WHITE PAPER–APRIL 2019

VMware AppDefense: Security Inside the Perimeter

vmware®

As an industry, cybersecurity focuses on threats. It's easy to see why. Considering the ever-increasing number of cyberattacks in 2018 (up again in 2019) and the estimates of damages from these attacks—by some estimates, as much as $6 trillion by 2021[1]—it's easy to see why. But while this focus makes for great headlines, it doesn't make for great solutions. Although there weren't as many catastrophically large breaches in 2018 as there have been in recent years, there were enough to add hundreds of millions of people and thousands of businesses and organizations to the rolls of those who have been hacked.

Despite this focus, modest IT spending on cybersecurity, while rising, doesn't seem to have resulted in much of an improvement; that is, we aren't any safer today, and one could argue that we're actually much less safe than we've ever been. Here's why:

Threats constantly evolve

Attackers constantly adjust and change their techniques to find new ways to subvert any known controls to look for new ways to infiltrate new targets. For instance, in 2017, crypto-mining attacks increased by more than 8,000 percent, according to Symantec's latest threat report.[2]

The attack surface gets bigger every day

New applications are constantly being developed, and with new software comes new issues: new attack surfaces, new platforms, all with new security holes and new opportunities for attacks. There will be more mobile attacks, for example. And with more people online than ever before, and more enterprise-borne vulnerabilities such as containers and cloud-native apps, attackers have the most target-rich environment. Those Internet of Things (IoT) devices we're beginning to take for granted, for instance, will grow to more than 200 billion by 2020.[3] All of this presents a new attack surface.

Attackers control the game

Threat-based protection has stopped billions of attacks every year and will continue to do so. But, enough attacks also get through to wreak havoc. The fact is that even if threat-based protection could be designed to recognize and stop every attack that is ever going to be launched, you would still be looking backward at what has been done and not forward at what might or could be done. In this ever-evolving game of cat and mouse, there's little doubt as to who the cat is and what happens to the mice.

The reality is that if your organization relies solely on threat-based protection, you're employing an incomplete solution. And as such, attackers will consistently take advantage of you by designing threats that elude your security posture.

That's because, traditionally, threat-based security is architected with a hard exterior and a soft interior. The idea here is that trying to stop infiltration at the figurative front door—the infiltration point—is the best way to stop attacks from being successful. But that reasoning, while popular, obviously doesn't solve the problem. Attackers can launch attacks over and over again, and they only need a small percentage to get through, but to defend your organization, you have to catch 100 percent of their attempts, an obvious impossibility. This approach also begs the question that the infiltration point—the point at which the attacker initially penetrates your environment—is the best place to blunt the attack. It is not.

1 Cybersecurity Ventures. "Cybercrime Damages $6 Trillion By 2021." Steve Morgan. December 7, 2018.

2 Symantec. Internet Security Threat Report, Volume 23. April 2018.

3 Intel. A Guide to the Internet of Things.

Even though the perimeter, with its many infiltration points, is where everyone spends their money, it's also the most porous, most difficult thing to secure. Infiltration happens because end users open emails they shouldn't. It happens through open ports in your firewalls that your company needs to keep open to do business. It happens from people browsing the Internet while inside the enterprise. So the idea that you're going to be able to stop all these attempts at the perimeter is obviously wrong. And when you put these two ideas together—focusing on a threat-based approach and trying to stop attacks at the perimeter before they get in—you're simply bound to fail. It's an asymmetric battle, one that attackers are going to win 99.9 percent of the time.

Taking up a new line of defense inside the perimeter

VMware is initiating a new approach: layering security, protection, visibility, and control at points farther down the kill chain. Assuming attackers can and will get past the infiltration point on your perimeter, layering protection inside that just makes sense. Even if an attacker breaches your perimeter, they still have to propagate the attack by moving laterally through your environment to get to your data and exfiltrate it for it to be considered a data breach. Moving robust protection into these later stages of the kill chain gives you the chance to stop the attack before the real damage occurs. In this case, even if your perimeter defenses are penetrated, having another layer of security inside them lets you change the game on the attacker. Inside the perimeter, you have the advantage because you know more about your data center than the attacker does. From this fortified position, you can lock down your environment and protect yourself from a full-on data breach.

By retreating to a more defensible position, you can stop playing on the attacker's turf. By focusing your interior defenses to protect the things you know best—your data center applications—you can stop attackers in the one place you have control, as opposed to out in the wild where you're just another bug waiting for a windshield. By taking back the advantage in this way, you control the battlefield, not the attacker.

By flipping the script of an attack in this way, VMware believes we have solved a major hole in the security industry by replacing the uncertainty, but inevitability, of a data breach with security and protection of a known good, least privilege environment in which attackers simply can't hide. In this environment, when they do the things attackers do, they flag their own activity, stand out, and give themselves away. The moment they do that, the game is over, and they've lost.

A full, zero-trust environment is no longer beyond your capabilities

Few people would argue that a full, zero-trust environment would not be ideal. If only it weren't so impossibly difficult to implement. There's so much to do, so much operational overhead involved, it's just not efficient or economical, and it simply gets in the way of the business. But if you could remove the operational complexity, eliminate the cumbersome and time-consuming process of manual policy creation, and automate the whole process, that would make creating this powerfully protective environment finally worthwhile. That's exactly what VMware has done.

Introducing VMware AppDefense: Implement a least privilege environment with the least effort

The first benefit VMware AppDefense™ delivers is to remove the operational burden involved in implementing a least privilege, zero-trust environment. You don't have to worry about manual policy creation because that isn't how policies are created. The AppDefense system does that for you. You don't have to inspect the data center to profile what behavior is normal across the entire range of potential events.

AppDefense does that for you, as well. You don't even have to figure out how to properly classify change events in your applications. AppDefense does that for you, too.

The VMware App Verification Cloud: Where AppDefense gets its smarts

The overridingly important capability of a least privilege environment is its ability to classify events. This involves determining whether a particular event occurs normally in the operation of a particular application and, if that event isn't deemed to be normal, figuring out what to do about it. It sounds fairly straightforward but, as it turns out, it's difficult to get the appropriate visibility to ascertain what even happened, let alone whether that event is normal or potentially malicious. With a threat-based approach that looks at all previous threats, when a new threat approaches, it will look normal, so it's likely that it is going to be ignored, allowing attackers to get into the environment. The opposite approach—a pure all-deny approach—means that everything that's potentially good is going to look bad, resulting in a raft of false positives and a policy that won't let anything through.

Nobody knows more about how applications behave than VMware, because nobody sees more running applications. We actually pull all the behavior of every application from all AppDefense customers across the entire VMware population of nearly 65 million servers. We compare the behavior we see in one customer and aggregate it across the entire VMware universe. That lets us model the behavior of your applications against the behavior of the population. When we see enough of the population exhibiting similar behavior, we'll start to verify that as accepted known good behavior.

Under the hood, there are machine learning and distributed consensus algorithms, along with an omnipresent security analyst sitting above your data centers looking at everything and helping you classify everything as good or not good, based on what we know and constantly learn. When we see an application attempt to execute, we will leverage the consensus to verify it as reputable. When the application begins to communicate over the network, we will take that data and run it against a machine learning algorithm to see how closely it matches the rest of the population. Classic clustering algorithms identify outliers in population analysis and tell us if there is something too far out of the norm, and if so, we'll flag that as suspicious. Armed with that analysis, we then give you a range of choices for how to respond, which includes blocking, alerting, and a visibility mode that lets you monitor everything. You also have the ability to change those responses depending on the application you want to protect.

Continuous analysis + continuous verification = continuous assurance

Every other method of creating a zero-trust environment requires that you write rules and create policy manually. That's fine if there are only a few moving parts. But with today's applications, you can have thousands and thousands of events. And you simply can't write rules for them all.

But we can because the App Verification Cloud lets the actions of thousands of applications running across the VMware population tell us what is and isn't good.

For a competitor to offer this level of visibility and control, they would need to have a fleet of agents collecting the data and feeding it back. But agents are difficult to deploy and eat up resources. You would need more of them than could be easily managed, each one creating a new attack surface.

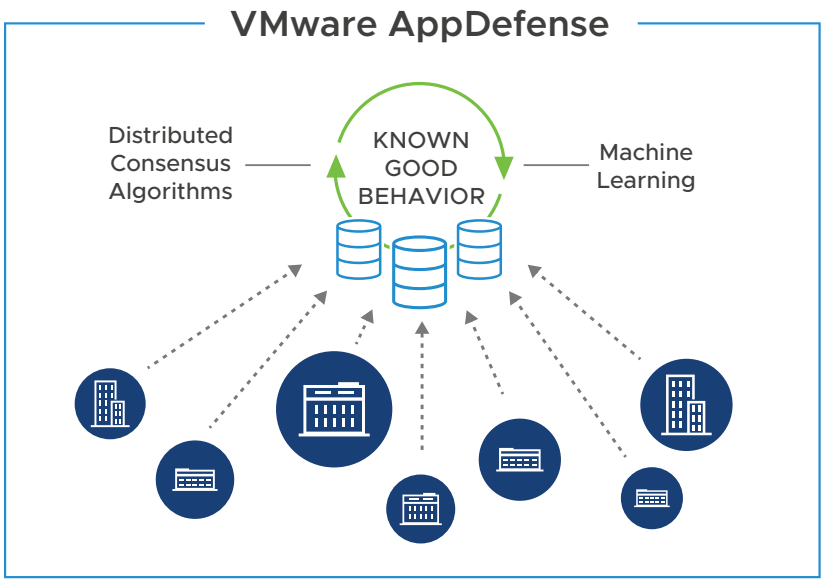

By comparison, we run this entire process in an agentless way with VMware Tools™, without having to run anything else. This not only solves the operational downside agents typically present by having to manage their lifecycle, but it virtually eliminates the performance overhead typically associated with creating zero-trust environments.

You can believe in it, too, because the VMware ESXi™ hypervisor runs beneath these virtual machines (VMs), so we can actually operate AppDefense from a higher degree of integrity. This is because ESXi uses virtual memory inspection to create a kind of tamper-evident seal that tells us if someone is trying to turn us off or manipulate AppDefense running in the guest VM. And if they are, we'll both know about it. That kind of automation creates visibility, control, and peace of mind.

Even so, we haven't turned everything over to the machines just yet. AppDefense is delivered as a service, so there's always human oversight from our side. We continually evaluate what we see across tenants and share it with you to make your particular deployment better. Everything that gets flagged goes into a forensic log you can examine. We'll tell you the score we gave it and alert you as to how dangerous or harmless it is.

All operation and management is done by us. We provision your account in the cloud. Your team will need to install the AppDefense software, but typically, you will be up and running in an hour.

Your secrets are safe with us

We're not performing full packet capture, so none of the contents of your applications are stored within AppDefense. That means there are no passwords, no credit card numbers, and no personal information exposed. We use metadata to make security decisions, such as process hash, file paths, machine names, and port/protocol info. Process info is encrypted within the cloud, and it's strictly a one-way communication channel from your operation to ours.

See how good your security can be

For more information about how VMware AppDefense can solve your security issues, have your team contact your VMware representative. And in the meantime, take heart: Fortifications have arrived.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com Item No: 225494aq-wp-appdefense-summary-uslet 4/19