

10

Things Your Next Firewall Must Do



Introduction: 10 Things Your Next Firewall Must Do

Much has been made about bringing application visibility and control into network security. The reason is obvious: applications can easily slip by traditional port-based firewalls. And the value is obvious: employees use any application they need to get their job done—often indifferent to the risk that use poses to the business. Nearly every network security vendor has acknowledged that application control is an increasingly critical part of network security. While the next-generation firewall (NGFW) is well defined by Gartner as something new, enterprise-focused, and distinct, many network security vendors are claiming NGFW is a subset of other functions (e.g., UTM or IPS). Most traditional network security vendors are attempting to provide application visibility and control by using a limited number of application signatures supported in their IPS or other external database. But underneath, these capabilities are poorly integrated and their products are still based on legacy port-blocking technology, not NGFW technology. Perhaps most importantly, these folks are missing the point – it's not about blocking applications, but safely enabling them. Unfortunately, the products proffered by traditional network security vendors ignore much of what enterprises do with applications today – they use them to enable their business – and as such, need to make sure that those applications run securely. It is obvious that a next-generation firewall is a different and revolutionary class of product, but the interest from enterprise customers is so strong

that vendors of traditional products are trying to subvert the interest of enterprise network security team by attempting to look like an NGFW.

Definition: Next-generation firewall.

5 Requirements:

1. Identify applications regardless of port, protocol, evasive tactic or SSL
2. Identify users regardless of IP address
3. Protect in real-time against threats embedded across applications
4. Fine-grained visibility and policy control over application access / functionality
5. Multi-gigabit, in-line deployment with no performance degradation

For enterprises looking at NGFWs, the most important consideration is: Will this new technology empower security teams to securely enable applications to the benefit of the organization? Key questions to ask include:

- Will it increase visibility and understanding of application traffic?
- Will it expand traffic control options beyond blunt allow/deny?
- Will it help prevent threats?
- Will it eliminate the need to compromise between performance and security?
- Will it reduce costs for my organization?
- Will it make the job of risk management easier or simpler?

If the answers to the above questions are “yes,” then transition is easy to justify.



There are substantial differences between NGFWs and UTM-style devices – in terms of the kinds of organization each targets, and in terms of architecture and security model. These differences have dramatic impacts on real-world functions/features, operations, and performance – as we’ve attempted to capture in the “ten things” section below.

Architecture and Security Model: Traffic is Best Classified in the Firewall

In building “next-generation firewalls,” security vendors have taken one of two architectural approaches:

1. Build application identification into the firewall as the primary classification engine
2. Add application signatures to an IPS or IPS-like pattern matching engine which is then added to a port-based firewall

Both can recognize applications – but with varying degrees of success, usability, and relevance. Most importantly, these architectural approaches dictate a specific security model for application policies – either positive (default deny), or negative (default allow).

Firewalls use a positive security model. Another term for it is default deny. Which means that administrators write policies to ALLOW traffic (e.g., allow WebEx)...and then everything else is denied or blocked. Negative policies (e.g., block Limewire) can be used in this model, but the most important fact is that the end of the policy in a positive security model says, “all else deny.” One of the key implications of this approach is that all traffic must be classified in order to allow the appropriate traffic. So visibility of traffic is easy and complete. Policies enable applications. Another key result of this approach is that any unknown traffic is, by default, denied. In other words, the best next-generation firewall is a firewall.

Intrusion prevention systems (IPS) typically employ a negative security model, or default allow. Which means that IPS identifies and blocks specific traffic (traditionally threats)...and everything else is passed through. Traditional network security vendors are adding application signatures to an IPS-style engine and bolting it onto a traditional port-based firewall. The result is an “application prevention system.” The application control is in a negative security model – in other words, it’s not in a firewall. Implication: one only sees what is expressly looked for, and unknown traffic is, by default, allowed.

While this paper is focused on the 10 specific things your next (generation) firewall must do, knowledge of the architecture and model as outlined above are prerequisites to understanding the different capabilities of the different products on the market and their ability to deliver these functions.

The “ten things” discussed below represent some of the critical, specific requirements we’ve gathered from thousands of IT organizations since we began selling next-generation firewalls in 2007. These are all real-world examples of requirements that make the job of securing enterprise networks easier, better, or simpler – marketing hype aside.

The 10 Things Your Next (Generation) Firewall Must Do

There are three areas of difference – security functions, operations, and performance. The security functional elements correspond to the efficacy of the security controls, and the ability for enterprises to manage risk associated with network traffic. From an operations perspective, the big question is, “where does application policy live, and how hard or complex is it to manage?” The performance difference is simple: can the firewall do what it’s supposed to do at the throughput it’s supposed to do it? The Ten Things Your Next (Generation) Firewall Must Do are:

1. Identify and control applications on any port
2. Identify and control circumventors
3. Decrypt outbound SSL
4. Provide application function control
5. Scan for viruses and malware in allowed collaborative applications
6. Deal with unknown traffic by policy
7. Identify and control applications sharing the same connection
8. Enable the same application visibility and control for remote users
9. Make network security simpler, not more complex with the addition of application control.
10. Deliver the same throughput and performance with application control active

1

Your next firewall must identify and control applications on any port, not just standard ports (including applications using HTTP or other protocols)

Business case: Application developers no longer adhere to standard port/protocol/application mapping. More and more applications are capable of operating on non-standard ports or are can hop ports (e.g., instant messaging applications, peer-to-peer file sharing, or VOIP). Additionally, users are increasingly savvy enough to force applications to run over non-standard ports (e.g., MS RDP, SSH). In order to enforce application-specific policies where ports are increasingly irrelevant, your next firewall must assume that any application can run on any port. This is one of the fundamental changes in technology that made the NGFW an absolute necessity. It was this change to applications that made the positive control of traditional port-based firewalls obsolete. It also underscores why a negative control model can't solve the problem. If an application can move to any port, a product based on negative control would have to run all signatures on tens of thousands of ports.

Requirements: This one is simple – if any application can run on any port – your next firewall must classify traffic, by application, on all ports – all the time (see #4 and #7). Otherwise, security controls will continue to be outwitted by the same techniques that have plagued them for years.

2

Your next firewall must identify and control circumventors: proxies, remote access, and encrypted tunnel applications

Business case: Most organizations have security policies – and controls designed to enforce those policies. Proxies, remote access, and encrypted tunnel applications are specifically used to circumvent security controls like firewalls, URL filtering, IPS, and secure web gateways. Without the ability to control these circumventors, organizations cannot enforce their security policies, and expose themselves to the very risks they thought their controls mitigated. To be clear, not all of these types of applications are the same – remote access applications have legitimate uses, as do some encrypted tunnel applications. But external anonymous proxies that communicate over SSL on random ports, or applications like Ultrasurf and Tor have only one real purpose – to circumvent security controls.

Requirements: There are different types of circumvention applications – each using slightly different techniques. There are both public and private external proxies (see proxy.org for a large database of public proxies) that can use both HTTP and HTTPS. Private proxies are often set up on unclassified IP addresses (e.g., home computers) with applications like PHPProxy or CGIProxy. Remote access applications like MS RDP or GoToMyPC can have legitimate use – but due to the associated risk, should be managed. Most other circumventors, (e.g., Ultrasurf, Tor, Hamachi) don't have business uses. There are, of course, unknown circumventors – see #6 below. Regardless

of the policy stance, your next firewall needs to have specific techniques to deal with all of these applications, regardless of port, protocol, encryption, or other evasive tactic. One more consideration: these applications are regularly updated to make them harder to detect and control. So it is important to understand not only that your next firewall can identify these circumvention applications, but also how often that firewall's application intelligence is updated and maintained.

3

Your next firewall must decrypt outbound SSL

Business case: Today, more than 15% of network traffic is SSL-encrypted (according to more than 2,400 enterprise network traffic samples – see Palo Alto Networks' Application Usage and Risk Report for details). In some industries (e.g., financial services), it's more than 50%. Given the increasing adoption of HTTPS for many high-risk, high-reward applications that end-users employ (e.g., Gmail, Facebook), and users' ability to force SSL on many websites, network security teams have a large and growing blind spot without decrypting, classifying, controlling, and scanning SSL-encrypted traffic. Certainly, an NGFW must be flexible enough that certain types of SSL-encrypted traffic can be left alone (e.g., web traffic from financial services or health care organizations) while other types (e.g., SSL on non-standard ports, HTTPS from unclassified websites in Eastern Europe) can be decrypted via policy.

Requirements: The ability to decrypt outbound SSL is a foundational element – not just because it's an increasingly significant percentage of enterprise traffic, but also because it enables a few other key features that would end up incomplete or ineffective without the ability to decrypt SSL (e.g., control of circumventors - #2, application function control - #4, scanning allowed applications - #5, and control of applications sharing the same connection - #7). Key elements to look for include recognition and decryption of SSL on any port, policy control over decryption, and the necessary hardware and software elements to perform SSL decryption across tens of thousands of simultaneous SSL connections with good performance and high throughput.

4

Your next firewall must provide application function control (e.g., SharePoint Admin vs. SharePoint Docs)

Business case: Many applications have significantly different functions, presenting different risk profiles and value to both the user, and the organization. Good examples of this include WebEx vs. WebEx Desktop Sharing, Yahoo Instant Messaging vs. the file transfer feature, and regular Gmail vs. sending attachments. In regulated environments, or in organizations heavily dependent on intellectual property this is a significant issue.

Requirements: Continuous classification and fine-grained understanding of each application. Your next firewall has to continually evaluate the traffic and watch for changes – if a different function or feature is introduced in the session, the firewall should note it and perform a policy check. Understanding the different functions of each application and the different associated risks is equally important. Unfortunately, many firewalls classify a traffic flow once, and then “fast path” it (read: never look at that flow again) for better performance. This method pre-dates modern applications and prevents those firewalls from meeting this requirement.

5

Your next firewall must scan for threats in allowed collaboration applications – e.g., Sharepoint, Box.net, MS Office Online

Business case: Enterprises continue to adopt collaborative applications hosted outside their physical locations. Whether it's hosted Sharepoint, Box.net, Google Docs, or Microsoft Office Live, or even an extranet application hosted by a partner, many organizations have a requirement to use an application that shares files – in other words, is a high-risk threat vector. Many infected documents are stored in collaboration applications, along with some documents that contain sensitive information (e.g., customers' personal information). Furthermore, some of these applications (e.g., Sharepoint) rely on supporting technologies that are regular targets for exploits (e.g., IIS, SQL Server). Blocking the application isn't appropriate, but neither is allowing a threat into the organization.

Requirements: Part of safe enablement is allowing an application and scanning it for threats. These applications can communicate over a combination of protocols (e.g., Sharepoint – HTTPS and CIFS, see requirement #3), and require a more sophisticated policy than “block application.” First step is to identify the application (regardless of port or encryption), allow it, and then scan it for any of the appropriate threats – exploits, viruses/malware, or spyware...or even confidential, regulated, or sensitive information.

6

Your next firewall must deal with unknown traffic by policy, not by just letting it through.

Business case: There will always be unknown traffic and it will always represent significant risks to any organization. There are several important elements to consider with unknown traffic – minimizing it, easily characterizing custom applications so they are “known” in network security policy, and having predictable visibility and policy control over traffic that remains unknown.

Requirements: First, by default, your next firewall should attempt to classify all traffic – this is one area where the earlier architecture and security discussion becomes very important. Positive (default deny) models classify everything, negative (default allow) models classify only what they’re told to classify. Second, for custom developed applications, there should be a way to develop a custom identifier – so that traffic is counted among the “known.” Third, the security model plays into these requirements again – a positive (default deny) model can deny all unknown traffic – so what you don’t know can’t hurt you. A negative (default allow) model allows all unknown traffic – so what you don’t know will hurt you. For example, many botnets will use port 53 (DNS) for communication back to their control servers. If your next firewall lacks the ability to see and control unknown traffic, bots will be able to drive right through, unimpeded.

7

Your next firewall must identify and control applications sharing the same connection

Business case: Applications share sessions. To ensure users are continuously using an application “platform,” whether it’s Google, Facebook, Microsoft, salesforce, LinkedIn, or Yahoo, application developers integrate many different applications – which often have very different risk profiles and business value. Let’s look at our earlier example of Gmail – which has the ability to spawn a Google Talk session from within the Gmail UI. These are fundamentally different applications, and your next firewall should recognize that, and enable the appropriate policy response for each.

Requirements: Simple classification of the platform or website doesn’t work. In other words, “fast path” is not an option – “once and done” classification ignores the fact that applications share sessions. Traffic must be continuously evaluated to understand the application, its changes (see #5), when the user changes to a completely different application using the same session, and enforce the appropriate policy controls. Looking briefly at the technical requirements using our Gmail/Google Talk example: Gmail is by default HTTPS (see #3) so the first step is to decrypt – but it has to be continuous, as does the application classification, because at any time, the user can start a chat...which may have a completely different policy associated with it.

8

Your next firewall must enable the same application visibility and control for remote users as for on-premise users

Business case: Users are increasingly outside the four walls of the enterprise. Once the domain of road warriors, now a significant portion of the enterprise user population is capable of working remotely. Whether working from a coffee shop, home, or a customer site, users expect to connect to their applications via WiFi, wireless broadband, or any means necessary. Regardless of where the user is, or even where the application they're employing might be, the same standard of control should apply. If your next firewall enables application visibility and control over traffic inside the four walls of the enterprise, but not outside, it misses the mark on some of the riskiest traffic.

Requirements: Conceptually, this is simple – your next firewall must have consistent visibility and control over traffic regardless of where the user is – inside or outside. This is not to say that enterprises will have the exact same policy for both – some organizations might want employees to use Skype when on the road, but not inside headquarters, where others might have a policy that says if outside the office, users may not download sales-force.com attachments unless they have hard disk encryption turned on. This should be achievable on your next firewall without introducing significant latency for the end user, or undue operational hassle for the administrator, or significant cost for the organization.

9

Your next firewall must make network security simpler, not more complex with the addition of application control.

Business case: Many enterprises struggle with incorporating more information feeds and more policies, and more management into already overloaded security processes and people. In other words, if teams cannot manage what they've already got, adding more management, policies, and information doesn't help. Furthermore, the more distributed the policy is (e.g., port-based firewall allows port 80 traffic, IPS looks for/blocks threats and applications, secure web gateway enforces URL filtering) – the harder it is to manage that policy. Where do admins go to enable WebEx? How do they resolve policy conflicts across these different devices? Given that typical port-based firewall installations have rulebases that include thousands of rules, adding thousands of application signatures across tens of thousands of ports (see #3 above) is going to increase complexity by several orders of magnitude.

Requirements: Firewall policy should be based on user and application. Subsequent content analysis can be performed on allowed traffic, but fundamental access control should be based on relevant elements (i.e., application and user or group). This can have a significant simplifying effect. Firewall policy based on port and IP address, followed by subsequent analysis to understand the application makes things more complicated than they are today.

10

Your next firewall must deliver the same throughput and performance with application control fully activated

Business case: Many enterprises struggle with the forced compromise between performance and security. All too often, turning up security features in the network security realm means turning down throughput and performance. If your next-generation firewall is built the right way, this compromise is unnecessary.

Requirements: The importance of architecture is obvious here too – in a different way. Cobbling together a port-based firewall and other security functions from different technology origins usually means there are redundant networking layers, scanning engines and policies – which translates to poor performance. From a software perspective, the firewall must be designed to do this from the beginning. Furthermore, given the requirement for computationally intensive tasks (e.g., application identification) performed on high traffic volumes and with the low tolerance for latency associated with critical infrastructure, your next firewall should have hardware designed for the task as well – meaning dedicated, specific processing for networking, security (including SSL termination – see #3), and content scanning.

Conclusion: Your Next Firewall Should Safely Enable Applications – and Business

Users continue to adopt new applications and technologies – and the threats carried by them. In some organizations, obstructing the adoption of new technologies can be a career-limiting move. Even when it isn't, applications are how employees get their jobs done, or maintain productivity in the face of competing personal and professional priorities. Because of this, safe enablement is increasingly the correct policy stance. But to safely enable these applications and technologies, and the business that rides atop them, network security teams need to put in place the appropriate policies governing use, but also controls capable of enforcing them. The ten things described here are critical capabilities for putting the necessary controls in place – especially in the face of a more varied and rich application and threat landscape. Without the network security infrastructure to cope with that variety and depth, security teams cannot safely enable the necessary applications and manage risk for their enterprises.