Analysis of Virtualization Technology in the DMZ

Randy Armknecht, *Computer Security Graduate Student, DePaul University*
rarmknec@students.depaul.edu

Abstract

Virtualization technologies are constantly evolving as more uses are found and the technologies are further vetted and refined. An overview of currently known security issues with virtualization, in particular VMware, is conducted. Through analysis of related works and a basic application of the Factor Analysis of Information Risk methodology it is determined that virtualization can be used in a DMZ environment if trust boundaries remain physical and are not collapsed into the virtualized environment.

I. INTRODUCTION

Over the last several years implementation of virtual machine emulation, or VME, technologies has been growing substantially within enterprises. Major players in the market include the likes of VMware owned by EMC, Xen owned by Citrix, and VirtualPC owned by Microsoft. The core concept of all VMEs is the same: one operating system running a hypervisor is the host, multiple other operating systems of the same or different types can then be installed inside the host as a guest. Thus multiple systems can coexist on a shared set of hardware while remaining "isolated."

The benefits of VME technology has been seen and proven first in development labs and then in production data centers. Some of the most common benefits are in regards to efficiency of operations and utilization of existing resources. VMware's ESX product is one that the author has first hand experience with and helps immensely to achieve both of these goals. For one, the product is able to be clustered on multiple powerful servers. To achieve better utilization of resources the virtualization software will automatically spread virtual machine guests across the nodes in the cluster to maximize hardware utilization. Instead of dozens of machines each dedicated to a specific application with minimal hardware utilization, they can all be aggregated into the ESX cluster where fewer hardware resources exists, but will be better capitalized on. Additionally, having fewer physical machines will reduce the overall amount of power consumed, heat produced, cooling required, and rack space required. These are all important factors when considering the cost of resources in a data center. The operational efficiency comes into play with the ability to create templates for server operating systems. With templates, the time for deployment of a new server involves nothing more than a verification that the cluster has enough resources to support another server, and a few clicks by the ESX administrator. The ability to quickly deploy a standard vetted operating system image is a huge cost savings for the system administrators. With such benefits being reaped in the production data center its a natural evolution for system administrators to encourage the adaptation of this technology in the DMZ.

The DMZ, short for Demilitarized Zone, is an area of a network that is assigned a different trust level than the internal network. Typically, it's located between an internal network (intranet) and the internet. DMZs can be designed in numerous ways, two of which will be discussed in this paper. The first is the simple DMZ, it consists of a single subnet between the internet and intranet, and contains all the servers that have some public exposure. An example is seen below in Figure-1. The second is a tiered DMZ consisting of 3 tiers (web, application, database) and can be seen in Figure-2.

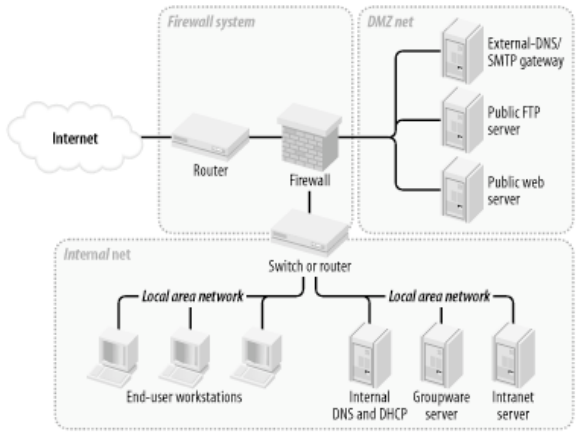

Figure 1 - Simple DMZ

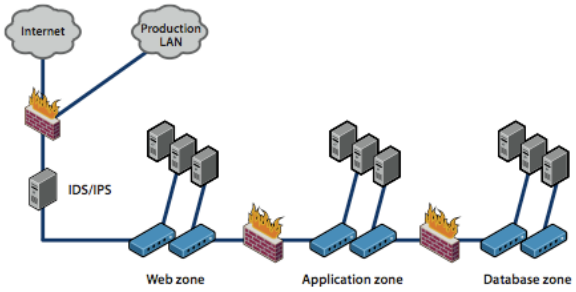

Figure 2 - Three Tier DMZ

II. Separation of Duties

The separation of duties among teams for the handling of any risk sensitive activity has long been a tenant of any quality security program. In terms of the DMZ the separation is that the servers and applications are generally managed by one team, the system administrators, and the security devices that control access to and from the servers between different zones of trust are controlled by another team, the security team.

If the virtualization is applied to each trust level individually, there is no difference in separation of duties. The difference is when multiple trust levels are collapsed into one ESX environment. In this scenario the firewalls between the multiple trust levels are themselves virtual. According to the Gartner report *Limited Choices Are Available for Network Firewalls in Virtualized Servers* this lack of major players has resulted in solutions with limited capabilities. Even when a solution is in place its heavily dependent on the configuration of the virtual environment. Typically, two vSwitches (virtual switches) are created, machines of one trust level are connected to vSwitch A and machines of a different trust level are connected to vSwitch B. The virtual firewall is then built and connected to both vSwitch A and B, and is configured to police the traffic between the zones [2].

The issue with such a configuration is that separation of duties can fall apart rather quickly. It's trivial for an administrator of the virtual system to cause harm in a number of ways both maliciously and through configuration mistakes. One example is that of a Denial of Service to traffic between zones. In the virtualized world the system administrators have the ability to power down systems with a few clicks. They could

maliciously power down the virtual firewall or mistakenly power it down when they meant to power down a different system within the virtual environment. Another is the bridging of two networks of different trust levels. Since the network a machine is connected to can be configured with a drop-down menu in the administrative software it's possible for a system administrator to maliciously or mistakenly assign a system to two networks, thus bridging trust levels without a policy enforcement device in-between. A third way that a system administrator could potentially cause harm is by maliciously or mistakenly placing a system in the wrong zone of trust.
It's important to note that in the above three examples each could be caused by either malicious actions or simple mistakes. In the physical world, the mistakes are not as simple and less likely to occur. In regards to the example of assigning a server to the wrong switch: in ESX it's a click; in the physical world a cable must be run to, plugged in, and enabled on the wrong physical switch. To do so maliciously in the physical world is possible, but due to separation of duties would require the less likely collusion of multiple people. In the virtual environment separation of duties has disappeared and makes the malicious attack possible to be carried out by a single individual.

The risk posed by this loss of separation of duties is still relatively small. Other controls, particularly those available to management exist to ensure that only trusted individuals have such access in the first place. The risk can be further mitigated by ensuring that proper change control and configuration audit programs are in place. These reviews can be used to uncover configurations that violate security policy regardless of intent.

III. Software Vulnerabilities

The ability to isolate guests from each other and the host system such that no guest can have an impact other than the typical consumption of resources is key to a secure virtualization implementation [6]. In order to accomplish this the host hypervisor must be sufficiently secure in order to provide that protection [1].

There are a multitude of methods for detecting if code is running inside a VME, some of which have started to appear in current malware [3]. So far the behavior of malware that detects the presence of a VME is to disable malicious functionality and appear benign. The use of virtualization for the study of malware has become common [3] and is being specifically detected in the hopes of avoiding reverse engineering by antivirus companies.

At first glance this can appear to be a benefit for virtualized systems in the DMZ. If for instance, a machine is compromised and loaded with malicious code the malware may detect that it's running in a VME thinking that its on an analyst's machine or honeypot [7] and cease to perform it's true purpose. However, VME detection is significant and required if the goal of the malware is to escape the guest system and control the host or to cause a DoS against the host machine [1].

In 2005 Tim Shelton discovered a flaw in the vmnatd service of VMware that allowed a malicious attacker to send specially crafted EPRT and PORT commands to a virtual machine running an FTP server. The flaw allowed for the attacker to escape from the virtual machine and execute arbitrary code on

the host system [8]. Fortunately, this flaw only existed on the Windows based VMware products and did not exist in the ESX based variants of VMware's products.

While research into VME escape techniques continues [1] [6] there are currently no known attack vectors in VMware's ESX line of products (the recommended VME for production systems and data centers). There have however been bugs in the software, so it is important that the hypervisor and related host software be maintained and kept up to date.

It's then worth noting that the addition of a VME into the DMZ environment adds an additional layer of software that must be routinely patched. The patching of the guest operating systems and their services also become more important. Several of the detection and escape techniques require the use of elevated privileges within a guest system. If the guest services and OS are appropriately patched escape to the host is less likely to be successful.

The introduction of a VME into the DMZ adds the risk that if a machine is compromised rather than it being a stepping stone to launch attacks against other machines it can be used to compromise the host, effectively gaining control over all of the machines in the DMZ through a single attack.

IV. Separation of Trust Levels

As mentioned previously there are multiple levels of trust within an enterprise network, and the DMZ typically consists of machines on a lower trust level.

If the DMZ to be virtualized is of the simple architecture there are only two possible VME solutions. The first is to setup a new VME environment within the DMZ's physical boundaries. If using the VMware ESX solution a cluster of hardware is required and the only ports open to the clusters are 22/tcp and 902/tcp, both of which are fully encrypted and should only be allowed from the system administrator's machines on the intranet [2]. In this scenario the current trust boundaries are maintained. If a virtual machine is compromised it must still contend with the firewall's traffic enforcement policies, regardless of whether or not it was able to escape and compromise the host system.

The second option is to fully collapse the DMZ into the VME thats currently in place in the production data center. This is an option proposed in the paper *DMZ Virtualization with VMware Infrastructure* written by VMware. This is deemed safe with the application of virtual security devices (firewalls, IPS), proper configuration and regular auditing of configurations both of the VME and the virtual devices within. The interesting item is that VMware requires virtual security devices for this configuration to be considered secure, yet Gartner has stated that their is a current lack of just such devices [10]. Also, all of the separation of duties issues identified previously in this paper exist with such a configuration. One of the more severe risks, though less likely to occur is that of VM escape. In a mixed trust environment if a DMZ VM is successfully compromised and subsequently escaped to compromise the host, the attacker is no longer subject to network traffic enforcement of the firewall. Essentially, the attacker can gain network access in the production

environment, not to mention, the ability to shut down a production node of virtual machines in the data center.

In the three tiered DMZ architecture there are three possible ways to deploy VME technologies. The first is virtualization of each trust level as described for the simple DMZ that continues to rely on the physical boundaries and enforcement points already in existence. The second is a fully collapsed DMZ that virtualizes all the networks and enforcement points as described above. The third is to partially collapse the DMZ.

A partial collapse of the DMZ means that a single VME is setup to contain the web, application, and database zones of the 3-tier DMZ but still utilize the existing physical enforcement points to reach the internet and intranet. The tiers within the VME are each assigned to virtual switches with virtual firewalls between each. As in the fully collapsed DMZ more trust is being placed in the hypervisor's resistance to escape, the virtual firewall's effectiveness, and in the correctness of each piece's configuration. The risk of mixing these trust levels is significantly reduced when it becomes apparent that if the enforcement points within the VME where to be evaded, the environment would be essentially the same as a simple DMZ, found to be risk acceptable at many corporations.

V. Regulatory Issues

In all the research conducted there was only one mention [5] of regulatory issues. Essentially, organizations that choose to implement virtualization technologies within the DMZ will come under greater scrutiny from auditors than those that don't. One thought is that this is related to Sarbanes-Oxley regulations that require adequate controls to protect systems that could have an impact on financial statements.

No matter which design for virtualizing the DMZ is chosen it will be up to the implementing organization to provide thorough documentation that they understand the risks of using the technology and have implemented controls to mitigate those risks to a level acceptable by the business. Although the use of virtualization in the DMZ is relatively new (10% of large organizations in 2006) it is expected to grow and become more pervasive (70% of large organizations by year-end 2010) [5].

A sample risk analysis will be conducted at the end of this paper including a list of recommended controls. Such an analysis is important to ensure that auditors will not take issue with the virtualization of the DMZ.

VI. Environmental Risks

There is currently an apparent lack of research regarding the physical/environmental risks that come into play when considering a virtualization technology solution. The Gartner document titled *Server Virtualization Can Break DMZ Security* mentioned the increased risk of a denial of service which was covered previously under Section II. The denial of service that it did not cover is that of power failure and hardware failure.

A basic ESX cluster will consist of two nodes each equipped with multiple hard drives in an RAID array and dual power supplies. If the environment being virtualized consisted of 12 servers, each serving a different application, they could easily be

hosted on the 2 node cluster. The issue that arrises is when failure strikes one of the nodes. Performance may have been at an acceptable quality when spread across multiple servers, however, all 12 servers would experience reduced performance levels if a node in the cluster fails. In a traditional DMZ design the failure of a single server will affect a single application. In a virtualized environment the failure of a single server will affect the entire group.

Basic controls can be implemented to limit the affect of hardware failures. For instance, NICs can be teamed; this will not only increase the amount of bandwidth available to the ESX node but will provide redundancy if a switch port or NIC port were to go bad. As with any server in the data center the dual power supplies should be connected to separate power grids. Finally, the cluster for a critical tier such as the DMZ that controls the public presence of an organization should consist of no less than a three node cluster. This will ensure that performance degradation is not significant should a node from the cluster fail.

Virtualization has also been shown to be a useful technology in disaster recovery scenarios [9]. The architecture of how it's recovery capabilities can be applied to restoring a DMZ environment are tightly joined to an organization's network design and business continuity plans. As such it is not worth studying further in this paper and will instead be left as an exercise for the reader.

VII. Quick FAIR Analysis

Throughout this paper a series of risks have been presented under each section. In order to understand at a high level the magnitude of these risks and how they affect the decision to implement virtualization technologies in the DMZ several will be applied to the Factor Analysis of Information Risk (FAIR) method [4]. A chart of the analysis is provided in Appendix A.

For the asset of a DMZ Guest Machine that exists within the VME five risks have been analyzed. The first three take the system administrator into account as the threat agent and the last two take a malicious attacker into account as the threat agent.

For the action of accidently powering a system down a Threat Event Frequency (TEF) of Low has been selected. System administrators are often highly trained and cautious when handling production systems. As such it can be reasonably argued that they would not make that mistake more than once in a year, but will likely make it at least once within a span of 10 years. A system administrator's Threat Capability (Tcap) can also be reasonably argued as High (top 16%) due to training and experience. The controls for this are a proper change control environment and daily auditing of changes made to the VME. Such controls exists as part of VMware's best practices and would certainly catch any inconsistencies in the environment. This controls strength (CS) was rated as Very High (catches all but top 2%).

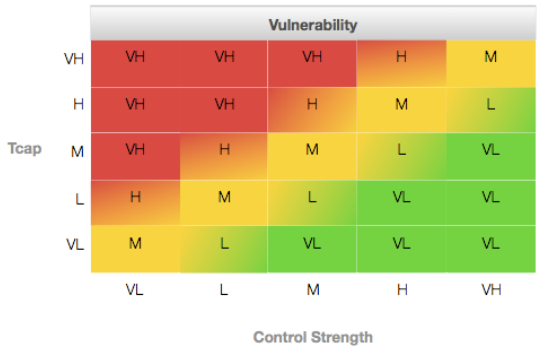

Using the above chart from FAIR it can be concluded that the vulnerability to such an incident is Low. From here the Vulnerability and TEF were compared to arrive at the Loss Event Frequency (LEF) of Very Low. It's chart is shown below.

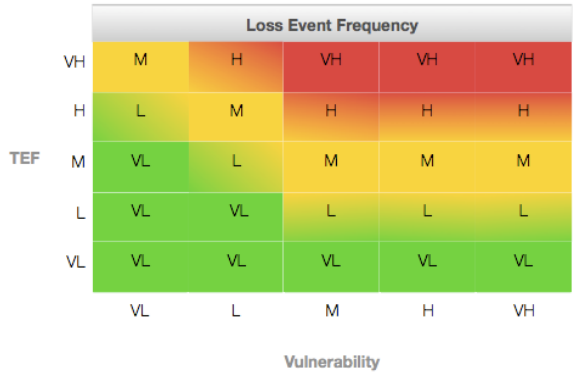

Continuing on, a Probable Loss Magnitude (PLM) was estimated to fall within the Moderate range of $10,000 to $99,999 if a system were to be powered off. This value will depend heavily on the specific system affected and should be calculated appropriately for each system by the organization conducting the analysis. For the purposes of this paper a conservative middle value was chosen. Lastly the PLM and LEF are evaluated against the below chart to arrive at an overall low risk.

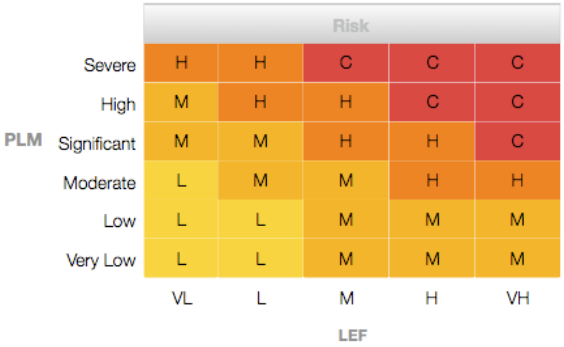

Now that a walk through of the general analysis process has been covered only interesting choices made will be further discussed. As previously mentioned the full chart can be seen in Appendix A.

For a malicious attacker against a DMZ Guest Machine the control strength in a non-collapsed environment is Very High and High in a collapsed environment. This is due to the fact that in a non-collapsed environment the attacker will still find themselves subject to the enforcement devices that separate trust zones outside of the VME. In a collapsed system, if VM escape is possible, virtual enforcement controls can be evaded, thus the virtual control is slightly weaker than the physical control. The same reasoning applies to the difference in PLM for these two cases.

In the case of malicious attackers going after the host system, the control strength against escape is Very High while the control strength against detection is Very Low. The ESX code base is extensively tested by VMware, has never suffered an escape vulnerability to date, and is under constant research, so for an attacker to successfully escape from it they would likely be in the top 2% of attackers. Detection of VMware on the other hand is trivial. If an attacker is unable to detect they are executing inside a virtual machine, they are likely in the bottom 2% of attackers given the extensive documentation on detection methods. The PLM is also of interest with these two cases. If an attacker is able to gain root control of the host machine the damage they would be capable of causing is most certainly in the Significant range. If an attacker is able to discover that they are in a virtual machine there is no direct damage caused or cost incurred.

VIII. Conclusion

The information presented in this paper has clearly demonstrated that virtualization technologies are highly active both in terms of development and research. While the current technology is not as secure as one might wish, especially when virtualizing zones of varying trust levels, mitigating controls exist that may make the risk acceptable. When virtualization is used within a single zone of trust little additional risk in incurred, making this architecture the most palatable to risk adverse organizations.

IX. References

1. Carpenter, Matthew, Tom Liston, and Ed Skoudis. "Hiding Virtualization from Attackers and Malware." IEEE Security & Privacy (2007): 62-65.

2. Dell'Era, Ivan. "Testing New Applications In The DMZ Using VMware ESX." IBM, VMWORLD 2006, 2006. 24 Oct. 2008 <http://download3.vmware.com/vmworld/2006/dvt0026.pdf>.

3. Ferrie, Peter. Attacks on Virtual Machine Emulators. Symantec Advanced Threat Research. 24 Oct. 2008 <www.symantec.com/avcenter/reference/virtual_machine_threats.pdf>.

4. Jones, Jack A. An Introduction to Factor Analysis of Information Risk (FAIR). Rep.No. Risk Management Insight. Sept. 2008 <http://www.riskmanagementinsight.com/media/documents/fair_introduction.pdf>.

5. MacDonald, Neil, and Greg Young. Server Virtualization Can Break DMZ Security. Rep.No. G00147785. Research, Gartner.

6. Ormandy, Tavis. An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments. Research. Google, Inc. 27 Oct. 2008 <http://taviso.decsystem.org/virtsec.pdf>.

7. Provos, Niels. Honeyd: A Virtual Honeypot Daemon (Extended Abstract). University of Michigan. 24 Oct. 2008 <http://www.physnet.uni-hamburg.de/provos/papers/honeyd-eabstract.pdf>.

8. Shelton, Tim. "[Full-disclosure] [ACSSEC-2005-11-25-0x1] VMWare Workstation 5.5.0." Full-disclosure. 21 Dec. 2005. 24 Oct. 2008 <http://lists.grok.org.uk/pipermail/full-disclosure/2005-december/040442.html>.

9. VMware. "Disaster Recovery Virtualization." Disaster Recovery Virtualization. 30 July 2007. VMware. 27 Oct. 2008 <www.vmware.com/files/pdf/DR_VMware_DoubleTake.pdf>.

10. Young, Greg, Neil MacDonald, and Joe Pescatore. Limited Choices Are Available for Network Firewalls in Virtualized Servers. Rep.No. G00154065. Research, Gartner.

Appendix A - Chart of Quick FAIR Analysis

Asset	Threat Community	Threat	Threat Event Frequency (TEF)	Threat Capability (Tcap)	Control Strength (CS)	Vulnerability (Vuln)	Loss Event Frequency (LEF)	Probable Loss Magnitude (PLM)	Risk
DMZ Guest Machine	System Administrator	Power Down System	Low	High	Very High	Low	Very Low	Moderate	Low
		Bridge Networks	Low	High	Very High	Low	Very Low	Moderate	Low
		Placed in Wrong Zone	Moderate	High	Very High	Low	Low	Low	Low
	Malicious Attackers	Compromise Guest - Non-Collapsed Architecture	Low	Moderate	Very High	Very Low	Very Low	Moderate	Low
		Compromise Guest - Collapsed Architecture	Low	Moderate	High	Low	Very Low	Significant	Medium
Host Machine	Malicious Attackers	Escape from VM	Low	Moderate	Very High	Very Low	Very Low	Significant	Medium
		Detect VM	Low	Moderate	Very Low	Very High	Low	Very Low	Low
	Malfunction	Hardware Failure	Low	Very High	High	High	Low	Very Low	Low