send copies of those packets to a network monitoring application (also called an analyzer application). The purpose of a network monitoring application is to perform security analysis, network diagnostics, and network performance metrics generation. Configuration options are available in hypervisors to turn on port mirroring functionality. Depending upon the hypervisor offering, this configuration option may exist as either a VM-configuration feature or a virtual switch port configuration feature.

Based on analysis of hypervisor offerings and ways of configuring a virtual switch, the publication provides security recommendations for VM traffic monitoring. Each recommendation has a unique identifier of format VM-TM-Rx, where VM stands for virtual machine, TM stands for traffic monitoring, and Rx for the recommendation sequence.

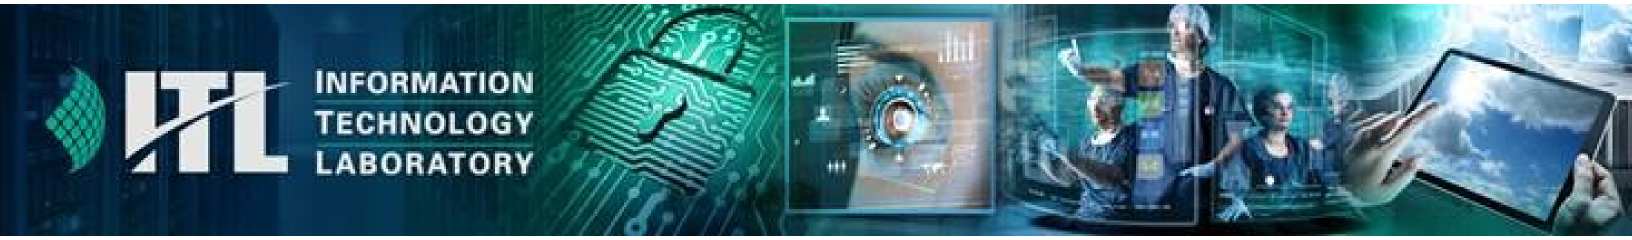

Conclusion

As the virtualized infrastructure in enterprise data centers increases, the VMs hosting mission-critical applications become a critical resource to be protected. VMs, just like their physical counterparts (i.e., physical servers), can be protected through host-level and network-level security measures. In the case of VMs, since they are end nodes of a virtual network, the virtual network configuration is a critical element in their protection.

Four virtual network configuration areas have been considered in SP 800-125B: network segmentation, network path redundancy, traffic control using firewalls, and VM traffic monitoring. Each area has been analyzed and corresponding security recommendations have been provided.

NIST will continue to work with stakeholders to identify additional ways to improve the security of VMs and the applications hosted on them, such as through host-level protection and VM data protection.

ITL Bulletin Publisher: Elizabeth B. Lennon
Information Technology Laboratory
National Institute of Standards and Technology
elizabeth.lennon@nist.gov

Disclaimer: Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by NIST nor does it imply that the products mentioned are necessarily the best available for the purpose.