FIREMON

5 STEPS TO KEEP NETWORK SECURITY ENFORCEMENT POINTS SECURE AND UP-TO-DATE

By Randy Franklin Smith, Ultimate Windows Security

A basic truth in network security is that it's never static. Changes in the network, applications, business strategy, cyberattack vectors and methods, and even the security technologies used all impact your network security configuration. The goal is always to reduce the potential threat surface to as close to zero as is possible, which means you must constantly update firewall rues as you *deploy, move, or retire workloads*, allowing necessary access but nothing more.

In principal, your network security should aim to limit the number of ways workloads can communicate while still allowing them to properly function. Security models such as *Least Privilege* or *Zero Trust* take the default stance of reducing access down to bare bones—opening access based only on business need—and reflect this in firewall rules. Technologies like microsegmentation can assist with this restricted form of access, but you want to have the same limitations in play for both virtual and physical network devices.

Ensuring policy consistency requires having detailed context around who needs to talk to who, where they are coming from, what they need access to, over what ports, etc. These details help to create enforceable points necessary to keep the organization both secure and productive.

So, where should you start?

In this whitepaper, we'll discuss five steps you can take to understand the current state of your network connections and ensure that the enforcement points designed to protect the environment remain up-to-date and carried out.

FIREMON: FOCUSED ON NETWORK SECURITY

Businesses today require a complex mix of disparate environments, networks, and devices, making network security a difficult task. At the same time, the threat of cyberattack, insider threats, and compliance violations require organizations to ensure that every part of the business environment is secure. FireMon's *Global Policy Controller* enables organizations to visualize the changing state of workload connectivity and orchestrate the enforcement of continual security across a hybrid infrastructure.

Step 1: Identify New Network Connections

A key challenge in network security is maintaining awareness of the changing network environment. For firewalls, this includes changes in logical network connections between two systems with at least one of them existing logically in your environment. New network connections arise whether a security team has identified them or not. Recognizing when a new connection is required (or occurs without sanction) is an important part of implementing access restrictions in the form of firewall rules.

Some examples of new network connections include:

- **Dedicated and Site-to-site VPNs** – on-premises networks may require connectivity to subsidiaries, business partners, or virtual environments in the cloud, as well as between the workloads in each environment; one example is the use of virtual machines in Azure via Azure ExpressRoute
- **Cloud Application Gateways** – using the Azure example, an Azure Application Gateway helps to manage and optimize your web traffic to applications in Azure
- **Web Application Gateways** – each time that you implement a new web application, you may also be installing an Internet-facing gateway that can put the internal environment at risk
- **Cloud-Based Solutions Using Local Appliance**s – software that uses a virtual or physical appliance that connects to your network and calls home through your firewall can also create risk

In each case, the network security challenge is that you have extended your network. Connecting to any external environment requires continuous adjustments to security controls. For example, you must adjust firewall rules to allow a local workload to connect to one of the previously-mentioned external environments, but also adjust the rules when you retire the workload. Leaving a port open after you retire the workload that required the open port creates a vector to attack the network.

But it's not just your end of the connections you need to be thinking about; changes on the other end of a connection should be of concern. Take the example of a new virtual workload in Azure that has an application published to the Internet. You must ensure that the correct security policies are in place so that, should the web server become compromised by an attacker, the attacker will not gain access to the Azure virtual network or to your on-premises network.

Security controls should align with proven effective security models – such as *internal segmentation, least privilege,* and *zero trust* – that can address new connections that were previously unaccounted for and, to varying degrees, require the intervention of security or firewall teams. Your default assumption should be that *you don't know what you don't know* and that it is imperative to identify new connections in real-time.

So, how do you find out about these new connections?

There are a number of practical ways to identify new connections caused by changes in workloads. These include:

- **Network Monitoring** – this simply entails recognizing new applications, protocols used, and source/destination IP/port combinations
- **Network Footprintin**g – this requires analyzing the network (similar to activities performed in pen testing) to gather information on network activity and subsequently identify devices on the network
- **Vulnerability Scanning** – this function, normally performed to find potential points of exploit on systems and applications, can also help you spot new workloads and their connections
- **Change Control** – when a business process changes, you can infer the corresponding change to a workload, such as the workload being new, moved or retired

FIREMON INSIGHTS: MANAGING DISPARATE DEVICES

For every new network connection, location, data center, or cloud instance, there's yet *another* firewall or network device to manage. Only the savviest of network engineers are up to speed on every management console, making it difficult to ensure consistent security across your entire environment.

Global Policy Controller utilizes an abstraction layer that allows you to manage enforcement points by focusing on the *intent* (for example, allowing Antivirus to scan, as shown below), without requiring you to be an expert on all the disparate firewalls and network devices in line between the Antivirus application and the endpoints it manages.

Global Policy Controller automatically distributes changes to enforcement points, eliminating the need for manual, one-off management.

In every case, you'll need to assess the security needs of the new connection. That brings us to Step 2.

Step 2: Understand the Protocol Requirements

New workloads can require new network connections and application gateways, which in turn can change your organization's threat surface. So, you need to know which protocol

each workload requires. You can gather specific protocol requirements in a number of ways, such as:

- **Knowledge of the Workload** – familiarity with the applications involved is a reasonable start; for example, you already know DNS is on TCP port 53
- **Documentation** – vendors are getting more transparent about network protocol requirements for their applications, so consulting the documentation is a must
- **Network Monitoring** – sniffing packets to see what is going across the wire provides insight into the traffic patterns and frequency of a given workload

Using all three of these methods provides the most complete picture. For example, you may monitor the network traffic of a given application and find that it looks consistent. But the documentation may indicate that once every 24 hours a process kicks off that may perform a sync, a heartbeat test, etc., using a different port that your monitoring may not find. This is why sniffing should cover a full business/usage cycle of the application.

You should also leverage multiple sources to confirm your findings so that you can be certain your list of protocol requirements is accurate and complete.

FIREMON INSIGHTS: WHEN PORT 443 IS MORE THAN JUST HTTPS

Many applications, regardless of the technology, use TCP port 443 for encrypted traffic. On the plus side, it's easy to monitor a single port. But the encrypted traffic all looks the same going across the wire – even if it has nothing to do with HTTPS traffic. You still have to monitor known ports like 443 because, despite encryption being common, threat actors still scan publicly-accessible IP addresses looking for exposed connections to services, databases, etc. that are unencrypted over this port.

Step 3: Determine Who Needs to Communicate with the Workload

With a solid understanding of the networks you're connected to, the workloads that need to communicate, and the specific protocols in use, the next step is to understand the specific endpoints on either end of given connections. The goal is to avoid opening up the entire IP address range on either end and limit connectivity to just those systems deemed necessary.

This is especially important with inbound traffic – you don't want to allow unnecessary communication. For example, if a single endpoint within your network communicates with a cloud-based workload, it makes sense to facilitate just that one connection. If you were, instead, to allow traffic from any system in the cloud to any system in your network over a given workload's defined set of protocols, you'd be providing a means for an attacker to gain access to your network. Rather than focus on restricting access to one *internal* endpoint, it would benefit the organization's security stance to restrict the specific workloads on the externally-trusted network as well.

FIREMON INSIGHTS: WHO SHOULD BE DOCUMENTING ALL THIS?

Should security and/or network teams be keeping updated documentation of the allowed network connections for auditing purposes? Global policies, one-off configurations, and everything in between should, in concept, be documented to demonstrate that unsanctioned communications aren't happening and that the right policies are in place to keep it from happening. Teams should always be prepared for the "audit from hell," but building documentation is likely the *last* thing a security or network engineer wants to do.

In most cases, organizations are using the basics – Word or Excel documents, a homegrown database, etc. This is a step in the right direction, but it relies on the human element to keep it updated. It's just not realistic to keep this form of documentation current, given the velocity of changes made within an organization. What's needed is automatic documentation to ensure timeliness and accuracy.

FireMon's Global Policy Controller automatically documents all parts of the network involved in the communications of a given application or workload. In the example below, the enforcement points are visualized to show how traffic will be allowed or denied, the devices in line that need to be managed, and the current state of enforcement.

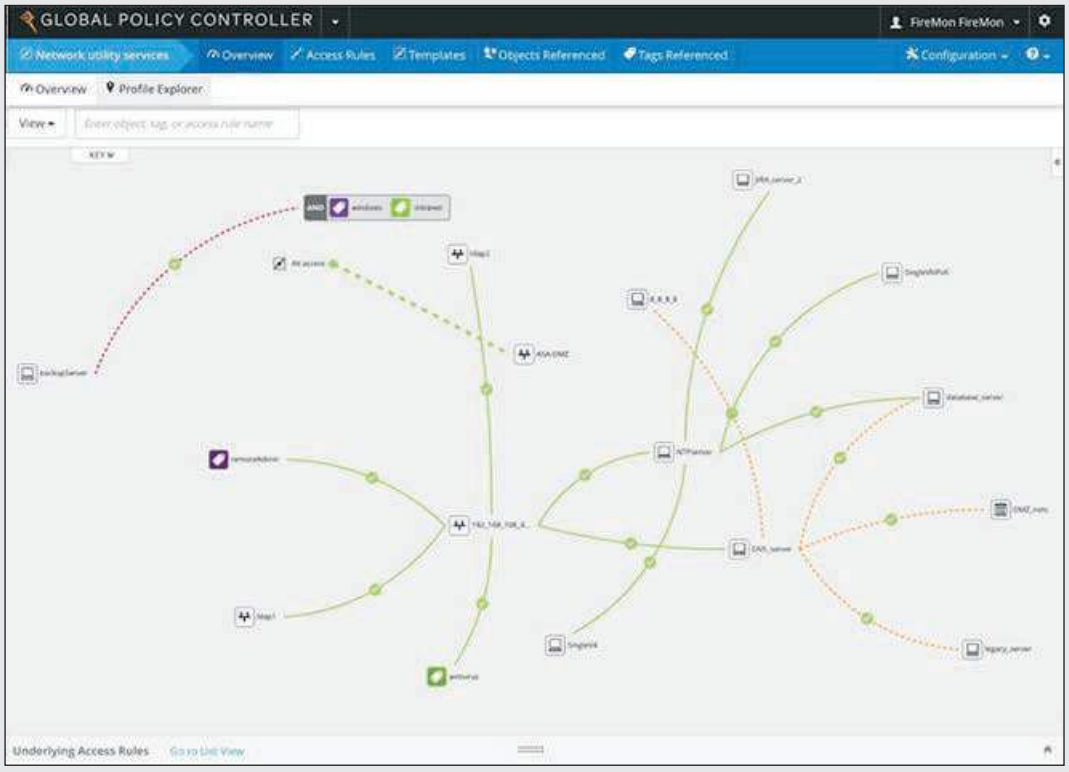

Information sources that can help you identify the specific endpoints that will be communicating are similar to those that aided your discovery of protocol detail. The difference is that here, you investigate at the system/service/client application level. Sources include:

- **Knowledge of the Workload Interaction** – institutional knowledge within IT can attest to the need for specific endpoints to communicate with certain workloads, both locally and on externally-trusted networks

FIREMON

- **Documentation** – documentation won't tell you which systems in your environment need to talk, but it can provide context of the various servers, services, and applications involved, which you can use as the basis for identifying specific system-to-system interactions
- **Network Monitoring** – as with protocol requirements, watching interactions between specific endpoints can provide useful details; using DNS names is better than using IP addresses because it provides better context of who and why (with DNS names serving as a form of documentation) and can more easily adapt to changes in configurations on the external network without impacting local firewall rules

Be sure to consider both ends of the communication; in some cases, your network may be providing the "server" side of an application and the externally-trusted network may have many clients. You should identify these individually to ensure that only specific endpoints can communicate to resources on your network.

FIREMON INSIGHTS: DON'T FORGET TO INCLUDE PEOPLE

In some cases, determining the specifics around the "who" may need to be less a technical discussion and more a business discussion. Remember, someone within the organization owns the use of a given data set, application or workload – a line of business owner, stakeholder, etc. who is personally interested in it. Your diligence gathering should include conversations with these individuals, as they can provide proper context around which machines on both ends are involved.

Step 4: Identify the Risk Differentials

Think of the overarching operational network environment – your network, any public cloud virtual environments in play, externally-trusted networks of partners, etc., and anything else that is part of your network environment. Now, logically think about those separate parts of the operational environment as *zones* – for example, your network, Azure, your partner's networks, and your production and development networks. Once you've done this, you need to analyze whether there is a risk differential between each zone. Ideally, each zone of your network is protected by some set of least privilege or zero trust firewall rules. But, realistically, zones are rarely equally restrictive.

These differentials in security create risk for your organization. So, it's important to determine how much risk the current configuration creates versus one that's more restrictive, regardless of whether the zones involved are at the same location or across the globe. Take the following example:

> You have one internal endpoint talking to a cloud-based endpoint with only the necessary ports open. The cloud endpoint may be owned by a partner with fewer safeguards in place, causing that external endpoint to be more susceptible to attacks (e.g., there's no antivirus in place) from another endpoint on the same externally-trusted network. Therefore, the risk level of the external endpoints is greater than that of your internal endpoint. Thus, the *risk differential*.

Identified risk differentials may require additional firewall rules that are a subset of available network controls. Even with a perfect Zero Trust firewall rules in place, the risk differentials can still warrant additional types of controls, such as malware protection, intrusion detection, or performing packet inspection for malware over an approved port.

When you find a risk differential, you must ensure that there is a compensating control in enforcement points.

Step 5: Identify any Required, Additional Enforcement Points

For each risk differential that requires addressing, determine whether you need to strengthen the security stance beyond just firewall policies that restrict traffic. You can put a number of enforcement points – security-focused safeguards that exist at the intersection between two security zones, endpoints, or workloads – in place to augment the security stance created by firewall rules. Examples of the use of additional enforcement points include:

- Placing a new cloud application gateway appliance *in front of an internal firewall* to ensure it can only communicate with select endpoints on your network
- Putting *network segmentation* in place between a new Internet-facing application deployed on a cloud-based virtual network and downstream resources
- Ensuring that an email connection that routes directly to a business partner (and, therefore, bypasses your email security gateway) *goes over a VPN* so that all communications are encrypted

FIREMON INSIGHTS: MIGRATING WORKLOADS TO THE CLOUD

Everything mentioned in this paper applies when you move an on-premises workload to the cloud. But you must update rules placed on existing enforcement points that allowed traffic to and from the old workload to reflect the (potentially) new network connection and the endpoints involved in communicating with the now cloud-based workload. Additionally, you should retire rules reflecting the old configuration, as leaving them in place slows down firewalls and serves as a distraction when security or firewall teams are trying to understand current policies.

The more difficult it is to manage, the more mistakes and risk. Leaving old rules in place creates a direct security risk when you deploy new workloads that reclaim the old address space; old rules designed to permit access for an old workload are likely inappropriate for the new workload.

Global Policy Controller leverages the concept of tags (shown below). Rather than requiring you to manage specific devices or applications, you can simply tag objects on the network, making global changes a simple task.

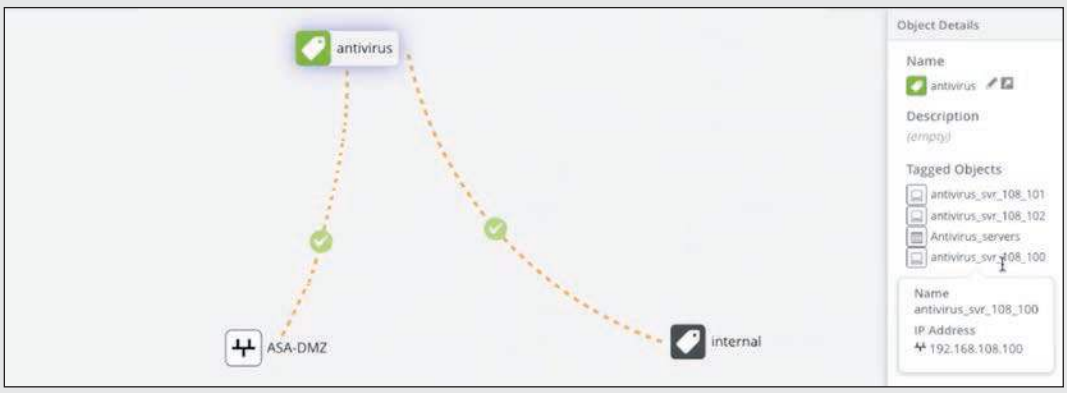

FIREMON

ABOUT RANDY FRANKLIN SMITH

Randy Franklin Smith is an internationally recognized expert on the security and control of Windows and AD security. Randy publishes **www.UltimateWindowsSecurity.com** and wrote The Windows Server 2008 Security Log Revealed—the only book devoted to the Windows security log. Randy is the creator of LOGbinder software, which makes cryptic application logs understandable and available to log-management and SIEM solutions. As a Certified Information Systems Auditor, Randy performs security reviews for clients ranging from small, privately held firms to Fortune 500 companies, national, and international organizations. Randy is also a Microsoft Security Most Valuable Professional.

ABOUT FIREMON

FireMon is the #1 network security management solution for hybrid enterprises. FireMon delivers continuous security for multi-cloud enterprise environments through a powerful fusion of vulnerability management, compliance and orchestration. Since creating the first-ever network security policy management solution, FireMon has continued to deliver real-time visibility into and control over complex network security infrastructures, policies and risk postures for more than 1,700 customers located in nearly 70 countries around the world.

FIREMON

Staying Secure in an Ever-Changing Environment

Your organization's network will never remain static; digital transformation, business innovation, and cloud adoption are all driving big changes in the way you operate. Disparate network environments elevate organizational risk, requiring centralized enforcement of security policies throughout every part of the network – regardless of whether it's virtual or physical, on-premises or in the cloud, in your organization or in that of an externally-trusted environment.

To ensure enforcement points are up-to-date, follow the five steps mentioned in this paper. You must identify new network connections and new application workloads in real-time, including specific protocols and endpoint requirements. Risk differentials help provide context around what enforcement points are necessary – whether simple firewall rules or other compensating controls.

In this chaotic environment, organizations need a global view of, and the ability to manage, enforcement points that meet the changing needs of the business without regard to the technologies in place. Given the variety of firewalls and other enforcement point solutions in play today, the best answer is to look to use of a third-party solution designed to help you implement the steps outlined in this paper and manage the necessary security changes they require.