10 Things to Test in Your Future Next-Generation Firewall

In the market for your next firewall? How do you navigate the risks and opportunities cybersecurity presents to your organization? How can you determine if the features of your new next-generation firewall are what your organization needs to grow and move forward?

The answer is simple: you test it.

Security professionals agree that organizational security should not be approached with a one-size-fits-all mindset. Every organization has unique needs, and their security architectures should reflect that. Security tools, services, and features should be flexible enough to address these individual needs while remaining true to the capabilities advertised.

This paper discusses 10 points to consider and actively test in your current security infrastructure as well as your future next-generation firewall. Using these as guidelines for cross-functional conversations, you can widen the lens through which you view next-generation firewalls to determine if your potential security investments will be easy to implement, alleviate operational burdens, and offer your organization the best protection and value, today and in the future.

1. Prevent Credential Theft

Users and their credentials are among the weakest links in any organization's security infrastructure. As such, the majority of breaches involve credential theft at some point in the attack lifecycle, and attacks based on credential theft have a high probability of success. The key to stopping them is preventing the theft in the first place.

Why Should You Consider and Test This Capability?

Preventing credential theft and stopping phishing attacks will reduce exposure to one of the most prevalent forms of targeted attacks on organizations. These measures are crucial when dealing with targeted phishing attacks, which typically go after nontechnical employees who follow links to previously unknown phishing sites.

Move Beyond the Status Quo

Most organizations work to stop these attacks primarily by educating their employees—but this isn't nearly effective enough.

Traditional products commonly rely on identifying known phishing sites and filtering email, but these methods are easily bypassed. Checking for known bad sites will miss new ones, and attackers can evade email filtering technology by sending links through social media. A next-generation firewall with machine learning-based analysis can accelerate protection. If the analysis identifies a site as malicious, your firewall should be updated and block it.

Still, there will always be never-before-seen phishing sites that are treated as "unknown." To protect your network and users, it's critical to prevent submission of credentials to such sites. By using credential filtering, organizations can allow authentication to authorized applications and block credential submission to unknown sites.

Recommended RFP Questions

- Can the next-generation firewall prevent use of credentials on unknown websites?
- Can the next-generation firewall block users from submitting corporate credentials without storing a copy of the hash in the firewall?
- Does your next-generation firewall use machine learning to identify and stop credential phishing as well as JavaScript-based attacks inline as they try to enter the network?
- How quickly does the next-generation firewall analyze previously unseen phishing sites and update its protections?
- Does the next-generation firewall log user attempts to submit credentials in HTTP post request?

2. Prevent Credential Abuse

Attackers can obtain stolen credentials in many ways: phishing, malware, social engineering, brute force, or black-market purchase. Once attackers have the stolen credentials, they can abuse them to enter an organization, move laterally, and escalate privileges for unauthorized applications and data.

Why Should You Consider and Test This Capability?

Implementing multi-factor authentication (MFA) on your firewall helps prevent attackers from moving laterally with stolen credentials. MFA allows your organization to protect all types of applications, including legacy applications and client servers. Also, authentication that occurs at the firewall, before users connect to applications, moves the line of exposure farther away.

Move Beyond the Status Quo

Many organizations have an MFA solution, but it is often challenging and time-consuming to integrate with all applications. As a result, most only use MFA with a handful of applications, such as VPN gateways or a few cloud applications, leaving many others vulnerable to credential abuse.

Protection at the Network Layer

MFA is a great tool, but it must be implemented in a way that protects all critical applications. Rather than modifying the applications themselves, use your firewall with MFA policy to control traffic to specific applications. The firewall should be able to control access and require authentication before allowing traffic to pass. Attackers operating inside the organization, even from a compromised endpoint, would not be able to complete the MFA.

Tailored User Experience

Policies for MFA that are created within the firewall should be granular, both in terms of the user and the security needs of the application. For example, policy should match how often users must reauthenticate based on a given application's level of sensitivity.

Accelerate the Time to Protection for Applications

MFA as a part of firewall policy improves the speed of implementation, as you only need to integrate MFA within network policy, rather than changing applications themselves. This allows for quicker deployment of protections to meet compliance. MFA on the firewall stops attackers from using stolen credentials or moving laterally within an organization, and protects all types of applications, including legacy applications and client servers.

Recommended RFP Questions

- Does the next-generation firewall support MFA as part of the access-control policy based on the sensitivity of the resource accessed?
- Does the next-generation firewall provide choices for a variety of MFA integration with partner technologies?
- Can the next-generation firewall support RADIUS and API integrations with MFA partner technologies?
- Does the next-generation firewall support MFA policy for any type of application, including web, client-server, and terminal applications?
- Is the next-generation firewall's MFA capability limited to certain protocols?

3. Provide Dynamic Security Policies for Dynamic Virtual Workloads

When security policies for data center environments are first created and deployed to firewalls, it's assumed that the assigned IP address will remain the same throughout the life of the policy. These policies are static, blanketed, and applied in a generic fashion. With data centers transitioning to virtualized environments, workloads are no longer fixed to a particular location or networking schema.

Why Should You Consider and Test This Capability?

To address security in the virtualized data center, your firewall security policies should be based on the attributes of the workloads rather than tied to static IP addresses, since the data center environment is highly dynamic. This can be done through dynamic address groups on a next-generation firewall.

Move Beyond the Status Quo

Transient workloads are frequently spun up and down to optimize compute resources, repeatedly acquiring new IP addresses. This makes it difficult to manage access control policies on the firewall when dealing with hundreds or thousands of address groups—each with their own address objects—with constant additions, deletions, or changes.

Your firewall should support policies that automatically adapt to the dynamic nature of today's data center, which involves constantly adding, moving, or removing workloads for optimal use of compute resources. Adaptive policies help enforce consistent security across your dynamic virtual machines and applications.

Dynamic address groups decouple security policies from IP addresses and instead build granular security policies based on the attributes of your virtual workloads. Policies on the firewall use tags mapped to workload attributes. For example, the tag on the firewall may be "App-Server," which can be mapped to attributes that identify the specified application server regardless of its IP address. The attributes will continue to place the workload in the desired security policy even if the workload gets relocated.

This helps you build security policies bound to workloads, enhancing your security posture. Dynamic address groups lower operational overhead by reducing dependence between applications and security teams.

Recommended RFP Questions

- How does the next-generation firewall create security policies based on VM attributes of workloads?
- Can the next-generation firewall create security policies for dynamic workloads in both private and public clouds?
- Can the next-generation firewall ensure consistent security policies for workloads even when their IP addresses or locations change in the data center?

4. Manage Your Next-Generation Firewall with Simple and Effective Tools

To be responsive to business needs, security teams need the flexibility to make firewall changes both from a centralized tool and on-site in real time. If a firewall manager allows local administrators to make changes only to a limited set of features, the local team must heavily rely on global teams, potentially located in other regions, to make changes. This results in delays, gaps, limited visibility, and granular administrator access.

Why Should You Consider and Test This Capability?

To minimize the delay in making changes locally and keep your security aligned with your organization's guidelines, your firewall should support complete management of all firewall features and offer role-based access control (RBAC) for multiple administrators. Your local firewall managers' tools should support the full feature set on the centralized tool for local administration, allowing local teams to accomplish their respective tasks on time. Your central management tool should augment local data with overarching visibility into the actions of local administrators and, if required, alert and allow for remote override changes to keep the firewall in line with organizational guidelines.

Look for a next-generation firewall that provides a single pane of glass from which to manage all your firewalls irrespective of their form factors and locations. This reduces complexity by simplifying the configuration, deployment, and management of your security policies. You want a tool that correlates firewall logs to provide network and security insights as well as surface malicious behavior, which can often get buried in the noise.

Move Beyond the Status Quo

Ensure Granular Control While Deploying Configuration Changes

In a multi-firewall environment, it's not unusual for multiple administrators to make configuration changes at the same time. It's likely that one will want to commit recent changes before another is completely done making his or her own, and if your firewall manager doesn't allow for selectively committing changes, those incomplete changes will also be deployed. This can have serious security implications, such as allowing users to access blocked sites or blocking their access to business-critical applications. When selective configuration deployment and rollback isn't possible, administrators must manually undo half-baked changes, and then redo and redeploy them. This adds to operational overhead and delays improvement of security posture.

Manage Logs Effectively at Scale

A central manager acts as a single pane of glass for the organization's security and network, providing a holistic view and context for analyzing security events. In many cases, central firewall managers collect and consolidate firewall logs in multi-firewall deployments. An incoming log rate (generally expressed in logs per second [LPS]) that exceeds the manager's capacity will impact its performance.

Performance impact on the central manager is generally seen through an unresponsive user interface or timed-out database queries. In our high-throughput digital world, it's common for a single high-end firewall to exceed the LPS capacity of the central manager acting as a log manager. The likelihood of running into capacity issues in a multi-firewall deployment is very high.

High-throughput log processing needs are generally addressed through a separate log management appliance. A firewall manager, in conjunction with a log manager, is the most appropriate solution for most enterprises. With this setup, the central manager is relieved of log management responsibility and can focus solely on firewall management. When provisioned, the central manager queries the data on the log managers to provide centralized visibility and brings raw logs to the central manager only when required, reducing performance impact.

Keep Your Security Posture Up to Date

Each of the many features of a next-generation firewall is built to address a specific network security needs and empower an organization's growth. In a multi-firewall environment, manual firewall configuration changes are inefficient and often result in security gaps and inconsistent prevention. Automation will provide faster, more accurate responses to ever-changing cybersecurity threats.

The preferred way to act on this is to leverage next-generation firewall application programming interfaces (APIs) to automate changes. This alleviates network security teams' operational overhead while reducing human error. For this to be possible, the APIs should be flexible enough to allow for automated changes to all firewall features.

Recommended RFP Questions

- Can local administrators work directly on the appliance and make configuration changes as needed, without having to log in to a central manager?
- Can central administrators monitor and view the changes made by local administrators?
- Can you choose which firewall administrator's configuration changes should be deployed on the firewalls?
- When deployments go wrong, can you quickly roll back changes from specific users and restore working configuration?
- Can the central firewall manager separate log management from core configuration management yet still act as a single pane of glass for unified visibility?
- Can your log managers ingest logs at high throughput (e.g., 50,000 LPS)?
- Does your firewall have APIs for every feature so that you can automate configuration changes?

5. Use Automation to Integrate Security, Prevent Fast-Changing Threats

Security as an add-on has become a legacy model. Today's security should be integrated with systems and processes at an early stage to avoid a buildup of complex, disconnected security tools and control systems. Automated detection and response tools as well as APIs integrate security devices with your organization's overall security ecosystems. This streamlines operations and turns detection data into an actionable response to mitigate attacks before they succeed—and before sensitive data leaves your organization.

Why Should You Consider and Test This Capability?

APIs can automate security workflows that need multiple security devices, often from different vendors, to work together. This moves security teams away from the cumbersome, error-prone processes of operating these workflows manually and increases the speed of effective enforcement. Automated security tools and services should be able to ingest alerts from multiple sources and execute standardized, automatable playbooks to speed up incident response as well as integrate with other tools to set off the next action in a workflow.

Move Beyond the Status Quo

According to the Verizon 2019 Data Breach Investigation Report, the time from an attacker's first action to the initial compromise of an asset can be measured in minutes.[1] That means your organization needs tools that can ensure all parts of your infrastructure can communicate effectively, respond quickly and automatically to identify known and unknown threats, and stop those threats more quickly than they can progress through the attack lifecycle. To do so, every step in the process, from discovery to full prevention, should be automated. Moreover, every infrastructure element must be able to communicate effectively with every other element to streamline operations and speed up incident response.

APIs offer a mechanism for the many elements in a data center, which may be from disparate vendors, to share data and kick off appropriate actions required in the workflow. As such, the APIs your security vendor uses must be able to integrate with a broad list of partners via documented and certified interoperability. Multi-vendor integration should also extend beyond the data center to vendors of endpoint security, email gateways, wireless security, and more. A firewall with natively integrated APIs allows firewall administrators to view, access, and change the entire feature set.

[1] "2019 Data Breach Investigations Report," Verizon, May 2019, https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf.

Security tools must be able to analyze and identify malicious behavior—ideally within a cloud environment, to take advantage of elastic compute and scalability—to prevent never-before-seen threats. New tools should be able to integrate with your existing cloud security tools to coordinate and automate response processes across cloud and on-premises environments. These tools must use automation to ingest alerts generated by multiple sources and execute standardized, automatable playbooks to accelerate incident response. Using automated data correlation, the tools should identify and surface hosts on your network exhibiting any of the same malicious behavior as a threat.

Some organizations may want to automate the immediate quarantine of potentially infected hosts. This can be done by moving a host to a policy that denies it access to all parts of the network while retaining connectivity for remediation efforts. Others may take a more nuanced approach by automatically applying MFA to a potentially infected host so that, if attackers gain access to it, they cannot access corporate data or applications.

Automation and APIs enable organizations to act against threats without waiting for human intervention, improving response time and—if implemented appropriately and in conjunction with the right tools—preventing successful attacks. A security vendor that offers automation and natively integrated APIs allows security teams to move away from basic operational tasks to focus on strategic efforts that directly benefit the organization. Reducing human intervention reduces avoidable errors, ultimately enabling a more secure security posture.

Recommended RFP Questions

- Can your firewall/manager create a ticket on a change management system based on a malicious event seen on the firewall?
- Can your firewall/manager trigger a quarantine action for an infected host on the wireless network?
- Can your firewall be completely programmed via API?
- Can your firewall collect user identification information via APIs from wireless controllers about hosts connecting to wireless networks?
- Does your security vendor support the ability to automatically generate prevention signatures across the attack lifecycle for all data relevant to attacks?
- Can your firewall correlate and identify infected hosts in your network, and then quarantine them to limit their access in the network?
- Can your firewall trigger MFA to prevent credential abuse and secure critical applications?

6. Protect Against Evasive and Never-Before-Seen Attacks

Millions of new malware samples and malicious websites are discovered every year. Traditional solutions mostly protect against newly discovered threats after the first victim in an organization has already been compromised. That's not enough when a single victim can turn into 10,000 in a five-minute window before protections are introduced.

Many attacks succeed because malware developers use various evasive techniques, such as wrapping malicious payloads in legitimate files, packing files to avoid detection, or extending sleep calls to wait out potential sandbox environments. Attackers have become acutely aware of the methods security teams use to analyze files for malicious activity. They also keep up on the virtual sandboxes organizations use for dynamic analysis, including scanning these environments for code used in known malware analysis tools; scanning for valid user activity, system configurations, or indicators of specific virtualization/emulation technologies; or observing hardware size for amounts of memory typically found in virtual machines.

Modern threats are frequently designed with knowledge gained from open source technology used in most malware analysis tools and hypervisors. With the availability and growth of the cybercrime underground, even a novice attacker can purchase plug-and-play threats designed to identify and avoid malware analysis environments. The ability to identify and protect against evasive malware is more crucial now than ever.

Why Should You Consider and Test This Capability?

Waiting to protect against new attacks only after the first victim has been compromised can lead to rapid lateral spread, putting your whole organization at risk. Some solutions try to prevent "patient zero" by holding files for analysis, but this ultimately doesn't work because it harms the user experience and delays business.

Most modern malware uses advanced techniques that can bypass traditional, common network security solutions to transport attacks or exploits through network security devices, firewalls, and sandbox discovery tools. Although we can't build individual tools to detect every piece of evasive malware, it's critical to utilize systems that can identify evasive techniques and automatically counteract them.

Move Beyond the Status Quo

Stop Never-Before-Seen Attacks with Inline Threat Prevention

In a world where new and unknown malware threats can spread exponentially in minutes, relying on offline analysis and periodic updates to firewall rules will always leave you one step behind. A next-generation firewall that offers inline, machine learning-based prevention at the network level to stop new threats across all applications can protect patient zero from credential phishing and other evasive threats without compromising business productivity.

Use Bare Metal Analysis

There are a number of ways to counter threats built to evade analysis environments, and a modern, effective security platform should combine multiple techniques. For example, combining dynamic analysis in a sandbox environment with bare metal analysis has proven effective in countering malware that assesses the environment to determine if it is being analyzed.

In bare metal analysis, suspicious files are sent to a real, racked-and-stacked hardware environment where they are detonated, and any response is observed for malicious behavior. The malicious activity of the file, which would otherwise have remained dormant in the virtual environment, will fully execute in the bare metal environment. Threats with virtual machine evasion techniques cannot evade this type of environment.

Avoid Open Source Hypervisors

A next-generation firewall that uses a proprietary hypervisor takes away the potential for adversaries to test their malware in well-known, virtual open source environments. Most adversaries test and refine their malware based on what they learn by executing in known open source analysis environments.

Fight Automation with Automation

Malware changes rapidly, and threat signatures that rely on specific variables—such as hash, filename, or URL—get one-to-one matches only against known threats. Attackers often make slight modifications to malicious code, resulting in malware variants and/or polymorphic malware. This "new" malware is considered unknown, as protections have only been created for the original malware, not its modified variants.

Hash-based threat signatures are particularly problematic. According to the Verizon Data Breach Investigations Report, "99% of malware hashes are only seen for 58 seconds or less."[2] If one bit is changed, the hash changes, and the signature no longer recognizes it as malware.

Rather than signatures based on specific attributes, next-generation firewalls should use predictive machine learning models resident on the firewall or content-based signatures to detect variants, polymorphic malware, or command-and-control (C2) activity. Content-based signatures detect patterns that allow them to identify whole families of malware, including known malware that has been modified. This results in signatures capable of automatically preventing tens of thousands of variants created from the same malware family, rather than trying to create signatures for individual variants.

C2 can pose a challenge, with malware authors creating C2 communications that automatically change the DNS or URL. Automated signatures based on these artifacts quickly become outdated and ineffective. C2 signatures based instead on analysis of C2 outbound communication patterns are much more effective protections that can scale at machine speed when created automatically.

Validate with More Than One Analysis Method

More determined, skilled attackers will create entirely new threats with purely new code, the costliest method for attackers. Any such threat will be treated as an unknown and go undetected.

When an entirely unknown threat enters an organization, the clock begins ticking. Protections must be created and distributed across all security products more quickly than a threat can spread. This can be accomplished by automating various aspects of the analysis, including static analysis with machine learning, dynamic analysis, and bare metal analysis. Implementing automation results in accurate identification of threats, enables rapid prevention, improves efficiency, makes better use of the talent of your specialized staff, and improves your organization's security posture.

Recommended RFP Questions

- Does your next-generation firewall deliver machine learning-based prevention of unknown malware files and variants, including both executables and fileless attacks leveraging scripts such as PowerShell®?
- Does your next-generation firewall deliver inline machine learning-based prevention of malicious website attacks, including JavaScript or credential phishing attacks?
- How quickly does your cloud-based malware analysis system distribute signatures after verdict generation?
- Does your next-generation firewall use signatureless technologies to prevent never-before-seen attacks?
- Does your cloud-based sandbox support bare metal analysis?
- Does your cloud-based malware analysis system use a custom-coded hypervisor to be effective against sandbox-aware malware?
- Does your malware analysis system, after analyzing malware, create threat prevention signatures, such as:
  - » Content-based AV signatures to prevent whole families of malware, including known and unknown variants?
  - » Pattern-based anti-spyware signatures to detect communications to known and unknown C2 infrastructure?
- Does your cloud-based malware analysis system support malware analysis for file types of Windows®, Android®, macOS®, and Linux operating systems?

[2] 2019 Data Breach Investigations Report," Verizon, May 2019, https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf.

7. Customize Your Next-Generation Firewall

Firewalls can import lists of predetermined rules and policies that, once consumed, allow the firewalls to act against the objects outlined in the lists. Afterward, firewall administrators are responsible for updating firewalls to reflect newfound threats, protections, and policy roles. With attackers employing more advanced methods, such as automation and evasion, having the most updated security posture possible requires moving at machine speed.

Why Should You Consider and Test This Capability?

Incorporating automation and dynamic lists into your next-generation firewall is the most effective and efficient way to improve your organization's security posture. Your next-generation firewall's vendor will often provide dynamic lists, which you can update manually or by integrating third-party threat intelligence. As a result, changes to rules and policies only need to be made to the dynamic list, and all firewalls tied to it will regularly and automatically import the most updated protections.

Move Beyond the Status Quo

Dynamic Lists

When new threats are identified, it falls to the firewall administrator to create a new rule or policy so that the firewall can respond appropriately. This must be done for each risk object and potentially each firewall in the network—a labor-intensive, often error-prone, and manual process.

Working with dynamic lists dramatically reduces manual efforts and improves response time. Modern dynamic lists include protections against known and high-risk malicious IP addresses validated by your next-generation firewall vendor. They also protect against high-risk IP addresses drawn from correlated third-party data that has not been validated by your vendor, which you can opt into at the level of policy enforcement appropriate for your organization.

Dynamic User Groups

Dynamic user groups help you to create policy that provides auto-remediation for anomalous user behavior and malicious activity while maintaining user visibility. After you create the group and commit the changes, the next-generation firewall registers the users and associated tags, and then automatically updates the dynamic user group's membership. Updates to dynamic user group membership are automatic. Using dynamic user groups instead of static group objects allows you to respond to changes in user behavior or potential threats without making manual policy changes.

Dynamic user groups help automate policy enforcement and tighten security. You create groups on the next-generation firewall to immediately respond to changing needs, enforce security that matches the user's behavior, and eliminate over-provisioned access. Dynamic user groups allow an administrator to change a user's group membership on the fly on the next-generation firewall without waiting for changes to be applied in the directory. Now, you can dynamically change user access based on changes in activity.

Third-Party Threat Intelligence Feeds

Organizations subscribe to third-party threat intelligence feeds for access to continuously updated data on potential threats and attack sources. These feeds provide a massive amount of data on raw indicators of compromise (IOCs), which can be used to turn unknown threats into known before attackers have a chance to compromise an organization.

Turning threat intelligence into actionable protections, much like creating new rules and policies based on activity seen on the firewall, is a time-consuming, manual process that many security teams struggle to manage. The data from threat intelligence feeds must be current and formatted, potentially requiring the data format to be changed to a consumable form. The data must also be correlated to validate whether a given IOC is malicious, reveal larger patterns of malicious behavior with multiple IOCs, and add necessary context, such as the priority and relevance of newly identified threats. Once the data has been validated and enriched with context, security teams can much more efficiently create and distribute protections to address specific threats across various enforcement points. Alternatively, vendors can push protections out to enforcement points, but consolidation with local traffic isn't as effective. Without completing these steps, threat feeds remain inert reams of data.

Automation is necessary to rapidly improve your security posture with the latest threat intelligence, alleviate manual intervention, and eliminate human error. Based on context collected from outside your organization, automation can turn unknown threats into known protections more quickly than attackers can successfully complete the attack lifecycle.

Recommended RFP Questions

- Does your next-generation firewall support dynamic lists and dynamic user groups to automate policy enforcement and tighten security?
- Can your next-generation firewall dynamically incorporate third-party or custom threat intelligence feeds in the firewall without policy commits?
- Does your security architecture support threat feed aggregation, consolidation, and deduplication of threat feeds before pushing the indicators to your firewall?
- Does your security architecture integrate with your next-generation firewall to automate timeout of expired threat indicators and avoid using stale threat intelligence?
- Does your security architecture allow you to target threat indicators from recent advanced persistent threat campaigns and incorporate threat feeds proactively on your next-generation firewall?
- Does your security architecture allow you to enrich threat intelligence based on a confidence rating to reduce the operational overhead from dealing with false positives?

8. Prevent Modern Attacks

There are more types of attacks and more devices at risk than ever. Phishing, credential abuse, social attacks, denial of service, ransomware, and backdoor or C2 malware all pose real threats. The internet of things (IoT) presents a new type of security frontier, one particularly appealing to adversaries. According to the 2020 Unit 42 IoT Threat Report, 57% of IoT devices are vulnerable to medium- or high-severity attacks, making them low-hanging fruit for attackers. Many attacks use IoT devices as stepping stones in lateral movement to attack other systems on a network.

Why Should You Consider and Test This Capability?

No single security product can successfully combat every type of attack on its own. As there are multiple stages in the attack lifecycle, there should be multiple layers of defense to prevent modern attacks. Your organization's ability to effectively defend against attacks like malware and ransomware lies in the natively engineered automation and integration among your security products to proactively detect and prevent ransomware. A multilayered defense is the most effective way to disrupt possible attacks, and new additions to the security architecture should complement protections throughout the network.

According to Unit 42, the Palo Alto Networks threat intelligence team, 30% of devices on enterprise networks today are IoT devices, excluding smartphones. These devices pose a huge cybersecurity risk as they are often introduced and deployed with vulnerabilities, are difficult (or impossible) to patch, and can have unfettered access to an organization's network.

Move Beyond the Status Quo

There Is No Silver Bullet

Protecting against modern attacks requires visibility into network traffic and enforcement of applications, as well as user- and content-based policies. It also requires security products to protect against known and unknown exploits, malware, and C2 traffic, as well as prevent access to known malicious and phishing URLs.

Attacks Are Time-Sensitive

Automation is the only way for prevention to move more quickly than an attack can transition through the full attack lifecycle within your organization. To identify and block unknown threats, malicious files and URLs must be detonated, analyzed, and observed for malicious activity. Once a file or URL is identified as malicious, protections must be created and automatically distributed throughout the security infrastructure—across the network, cloud, and endpoint. This ensures all points of entry are informed and capable of protecting against the latest version of the ransomware.

Combine Preventive Efforts

For effective prevention, you must employ automation and share information among various security tools. These tools then work together to identify known and unknown malware and exploits in your environment, and subsequently identify and quarantine any infected host, preventing the attack from spreading.

Threat intelligence should always be a component of your organization's threat prevention efforts. Your firewall should be capable of dynamically updating preventions against malicious IPs, domains, and URLs based on information gathered from the threat intelligence cloud and IOCs.

Reduce Exposure to IoT Security Risk

With so many network-connected devices in an organization, it can be a challenge for security teams to keep pace and mitigate risk as each new device type presents a possible new threat vector. Properly classifying IoT devices, keeping software up to date with the latest patches, segmenting your network, and enabling active monitoring ensures IoT devices are granted access to appropriate resources and placed in the right network segments. This effectively lessens the risk of threats to other resources and networks in addition to reducing your overall attack surface.

Recommended RFP Questions

- Can your next-generation firewall block executables and other risky file types from unknown applications and URLs to prevent ransomware attacks?
- Can your next-generation firewall prevent successful cyberattacks targeting IoT devices?
- Can your next-generation firewall automatically and dynamically import all known IOCs (i.e., IPs, domains, and URLs) into the block list to be proactive against all known ransomware families?
- Does threat intelligence integration with the next-generation firewall support dynamic updates for malicious URLs related to ransomware in the malware category of the URL filtering database?
- Does threat intelligence integration with the next-generation firewall support dynamic updates for malicious domains related to ransomware as DNS signatures to be automatically added to a block list or sinkholed?
- Can your next-generation firewall learn about threats and suspicious behavior from your endpoint protection software and vice versa?
- Does your next-generation firewall provide visibility into all IoT devices in your network, including those never seen before?
- Can your next-generation firewall segment IoT device traffic from the rest of the network to prevent compromised IoT devices from becoming beachheads for attacks?
- Does your next-generation firewall offer policy recommendations, based on risk assessment results, that can be automatically enforced?

9. Offer Consistent Protection, Wherever Users or Applications Are Located

Users are becoming more mobile and distributed, and they need access to applications from remote locations around the globe. As applications move to the cloud, branch locations need the same connectivity and security as your headquarters. However, many organizations lack visibility into traffic when users access the internet and cloud applications off-premises, which can compromise security.

Why Should You Consider and Test This Capability?

Your organization should be able to protect all users in the same way, without having to create different security profiles depending on users' locations. Since security policies are more effective when they can be administered in a consistent manner, a single set of tools and a common policy framework will give security teams greater control.

Move Beyond the Status Quo

Software-defined wide area networking (SD-WAN), an approach that uses commodity links and allows you to intelligently manage as well as control connectivity between branches and cloud instances, is now a necessity for distributed enterprises. A next-generation firewall must be able to operate as an SD-WAN edge device in a branch location and as an SD-WAN hub in a central location. SD-WAN features should be simple to enable and centrally managed on a single network security management interface.

With SD-WAN, each device is centrally managed, with routing based on application policies, so WAN managers can create and update security rules in real time as network requirements change. In addition, combining SD-WAN with zero touch provisioning—a feature that helps automate deployment and configuration—can further reduce the complexity, resources, and operating expenses required to spin up new sites. In addition, SD-WAN allows efficient access to cloud-based resources without the need to backhaul traffic to centralized locations, so you can provide a better user experience.

Your organization should maintain the same protection whether your users work on-premises or off. Deployment options should provide flexibility to support consistent coverage for all users and locations. This way, no matter where users are, they can easily connect to a cloud service or firewall for security and receive the same protection against known and unknown threats.

Recommended RFP Questions

- Can your next-generation firewall provide a consistent security policy for mobile users?
- Can you protect users who are not behind a next-generation firewall?
- Can your next-generation firewall operate as an SD-WAN edge device?
- Does your next-generation firewall offer zero touch provisioning?
- Can your next-generation firewall use multiple physical/virtual firewalls to support an always-on VPN connection?
- Can your next-generation firewall utilize the cloud to bring protection closer to your users?

10. Adopt a Zero Trust Approach

The old network security approach of classifying users as "trusted" or "untrusted" doesn't work anymore. Anyone who accesses your network—whether they're a malicious actor, an insider with ill intent, or a user with more access than they should have—is free to move laterally throughout your network unless the network is segmented or policy enforces access control. This leaves your network open to all kinds of threats, from credential and data theft to intellectual property exfiltration and deployment of malware. Zero Trust is not about making a system trusted, but instead eliminating trust. Its central tenet is "never trust, always verify."

Why Should You Consider and Test This Capability?

The most effective way to secure modern networks is to adopt a Zero Trust approach that enforces least-privileged access and inspects all users, devices, content, and applications across all locations. The Zero Trust methodology gives you a roadmap to building a segmented network. Zero Trust is a powerful prevention strategy when implemented across an entire enterprise—from the network to endpoints and clouds. With an integrated approach, security teams can automate and streamline Zero Trust policy management across the enterprise, from creation and administration to deployment and maintenance.

One criterion of Zero Trust is Layer 7 policy and enforcement via a next-generation firewall, which serves as a segmentation gateway. This policy is written to reflect your business, based on the way traffic traverses the network and the interdependencies of your data, users, and applications.

Move Beyond the Status Quo

In Zero Trust, you identify a "protect surface," which is made up of the network's most critical and valuable data, assets, applications, and services (DAAS).

The next-generation firewall enables microsegmentation of perimeters and acts as an enforcement point as well as a network of sensors for comprehensive visibility into your traffic. While it's important to secure the traditional perimeter, it's even more crucial to gain the visibility and precise control necessary to verify traffic as it moves north-south and east-west. Tools like MFA and other verification methods improve your ability to confirm user identities.

A layered defense is a major component of Zero Trust. To that end, it's important to incorporate tools that can block advanced threats and provide insight into what's happening across your network. Not only should you use policy to prevent the wrong people from accessing things they shouldn't; you should use the visibility you gain to identify anomalies—such as changes in user behavior due to job function and/or malicious activity—and to update policies. You should also be able to use that visibility to quickly identify and act when malicious activity is identified.

Recommended RFP Questions

- Can your next-generation firewall function as a segmentation gateway?
- Does your next-generation firewall enable visibility, decryption, and inspection into all traffic?
- Are you able to write and enforce context-based Layer 7 policy on your next-generation firewall?

Bonus: Provide Flexible Deployment Options, Including Containers

For many organizations, physical appliances remain key control points for data center ingress and egress in addition to providing the required horsepower for high-performance networks. At the same time, virtual appliances provide visibility into east-west traffic in software-defined data centers and enable security to keep pace with the dynamic nature of these environments.

Container adoption, however, has added new dynamics to network security. As cloud native applications are connected to legacy applications, network security teams must ensure that cloud native applications are secured and compliance is maintained with the same rigor as the rest of their environment.

Why Should You Consider and Test This Capability?

Traditional solutions are limited. While some of them integrate with container orchestration tools, the visibility is often limited to the node or cluster and focused on basic port and protocol filtering. Cloud native firewalls offer only basic port blocking to enforce container microsegmentation, without the ability to prevent threats. Conventional next-generation firewalls lack the native container context and orchestration framework integration needed to provide seamless advanced security in these agile, dynamic environments.

Move Beyond the Status Quo

To properly address cloud native container environments, next-generation firewall visibility and enforcement must extend to inter-pod and container-to-container traffic to apply threat inspection within a cluster and improve protection within containerized data centers. Further, containers often access web resources such as GitHub® repositories to pull down source code. This becomes an additional driver to have visibility and control over web application traffic to ensure containers only communicate with the correct repository and that the resulting traffic is not spoofed or malicious. Next-generation firewalls that can offer support across physical, virtual, and containerized environments help drive consistent security across increasingly hybridized data centers.

Recommended RFP Questions

- Does your firewall deliver consistent network security and threat prevention aligned to applications hosted on-premises as well as in virtualized and container environments?
- Does your firewall natively deploy within Kubernetes® to integrate firewall provisioning into a continuous integration/continuous development (CI/CD) processes?
- Does your firewall integrate into software-defined networking (SDN) solutions to extend security protections to remote locations for branch segmentation and meet PCI compliance?

Comprehensive Security

Sophisticated attackers and techniques enable advanced attacks that are targeted, automated, evasive, and capable of spanning multiple environments.

Your new next-generation firewall and the various security products that make up your security infrastructure should be comprehensive, including all of the following.

The Best Technology

Utilize industry-leading technology with the ability to rapidly, automatically stop known and unknown threats at every step of the attack lifecycle. These products should deliver consistent, risk-appropriate protection for data and users regardless of location. Your security ecosystem should offer agile, flexible updates, allowing you to adapt to changing risks and workloads.

Operational Efficiency

Automated delivery and API integrations reduce time spent on error-prone, manual tasks. Security should be operationalized over various environments without straining resources or budget, and without adding complexity, allowing security teams to focus on strategic efforts that are more critical to the organization.

Knowledgeable, Responsive Service

Service and support teams who are knowledgeable and responsive minimize your learning curve and help you keep improving your security posture long after the initial migration. You should be able to maximize your investment over time and achieve higher levels of security.

When planning your next purchase or assessing your current firewall deployment, it's important to test the capabilities and features of your firewall with all security teams in your organization.

The 10 points of consideration discussed in this paper, when tested, will help determine if your next firewall purchase matches the needs of your organization in its current and planned states by keeping future innovation in mind.

Are you ready to evaluate your next firewall? Take an Ultimate Test Drive.

paloalto®
NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com