From the Physical to the Virtual – Threat Modeling the Landscape of Virtualization

Elliot Lewis
Principal Security and Networking
Architect, DVS, Dell Corporation

Legacy security can be viable as long as the virtual desktop solution can use it effectively.

Introduction

While today's forward-looking IT departments have a good sense of the threat landscape their company's face, the mechanics and necessity of VDI security remains a problematic issue. Indeed, regardless of whatever assurances the CIO may have received or actually read, customers considering desktop virtualization still want to know "Can my virtual environment be as secure, or even more secure, than the physical environment I have today?" While the physical and virtual environments are implemented on the same IT environment, many of the threat vectors that exist in the physical implementation have evolved to threaten the virtualized space.

One common misconception is that simply transposing well established security solutions from the physical environment into the virtual environment will maintain the security profile of the environment. Another is that virtualized end points running virtual operating systems and virtualized applications are impervious to virus attack. Indeed, although many threat vectors become irrelevant in virtual environments, customers should acknowledge that simply adopting new reference architecture will not solve every data center security issue going forward. The key is for companies to understand how their transition from the physical to virtual worlds will reshape the threat landscape and mitigate their data centers' exposure to a constellation of vulnerabilities.

The truth is virtual desktop environments still require a full complement of security mitigation controls and strategies for ensuring their efficacy while not degrading solution performance. Dell's Desktop Virtualization Solutions (DVS) offerings allow customers to easily deploy, manage, and integrate security mitigations into the seemingly complex architectures that enable virtual desktop infrastructure (VDI) with two compelling offerings: DVS Simplified and DVS Enterprise. These solutions, sized to fit customer deployments by Dell's blueprinting and benchmarking process, can greatly increase the speed of deployment to a successful virtualized operation. In this white paper, we will present a threat model for a typical physical environment of a generic industry vertical and then analyze the changes in the security posture and threat model that typically occur during a shift to a virtualized environment.

We will cover two threat models focused on the same enterprise environment: one for the physical "controlled" environment and the other for the newly-adopted, on-premises virtualized environment, where the company owns and manages the endpoints with network attached services. After reviewing these threat models we will cover the top five security changes customers need to consider when migrating to virtualized environments. Although "consumerization' is certainly a hot topic, the risks inherent in "Bring Your Own Device" BYOD programs will be examined in a future white paper. The bottom line is that in recognizing the critical nature of data and system security Dell has been taking several pro-active and aggressive steps to address a myriad of threats that jeopardize enterprise data centers and end points.

What are Threat Vectors and what is "Inherent Risk"?

Before we dive into the actual threat modeling, it is important to understand the two basic concepts of "inherent risk" and "control weakness" in the context of desktop virtualization. These are the key parameters that we will be mapping our models on and how each threat vector on the model is rated against others. A "threat vector" is a security function or security issue that a given company needs to be aware of in order to protect its business, data, revenue, and/or operation lines. Threat vectors identify how a company’s critical assets can be accessed or compromised and are measured against the inherent risk that a particular vector represents against the mitigation controls that allow the company to reduce that particular threat. To understand the threat models discussed below, it is important to begin any discussion of security with an understanding of how “inherent risk” is formulated against a given threat vector. In Figure 1, we can see a methodology to understand how “inherent risk” is analyzed:

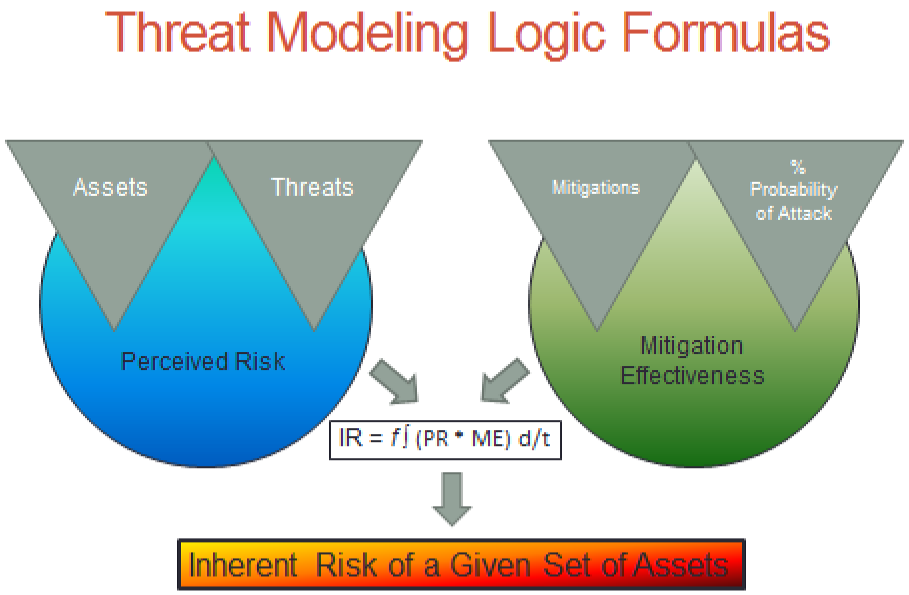

Figure 1: Inherent Risk Model

When visualizing any given threat model for a company’s operations, the first input factor should be the assets that the company possesses. These assets can be physical, virtual, or intangible, meaning IT departments have to deploy strategies to protect devices, data sets, IP, and services – anything that represents a critical business component to the company’s operation and future success. Once these assets are identified, the analysis shifts to assessing what kinds of threats can put any given asset at risk. As an example, databases containing customer personal data, can face threat vectors such as:

- Escalated privileged access to the database
- Web-front end SQL injection
- Corruption of the database tables
- Theft of the database files by external entities

Unpacking The Equation

The equation in Figure 1 may look complex – but it is just a mathematical representation of the business model of any given business’s daily operations. Every company has a unique set of working parameters that determine how they handle data, how they generate revenue, and how they leverage or monetize their customer base – and the graphic represents is a visual representation of this operational modeling. In the daily workings of the given company, and depending on its operational models and defenses, a constellation of threat vectors may or may not pose a risk. The next step is to evaluate these vectors for inherent risk in each environment where they could be disruptive and rank the vectors and their corresponding threat level.

To get a more accurate reading of the potential threats against a company's assets, we then incorporate two more factors: the "mitigation control efficacy" and "probability of attack" for each vector. Mitigation control efficacy is a measure of how well a given security mitigation actually limits exposure to a threat vector. For instance, looking back at our database problem:

- Theft of the database files can be mitigated by deploying an encryption solution for the database files. A "weaker" mitigation control would be to not to encrypt the actual data, but to rely solely on an upstream firewall to protect access to the database.

The next step in the process for IT department is to map out the threat mitigations available to the company and assess them for how well that control can mitigate that particular threat – also known as the "control weakness."

The other factor in assessing the inherent risk for a threat vector is the probability of attack for the given vector – i.e. "What is the chance, given my environment and how it is configured and operates, that this threat vector can be used to attack my company?" As an example, if the database above were actually located in an "air-gapped" network segment and not accessible online in any given way, then the probability of attack would be drastically reduced. If the database is exposed to the public for searching and data entry, then the probability of exposure or of incursion is much greater and requires more mitigation controls to be put in place. Probability of attack allows us to judge how much a vector adds to the inherent risk – and measures the potential weakness of the mitigation controls put in place to address the probability of attack.

When analyzing assets, threats and threat vectors, we map each vector against the mitigations available to control the threat and the probability of attack against the assets that are exposed by that threat vector. The complete analysis is then represented in the overall threat model. We will now summarize and compare two threat models regularly found in both the physical to the virtual operational environments.

An Example Threat Model for Financial Services

Threat model designs will vary by business vertical and by business operations models for a given company. For the purposes of this discussion we will focus on a typical financial services firm for two reasons:

1. Financial Services is an industry vertical that is highly regulated, so it is already clearly defined in many operations what is considered "risky" and what needs to be protected.

2. In Financial Services, there is often a clear operational example of how virtualization provides a highly positive security mitigation due to the fact that many financial traders have embraced consumerization and BYOD and currently work on various mobile and laptop devices from disparate network environments and require 24/7 access to critical financial data to perform their functions effectively.

Figures 2 and 3, we present two threat models based on the operations of a typical enterprise-level Financial Services company.

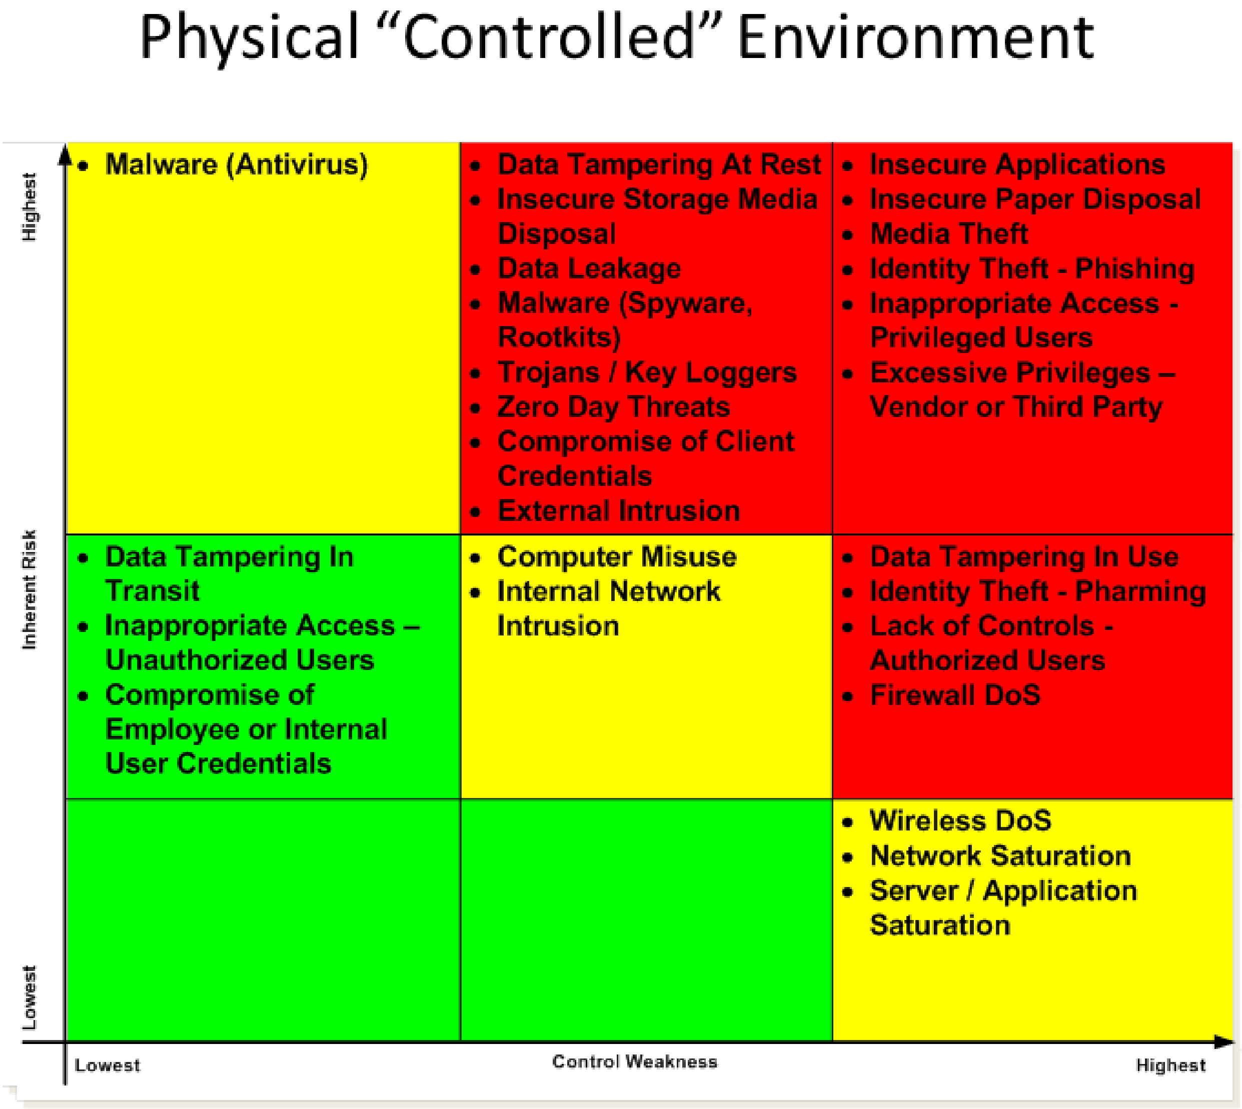

Physical "Controlled" Environment

	Lowest	Control Weakness	Highest
Highest	🟨 • Malware (Antivirus)	🟥 • Data Tampering At Rest • Insecure Storage Media Disposal • Data Leakage • Malware (Spyware, Rootkits) • Trojans / Key Loggers • Zero Day Threats • Compromise of Client Credentials • External Intrusion	🟥 • Insecure Applications • Insecure Paper Disposal • Media Theft • Identity Theft - Phishing • Inappropriate Access - Privileged Users • Excessive Privileges – Vendor or Third Party
Inherent Risk	🟩 • Data Tampering In Transit • Inappropriate Access – Unauthorized Users • Compromise of Employee or Internal User Credentials	🟨 • Computer Misuse • Internal Network Intrusion	🟥 • Data Tampering In Use • Identity Theft - Pharming • Lack of Controls - Authorized Users • Firewall DoS
Lowest	🟩	🟩	🟨 • Wireless DoS • Network Saturation • Server / Application Saturation

Figure 2: Financial Services Physical "Controlled" Threat Model

When looking at the "physical" threat model in Figure 2 above, many of the threat vectors under consideration are based on physical media vectors such as local storage or SAN solutions, individualized end-user client computing on traditional endpoints, "off-network" capabilities that users obtain when not attached directly to the corporate resources, and the loss of control – even in a corporate environment – that new BYOD physical assets represent.

When an enterprise migrates to a virtualized solution, the outcome of the same threat model shifts. The virtualization of the system mitigates several of the threat vectors. As a result of the reduction of the inherent risk associated with the threat vector, the new virtualized solution enhances the mitigation options against that vector.

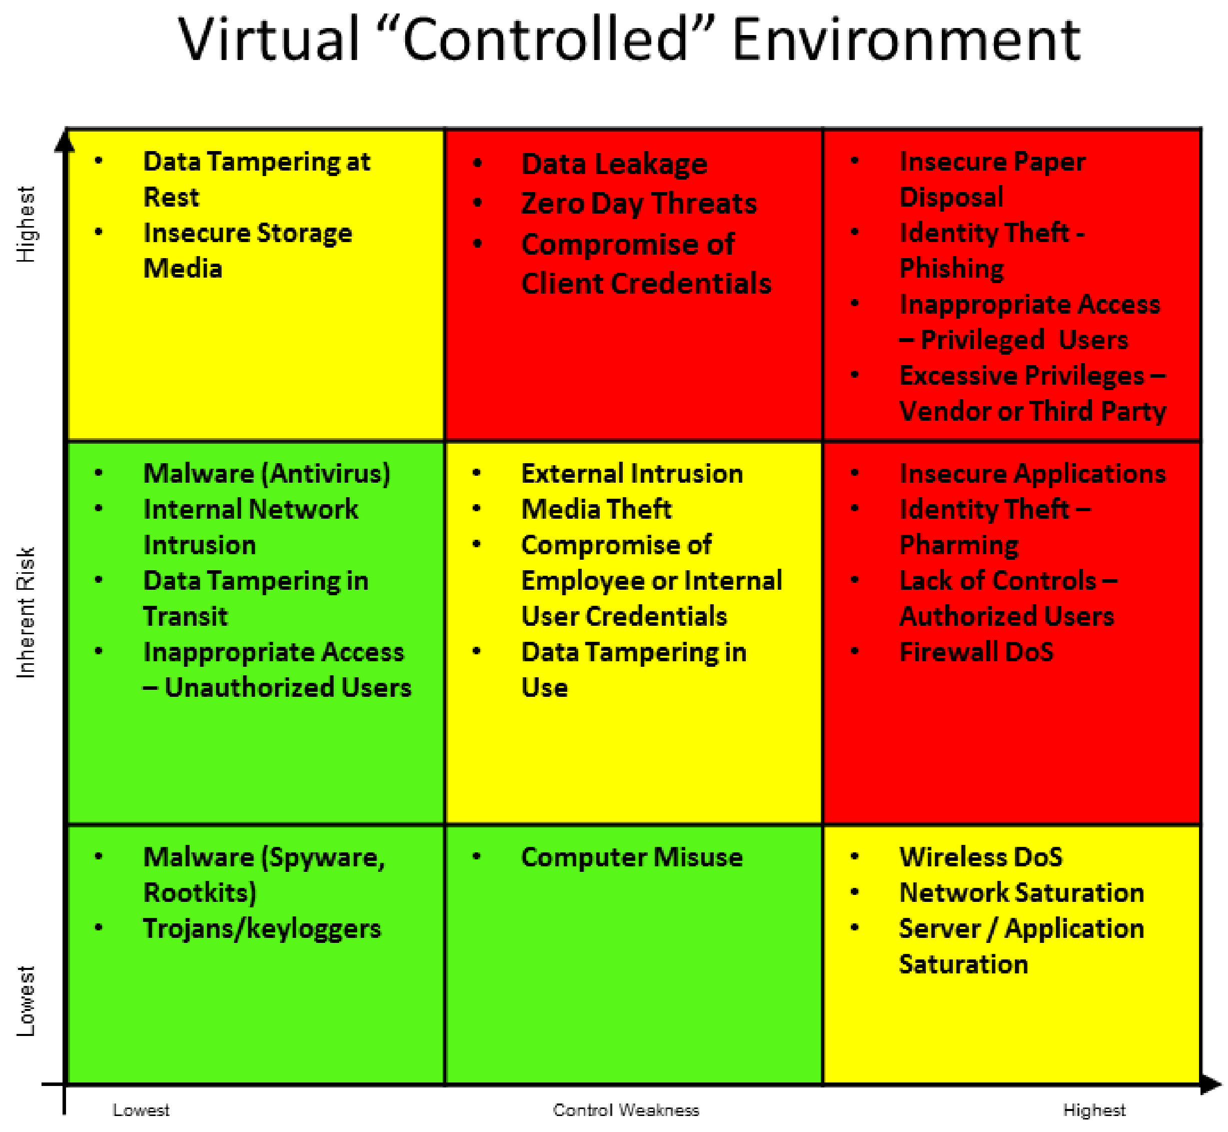

Virtual "Controlled" Environment

Inherent Risk / Control Weakness	Lowest	Control Weakness	Highest
Highest	🟨 • Data Tampering at Rest • Insecure Storage Media	🟥 • Data Leakage • Zero Day Threats • Compromise of Client Credentials	🟥 • Insecure Paper Disposal • Identity Theft - Phishing • Inappropriate Access – Privileged Users • Excessive Privileges – Vendor or Third Party
Inherent Risk	🟩 • Malware (Antivirus) • Internal Network Intrusion • Data Tampering in Transit • Inappropriate Access – Unauthorized Users	🟨 • External Intrusion • Media Theft • Compromise of Employee or Internal User Credentials • Data Tampering in Use	🟥 • Insecure Applications • Identity Theft – Pharming • Lack of Controls – Authorized Users • Firewall DoS
Lowest	🟩 • Malware (Spyware, Rootkits) • Trojans/keyloggers	🟩 • Computer Misuse	🟨 • Wireless DoS • Network Saturation • Server / Application Saturation

Figure 3: Financial Services Virtual "Controlled" Threat Model

What Changed in the models from the Physical to the Virtual?

Now that we have the physical vs. the virtual mapped out in comprehensive models, we'll go over the five top threat vectors that have the greatest differential in inherent risk and control efficiency and discuss why each one moved position on the models. We will also cover what specifically has not changed in the models and the reasons for its static positioning.

Data at Rest/In Transit/In Use

The threats against the control of data and storage of data are reduced on several levels due to the decreased attack surface that desktop virtualization and thin client deployment provides. The inherent risk is still at an equal level for each area of data at rest, in transit, and in use but the control weakness ratios Figure 3 are modified because data is no longer exposed to the increased attack surface that physical endpoints provide.

Data at Rest: the physical threat here is directly related to data sitting on physical devices that are subject to the potential loss/theft/destruction of the physical device. By virtualizing the application/desktop and providing storage for that data back at the datacenter or cloud, the data is no longer exposed by sitting on an exposed device in a "resting" state. This new level of control on data at rest also applies to other threat vector modifications discussed below in the section titled "Insecure Media Storage/Media Theft."

Data in Transit: the physical threat here is related to the multitude of networks that need to be transitioned from the storage source to the applications/services where the data is utilized locally. When using a laptop or mobile device with non-virtualized applications, the data will have to traverse several network segments (such as the local network, the internet provider, and all routing/switching and far end networks) before interacting with corporate resources. (VPN, as an example, is an effective mitigation source for this kind of threat vector.) With virtualization, the potential exposure level and corresponding attack surface are drastically reduced.

NOTE: The data in transit vector is specifically related to networked data delivery from one corporate owned asset to another within the virtualized environment. The threat vectors of data being sent externally in an authorized manner to third parties (i.e. – via email, et al) are related to "Data Leakage" in the above models. Note that "Data Leakage" has not been moved in the models. Virtualization of desktops and applications do not affect the scenarios of data leaving the corporate controlled assets via authorized means unless data loss prevention (DLP) is used as a corporate data breach / data loss mitigation tool.

Data in Use: Corporate data can be corrupted/lost/stolen via compromised systems (i.e. – those that have been root kitted, infected with malware, or otherwise breached (see section below "Malware/Trojans/Keyloggers" for more in-depth information) that are exposed to external pathways and not kept up to date or protected upstream/downstream by corporate security measures. By using virtualized desktops and applications, the potential for exposure is drastically reduced mitigating the effectiveness of common attack scenarios. Further, if a compromise does occur in the regular use of the data, active corporate controls can more readily identify and mitigate the exposures. This threat vector is also positively impacted by the use of virtualization deployment models called "non-persistent virtual desktops" where the VM's are launched from "clean" images every time a new session is established. Additionally, the use of virtualization layering by Dell in its architecture for its DVS

Enterprise solution, allows for dissemination of "personal" settings from the corporate OS images, thus ensuring that – while the personal user experience is maintained – the processing and data controls from the base OS are kept in a "clean" manner.

Malware/Trojans/Key-loggers

When using virtualization techniques to deploy corporate computing solutions, such as virtual desktops and application streaming, one of the deployment models widely used is "non-persistent" desktop provisioning. In this model, users accessing virtualized desktops receive a desktop image that is brought made available using a "clean" image that has been virus scanned, malware scanned, and is free from common threats such as contagions, rootkits, and keyloggers. When using "persistent" deployment models, the virtual machine states are maintained for the user and the image is not reset – but simply saved in state – thus providing an experience more like the traditional usage scenarios of individual laptops in a corporate environment.

While the user experience is more akin to what the user expects today from a thick client physical laptop , it is not as effective in reducing the attack surface as "non-persistent" VDI. Traditional persistence can create image sprawl which drives up the time to update/patch inconsistencies, potentially lengthening vulnerability exposure and raising the vulnerability quotient in the threat model. OS layering can shorten cycle times leveraging many patching/updates due to single gold virtual image management across the enterprise. Careful consideration should be applied when choosing the benefits of "non-persistent" vs. "persistent" VDI architectures and the balance between user experience and productivity with security.

In the area of malware on VDI deployments, there are varying technology developments and strategies how to deploy malware protection on virtual machines. Anti-malware is still a necessity. Even in "non-persistent" deployment scenarios, active malware contagions acquired during a live session and can impact the corporate resources while a given virtual machine is "active" prior to periodic resets and/or log-off resets by the user. In other words, users in "non-persistent" virtual desktop architectures are still vulnerable to virus and malware attack and IT departments should never feel 100% secure. Anti-virus/anti-malware software is always required no matter how minimal the risk – but the deployment methodologies can change – shifting the weighting of these threat vectors in our threat models over to a more secure position as we move from the physical to the virtual. The primary operational differences, from the physical to the virtual deployments, are in two areas:

1. Boot storms –In physical deployments, where the boot process of each individual system is isolated using its own processor and memory stores to bring up the OS, and each system will use scanning systems for anti-malware. Virtualized deployments on the other hand will have many virtual machines accessing one set of processor and memory store spreading the scan or the IOPS load from many virtual machines using the same hypervisor. When many virtual machines start up on the same hypervisor simultaneously, this can cause a significant degradation of the user experience during boot processes, which can lead to user revolt. Anti-malware solutions installed to run during boot processes actually exacerbate the time and processor cycles for this event as well.

2. Scanning storms – while many virtual machines remain active on a hypervisor, the IT/DVS administrator needs to be careful not to use the default settings of traditional "physical" anti-malware agents that are designed for individual physical machines. For instance, if all agents are set to kick off a full virus scan at 2 a.m. every morning, this will cause a "scan storm" to

occur on the hypervisor cluster and potentially lock up most of the virtualized machines associated with the hypervisor until long after 8-9 a.m. and the start of the traditional work day.

To mitigate these issues, automatic VM startups can be staggered over a period of time for the general population use. Full volume scanning can be staggered or potentially bypassed, or in the case of "non-persistent" virtual machines set to commence only during specific, scheduled reboots. In the case of anti-malware deployment options, reduced stand-alone anti-malware agents can be deployed throughout the virtualized environment and launched from a single dedicated virtual machine. This allows the IT department to sandbox the VM's and limit scanning to a specific hypervisor cluster instance, thus offloading the processing and memory usage constraints from the related virtual machines.

Computer Misuse

Corporate owned laptops or mobile devices – where administrative privilege controls and security features are active on the device – can still become compromised while the device is off the corporate network and susceptible to external access and threats. This can happen when accessing an enterprise network through relatively insecure airport wi-fi, or when a laptop or mobile device is being used on a home LAN or hotel LAN, In these instances, the efficacy of corporate mitigation tools positioned downstream from the device while it is on a corporate LAN are undercut, leaving firewalls, URL filters, peer-to-peer agent controls, data loss prevention controls, and router access control list solutions are less effective against various techniques often employed by bad actors. Another threat vector to consider regarding data breach comes from studies indicating that nearly 70% of most organization's data resides on client systems.[1] Dell DVS solutions can mitigate the fallout from the loss of physical devices either through theft, carelessness, or inadvertent erasure. Finally, virtualization allows for central policy controls to mitigate data loss from end users copying corporate data onto USB drives and similar devices which are often misplaced.

Insecure Media Storage/Media Theft

Due to the threat vector of the "stolen device" scenario, many corporations deploy full disk encryption solutions to maintain the corporate data and access to corporate resources. However, desktop virtualization solutions are far more effective. By virtualizing the desktop or application, this threat vector is no longer applicable – the theft would have to occur by breaking into the "defense in depth" controlled datacenter, by bypassing the storage administration controls, and then performing a large download of gigabytes of data to steal a virtual machine's image files and storage. By keeping the data on premises, and by deploying desktop virtualization and defense-in-depth controls on the datacenter, this threat vector is vastly reduced due to the reduced opportunity, extreme inconvenience, and low probability of success with this approach.

This being said, data protection scenarios do still need to be deployed in the datacenter and across the storage solutions to ensure the sanctity of personal data, to protect sensitive corporate assets and non-public data, and to keep trade secret data from unauthorized data extrusion or exfiltration. The benefit from using desktop virtualization is that encryption solutions can now be centrally deployed and IT-controlled. Additionally, processing can be increased by providing encryption on server grade datacenter processing resources. Finally, key management can be tightly controlled rather than widely

[1] "Protecting Critical Data in your Organization," Micro Strategies, Page 4, http://bit.ly/Noxvcs

disseminated by reducing the expansion, replication, and access to the key management solutions to the datacenter where the images reside.

External Intrusion

Prolific hacker techniques take advantage of corporate devices being out "in the wild" and subsequently vulnerable. Hacker collectives today work to find ways to compromise that individual device in order to gain access to corporate data centers and the assets within. Hackers know that datacenters are highly monitored, controlled and are typically safeguarded by defense-in-depth mitigation controls. By targeting the individual devices, the probability of a successful attack is increased significantly, and the hackers can potentially leverage the following vectors:

1. User identity: keyloggers, rootkits, and other various malware target the individual device and can be a high-yield attack vector for hackers. Users are susceptible to phishing attacks, insecure network environments, and access to insecure websites – all of which are specifically designed to grab personal and corporate identification material.

2. Machine identity: if a system becomes root-kitted, the certificate stores of the individual device can be exposed, thus allowing the hacker to obtain commonly used corporate PKI certificates used for network access, application access, VPN access, et al. Once these certificates are exposed, the hacker can often sign in to the corporate data center as an "authorized" user.

Fortunately, many corporate controls, when properly used, can drastically mitigate these threats. Indeed, the proper use of TPM technologies can reduce this threat in the physical devices. What IT may not know is that VDI solutions can significantly reduce the probability of attack for this threat vector instead of having to manage such solutions on individual devices. When virtualizing the desktop and applications, the ability to use the corporate upstream and downstream security mitigations greatly reduces the attack surface that bad actors can use as a jumping off point to grab the unique identity materials of the user or device. Security measures inherent in the architecture of desktop virtualization leave hackers relegated to attempting to breach the far more protected datacenters rather than the exposed physical device "in the wild."

Conclusion

Desktop virtualization can be a powerful tool to enhance the security of the company's security profile and, if implemented properly, can significantly reduce the overall attack surface for the company. Dell DVS Solutions are an incredibly powerful way for customers to easily deploy and manage the complex infrastructures and architectures required to enhance data security in a dangerous world. Of course, customers will still need to evaluate the current security tools being used in their physical environments and correctly deploy them into virtualized environments.

As more enterprise IT departments consider migrating from physical environment operations to virtualized operations, many wrongly assume that legacy security solutions and mitigation tools deployed in the physical realm will translate to the virtual world. In several respects, this is not the case. Because threat vectors and mitigation effectiveness change in many ways from the physical to the virtual companies need to rationalize the transition and maintain a vigilant corporate security posture during and after that transition.

Virtualization can be very effective in reducing the attack surface of the corporate environment by streamlining operations, reducing systems access and exposure, allowing for stronger control of company assets and resources, and simplifying management and deployments. The underlying message in producing these threat models is that, by deploying compelling desktop virtualization solutions such as Dell's DVS Enterprise and DVS Simplified, the corporate based mitigations and controls that are available on the on premises corporate network remain in effect. This improves the company's overall threat model while increasing mitigation effectiveness and ease of control even with the march of consumerization and the use of unauthorized mobile devices – while significantly reducing the probability of attack from known threats.

Because desktop virtualization shifts the processing model, certain security mitigation controls such as anti-malware need to be adjusted. Encryption strategy, tools and techniques also need to be modified for the new operational models. Corporate security controls upstream and downstream from the end-user clients may need to be adjusted in order to accommodate the new levels of IOPS, data usage, and processing paradigms.

Using our comparative threat models, five key areas in security mitigation controls have been identified for close scrutiny:

- Data Controls (at rest, in transit, and in use)
- Malware/Trojans/Keylogger Threats
- End User Computer Misuse
- Insecure Media Storage/Media Theft
- External Intrusion threat vectors

While these leading threat vectors may change in our threat models over time during the migration from physical to virtual enterprise desktops, new IT deployment and usage models with new threat vectors still need to assessed and rationalized. When considering new deployment models such as Bring Your Own Device (BYOD), new threat vectors can emerge. (In a subsequent white paper from Dell, we will analyze the BYOD trend and the new threat models it represents.) Yet however the security situation evolves, enterprise customers can rest assured that Dell will continue to stay focused on the virus and malware threat landscape and incorporate best of breed technologies and best-practices for keeping sensitive data safe into its solutions.

To get insight on the array of Dell's desktop virtualization solutions, visit: **www.dell.com/desktopvirtualization**

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND/OR TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.