

Next-Generation Firewall Buyer's Guide

The Definitive Guide for Evaluating Enterprise Network Firewalls

The rapid evolution of IT has changed the face of the network perimeter. Data and users are everywhere. New devices are joining the network and proliferating more quickly than most organizations can keep up, sometimes without notice. At the same time, IT teams are adopting the cloud, big data analytics, machine learning, and automation to accelerate delivery of new applications to drive business growth. Meanwhile, applications are increasingly accessible. The result is an incredibly complex network that introduces significant business risk. Organizations must minimize this risk without slowing down their business.

Traditional cybersecurity is falling behind as attacks are constantly and automatically morphing. Spending on security feels endless, and the reduction of risk is unclear. Deploying disparate, non-integrated tools and technologies leaves your business exposed to threats. Security tools that weren't designed for automation require analysts to manually stitch together insights from many disconnected sources before acting. We need a new, disruptive approach.

It starts with a next-generation firewall platform as the cornerstone of an effective network security strategy. With a prevention-focused architecture, security teams can easily adopt best practices to proactively prevent successful attacks, use automation and analytics to reduce manual effort, replace disconnected point products, and deploy tightly integrated innovations that strengthen and simplify security.

This paper describes the evolution of the firewall to “next-generation,” highlights critical capabilities a next-generation firewall must have to secure your network and your business, and provides key questions you should ask during the request-for-proposal (RFP) process as you formally evaluate your next next-generation firewall.

The Evolution of the Next-Generation Firewall

Early on, stateful inspection firewalls classified traffic by looking only at the destination port, such as TCP port 80 for HTTP. As the need for application awareness arose, many vendors added application visibility and other software or hardware “blades” into their stateful inspection firewalls, which they subsequently sold as unified threat management (UTM) offerings. However, since their functions were retrofitted—not natively integrated—UTMs did not improve security.

Gartner predicted that, by the end of 2019, 90% of enterprise internet connections for the installed base would be secured with next-generation firewalls.¹

Unlike UTM offerings, **next-generation firewalls** are application-aware and make decisions based on application, user, and content. The integrated design improves security and simplifies operations. Given the model's success, the term “next-generation firewall” is now synonymous with “firewall.”

Next-Generation Firewall Requirements

- Identify applications regardless of port, protocol, evasive tactics, or encryption.
- Identify users regardless of device or IP address.
- Decrypt encrypted traffic.
- Protect in real time against known and unknown threats embedded in applications.
- Deliver predictable, multi-gigabit, inline throughput.

Over the last decade, many security vendors have focused on reducing the time it takes to react to new attacks, new devices, and a need for policy changes—from weeks to days, hours, and in some cases, even minutes. However, only shortening reaction times is becoming unproductive given how rapidly attacks, devices, and environments are changing. The next evolution of the next-generation firewall will require network security solutions that proactively deliver protection against tomorrow's threats.

Next-generation firewall selection criteria typically fall into three areas: security functions, operations, and performance. The security functions correspond to the efficacy of the security controls and your team's ability to manage the risk associated with the applications traversing your network, without slowing down the business. From an operations perspective, application policy should be accessible and simple to manage, applying automation to reduce manual effort so security teams can focus on high-value activities. Performance criteria are simple: the firewall must do what it's supposed to do at the required throughput for your business needs. As part of this, innovations should be tightly integrated and easy to adopt. Although requirements and priorities will vary within these criteria, there are certain capabilities your next firewall must have.

15 Things Your Next-Generation Firewall Must Do

1. Identify Users and Enable Appropriate Access

The Problem

Employees, customers, and partners connect to different repositories of information within your network, as well as to the internet. These people and their many devices represent your network's users. It's important for your organization's security posture that you're able to identify your users beyond IP address as well as grasp the inherent risks they bring based on the devices they're using—especially when security policies have been circumvented or new threats have been introduced to your network.

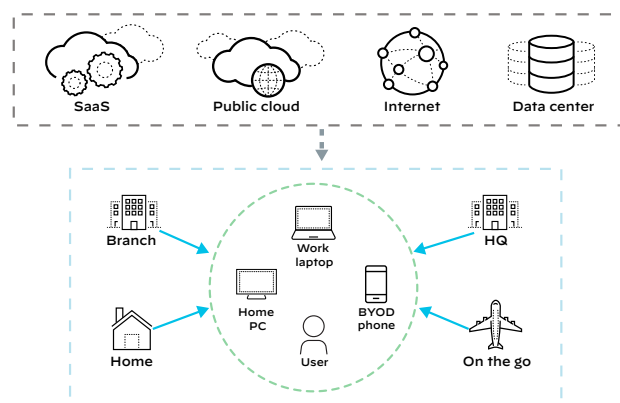


Figure 1: Users access data from different devices and locations

1. Adam Hills, Jeremy D'Hoinne, Rajpreet Kaur, “Magic Quadrant for Enterprise Network Firewalls,” Gartner, July 10, 2017, <https://www.gartner.com/en/documents/3757665/magic-quadrant-for-enterprise-network-firewalls>.

In addition, users are constantly moving to different physical locations and using multiple devices, operating systems, and application versions to access the data they need. IP address subnets are mapped only to physical locations, not individual users, meaning that if users move around—even within the office—policy doesn't follow them. Finally, user directories don't include behavioral characteristics, so a user's risk profile can change or credentials can become compromised, yet the user retains the same level of access simply based on role. Since user directory changes take time, risky or malicious activity can go unchecked and increase the risk to which an enterprise is exposed.

Addressing the Problem

User and group information must be directly integrated into the technology platforms that secure modern organizations. Your next firewall must be able to pull user identity from multiple sources, including virtual private networks (VPNs), wireless local area network (WLAN) access controllers, directory servers, email servers, and captive portals. Knowing who is using the applications on your network, and who may be transmitting a threat or transferring files, strengthens security policies and improves incident response times.

Your firewall must allow policies to safely enable applications based on users or groups of users, outbound or inbound—for example, by allowing only your IT department to use tools such as SSH, telnet, and FTP. User-based policies follow users no matter where they go—at headquarters, branch offices, or home—and on whatever devices they use. However, developing policies based on user information in the directory is not sufficient. You should be able to dynamically change user access based on changes in circumstances, whether the change is due to new indicators of compromise (IOCs) or a business need, such as granting temporary access to a set of users.

2. Prevent Theft and Abuse of Corporate Credentials

The Problem

Users and their credentials are among the weakest links in an organization's security infrastructure. According to the 2019 Data Breach Investigations Report by Verizon, in the 12-month period covered in the report, 29% of hacking-related breaches took advantage of stolen and/or weak passwords.² With stolen credentials as part of their toolset, attackers' chances of successfully breaching go up, and their risk of getting caught goes down. To prevent credential theft, most organizations rely on employee education, which is prone to human error by nature. Technology products commonly rely on identifying known phishing sites and filtering email.

However, these methods can sometimes be bypassed—checking for known bad sites misses newly created ones, and attackers can evade mail filtering technology by sending links through social media. Attackers can easily steal credentials through phishing, malware, social engineering, or brute force, and can even buy them on the dark web. Attackers use these credentials to gain access to a network, move laterally, and escalate their privileges for unauthorized access to applications and data.




Figure 2: Verizon 2019 DBIR finding on stolen credentials

Addressing the Problem

Organizations should look for a firewall with machine learning-based analysis to identify websites that steal credentials. If the analysis identifies a site as malicious, the firewall should be updated and block it. Still, there will always be new, never-before-seen phishing sites that are treated as "unknown." Your next firewall must allow you to block submission of corporate credentials to unknown sites. The firewall must also allow you to protect sensitive data and applications by enforcing multi-factor authentication (MFA) to prevent attackers from abusing stolen credentials. By integrating with common MFA vendors, your firewall can protect your applications containing sensitive data, including legacy applications.

3. Safely Enable All Apps and Control Functions

The Problem

More and more applications, such as instant messaging applications, peer-to-peer file sharing, or voice-over-IP, are capable of operating on nonstandard ports or hopping ports. Additionally, users are accessing diverse types of apps, including software-as-a-service (SaaS) apps, from varying devices and locations. Some of these apps are sanctioned, some tolerated, and others unsanctioned, and users are increasingly savvy enough to force applications to run over nonstandard ports through protocols such as RDP and SSH.

Further, new applications provide users with rich sets of functions that help ensure user loyalty but may represent different risk profiles. For example, WebEx® is a valuable business tool, but using WebEx desktop sharing to take over an employee's desktop from an external source may be an internal or regulatory compliance violation. Gmail® and Google Drive are other good examples. Once users sign in to Gmail, which may be allowed by policy, they can easily switch to YouTube® or Google Photos, which may not be allowed. Security administrators want complete control over usage of these apps and set policy to allow or control certain types of applications and application functions while denying others.

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

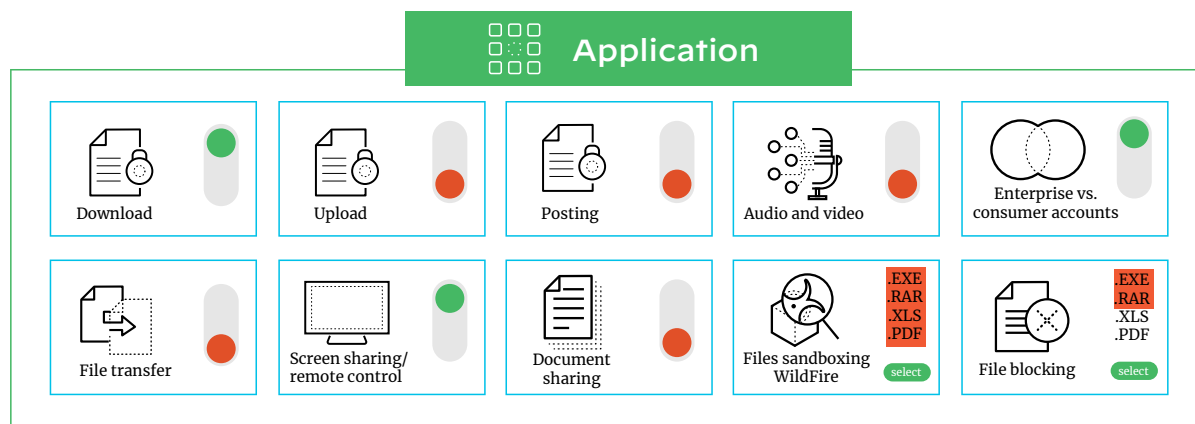


Figure 3: Control application usage in policy

Addressing the Problem

Your next firewall must classify traffic by application on all ports, all the time, by default—and it should not burden you with researching common ports used by each application. The firewall must provide complete visibility into application usage along with capabilities to understand and control their use (see figure 3). For example, it should understand usage of application functions, such as audio streaming, remote access, and posting documents, and be able to enforce granular controls over that usage, such as upload versus download permissions, chat versus file transfer, and so on. This must be done continuously. One-and-done traffic classification is not an option as it ignores the fact that these commonly used applications share sessions and support multiple functions. If a different function or feature is introduced in the session, the firewall must perform a policy check again. Continuous state tracking to understand the functions each application may support—and the different associated risks—is a must for your next firewall.

4. Close Dangerous Policy Gaps

The Problem

Legacy firewalls allow and block traffic based on ports and IP addresses. This approach is inadequate as port-based rules allow both good and bad applications through the firewall. Applications can easily go through a port-based firewall by hopping between ports, using SSL and SSH, or using well-known open ports such as 80 and 443. Over time, customers accumulate thousands of port-based rules on their firewalls, and often migrate these rules as-is to their next-generation firewalls. These rules leave dangerous policy gaps. Customers realize that they must migrate to application-based rules for effective security, but this requires significant manual effort—and due to the cybersecurity skills shortage, most organizations do not have the resources. This becomes a high security risk that may cause a business disruption. In fact, according to Gartner, through 2023, 99% of firewall breaches will be caused by firewall misconfigurations, not firewall flaws.³

Addressing the Problem

When evaluating your next firewall, look for one that reduces the complexity of rule and policy management. This begins with showing you what applications are running on your network, mapping them to the legacy rules, and helping replace the legacy rules. A next-generation firewall should help your security team easily replace legacy rules with intuitive, application-based policies. Because rules based on application identification are easy to create, understand, and modify as business needs evolve, they minimize configuration errors that leave you vulnerable to data breaches. These policies strengthen security and take significantly less time to manage. Finally, your next firewall should aggregate telemetry information and apply machine learning to automatically identify required policy and configuration changes. These capabilities can improve security policy optimization to eliminate breaches due to misconfiguration.

5. Secure Encrypted Traffic

The Problem

Most enterprise web traffic today is encrypted, and attackers exploit encryption to hide threats from security surveillance. This means even businesses with mature, comprehensive security measures in place can be breached if they are not monitoring encrypted traffic. Additionally, TLS/SSL encryption is used nearly universally, and end users can easily configure it to hide non-work-related activity.

Addressing the Problem

The ability to decrypt TLS/SSL-encrypted traffic is a foundational security function. Key elements to look for include recognition and decryption on any port, inbound or outbound; policy control over decryption; and the necessary hardware and software elements to perform decryption across tens of thousands of simultaneous SSL connections without compromising performance. However, your next firewall must be flexible enough to easily decrypt certain types of encrypted traffic (e.g., HTTPS from unclassified websites) via policy,

3. Rajpreet Kaur, Adam Hills, John Watts, “Technology Insight for Network Security Policy Management,” Gartner, February 21, 2019, <https://www.gartner.com/doc/3902564/technology-insight-network-security-policy>.




Figure 4: Google 2019 findings on encrypted traffic⁴

while other types (e.g., web traffic from known financial services organizations) are left alone in compliance with privacy standards. A next-generation firewall should apply security and load balancing to decrypted flows across multiple stacks of security devices for additional enforcement. This eliminates dedicated SSL offloaders, reducing network complexity and making decryption simpler to operate. It must also offer support for decryption of modern protocols like TLS 1.3 and HTTP/2 that are gaining widespread adoption. Read Decryption: [Why](#), [Where](#), and [How](#) for a detailed overview of this important capability.

6. Stop Advanced Threats to Prevent Successful Cyberattacks

The Problem

Most modern malware—including ransomware variants—uses advanced techniques, such as wrapping malicious payloads in legitimate files or packing files to avoid detection, to transport attacks or exploits through network security devices and tools. As organizations have increasingly deployed virtual sandboxes for dynamic analysis, attackers have evolved to focus on ways to evade them. They employ techniques that scan for valid user activity, system configurations, or indicators of specific virtualization technologies. With the growth of the cybercrime underground, even a novice attacker can purchase plug-and-play threats designed to identify and avoid malware analysis environments.

Addressing the Problem

Your firewall, using integrated security services, should automatically prevent known threats. Unknown threats need to be automatically analyzed and countered, too. Where possible, unknown file- or web-based threats should be blocked inline. In the remaining cases, your organization needs a service that can contextualize behaviors through robust analytics and integration with other parts of your security infrastructure to identify threats at all points within the cyberattack lifecycle, not just when threats first enter your network. Blocking known risky file types or access to malicious URLs before they compromise your network reduces your threat exposure.

Your firewall should protect you from known vulnerability exploits, malware, and command-and-control (C2) activity without requiring you to manage or maintain multiple single-function appliances. Signatures should be updated automatically and distributed to all enforcement points as soon as new malware is encountered, keeping you protected while allowing your security and incident response teams to focus on the things that matter.

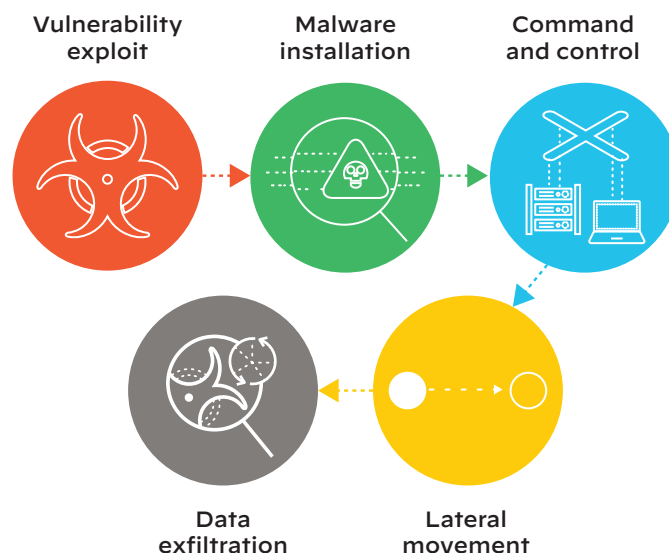


Figure 5: Disruption at every step to prevent successful attacks

A next-generation firewall that utilizes inline machine learning-based prevention and multiple methods of cloud analysis to detect unknown threats, including static analysis with machine learning, dynamic analysis, and bare metal analysis, is capable of high-fidelity, evasion-resistant discovery. Rather than signatures based on specific attributes, firewalls should use predictive machine learning models resident on the firewall or content-based signatures to detect variants, polymorphic malware, or C2 activity. In addition, C2 signatures based on analysis of outbound communication patterns are much more effective protective measures that can scale at machine speed when created automatically. Finally, cloud-delivered security infrastructure is critical for security enforcement. It supports threat detection and prevention at massive scale across your network, endpoints, and clouds in addition to allowing you to tap into an open ecosystem of trusted innovators.

7. Stop Attacks That Use DNS

The Problem

DNS is a massive, often overlooked channel that can be used for malware delivery, C2, and data exfiltration. Adversaries take advantage of the widespread nature of DNS to abuse it at

4. "Google Transparency Report: HTTPS encryption on the web," Google, Inc., accessed February 5, 2020, <https://transparencyreport.google.com/https/overview?hl=en>.

multiple points of an attack. According to Palo Alto Networks Unit 42 threat research team, almost 80% of malware uses DNS as a way to establish communication with a C2 server. Attackers establish reliable command channels that are difficult to take down or identify since DNS is such a reliable way to maintain a connection to DNS servers. Once a connection is established, attackers can use DNS traffic to deliver malware into a network or tunnel data out. Additionally, attackers develop domain generation algorithms (DGAs), which automatically create thousands of malicious domains that can be used for C2. As adversaries increasingly automate their attacks, it becomes almost impossible to identify and stop these threats.

Addressing the Problem

Your organization cannot simply add attacks that use DNS to the block list as this tactic often relies on relatively static threat feeds that work off known bad domains. Without analytics, it is impossible to predict highly dynamic malicious domains. Stopping attacks that use DNS requires a next-generation firewall that can apply predictive analytics and machine learning to identify unknown bad domains dynamically.

8. Protect Your Growing Mobile Workforce

The Problem

The mobile workforce continues to grow along with the use of mobile devices to connect to business applications, often through public networks and devices that are open to advanced threats. This increases risk when users are off-premises because there is no network firewall to stop attacks, and the issue becomes even more complex when considering the effects of cloud and bring-your-own-device (BYOD) practices. In addition, remote locations and small branch offices often lack consistent security because it is operationally inefficient and costly to ship firewalls to them or backhaul traffic to headquarters.

Addressing the Problem

The mobile workforce and remote locations need access to applications from places far beyond your network. They also need protection from targeted cyberattacks, malicious applications and websites, phishing, C2 traffic, and other unknown threats. This requires consistent security. Your next firewall must enable the required levels of visibility, threat prevention, and security policy enforcement to protect your distributed users and locations by delivering next-generation firewall capabilities from the cloud, securing them without the need to deploy physical hardware.

Holistic Coverage for All Operating Systems

Given the onslaught of BYOD initiatives and an increasingly mobile workforce, holistic coverage across Windows®, macOS®, Android®, and Linux environments and workloads is critical. Holistic coverage allows organizations to confidently prevent known and unknown malware regardless of which operating systems their users prefer.

9. Extend Security to Your Evolving Cloud Environments

The Problem

Data and applications reside everywhere—in your network and in the cloud. According to the RightScale 2019 State of the Cloud Report™, 84% of enterprises use multiple public, private, and/or hybrid clouds—five different clouds on average.⁵ Compounded with SaaS environments, organizations must now secure sensitive data in the network and a variety of cloud environments. In addition, legacy security tools and techniques designed for static networks weren't designed to work with cloud native tools or capabilities. Moreover, native security services from the cloud providers themselves, such as Google Cloud Platform (GCP®), Amazon Web Services (AWS®) and Microsoft Azure®, typically provide only Layer 4 protections and are specific to that cloud provider.



Figure 6: RightScale findings on multi-cloud strategy

Addressing the Problem

To succeed, your organization needs cloud security that extends policy consistently from the network to the cloud, stops malware from accessing and moving laterally (east-west) within the cloud, simplifies management, and minimizes the security policy lag as cloud workloads change. Your next firewall must protect the resident applications and data with the same security posture you may have established on your physical network. To secure multi-cloud deployments, the firewall must support a variety of cloud and virtualization environments, including all major public clouds and virtualized private clouds. The firewall must integrate with native cloud services, such as AWS Lambda and Azure, and automation tools, such as Ansible® and Terraform®, to integrate security into your cloud-first development projects.

10. Use a Zero Trust Strategy

The Problem

Conventional security models operate on the outdated assumption that everything inside an organization's network can be trusted. These models are designed to protect the perimeter. Meanwhile, threats that get inside the network go unnoticed and are left free to compromise sensitive, valuable business data. In the digital world, trust is nothing but a vulnerability.

5. "RightScale 2019 State of the Cloud Report," Flexera, February 2019, <https://resources.flexera.com/web/media/documents/rightscale-2019-state-of-the-cloud-report-from-flexera.pdf>.

Addressing the Problem

When evaluating a next-generation firewall, consider a firewall that can act as a segmentation gateway to enable a Zero Trust architecture. **Zero Trust** is a cybersecurity strategy that eliminates the notion of trust. In a Zero Trust world, there are no trusted devices, systems, or people. You identify the data, assets, applications, and services most critical to the business, determine who or what should have access based on their specific job function, and enforce a least-privileged access model through network segmentation, granular Layer 7 security policy, user access control, and threat prevention.

Your next firewall should directly align with Zero Trust, including enabling secure access for all users irrespective of location, inspecting all traffic, enforcing policies for least-privileged access control, and detecting and preventing advanced threats. This significantly reduces the pathways for adversaries, whether they are inside or outside your organization, to access your critical assets. [Watch this webinar](#) to get insight into effectively implementing Zero Trust.

11. Maintain Consistent Access Control and Policy Across Clouds and On-Premises, Remote, and Mobile Networks

The Problem

Organizations have adopted a wide range of point products to address different network and security requirements for applications hosted on-premises, in cloud environments, or both. However, for each product comes a separate policy and interface to manage, creating extra costs, complexity, and gaps in security. Additionally, these products are not integrated and cannot share insights into network access, application access, or policy violations, nor can they provide consolidated logs. Organizations also find it challenging to onboard new firewall appliances at scale, maintain consistent security policies, and deploy policy changes across thousands of firewalls. This approach causes gaps in security and network performance, leading to staff and cost shortages.

Addressing the Problem

According to ESG Research, 66% of organizations are actively consolidating the number of cybersecurity vendors with which they do business.⁶ To be successful, firewall solutions must deliver security capabilities in a variety of form factors—hardware, software, and containerized—to integrate security protections into the optimal parts of the environment. You must be able to operationalize the deployment of consistent, centralized security policies across tens of thousands of firewalls spanning on-premises and cloud deployments—including remote locations, mobile users, and SaaS applications—through centralized management, consolidated core security tasks, and streamlined capabilities. For example, you should be able to use a single console to view all network traffic, manage configurations, push global policies, and generate reports on traffic patterns or security



Figure 7: ESG Research findings on cybersecurity vendor consolidation

incidents. Your reporting capabilities must let your security personnel drill down into network, application, and user behavior for the context they need to make informed decisions.

When these capabilities are delivered from the cloud, your teams can get the networking and security needed in an architecture designed for everything: traffic, applications, and users, no matter their location. In today's constantly changing threat landscape, using a single security vendor to address the vast spectrum of your security and business needs may not be practical. In this case, the ability to integrate with and consume third-party insight and innovation is critical. When evaluating security vendors, be sure to evaluate the flexibility, extensibility, and programmability of what they offer. Read this [e-book](#) to learn about a new approach to securing cloud-enabled organizations as well as delivering speed and agility to enterprise networking and security.

12. Automate Routine Tasks and Focus on the Threats That Matter

The Problem

In a 2018 report, Cybersecurity Ventures estimated that there would be 3.5 million unfilled cybersecurity jobs by 2021.⁷ This is compounded by a dependency on too many manual processes for day-to-day security operations, such as chasing down data, investigating false positive alerts, and managing remediation. Manually analyzing and correlating the vast number of security events slows mitigation, increases the chance for error, and is difficult to scale. Security teams can easily drown in the volume of alerts and miss the critical, actionable ones. The looming shortage of skilled cybersecurity professionals only exacerbates this. Although big data analytics uncovers hidden patterns, correlations, and other insights to provide security teams with actionable intelligence, you still need the right data. That data must be sourced from everywhere—networks, endpoints, SaaS applications, public clouds, private clouds, data centers, and so on—and be ready for analytics.

6. "The Cybersecurity Technology Consolidation Conundrum, March 26, 2019, <https://www.esg-global.com/blog/the-cybersecurity-technology-consolidation-conundrum>.

7. "Cybersecurity Jobs Report 2018–2021," Cybersecurity Ventures, November 28, 2018, <https://www.herjavecgroup.com/wp-content/uploads/2018/11/HG-and-CV-Cybersecurity-Jobs-Report-2018.pdf>.




Figure 8: Cybersecurity Ventures findings on cybersecurity jobs

Addressing the Problem

By using precise analytics to drive automation, you can easily operate security best practices like Zero Trust, streamline routine tasks, and focus on business priorities—such as speeding up application delivery, improving processes, or hunting for threats. There are three ways to think about automation:

1. **Workflow automation:** The firewall must expose standard APIs so it can be programmed from other tools and scripts you may be using. In the cloud, it must integrate with tools like Ansible and Terraform. In addition, the firewall must be able to kick off workflows on other devices in your security ecosystem, using their APIs, without manual intervention. This automation should extend to operationalization of rules and threat information in open source formats, such as Snort or Suricata®.
2. **Policy automation:** The firewall must be able to adapt policies to any changes in your environment, such as movement of applications across virtual machines. It must also be able to ingest threat intelligence from third-party sources and automatically act on that intelligence.
3. **Security automation:** Your environment must be able to uncover unknown threats and deliver protections to the firewall so new threats are blocked automatically.

Some threats remain hidden in data. By looking deeper into that data across locations and deployment types, you can find threats that may be lurking in plain sight. With automation, you can accurately identify threats, enable rapid prevention, improve efficiency, better utilize the talent of your specialized staff, and improve your organization's security posture.

13. Coordinate Detection and Analytics with Other Security Tools

The Problem

Advanced adversaries don't limit themselves to one part of your architecture. Instead, their goal is to move laterally from endpoints to your network, clouds, and other data structures to access and exfiltrate valuable data. With this in mind, siloed security approaches that can only see and

understand one slice of your infrastructure produce sub-optimal results. They limit the application of analytics and force security analysts to bounce between interfaces to try to manually piece together attacks—a process that is both time-consuming and prone to error.

Addressing the Problem

As the number of needed security functions increases, so does the potential value of platforms/devices that can provide meaningful integration between them. If your firewall can act as a sensor and enforcement point for a more comprehensive, machine-learning driven analytics platform (such as an [XDR](#) solution), your security team will gain both efficacy and efficiency in uncovering, remediating, and preventing sophisticated attacks. Your next firewall should integrate with XDR to allow both your network and security teams to understand the full scope of an attack, share threat context and intelligence, and drive automated response as well as enforcement between the firewall and other enforcement points.

14. Consolidate Connectivity with Security

As more enterprises embrace digital transformation and move applications to the cloud, IT teams are challenged to quickly, reliably, and securely connect corporate locations and branches to critical business resources. [Software-defined wide area networking](#) (SD-WAN) promises to increase bandwidth while improving connectivity and performance, and organizations are taking note. According to Gartner, SD-WAN will be included in 75% of WAN infrastructure refreshes by 2020.⁸ However, while SD-WAN offers many benefits, it also brings many challenges, such as degraded or bolted-on security, unforeseen architecture and deployment complexity, and unpredictable performance.

Addressing the Problem

Your next firewall should extend to your branches the same consistent security that protects your data center and cloud environments. Organizations can adopt SD-WAN safely by implementing a firewall that natively integrates with it, consolidating their connectivity and security. This can also




Figure 9: Gartner findings on SD-WAN adoption

help maintain consistent security policies from the network core out to branches. With SD-WAN configuration and monitoring as well as firewall user and application policy workflows available through a single pane of glass, organizations can avoid gaps in their security posture as well as benefit from improved security, simplicity, and efficiency. [Read this e-book](#) to learn how to achieve consistent security with SD-WAN.

15. Reduce Exposure IoT Security Risks

The Problem

According to Unit 42, [30% of devices](#) on enterprise networks today are internet of things (IoT) devices. These devices pose a huge cybersecurity risk as they are often shipped with vulnerabilities and are difficult (or impossible) to patch, yet have unfettered access to an organization's network. It can be challenging to secure these IoT devices for many reasons. IT is rarely notified of or involved in the purchase and management of IoT devices. Connected devices are incredibly diverse, running a variety of operating systems, firmware versions, and applications. IT teams lack governance over these devices and may have no means to monitor, manage, or track vulnerabilities. Security teams find it hard to define and enforce policies as they don't know how trusted behavior for a device is defined. Because the useful life of devices can exceed a vendor's supported life, many IoT devices with long lifecycles use legacy operating systems and outdated applications, and thus IT cannot put third-party agents on them to secure them.

Addressing the Problem

When evaluating your next firewall, consider a solution that can identify and classify all IoT devices on your network, including those never seen before. Your firewall solution should empower security teams to make decisions quickly with full context for each device, understanding device identity, risk level, and any behavioral anomalies. The firewall should then offer policy recommendations based on risk assessment results that can be automatically enforced natively so that users don't leave their devices insecure. [Read this article](#) to learn how to keep your organization secure in the age of IoT.

Using the RFP Process to Select a Next-Generation Firewall

Recommended RFP Questions and Considerations

Next-generation firewalls must deliver a wide range of capabilities to provide efficient, effective security while also integrating with other key prevention, detection, and response tools across the infrastructure. This section translates those requirements into a comprehensive RFP checklist to help you evaluate the quality of the firewall platforms you're consider-

ing. Use this checklist as a starting point, and tailor it to your company's needs to ensure you're able to identify vendors that can best protect your organization.

Identify Users and Enable Appropriate Access

Can your next-generation firewall...

- Provide consistent security policy for mobile users?
- Protect users who are not behind a next-generation firewall?
- Use multiple physical/virtual firewalls to support an always-on VPN connection?
- Utilize the cloud to bring protection closer to the user?

Prevent Credentials Theft and Abuse

Can your next-generation firewall...

- Prevent use of corporate credentials on unknown websites?
- Block users from submitting corporate credentials without storing a copy of the hash in the firewall?
- Quickly analyze previously unknown phishing sites and update its protections?
- Log user attempts to submit credentials in HTTP post?
- Support MFA as part of access-control policy based on the sensitivity of the resource accessed?
- Provide a variety of choices in MFA partner technologies?
- Support API integrations with MFA partner technologies?
- Support MFA policy for any type of application, including web, client-server, and terminal applications?
- Support MFA capability on any protocol, rather than be limited to certain protocols?

Safely Enable All Apps and Control Functions

- Can your firewall detect applications that can evade detection using nonstandard ports, port hopping, or by being configured to run on a different port?
- In traffic classification, is the first task the firewall executes based on application identity or the network port?
 - Are the application identification mechanisms part of the core firewall traffic classification (i.e., enabled by default)?
 - Do the application identification mechanisms depend on the application's standard port?
 - Can the signatures be applied to all ports? Is the process automatic or manually configured?
- When traffic first hits the device, is it first classified based on port (e.g., port 80, thus assumed to be HTTP) or application (e.g., Gmail®)?
- Describe in detail how the firewall can accurately identify applications:
 - Which mechanisms, besides signatures, are used to classify traffic?
 - What is the breadth of application and protocol decoder use?

8. Mike Toussaint, Ted Corbett, Andrew Lerner, "6 Critical Questions to Ask on SD-WAN," Gartner, June 6, 2018, <https://www.gartner.com/en/documents/3877766/6-critical-questions-to-ask-on-sd-wan>.

- How are SSL and SSH decryption and control implemented?
- Are the traffic classification mechanisms applied equally across all ports?
- Which mechanisms are used to detect purposely evasive applications, such as UltraSurf or encrypted P2P?
- Is application identification performed in the firewall, or is it performed in a secondary process, after port-based classification?
- Is application state tracked? If so, how is it utilized to ensure consistent control of the application and associated secondary functions?
- Is the identity of the application the basis of the firewall security policy, or is application control treated as a secondary policy element?
- How often is the application database updated? Is it a dynamic update or a system reboot upgrade?

Close Dangerous Policy Gaps

- Is stateful inspection traffic classification performed separately, prior to application identification?
 - Once an application is identified, how are changes in application state monitored, tracked, and used within policy?
- How does the application database hierarchy (flat, multi-level, other) expose functions within the parent application for more granular enablement policies?
- What levels of control can be exerted over individual applications and their respective functions?
- Can port-based controls be implemented for all applications in the application database so an administrator can enforce, by policy, the application and port relationship? For example:
 - Ensure IT personnel are the only ones allowed to use SSH and RDP?
 - Detect and block malware within the application, even on a nonstandard port?
- Which enterprise identity repositories are supported for user-based controls?
- Is an API available for custom or nonstandard identity-infrastructure integration?
- How are policy-based controls implemented by users and groups for terminal services environments?
- If any, what are the differences in application enablement options for hardware and virtualized instances?
- Does your firewall collect telemetry data for ML-based security policy optimization to eliminate breaches due to misconfiguration?

Secure Encrypted Traffic

- What is the process by which encrypted applications are identified on all ports, including nonstandard ports?
- What policy controls are available to selectively decrypt, inspect, and control applications that are using SSL?
- Are bidirectional SSL identification, decryption, and inspection supported?

- Is SSL decryption a standard feature or an extra cost? Is a dedicated device required?
- Is SSH control (a means of accessing remote devices) supported? If so, what is the depth of control?
- What mechanisms are used to identify evasive applications, such as UltraSurf and Tor?
- How does the product automatically identify a circumventor that is using a nonstandard port?

Stop Advanced Threats to Prevent Successful Cyberattacks

- Does your next-generation firewall deliver machine learning-based prevention of unknown malware files and variants, including executables as well as 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 cloud-based malware analysis system support multiple analysis techniques, including bare metal analysis, to detect evasive, sandbox-aware malware?
- 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 known and unknown variants of malware?
 - 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, and macOS operating systems?
- 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 automatically and dynamically import all known IOCs (i.e., IPs, domains, and URLs) into the block list to be proactive against known ransomware families?
- Does the threat intelligence cloud integration with the next-generation firewall support dynamic updates for malicious URLs related to ransomware in the malware category of the URL filtering database?
- Can your next-generation firewall learn about threats or ransomware behavior from your endpoint protection software and vice versa?

Stop Attacks That Use DNS

- Does cloud-based 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 the block list or sinkholed?

Protect Your Growing Mobile Workforce

- In detail, what are the available options for securing remote users, including all necessary components?
 - If a client component is included, how is it distributed?
- What are your sizing requirements? How many users can be supported simultaneously?
- Is the remote user security feature set transparent to the client?
- How is policy control over remote users implemented (in firewall policy, in a separate policy/device, etc.)?
- Which features and protections are provided by the remote capabilities (SSL, application control, IPS, etc.)?
- Can your firewall keep users connected to ensure consistent policy enforcement regardless of location?
- How do you address mobile device users? Will you be able to provide consistent policy enforcement whether users are on external networks or internal wireless networks?
- Can the firewall address BYOD issues, such as safely enabling corporate and personally owned laptops, phones, and tablets?

Extend Security to Your Evolving Cloud Environments

- 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?
- In virtualized environments, how is traffic classified throughout the virtual machine (east-west, north-south)?
 - What are the points of integration within the virtualized environment?
 - What does the process of building security policies for newly created virtual machines look like?
 - What features are available to track virtual machine moves, adds, and changes?
 - What features are available for integration with automation and orchestration systems?

Use a Zero Trust Strategy

- Does your next-generation firewall enable you to write context-based policy to determine who or what can access your protect surface?
- How does the next-generation firewall leverage network segmentation, prevent lateral movement, provide Layer 7 threat prevention, and simplify granular user access control?
- Does the next-generation firewall inspect all traffic for malicious content and unauthorized activity as well as log through Layer 7, both inside and outside, across the network and cloud environments?

Maintain Consistent Access Control and Policy Across Clouds as Well as On-Premises, Remote, and Mobile Networks

- Does your firewall solution deliver consistent network security and threat prevention aligned to applications hosted on-premises and in virtualized and container environments?
- Does your firewall solution 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 to meet PCI compliance?
- Can local administrators work directly on the appliance and change configurations as needed without logging in to a central manager?
- Can central administrators monitor and view 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?

Automate Routine Tasks and Focus on the Threats That Matter

- Does your security vendor support the capability to automatically generate prevention signatures across the attack lifecycle for all data relevant to attacks?
- Can your firewall correlate and identify infected hosts in the network and quarantine them to limit their access in the network?
- Can your firewall trigger MFA to prevent credential abuse and secure critical applications?
- Can your firewall correlate the threats seen in the network with information obtained from global threat intelligence?

Coordinate Detection and Analytics with Other Security Tools

Can your next-generation firewall or manager...

- Create a ticket on a change management system based on a malicious event seen on the firewall?
- Trigger a quarantine action for an infected host on the wireless network?

Can your next-generation firewall...

- Be completely programmed via API?
- Collect User-ID information via APIs from wireless controllers about hosts connecting to wireless networks?
- 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?
- Integrate with your next-generation firewall to automate timeout of expired threat indicators to avoid using stale threat intelligence?
- Allow you to target threat indicators from recent APT campaigns and incorporate threat feeds proactively on your next-generation firewall?
- Allow you to enrich cloud-based threat intelligence and IOCs with intelligence based on a confidence rating to reduce the operational overhead from dealing with false positives?

Consolidate Connectivity with Security

- What functionality does the next-generation firewall support to ensure the security of end-to-end communications?
- How does the next-generation firewall ensure the security of direct internet access (DIA)?

- How does the next-generation firewall enforce security policies for cloud services and applications delivered to the branch office?
- How does the next-generation firewall improve the security of cloud applications and services accessed from a branch office?

Reduce Exposure to IoT Security Risks

- Does your firewall solution provide visibility into all IoT devices in your network, including those never seen before?
- How does your firewall solution secure IoT devices against known and unknown attacks?
- Does your firewall offer policy recommendations, based on risk assessment results, that can be automatically enforced?
- Can your firewall segment IoT device traffic from the rest of the network to prevent compromised IoT devices from becoming beachheads for attacks?

Are you ready to evaluate your next firewall? Take an [Ultimate Test Drive](#).