

The Challenges of Securing Hosting Hyper-V Multi-Tenant Environments

by Brien M. Posey

In the not too distant past, VMware was the hypervisor of choice for hosting providers. More recently however, many providers have been adopting Microsoft's Hyper-V. Hyper-V offers core functionality that is very similar to that offered by VMware, but often does so at a substantially lower cost. In fact, a recent study by Gartner forecasts Microsoft and VMware sharing an equal percentage of new virtualization host deployments.

Although it might clearly be in a hosting provider's best interest to adopt Hyper-V as opposed to VMware, providers must exercise due diligence in making sure that their Hyper-V clouds are properly secured.

While the concept of server security might initially seem routine and mundane, there are two factors that hosting providers must keep in mind while choosing a security solution.

The first of these factors is that the solution of choice needs to be aligned with Hyper-V economics. Cost is one of the major reasons for choosing Hyper-V over competing solutions such as vSphere or XenServer. As such, a Hyper-V security solution should not erode the savings that were achieved by choosing Hyper-V in the first place.

Overpriced security suites make hosting providers less profitable, and may force a provider to pass costs on to customers. Doing so is especially undesirable when you consider that customers often turn to cloud hosting providers as a way of decreasing costs. When a provider's costs increase it risks losing customers to lower priced competitors.

The other high level concept that providers must keep in mind when choosing a hosting solution is efficiency. The solution must run efficiently on the hosting provider's servers, so as not to decrease performance or virtual machine density.

Likewise, the solution must be efficient to manage and it must provide scalability as the cloud networks grow. Cloud scale networks are challenging enough to manage without the security software getting in the way or adding to the complexity.

Cost and efficiency are easily two of the most important factors to consider when picking out a security solution for Hyper-V hosts, but cost and efficiency are meaningless unless the security solution is able to do its job. The bottom line is that whatever security software an organization adopts, it just has to work. A security suite must protect Hyper-V servers and the virtual machines running on them from a variety of threats – both known and unknown.

Securing a hosting provider's cloud is a tall order. Hosting providers by their very nature make big targets for those who have malicious intent. As such, the key to securing a public Hyper-V cloud is to anticipate and then mitigate various threats to security, all the while keeping an eye on costs and operational efficiency.

The Challenge of Public Cloud Security

Security must be a top priority for a public cloud provider. The negative press surrounding a security breach can potentially drive a provider out of business and may also subject the provider to civil litigation.

Unfortunately, the unique nature of a public cloud provider's infrastructure makes achieving the proper level of security more difficult than it might be in a private organization.

By far the biggest challenge when it comes to securing a public Hyper-V cloud is that of locking down a multi-tenant environment. In essence, public cloud providers must come up with a way to secure resources that they do not fully control. Providers typically control infrastructure components, but not virtual machine contents or configurations. However, this is only one of several major security challenges.

A second challenge stems from the fact that cloud providers typically try to maximize the return on their hardware investment by achieving the highest possible virtual machine density. The problem with this is that as the number of virtual machines running on a server increases, so too does the number of potential exposure points. If a virtual machine (or the host OS) is compromised, that virtual machine could potentially be used as a platform for launching attacks against other VMs.

Another challenge that public cloud providers must cope with is the complexity of Hyper-V virtual networks. Hyper-V hosting providers isolate customers from one another through the use of tenant level virtual networks.

The side effect to building virtual networks that provide tenant isolation is that tenant level virtual networks may create a blind spot for security software. Unless a firewall for example, has been specifically designed for Hyper-V, it will be unable to analyze virtual network traffic.

Finally, public cloud provider's networks are extremely dynamic. Tenants constantly delete unwanted virtual machines and provision new ones. Load balancing software live migrates virtual machines from one host server to another in response to shifting workloads. As such, any suitable security solution must be able to adapt in real time to constantly changing conditions.

As you can see, there are a number of unique challenges associated with securing a public Hyper-V cloud. The remainder of this paper explores specific security issues that public cloud providers must address on their Hyper-V clouds, and provides strategies that address the various security issues while also taking cost and efficiency into account.

Malware

One of the most important issues that public cloud hosting providers must contend with is that of malware prevention. Unfortunately, the traditional approach to malware detection and removal is unsuitable for public cloud environments.

Legacy malware prevention software typically relies on a scanning agent that must be installed on each server or on each virtual machine. It is this seemingly simple concept that leads to a number of different problems.

As previously mentioned, one of the major challenges in securing a public hosting environment is the fact that it is difficult to secure what you do not fully control. Hosting providers typically use templates as a mechanism for allowing customers to create virtual machines on an as needed basis. Although it is possible to include a malware scanning agent within a virtual machine template, there is nothing stopping a tenant from removing or disabling the agent after a virtual machine has been created.

The fact that customers can disable or even remove antivirus agents is disheartening enough, but there are other problems associated with legacy anti-malware techniques. One such problem is that of performance.

In a hosting environment large numbers of virtual machines must share a finite pool of physical hardware resources. That being the case, the only way to achieve a high level of virtual machine density is to take steps to ensure efficient hardware usage, and to put controls in place to prevent resource contention from becoming problematic.

Legacy antivirus software simply does not work well in virtual server environments. The scanning process is I/O intensive and may also consume significant CPU resources. While this may not be an issue for isolated scans, most antivirus software is designed to perform periodic scans on a scheduled basis. The problem with this is that the scanning process can lead to AV storms. Simultaneous scans across multiple virtual machines can consume an unacceptable level of hardware resources. It is possible to circumvent this resource consumption by putting controls in place, but doing so may lead to severely degraded performance for the virtual machine that is being actively scanned.

Malware Prevention Best Practices

Although there are clearly number of challenges associated with providing comprehensive malware protection in a hosted public cloud environment, there are some best practices that you can adhere to in an effort to provide malware protection without risking severely degraded performance, and without having to worry about end-users circumventing your security measures.

One such best practice is to avoid the use of anti-malware agents. The various advantages of going with an agentless solution will be discussed in more detail later on, but one of the main benefits is that by going agentless you remove the customer's ability to disable the agent. As such, you can ensure that every virtual machine is being protected against malware.

Of course this raises the question of how you can enforce malware protection if there are no anti-malware scanning agents running on the virtual machines. The trick is to focus on the hypervisor, rather than on the individual virtual machines. Anti-malware software should be moved to the virtualization stack rather than running inside of the virtual machines. Doing so allows the software to have a more comprehensive view of the system than it would if it were confined to the boundaries imposed by virtual machines.

One big caveat to this approach is that the anti-malware software that you are using needs to be truly Hyper-V aware. Moving the malware scanning process out of the virtual machines and into the hypervisor involves more than just running antivirus software on the host operating system. In fact, Microsoft recommends excluding folders containing Hyper-V virtual machines from being scanned by anti-malware software that is running in the host operating system. Similarly, Microsoft recommends excluding the VMMS.EXE and the VMSWP.EXE processes. Failure to exclude these folders and processes can result in virtual machine corruption.

Needless to say, this presents something of a security paradox. Failing to scan Hyper-V processes and virtual machines exposes the organization to malware threats. Conversely, scanning those virtual machines and processes risks corrupting the virtual machines. So what is a provider to do?

The solution is to adopt an anti-malware application that has been specifically engineered for Hyper-V. Such software should run on the host server and integrate directly with the virtualization stack. Doing so would allow the software to intercept malware before it is written to a virtual machine, and do so in a way that eliminates the risk of virtual machine corruption.

One of the really nice things about this type of approach is that it has the potential to greatly reduce the amount of time needed for performing anti-malware scans. The traditional approach to anti-malware scanning involves performing a comprehensive scan of the full contents of a virtual machine, and comparing each file to a signature database in an effort to detect malware infections. When malware scanning is performed at the virtualization stack level, then it becomes possible to perform incremental scans. In other words, the anti-malware software is able to keep track of the blocks that have been written since the last scan occurred, and then run the scan against only those blocks. Previously existing storage blocks are assumed to be clean since they have already been scanned.

This particular approach to antimalware scanning not only reduces scan times, but it also dramatically decreases the load that the scanning process exerts on the server, thereby minimizing performance degradation.

Although this approach to anti-malware scanning works extremely well for malware prevention, it is still theoretically possible for a malware infection to occur. Imagine for example that a tenant's virtual machine was infected by a previously unknown piece of malware. Because the malware was previously unknown, the antivirus software may not be able to detect it. Eventually however, the antivirus software would presumably be updated and consequently become aware of the previously unknown malware. The problem is however, that the malware is already infecting a virtual machine. If the security software were to rely solely on incremental scans then the malware infection might never be detected and the malware might never be removed. As such, anti-malware software needs a way to initiate a manual scan. Of course there are a few criteria that the scan should ideally adhere to.

First, the anti-malware software needs to be able to stagger the scanning process so that scanning does not occur simultaneously across all virtual machines. Otherwise, resource contention could become a huge issue. Similarly, the anti-malware software should impose scanning thresholds to prevent the scanning process from consuming excessive system resources.

A second criteria is that the antimalware software needs to have a centralized management and reporting console. A large-scale public cloud could potentially contain millions of virtual machines, and it would be next to impossible to track anti-malware scanning across all of those virtual machines without having a management console that was specifically designed for use in large-scale environments.

The third criteria is probably the simplest, but it is also one of the easiest to accidentally overlook. The anti-malware software that you are using needs to be compatible with any operating system that could potentially be used within a virtual machine. Public cloud hosting providers typically provide customers with predefined templates that allow them to deploy various types of virtual machines. Providers commonly offer a mixture of Linux and Windows VMs. In such an environment, an anti-malware solution that is not Linux compatible would result in unprotected exposure points.

Considerations for Anti Malware Agents

As previously discussed, it is in a public cloud hosting provider's best interest to avoid the use of anti-malware agents at the virtual machine level and focus instead on implementing malware protection at the virtualization stack. One of the reasons that was cited for doing so was that performing malware scanning at the hypervisor level eliminates the possibility of a tenant removing or disabling an anti-malware agent.

Of course this raises the question of why a tenant might disable anti-malware protection. Even if a tenant does not have malicious intent, there are any number of reasons why a tenant might choose to disable malware protection. Some for instance, might uninstall an anti-malware agent as a way of squeezing better performance out of a virtual machine.

There have also been situations in which tenants uninstalled anti-malware agents as a way of reducing costs. Public cloud hosting providers usually charge tenants based on the resources that they consume. Eliminating an anti-malware agent reduces resource consumption, potentially resulting in a lower bill.

Another common complaint from tenants is that certain types of agents are simply too intrusive. Some agents for example, attempt to install browser toolbars and modify the browser's homepage, among other things.

Similarly, a tenant might remove an anti-malware agent because they prefer to use a different anti-malware program, or because the agent is known to have compatibility issues with an application that they want to run inside of a virtual machine.

The point is that there are any number of reasons why a tenant might possibly disable or uninstall an anti-malware agent, thereby rendering themselves vulnerable to a malware attack. Going with an agentless solution is a great way to prevent this type of tenant activity. Even so, preventing anti-malware removal is not the only reason why an agentless solution is desirable.

At the beginning of this white paper, cost and efficiency were cited as to the factors that needed to be considered with regard to security. Both of these issues become factors in agent-based anti-malware products. Cost can become a factor because many vendors license their anti-malware products on a per agent basis.

Efficiency is also a factor with agent-based solutions. Not only do scanning agents place significant loads on the system and increase the problem of resource contention, but they also cause management inefficiencies. When an agent-based solution is used, administrators must perform a variety of agent related tasks, such as keeping the agents up to date or replacing missing agents. An agentless approach greatly reduces the administrative burden.

Firewall

Anti-malware scanning is extremely important, but it is only one part of a comprehensive security solution. Another critical part is a good firewall. While it is true that Windows operating systems include a built-in firewall, the Windows Firewall is inadequate for public hosting environments.

One of the main problems with the Windows firewall is that it takes a unidimensional approach to security. A firewall running on a host server for example, would be able to monitor traffic along internal and external network connections, but it would not analyze traffic on private networks. Cloud providers depend on a complex structure of interwoven physical, logical, Virtual network segments. Most firewall software is inadequate for protecting these types of environments because the firewall protects only the networks that are immediately visible to the operating system. Multilayer virtual network structures simply cannot be protected with such software.

The only way to adequately secure a Hyper-V based public cloud is to use firewall software that was specifically designed for Hyper-V, and that latches onto the Hyper-V extensible switch. Such a firewall should make use of the Windows Filter Platform, and to be able to inspect, drop, modify, or insert packets into a traffic stream. The bottom line is that unless a firewall knows how to make use of the Hyper-V extensible switch, it will inevitably have blind spots.

Another reason why full virtual network awareness is essential is because hosting providers must be able to guarantee complete isolation to tenants, even if those tenant's virtual machines reside on a common host. Unless this level of isolation can be guaranteed by the firewall and the underlying virtual network stack, the hosting provider will be forced to implement isolation through VLANs. This may require the manual creation of hundreds of separate VLANs, which greatly complicates and increases the cost of management.

Although it is critically important for a Hyper-V firewall to be virtual network aware, it is equally important for the firewall to perform all of the basic functionality that would be expected in a physical data center. Specifically, the firewall should be application aware and should be able to perform stateful packet inspections. Another important criteria is that the firewall needs to be manageable in a large-scale environment. The reporting engine needs to be able to consolidate thousands of individual firewall status messages into a concise report that is easy to read and that is fully actionable. At the same time however, the firewall needs to allow for granular action to be taken. Administrators should be able to isolate and secure virtual machines by name, ID, or even tenant, on an as needed basis.

Intrusion Detection Systems

Another security mechanism that every public cloud provider needs is a good intrusion detection system. Although many providers rely on hardware based intrusion detection, it is also a good idea to supplement that protection with a software-based Intrusion Detection System that has been specifically engineered for Hyper-V environment.

A Hyper-V Intrusion Detection System needs to be able to monitor physical, logical, and virtual networks, as well as individual virtual machines in a multitenant space. The software should include a database of known attacks, but should also be able to use heuristics to monitor for previously unknown attack methods.

A Hyper-V based Intrusion Detection System needs to support cloud level scaling. Not only is this important from a reporting standpoint, but also from a responsiveness standpoint. An Intrusion Detection System is only effective if it is able to detect threats in real time, and then block the threat, and alert an administrator. This means that the Intrusion Detection System software must be designed in a way that prevents any significant lag in its detection capabilities, even as the cloud grows.

Compliance

One last consideration that needs to be taken into account is that of compliance. Although a public cloud provider might not necessarily be subject to any sort of regulatory compliance, they may have customers who are. A healthcare organization that is subject to HIPAA compliance for example, is only allowed to use cloud providers who are also certified to be HIPAA compliant.

Although regulatory compliance goes beyond the scope of this white paper, it is worth mentioning that the security software that a public cloud provider uses must provide centralized management capabilities, and allow for the use of policies, rules, filters, and log analytics. In other words, the public cloud provider needs a way to perform a comprehensive security audit that will satisfy regulatory authorities should doing so become necessary.

Conclusion

The Internet has been described as the most hostile environment imaginable, and public cloud providers make easy targets. It is therefore critically important for public cloud providers to invest in security solutions that will keep their tenants safe. The only way for a provider to establish comprehensive security is to practice defense in depth by using a security package that is able to protect against numerous different types of threats across all layers of the cloud infrastructure.

About The Author



Brien M. Posey



@BrienPosey

Brien Posey is a freelance technical writer who has received Microsoft's MVP award eleven times for his work with Exchange Server, Windows Server, IIS and File Systems Storage.

Brien has written or contributed to about three dozen books and has written well over 4,000 technical articles and white papers for a variety of printed publications and web sites.

In addition to his writing, Brien routinely speaks at IT conferences and is involved in a wide variety of other technology-related projects. Before becoming a freelance technical writer, Brien worked as a CIO for a national chain of hospitals and healthcare facilities. He has also served as a network administrator for some of the nation's largest insurance companies and for the Department of Defense at Fort Knox



+1 617 982-1261
+44 (20) 7048-2021



info@5nine.com



www.5nine.com