WHITE PAPER–MARCH 2020

Five Critical Requirements for Internal Firewalling in the Data Center

Why traditional perimeter firewalls are becoming obsolete for protecting east-west traffic

vmware®

Table of contents

"[When] multiple [tactics, techniques and procedures] are utilized in concert… cybercriminals are able to gain and maintain access to a computer network, no matter their motives. Once they are inside a network their process is almost always the same: establish continued access, escalate or obtain administrator privileges, move slowly and quietly to map the entire network, look for open ports, locate the 'crown jewels,' and exfiltrate the data undetected for as long as possible."[4]

MICHAEL D'AMBROSIO
DEPUTY ASSISTANT DIRECTOR
UNITED STATES SECRET SERVICE

Introduction

No organization wants to see its name in the same headline as the words "massive data breach." Yet, day after day, companies of all sizes, as well as nonprofits and government agencies, continue to make the news as cybercriminals and malicious insiders breach their defenses to exfiltrate sensitive data. Research firm Forrester Consulting reports that 58 percent of companies faced a significant security incident in 2019 despite spending more to secure their networks.[1]

Clearly, traditional defenses such as perimeter firewalls aren't enough to thwart successful attacks. In fact, according to a Forrester survey commissioned by VMware, seven out of 10 enterprises are handicapped by an overreliance on perimeter firewalls.[2]

The perimeter has become highly permeable and, once breached, perimeter defenses can't stop an attacker from moving laterally inside the corporate network to reach and exfiltrate records. At the same time, attacks involving insiders, who are already within the perimeter, account for a growing percentage of breaches.

Instead of relying on perimeter-based security, organizations must focus on monitoring, detecting and blocking malicious internal traffic as a core component of their IT security strategy. This requires an internal firewall approach specifically designed to protect large volumes of internal data center traffic without sacrificing security coverage, network performance or operational agility.

This white paper explains the difference between traditional perimeter firewalls and purpose-built, software-based internal firewalls, and why the latter is best suited to protecting today's modern workloads.

The disappointing state of network security

In 2019, 15.1 billion records were exposed through more than 7,000 publicly reported breaches, making it yet another record-breaking year. This represented an increase in records exposed of more than 284 percent compared to 2018.[3]

In Verizon's 2019 data breach report, 69 percent of the breaches in its data set were perpetrated by outsiders.[4] These outside cyberattackers frequently employ tactics such as phishing to bypass perimeter firewalls and gain access to the internal network. They then move laterally to find and exfiltrate sensitive data.

External cybercriminals are also benefitting from an increased attack surface, courtesy of today's modern computing and application environments. As networks of workloads and microservices replace monolithic and three-tier applications, the attack surface, along with security complexity, expands exponentially.

To make matters worse, the percentage of breaches that involve internal actors has been steadily growing since 2015. In 2019, approximately 34 percent of the breaches on which Verizon reported involved internal actors.[4] These internal actors move through largely unmonitored network traffic within the data center to reach their targets.

The growing volume of east-west traffic

Network security controls created in the pre-DevOps, pre-distributed application era are simply inadequate for protecting today's workloads and microservices. Virtual machines (VMs) connect to other VMs, containers connect with other containers, workloads connect with other workloads and so on. All of this creates a great deal of network traffic within the enterprise.

1. Forrester Research. "Forrester Analytics Global Business Technographics® Security Survey, 2019." August 2019.
2. Forrester Consulting. "To Enable Zero Trust, Rethink Your Firewall Strategy." February 2020.
3. Risk Based Security. "Number of Records Exposed in 2019 Hits 15.1 Billion." February 10, 2020.
4. Verizon. "2019 Data Breach Investigations Report." May 2019.

FURTHER READING

For more information on data center traffic types and examples, read the white paper *"Knock, Knock: Is This Security Thing Working?"* from SANS.

To understand why the increased amount of internal traffic is an important factor for security, let's start by differentiating the two main types of traffic in the network today (see Figure 1):

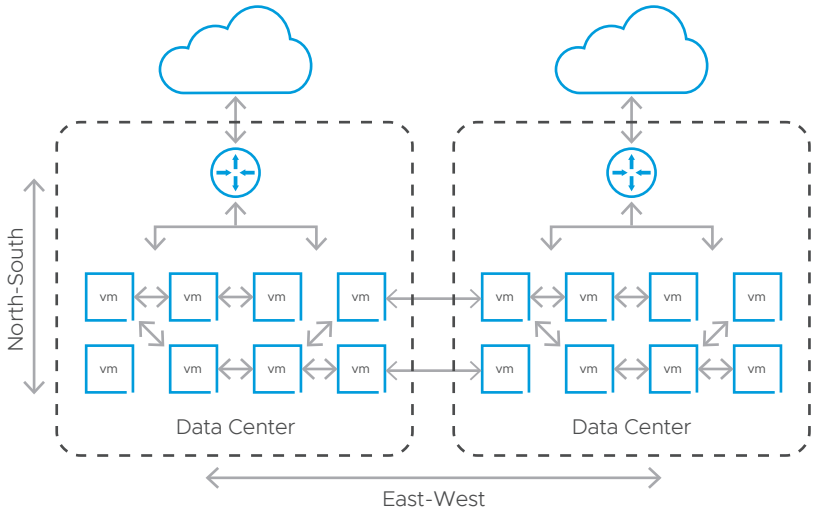

FIGURE 1: Data center traffic patterns.

- **North-south traffic** – This network traffic moves in and out of an organization's network; for example, to and from the internet. North-south traffic typically represents a much smaller percentage of the overall traffic on the network.
- **East-west traffic** – This traffic moves laterally (hence, east-west) across the data center, including workload-to-workload traffic (inter-data center, intra-data center, data center to public cloud, or public cloud to data center). As more monolithic applications are replaced with or rearchitected into distributed applications, the amount of east-west traffic (also known as internal traffic) has far surpassed that of north-south traffic.

A perimeter firewall only monitors north-south traffic. Yet, the lesson learned from the past decade of data breaches is that organizations cannot assume that east-west traffic can be trusted. Trusting all east-west traffic means that a cyberattacker who makes it through the perimeter firewall can then move undetected laterally within the network.

To properly defend against cyberthreats that breach the perimeter as well as malicious insiders, organizations should implement a distributed, internal firewall strategy. Internal firewalls proactively provide visibility and protection from internal threats, and minimize the damage from cyberattacks that make it past the traditional network perimeter.

The right firewall for the right type of traffic

As organizations realize they must focus greater attention, budget and efforts on improving network security, many make the mistake of using traditional perimeter firewalls designed to monitor north-south traffic to protect their internal networks. While it may be tempting to do so, provisioning perimeter firewalls for east-west traffic monitoring is not only expensive, it's highly ineffective in delivering the level of control and performance required to protect large numbers of dynamic workloads.

Distributed, granular enforcement

While both perimeter and internal firewalls enforce security policies by monitoring and blocking potential threats, the characteristics of east-west traffic and the network topology mean the enforcement approach must be different for an internal firewall.

For a perimeter firewall, it's acceptable to block traffic based on ports, protocols and IP addresses, or to identify traffic to or from a specific application, such as Skype.

On the other hand, an internal firewall needs to operate at a more granular level, that of individual workloads within an application. Using a three-tier application as an example, an internal firewall permits traffic between the web tier and the app tier of the application, and between the app tier and database tier of the same application. However, it blocks the traffic from the web tier to the database tier because this traffic should not exist in the normal course of operations.

Thus, the granularity of enforcement required of an internal firewall is much higher than that for a perimeter firewall. A typical perimeter firewall won't know that (in the example in the previous paragraph) the three tiers belong to the same application, but some traffic is permitted while other traffic is not within that application.

Scale and throughput

Centralized monitoring of north-south traffic using a perimeter firewall doesn't typically create performance bottlenecks because the volume isn't nearly as large as it is for east-west traffic. However, most enterprises have significantly more east-west traffic than north-south.

If an enterprise uses a perimeter firewall for east-west traffic and wants to inspect all (or most) of the traffic, it will have to deploy many perimeter firewalls to meet its throughput requirements. This can significantly increase the cost and complexity of the network security infrastructure. That's why, in practice, most organizations using perimeter firewalls to monitor east-west traffic don't inspect most of it—the cost and constraints to do so are simply too great.

For internal firewalls, a distributed enforcement approach is substantially more cost effective while delivering the scalability and performance needed. A distributed internal firewall is elastic and supports autoscaling as workloads are spun up or down. As the number of workloads expand, the internal firewall capacity expands automatically. As more servers are used to support workload expansion, a small portion of the server's capacity is then used for security controls, allowing the internal firewall to scale accordingly.

Infrastructure impact

If a perimeter firewall solution is used to monitor east-west traffic, the traffic is forced to and from a centralized appliance or capability. This creates a hair-pin pattern, which uses an inordinate amount of network resources in the process.

In addition to increasing latency, hair-pinning internal network traffic adds complexity, both from a network design as well as a network operations perspective. Networks must be designed to take into account the additional (hair-pinning) traffic routed through a perimeter firewall. From the operational side, the security operations team must adhere to the network design and be aware of constraints when sending additional traffic for inspection to the firewall.

Alternatively, a distributed internal firewall approach allows monitoring of large volumes of east-west traffic without creating a single chokepoint. A distributed architecture moves enforcement close to the data rather than the other way around, and secures all east-west traffic while maintaining a low impact on the network and server infrastructure. No hair-pinning of traffic occurs, which eliminates the complexity and latency issues involved in using perimeter firewalls to monitor the internal network.

OVERRELIANCE ON PERIMETER FIREWALLS

According to a Forrester Consulting survey, more than 75 percent of companies depend on virtual or physical perimeter firewalls to secure internal network traffic. However, 72 percent believe their overreliance on perimeter firewalls is a significant challenge to the security of their internal network.[2]

Intra-application visibility

Monitoring east-west traffic and enforcing granular policies requires visibility down to the workload level. Standard perimeter firewalls do not have clear visibility into the communication patterns between the workloads and microservices making up modern, distributed applications. This lack of visibility into application flows makes it extremely challenging to create (and enforce) rules at the workload or individual traffic flow level.

In comparison, an internal firewall should be able to automatically determine the communication pattern between workloads and microservices, make security policy recommendations based on the pattern, and check that traffic flows conform to deployed policies (i.e., enforce granular policies). A robust internal firewall solution can discover and visualize application topology, processes, acceptable state, application users and devices being used.

Policy lifecycle and mobility management

Traditional firewall management planes are designed to handle dozens of discrete firewalls but are not designed to support workload mobility with automatic reconfiguration of security policies. Therefore, when a perimeter firewall is used as an internal firewall, network and security operators must manually create new security policies whenever a new workload is created, and modify these policies when a workload is moved or decommissioned.

The management plane for internal firewalls is designed to manage tens of thousands of entities (including virtual switches and distributed firewalls) while accommodating policy lifecycle management and workload mobility. The internal firewall automatically adjusts security policies when a workload is created or decommissioned without manual intervention. It supports stateful workload mobility across the infrastructure with seamless forwarding of traffic to the new location and security policies that move automatically with the workload's VM.

Internal firewall must-haves

If traditional perimeter firewalls are not appropriate or effective as internal firewalls, what type of solution is best suited for monitoring east-west traffic? Summarizing the requirements from the previous section, an internal firewall approach must be able to support:

- Distributed and granular enforcement of security policies
- Scalability and throughput to handle large volumes of traffic without impeding performance
- A low impact on network and server infrastructure
- Intra-application visibility
- Workload mobility and automatic policy management

A perimeter firewall cannot deliver on these requirements without incurring exceptionally high costs and complexity while requiring too many security compromises. Instead, a distributed, software-defined approach is the most effective way to implement internal firewalls to monitor east-west traffic. The right software-defined, internal firewall approach delivers the scalability, cost-effectiveness and efficiency to secure tens of thousands of individual workloads across thousands of applications.

Yet, not all software-defined approaches can provide the level of internal network protection enterprises need to secure their sensitive workloads without sacrificing granular controls, consistency and flexibility. To achieve optimal security coverage, network performance and operational agility, organizations should seek out a purpose-built, internal firewall solution that offers intrinsic security, which is built into the infrastructure, distributed and application aware. To learn more about intrinsic security, read the white paper *"Knock, Knock: Is This Security Thing Working?"* from SANS.

Important use cases for internal firewalls

As more companies realize the limitations of perimeter-based security and the likelihood of malicious traffic moving undetected through the internal network, they're adopting a purpose-built, software-defined internal firewall approach to improve their overall security stance and protect against cyberthreats. Some of the most important use cases for an internal firewall strategy include the following:

- **Virtual security zones** – Internal firewalls can be used to support macro-segmentation of business units, partners, development from production environments and other security requirements. With a software-defined approach to internal firewalls, organizations can create and manage virtual security zones without the expense and effort of purchasing, configuring and maintaining physical appliances.
- **Lateral movement detection** – Inspecting all east-west traffic makes it possible to detect lateral movement early and limit its damage. Granular policies at the workload level help internal firewalls block cybercriminals' attempts to move laterally within the network to reach their targets.
- **Regulatory compliance** – To meet compliance requirements such as the Health Insurance Portability and Accountability Act (HIPAA), the Payment Card Industry Data Security Standard (PCI DSS) and the Sarbanes-Oxley Act (SOX), a distributed internal firewall approach helps companies achieve compliance by propagating regulation-specific security policies to all relevant workloads, and tracking traffic flows to and from sensitive applications. Software-based internal firewalls also eliminate the need to buy and deploy discrete appliances to support compliance.
- **Zero trust using micro-segmentation** – The zero trust approach assumes that all traffic should not be trusted until policy proves otherwise. Micro-segmentation is a core concept within a zero trust approach to isolate workloads and secure them separately. In support of a micro-segmentation approach, internal firewalls allow organizations to logically divide the data center into distinct security segments down to the individual workload level and then define controls for each unique segment.
- **Security tool consolidation** – Software-defined internal firewalls enable organizations to eliminate multiple security appliances and rein in appliance sprawl as applications become more distributed. Purchasing and managing fewer appliances reduces the cost of ownership and simplifies security operations.
- **Visibility** – Network and security operations teams need insight and context into all workload traffic to eliminate security blind spots, and accelerate incident investigation and remediation. The right internal firewall solution delivers 360-degree visibility into every workload, uses this visibility to determine expected behavior of applications and automatically generates security policies to enforce known good behavior.

Conclusion

To reverse the pace and volume of data breaches, enterprises must focus on securing all their east-west traffic. They can no longer afford to assume that perimeter defenses will be enough and that traffic within the network can be trusted.

A software-defined solution that is built into the infrastructure, distributed and application aware is the most effective way to improve security, reduce costs and simplify operations. The only solution built into the infrastructure, *VMware Service-defined Firewall* is designed to protect east-west network traffic across multi-cloud environments. By making security intrinsic to the infrastructure and virtualizing the entire security stack, the Service-defined Firewall enables security teams to mitigate risk, ensure compliance and simplify the operational model of firewalling every workload.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 vmware.com Item No: 486127aq-wp-five-reqs-intrnl-fw-dc-uslet-102 3/20