kaspersky BRING ON THE FUTURE

# Virtual machine RAM control: what you should know to build effective VDI cybersecurity

> 'The most important thing a business should do to adapt to the increasing presence of telework is to plan for it coherently.'
>
> Ars Technica, Oct 2020

During 2020, VDI has proved its worth as never before.  The sudden compulsory adoption of large-scale teleworking may well lead to a complete overhaul in strategic thinking about corporate IT - one in which virtualization will play ever more central role.

The concept of being able to just ship an image of a desktop from the datacenter to the user really is amazing. No data ever leaves the corporate perimeter, and if the user's computer, tablet or mobile phone is stolen, there's nothing for cybercriminals to get their hands on. Once VDI is implemented, it's also easier to procure new machines to users whenever needed – just clone a new virtual desktop and the employee is all set. No need to wait for hardware, load an operating system, install software – all the stuff that made providing an employee with a machine such hard work, particularly when remote working is involved, has essentially gone.

But, while virtual machines are indistinguishable to the user from physical PCs, there are significant differences in terms of data security and administrative effort. And there are several very important aspects that need to be taken care of if you're to extract full benefit from your VDI project.

One of these aspects is security.

## Why VDI infrastructure needs securing

> IT departments deliver the transparency that VDI users demand through providing virtual USB-ports, familiar browsers, etc. Unfortunately, these also provide opportunities for familiar malware threats.

Virtualization doesn't do away with the need for security. Many of the same vulnerabilities exist for virtual as for physical machines – the fact that an operating system is virtualized rarely means that a vulnerability can't be exploited. If anything, VDI administrators should be more cautious; there are known incidents where cybercriminals have taken advantage of the fact that they're attacking a virtualized infrastructure. Faced with same vulnerable operating systems and applications, the same threats, same users, and a potentially larger attack surface, a multi-layered defense for your virtualized desktop infrastructure remains a necessity.

## Different approaches to protection and the importance of RAM monitoring

The nature of virtual environments means that best-practice security approaches differ from those applicable to traditional physical infrastructures. It's worth briefly explaining why. Firstly, virtualized workloads need to be more power-effective than their physical counterparts.  The large number of machines on each host means that any simultaneous action by a number of machines (such as users logging on at 9 am) can create surges in the demand for power which products designed for physical endpoints struggle to meet. Then there are technology-based demands - or 'activity storms' - where every virtual machine attempts the same action, such as updating its security database, simultaneously. All these will result in slower overall response times if resource-hungry products are used.

One notable example of a fileless attack strategy that's almost impossible to detect without RAM access is the use of legitimate Mimikatz software by threat actors, to enable easy post-exploitation lateral movement throughout your network.

VDI solutions developers are always looking for new ways to minimize the load on virtual hosts. One approach is to centralize the security function to avoid the duplication of data and operations. The ideal scenario, from a systems performance perspective, is that no security agent whatsoever is installed on the secured virtual workload itself. A separate dedicated virtual machine is responsible for securing the entire virtual host. This approach does, however, have its own limitations, particularly in terms of the complexity of threats it's capable of seeing. This is due to lack of access to the internals of the secured virtual machines - including memory (RAM), storage, devices and the operating system. Even if RAM access is possible (as it can be under some conditions), the information available to the security product is very limited, e.g. only contains a memory dump.

As we'll see below, deep dynamic behavioral analysis is absolutely crucial when defending against complex threats, such as fileless malware, exploits and ransomware. That's why multi-layered security relies on much more than just executable file analysis, and access to virtual machine RAM is vital to the solid protection of the employee's virtual desktop.

## Why memory control is necessary, but not sufficient

Complex tactics such as fileless malware that leverages legitimate tools for an attack have been known to the information security industry for years, and the memory-scanning of physical and virtual machines is not new. We at Kaspersky have taken this approach further by developing a driver (also implemented in 'light agent' applications for virtual host protection) that scans the operating system kernel, other drivers, user space processes, etc. in the memory, using a number of different techniques. This creates context for events occurring on the virtual endpoint - absolutely crucial for understanding an attack leveraging a complex exploit, for example. All the technologies that you'd come to expect in a cutting edge security product contribute to this context, including File Threat Protection, Network Threat Protection, Exploit Prevention, Behavioral Protection and so on. Events recorded at file systems level, system registry manipulations, OS API calls and memory scanning are all different pieces of the puzzle, without which a security solution will either miss threats or – almost as bad - generate false positives.

## The technical bit

You'd think it would be good idea to just block any software that's packed or obfuscated. And some radical security products do just that. But there are very good reasons why legitimate software may use packing or obfuscation – to protect against piracy for a start. If you take a shortcut and just block any software that attempts to complicate its analysis, you're poised to generate a significant percentage of false positive verdicts.

One of the main reasons why memory scanning is so important is that the analysis of a packed malicious file is only possible in memory - that's the only place where it exists in its unpacked form.

Well-known packer programs have already been 'cracked' by security researchers, allowing packed malware to be analyzed by today's anti-virus engines while it's still packed and residing on the disk. But this becomes much more of a challenge when an attacker uses more advanced packers or obfuscators, which implement in-memory virtual machines, polymorphism and other technologies that are highly dynamic by nature, preventing any possibility for on-disk analysis. You have to be looking into the memory to even begin trying to differentiate this malware from legitimate software.

Basic memory threat scanning doesn't generate much analytic profit per se in the case of protected files, either. Without behavioral analysis, it's hard to judge the code assignment. And it's not always clear when to start scanning, to ensure that you've finished decrypting (de-obfuscating) a sample and found the correct entry point. And all this assumes the sample doesn't use a 'decrypt on the fly' mechanism that requires the continuous use of encryption routines.

The features of packers and obfuscators overlap with those of memory handling mechanisms in today's operating systems. In Microsoft Windows, many operations take place asynchronously. When writing data to a file, the data's initially buffered in memory, after which the OS writes the entire buffer to the disk - so the precise moment of writing to disk is unclear, as is the processing context in which it will occur.

Another issue – all operating systems now use 'swap files' to enhance performance. From time to time, virtual memory content is unloaded to the swap file. In the event of a security incident, the virtual memory can be loaded into the swap file just before a physical memory scan. This is a problem for agentless solutions in particular, as they only have access to physical memory. So dumping virtual memory into a swap file in itself interferes with memory scanning.

Kaspersky Hybrid Cloud Security supports a wide range of server virtualization and VDI platforms:

- VMWare vSphere
- NSX, Horizon
- Microsoft HyperV
- Citrix Hypervisor, Virtual Apps and Desktops
- KVM
- and others

Kaspersky Security for Virtualization Light Agent, a part of Kaspersky Hybrid Cloud Security, is a light-weight, linearly scalable solution for server virtualization and VDI. Highlights include:

- Patented architecture which efficiently centralizes security tasks. Shared caching and smart optimizations eliminate resource waste and maximize consolidation ratio
- Provides system hardening, vulnerability assessment and patch management, exploit prevention and behavioral analysis for desktop and server virtualization.
- Provides protection from ransomware for operating systems and shared data.
- Efficient Network Security provides firewall, deep traffic inspection, packet analysis, IDS and NAB.
- Single pane of glass management console ensures consistent visibility and control.

Finally, if scripted malware is used, a programming language interpreter is responsible for running a sample. The interpreter steps through a script and executes commands one by one, so only part of the sample may exist in the memory at any one time. Malicious code does not exist with a full context in the memory, but can still execute.

In terms of the specifics of virtual environment defenses, OS memory mechanisms also overlap with hypervisor features. Hypervisors are only allowed to take control during a limited number of strictly specified times, which in turn limits possible reactions to malware detection in the memory.

# What full, multi-layered VDI security looks like

To sum up, RAM threat scanning is very useful in fighting malicious code, but should operate in the context of detailed dynamic information about processes on the host. This data comes from the memory as well as from the file system and is aggregated by specialized security agents. An over-reliance on memory scanning can slow down security analysis with no significant benefits in terms of protection.

This is why Kaspersky Hybrid Cloud Security contains an highly-optimized application – Kaspersky Security for Virtualization Light Agent - which architecture is built around agents – a lightweight agent adds critical security capabilities while maintaining high VM density, providing a powerful multi-layered defense, capable of eliminating sophisticated malware and even the very latest threats.

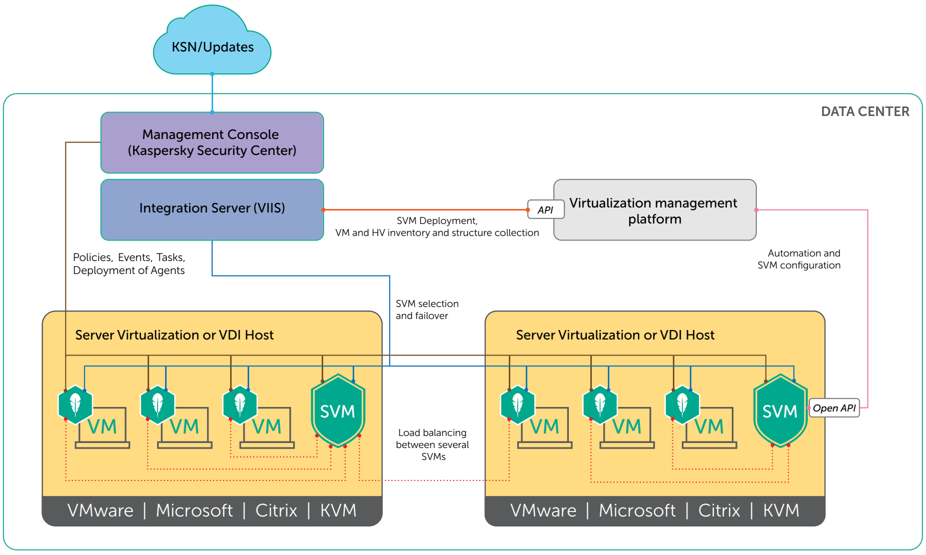
KSN/Updates
DATA CENTER
Management Console
(Kaspersky Security Center)
Integration Server (VIIS)
API
Virtualization management platform
SVM Deployment,
VM and HV inventory and structure collection
Policies, Events, Tasks,
Deployment of Agents
Automation and
SVM configuration
SVM selection
and failover
Server Virtualization or VDI Host
Server Virtualization or VDI Host
VM
VM
VM
SVM
VM
VM
VM
SVM
Open API
Load balancing
between several
SVMs
VMware | Microsoft | Citrix | KVM
VMware | Microsoft | Citrix | KVM

Users benefit from technologies like Behavioral Detection, HIPS (Host-Based Intrusion Prevention System), Network Threat Protection, Automatic Exploit Prevention, and a full set of endpoint controls. The solution architecture significantly reduces the attack surface, while saving precious computing resources.

## System hardening and multi-layered security

Rich and flexible system hardening drastically reduces the attack surface, eliminates arbitrary code execution on servers and blocks exploits. Memory and data control algorithms defuse fileless malware and ransomware attacks, both host and network-borne. Kaspersky Hybrid Cloud Security ensures a secure and responsive VDI environment, allowing users to focus on their work without risk of becoming a victim of cyberattack.

## Up to 30% reduction in virtualization resource usage

Light Agent technology significantly reduces the 'security tax' and, through a number of smart optimizations, such as shared caching and the elimination of redundant information, is able to cut the amount of data and number of operations needed. This dramatically reduces IOPS, CPU cycles, memory and disk footprints, all of which helps to achieve high consolidation ratios and protect investments in virtualization projects.

## Scalability

A patented architecture ensures linear scalability as the load on virtualization hardware increases. Simplified procurement and native support for Golden Image allows the dynamic spawning of virtual desktops as needed. Flexible licensing allows for gradual migration from physical machines with roll-backs as needed, making Kaspersky Hybrid Cloud Security an ideal companion for your VDI roll-out project and beyond.

## Single pane of glass management console

Kaspersky Security Center is the single pane of glass management console that takes the complexity out of security administration and IT systems management. Fully scalable, the console supports digital transformation and facilitates comprehensive security management, with easy separation of administrator responsibilities. The rich reporting tool enables a continuous audit of physical, virtual and public cloud infrastructures.

## Native integration with leading Virtualization Platforms

Native API integration simplifies deployment, streamlines daily administration routines and ensures consistent visibility and control. The single pane of glass management console provides the means for efficient orchestration of physical, virtual and public cloud workload security.

# It’s always a balance

Effective endpoint security relies in part on the continuous acquisition of dynamic information about processes on the machine. When the machine is a virtual one, the nature and footprint of the specialized security agent involved becomes critical to the effectiveness of machine performance. There’s a fine balance to be had between protection and performance, and you need to be sure you’re running a security solution which achieves that perfect balance for you.

For more about how Kaspersky Hybrid Cloud Security can help you build full VDI security, visit
www.kaspersky.com/small-to-medium-business-security/virtualization-light-agent

Cyber Threats News: www.securelist.com
IT Security News: business.kaspersky.com
Cybersecurity for SMB: kaspersky.com/business
Cybersecurity for Enterprise: kaspersky.com/enterprise

**www.kaspersky.com**

We are proven. We are independent. We are transparent. We are committed to building a safer world, where technology improves our lives. Which is why we secure it, so everyone everywhere has the endless opportunities it brings. Bring on cybersecurity for a safer tomorrow.

Known more at kaspersky.com/transparency

Proven.
Transparent.
Independent.