

13 THINGS YOUR NEXT FIREWALL MUST DO

The rapid evolution of IT has changed the face of the network perimeter. Data and users are everywhere. Devices are proliferating more quickly than most organizations can keep up. At the same time, IT teams are adopting the cloud, big data analytics, 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.

Cybersecurity is not keeping up as attacks continue to disrupt business. 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 different 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 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” and highlights the thirteen key things a next-generation firewall (NGFW) must do to secure your network and your business.

13 Things Your Next Firewall Must Do

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.

Unlike UTM offerings, NGFWs 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 “NGFW” is now synonymous with “firewall.”

NGFW 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, new innovations should be tightly integrated and easy to adopt. Although requirements and priorities will vary within these criteria, there are thirteen things your next 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 risk 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. 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.

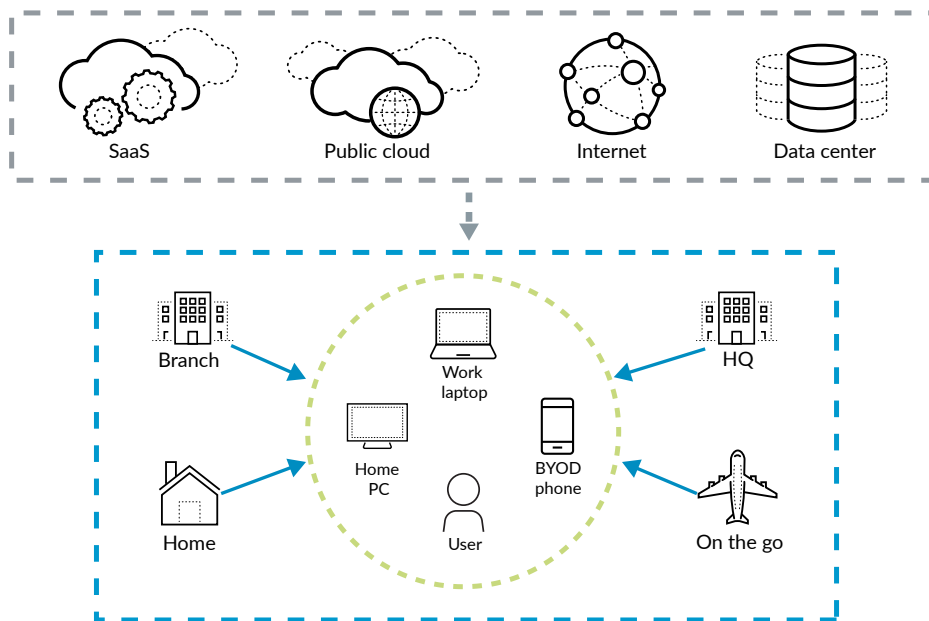


Figure 1: Users access data from different devices and locations

By the end of 2019, 90 percent of enterprise internet connections for the installed base will be secured using next-generation firewalls.¹

NGFW Requirements

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

1. Adam Hils, Jeremy D'Hoinne, Rajpreet Kaur, "Magic Quadrant for Enterprise Network Firewalls," Gartner, July 10, 2017.

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 VPN, 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. The 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, the issue of user identity goes beyond classifying users for policy reporting.

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 2017 Data Breach Investigations Report by Verizon, in the twelve-month period covered in the report, 81% 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 black market. Attackers use these credentials to gain access to a network, move laterally, and escalate their privileges for unauthorized access to applications and data.

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.

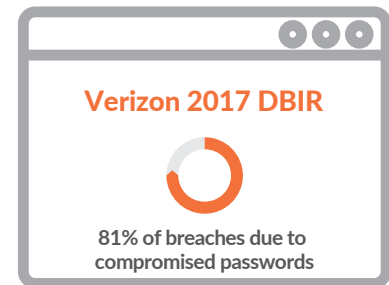


Figure 2: Verizon 2017 DBIR finding on compromised passwords

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 VoIP, 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 another good example. 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.

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

2. "2017 Data Breach Investigations Report," Verizon, 2017, www.knowbe4.com/hubfs/rp_DBIR_2017_Report_execsummary_en_xg.pdf.

usage, such as upload versus download permissions, chat versus file transfer, and so on. This must be done continuously. The concept of “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.

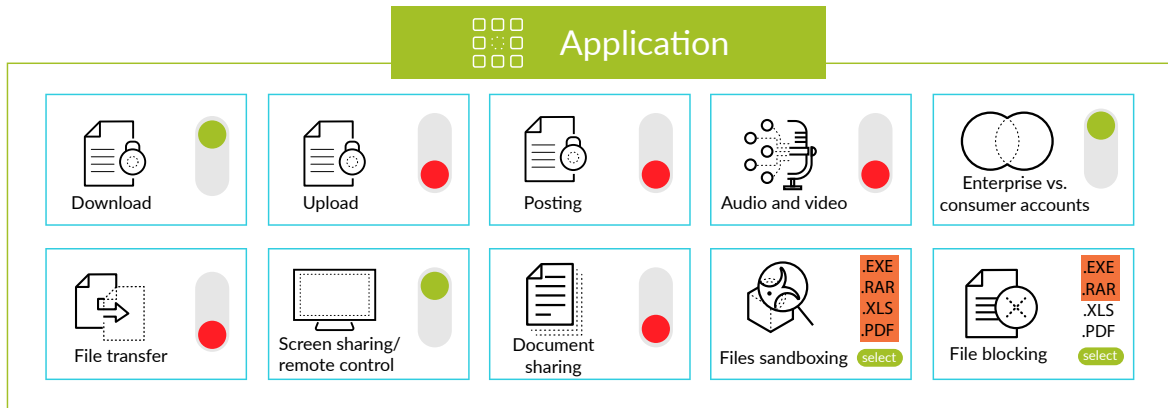


Figure 3: Control application usage in policy

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. An NGFW should help your security team easily replace legacy rules with intuitive, application-based policies. Because application-ID-based rules 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.

5. Secure Encrypted Traffic

The Problem

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

Addressing the Problem

The ability to decrypt SSL and SSH 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 with predictable performance. However, your next firewall must be




Figure 4: Google 2019 finding on encrypted traffic⁴

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

4. “Google Transparency Report: HTTPS encryption on the web,” Google, accessed March 8, 2019, transparencyreport.google.com/https/overview?hl=en.

flexible enough to easily decrypt certain types of encrypted traffic—such as HTTPS from unclassified websites—via policy, while other types—such as 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. 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, any attacker, novice or advanced, 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 block known threats. Unknown threats need to be automatically analyzed and countered, too. Your organization needs a service that looks for 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 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.

Attack Lifecycle



Figure 5: Disruption at every step to prevent successful attacks

A next-generation firewall that utilizes multiple methods of 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 use signatures based on specific attributes, firewalls should use 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 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 blacklist attacks that use DNS 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 2018 State of the Cloud Report™, 81% 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 do not 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.

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 virtual workloads change. Your next firewall must protect the resident applications and data with the same security posture that 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 cloud providers and virtualized private clouds. The firewall must integrate with native cloud services, such as Amazon Lambda and Azure, and automation tools, such as Ansible® and Terraform®, to integrate security into your cloud-first development projects.




Figure 6: RightScale findings on multi-cloud strategy

5. "2018 State of the Cloud Report," RightScale, 2018, www.suse.com/media/report/rightscale_2018_state_of_the_cloud_report.pdf.

10. Use a Zero Trust Strategy

The Problem

Conventional security models operate on the outdated assumption that you can trust everything inside your network. However, given the increased sophistication of attacks and insider threats, you need new security measures to stop them from spreading once they're inside. Because traditional security models are designed to protect your perimeter, threats that get inside your network are invisible to them and go uninspected, free to morph and move wherever they choose to extract sensitive business data. In the digital world, trust is nothing but a vulnerability.

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 strategy designed around the concept that users, applications, and data should never be trusted—that their actions should always be verified in an environment. The primary goal of the Zero Trust model is to eliminate trust in a system and prevent attackers from exploiting vulnerabilities hidden in trusted applications. The approach involves limiting the scope of an attack and blocking lateral movement by taking advantage of microsegmentation based on users, data, and location. A next-generation firewall platform should help with these steps, including enabling secure access for all users irrespective of location, inspecting all traffic, enforcing policies for least-privileged access control, and detecting as well as preventing advanced threats. This significantly reduces the pathways for adversaries to access your most critical data and applications, whether the adversaries are outside or inside your organization. [Watch this webinar](#) to get insight into effectively implementing Zero Trust.

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

The Problem

Individual security products typically come with their own management applications. To configure security for each one, security operators must work with different management devices. According to the 2017 U.S. IT Services Report from ResearchCorp, nearly 72% of organizations use products from three or more separate vendors to secure their network infrastructure.⁵ These products are disconnected and cannot share insights. Organizations also find it challenging to scale firewall onboarding, maintain consistent security policies, and deploy emergency changes across thousands of firewalls. This makes security complex and stretches IT teams to the limit.

Addressing the Problem

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 configuration, push global policies, and generate reports on traffic patterns or security 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 build out the right security architecture to prevent known and unknown threats at every corner of your extended network. In today's constantly changing threat landscape, using a single security vendor to address the vast spectrum of your security and business needs isn't always practical. In this case, the ability to integrate with and consume third-party insight and innovation is critical. When evaluating future security vendors, be sure to evaluate the extensibility and programmability of what they offer.

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

The Problem

A survey from the Enterprise Strategy Group found 51% of cybersecurity professionals feel their organization has a problematic shortage of cybersecurity skills.⁶ 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. This is exacerbated by a looming shortage of skilled cybersecurity professionals. 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. "2017 U.S. IT Services Report," ResearchCorp.org, 2017, www.fidelus.com/wp-content/uploads/2017/12/researchcorp-fidelus_us_it_servicesreport_full_report.pdf.

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 application delivery, improving processes, or hunting for threats. There are three ways to think about automation:

- **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.
- **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.
- **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. Consume New Security Innovations Easily

The Problem

Consuming cybersecurity innovation is arduous. Organizations waste time deploying additional hardware or software every time they want to take advantage of a new security technology. They invest more resources managing their ever-expanding security infrastructure instead of improving their security controls to stay ahead of attackers and prevent threats.

Addressing the Problem

As the number of needed security functions increases, there are two options: add more independent point products or use an existing device to support new capabilities. If your firewall can act as a sensor and enforcement point for third-party technology, you can rapidly adopt new security innovations without deploying or managing endless new devices. Your next firewall should enable teams to quickly discover, evaluate, and use new security technologies. Security teams should be able to collaborate between different apps, share threat context and intelligence, and drive automated response and enforcement with deeply integrated applications. This way, they can solve the most challenging security use cases with the best technology available, and they can do so without the cost or operational burden of deploying new infrastructure for each new function. [Watch this video](#) to learn how an open and integrated, AI-based continuous security platform can help you discover new innovative apps and capabilities.



Figure 7: ResearchCorp 2018 findings on multiple vendors for network security

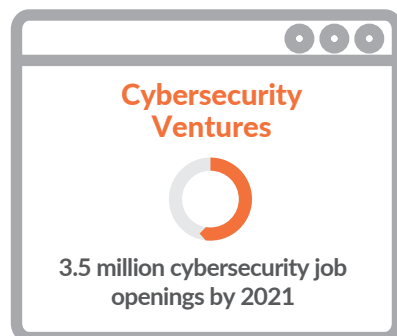



Figure 8: Cybersecurity Ventures findings on cybersecurity jobs

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



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

© 2019 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies. 13-things-your-next-firewall-must-do-wp-050719