



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Real-World Testing of Next-Generation Firewalls

Copyright SANS Institute
Author Retains Full Rights



Sponsored by Fortinet and Ixia

Real-World Testing of Next-Generation Firewalls

October 2013

A SANS Analyst Whitepaper

Written by Dr. Eric Cole

Today's Threat Landscape and Countermeasures *PAGE 2*

Verification and Test Design *PAGE 5*

Test Execution and Results *PAGE 10*

Introduction

Discussions of IT security that formerly took place in the server room have moved to the boardroom. Having invested significant amounts of money in securing their systems, board members and their IT departments see little return on their investment.

Traditionally, cyber attacks were visible, opportunistic and focused on the low hanging fruit; examples of such attacks include Code Red, Nachi and Conficker. Early firewalls addressed the middle layers of the networking stack and blocked IP addresses and protocol ports. Today's threats are much more stealthy, targeted and focused, preying on applications and their data.

Defending against these more complex attack methodologies requires a new generation of firewall that can also mitigate application-based attacks. As adversaries move up the protocol stack, security devices need to follow suit to be effective. By being aware of applications and their data, analysts—and their managers—can carefully examine the traffic and increase overall security.

This is the definition of a *next-generation firewall* or *NGFW*. Although many NGFWs claim they can properly inspect the traffic, not all offer equal protection, meaning IT security staff must carefully test potential solutions before recommending a purchase. The firewalls are not equally able to scale, integrate and detect advancement in attack methodologies.

This paper offers advice on what to expect from a next-generation firewall, features and business needs to consider, and a test methodology for IT and business professionals to use to enhance their investments in security through enhanced firewall capabilities.

Today's Threat Landscape and Countermeasures

Today's advanced threats are evading traditional firewalls with stealth, specificity and persistence. They are now, more than ever:

- **Stealthy.** They hide from traditional firewalls in legitimate traffic and ports, leaving no visible signs or traces. Once a system is compromised, hostile traffic is meant to blend in with normal network traffic—and can even turn off security monitoring if need be.
- **Targeted.** Specific individuals and systems are exploited rather than an entire network.
- **Data-focused.** In many cases victims do not know where their critical data is, making it easy for an adversary to succeed and hard for an organization to resist. Focusing on data flows is critical to proper security.
- **Persistent.** Today's adversaries do not merely want to compromise a system; they want persistent access. Multiple reports indicate that advanced attackers have been behind firewalls for months without detection.

What Next-Generation Firewalls Need

NGFWs must be able to analyze and inspect traffic, including the following data types:

- **Encrypted traffic.** NGFWs must be able to see into SSH or TLS/SSL traffic and not assume such traffic is authorized.
- **Packet headers.** NGFWs must look at more than just Layer 3 and Layer 4 headers to determine whether traffic should be allowed or blocked. They must be more application aware and not just assume that a particular TCP or UDP port is from a "known-good" application. A skilled adversary, having determined what outbound traffic an organization allows—typically TCP ports 25 (SMTP) and 80 (HTTP)—can route command-and-control traffic over any such port. So, these channels must be protected.
- **Use of nonstandard ports or "port hopping."** NGFWs must detect when nonstandard ports are being used because attacks often use random port hopping as a means to get beyond traditional firewalls.

Today's Threat Landscape and Countermeasures (CONTINUED)

A first-generation firewall's means of looking for known threat vectors effectively detects traditional attacks, but is less effective against advanced persistent threats (APTs), as illustrated in Figure 1.




Figure 1. Traditional defenses are ineffective against advanced threats.

New Countermeasures Required

Adversaries carefully examine how their targets secure their networks and have found many ways to bypass and defeat common security technologies, as outlined in Table 1.

Technology	How Advanced Threats Bypass	How NGFWs Protect
Antivirus	<ul style="list-style-type: none"> • APTs work around today's antivirus and antimalware on endpoints simply by changing their behaviors or hiding their activities in trusted systems and processes. 	<ul style="list-style-type: none"> • NGFWs can analyze the application traffic and potentially spot anomalies that are indicative of malicious activity. Traditional firewalls are limited in this area because they only filter on IP addresses and ports. With application awareness, an NGFW has additional details, leading to better-informed decisions, such as detecting malicious applications tunneling inside legitimate applications.
Legacy Firewalls	<ul style="list-style-type: none"> • APTs conceal malicious activity as ordinary HTTP traffic or encrypt their data. Legacy firewalls cannot penetrate such disguises. 	<ul style="list-style-type: none"> • NGFWs can identify the application-specific content within HTTP traffic and inspect encrypted traffic for malware by decrypting the packet stream.
Network Security Devices	<ul style="list-style-type: none"> • APTs planted internally open holes through firewall and network security. • With access to user accounts, attackers bypass internal network access controls. 	<ul style="list-style-type: none"> • Most of the holes that are created occur when command and control (C&C) channels mimic certain applications by using a well-known port. This activity can be detected by filtering on application data. Some NGFWs can also detect outbound C&C communications protocols used by botnets. The NGFW should receive regular updates to the database of botnet behaviors and blocked C&C servers.

Table 1. Advanced Threat Patterns

Today's Threat Landscape and Countermeasures (CONTINUED)

Advanced attack vectors, therefore, require a new approach to network security—one that is designed to thwart complex and customized threats.

As NGFWs evolved, they naturally moved up the stack. In addition to being “Layer 7 aware” they became able to inspect the content of the packet, searching for indicators of an attack within the data stream. This is typically where intrusion detection and prevention systems operate, so this represents an interesting convergence of previously dissimilar roles.

Verification and Test Design

Because needs vary from one organization to the next, a solution that fits perfectly in one environment might not suit another. This is one of the core reasons why testing an NGFW is so critical; the solution must meet an organization's specific needs. When preparing for testing and ensuring that the firewalls performed as intended, testers should work from a list of mandatory functionality for an NGFW.

Essential Capabilities

The core features offered by NGFWs that should be tested and verified for specific abilities include the following:

- **Performance.** Even though additional analysis is being performed, it is critical that the new inline NGFW devices (or services) do not become a bottleneck. Not only should the device be able to handle the current traffic on a network, but it should also account for the next several years of growth, ensuring that the device can properly scale to meet anticipated future needs. This includes scaling to various types of traffic and protocols, including encrypted traffic.
- **Application visibility.** Check the ability of the device under test (DUT) to verify the application through examination of Layer 7 headers instead of just looking at Layer 4 ports. Use various encoding methods (e.g., binary, EBCDIC and Unicode) in addition to ASCII to ensure that the device cannot be tricked into mislabeling traffic.
- **Content filtering.** Examine embedded content, making sure it aligns with the application data that the packet ostensibly contains. Use packet correlation and heuristics to determine the DUT's level of examination.
- **Backdoor detection.** Backdoors and command channels are common artifacts that attackers leave on targeted systems. Although meant to be covert, these "stay-behind" elements have distinct signatures that can be used to detect and block them. The most obvious of these is a fully encrypted outbound channel that the organization cannot decrypt using any known authorized keys. Other indicators of a backdoor or command channel include unusual TCP ports, traffic spikes at odd hours and unusual traffic destinations, such as the 175.45.176.0 to 175.45.179.255 range of IP addresses.
- **Visibility into encrypted packets.** Unauthorized encryption can enable an adversary to silently sneak past most security devices. Not only should evaluators gauge the DUT's ability to decrypt traffic, but in cases where it cannot decrypt the traffic, it must be able to block or shunt such traffic to a secure area and send an alert to the IT response team. The DUT needs to be able to do this without false positives and at line speed.
- **Integration with other functionality.** Some NGFWs incorporate functionality that brings them under the category of unified threat management (UTM), which combines IDS, firewall and other filtering techniques at one location. For small- to medium-sized networks, having a single device that can perform other roles can be of value, but be sure to thoroughly test the pros and cons. Test the firewall functionality of such a device and verify that other features of the DUT are not compromised.

- **Interoperation and integration with other devices.** When adding a NGFW to an existing environment, there are other security and nonsecurity devices with which it must cooperate. Make sure that the DUT does not conflict with or break the functionality of existing devices and can, in fact, work with them in making assumptions about good or bad traffic.
- **Advanced malware monitoring and intrusion protection.** Many NGFWs examine the entire packet and can look for indicators of both attempted attacks and successful ones. This is done through a wide range of techniques—from signature analysis to anomaly detection and heuristics, for example. These features should be fully tested to determine the DUT’s ability to detect malware samples thrown at them.
- **Inbound and outbound analysis.** Because adversaries can sneak into a network undetected and send commands and data outside the network, devices should perform both inbound and outbound analysis. Test inbound analysis to verify attacks are stopped and outbound analysis to verify a compromised system can be blocked.
- **User awareness.** In addition to monitoring and tracking applications, some NGFWs can also identify and track individual user activity. This enables the systems—and the IT teams managing them—to build user profiles and track anomalies associated with profiled users. This can also apply to devices. If a user is connecting to the network from a device that is not considered secure, additional security measures should be applied.
- **Self-learning.** The DUT should demonstrate the ability to learn traffic patterns, new methods of inspection and new methods of attacks that are detectable.
- **Prevention and detection.** The DUT must not only prevent attacks, but also detect them in a timely manner and without the same issue of false positives that has held administrators back from turning on the “prevention” part of their older IPS systems. Verify the device can stop known and unknown attack types.

Alignment with Needs

A determination of which NGFW is best for an organization must begin with three questions:

- What are the risks we’re protecting against?
- What are the priority risks?
- Is the purchase the most cost-effective way of reducing the risk?

Eventually, risk determines all security decisions. Organizations that take a risk-based approach succeed at IT security; those that ignore risk fail. When analyzing risk, always begin with the assets protected and their vulnerabilities to determine their relationship to the corresponding threat as illustrated in Figure 2.

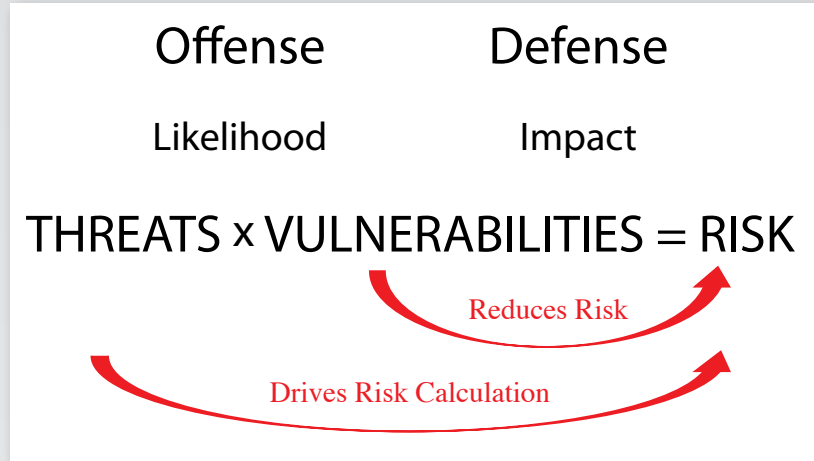


Figure 2. Core Components of Risk Calculation

Evaluators should draw up a checklist or requirements document based on the threats to an organization's critical assets and the known vulnerabilities. This document should be broken into two categories: requirements and options. "Requirements" are items that the solution must have; whereas "options" further distinguish the purchase candidates based on additional functionality that would enhance the requirements. Typically, the candidate with the most optional features to manage an organization's particular risk areas is the best choice, assuming that all are equally priced and features are integrated appropriately.

Prepare for Testing

The first step in testing an NGFW is to have an accurate, up-to-date network diagram of your devices, user accounts and their applications. This diagram will drive the testbed, and it is critical that the environment in which the device is tested be as accurate a representation as possible in terms of the infrastructure and the traffic on the production network.

Use existing network diagrams to identify holes in security that need to be closed and where an NGFW will do the job. The security holes the NGFW is meant to close must align with the internal and external threats, which are tied to the organization's risks. One example of this is dealing with C&C channels. Many adversaries sneak out of a network by tunneling IRC traffic over HTTP or sending unauthorized cryptography over SSL. Carefully examining traffic with an NGFW can minimize such threats.

Verification and Test Design (CONTINUED)

Establish the important criteria that are required from the NGFW. You'll use this information to create a scorecard for comparing the tested performance of the competing devices. For example, a scorecard similar to the example in Table 2 would help create quantifiable performance data for evaluating three vendors' equipment:

Metric	Goal	Vendor A	Vendor B	Vendor C
Transactions	10,000			
Concurrent Flows	30,000			
Average Latency (in microseconds)	5,000			
Attacks Blocked	90 percent			

Table 2. Scorecard for Rating Vendors

Although the most realistic testbed is the production environment, that is too dangerous for initial testing—especially because NGFWs are inline devices, and there is always concern that inline inspection devices could degrade the performance of the network. Eventually, you will use the production environment as the “acid test” of the best candidate, but this should only be done when you have a high level of confidence that the DUT works as advertised.

Create a Testbed

Once your organization has determined your needs, risks, assets and weaknesses, you should begin by testing an environment that replicates the network's key features, create representative traffic samples against the DUT, and observe the system over time.

Although traffic samples from the production environment can be useful in establishing a baseline of normal behavior, they lack the randomness and complexity of a live simulated network. The best option is to create a testbed that is able to simulate the intended environment at load. Traffic-generation solutions should have the ability to simulate a realistic mix of traffic at line rate. They must also simulate attacks that include continual updates of the latest threats as well as threats launched throughout history.

Consider Costs

The most expensive device is not always the best one. Having narrowed the choices to a few options (three to five candidates is an optimal field of contenders), an organization must compare the costs and features of each. This cost estimate should include:

- Variables in licensing
- Required manpower and training
- Support and maintenance costs
- Costs for expanding coverage as needed in the future

With this analysis in hand, the organization can determine if the additional functionality is worth the additional price.

Consider Payback

Ultimately, over time, the product should show a reduction in administrative costs for overhead and updating of firewall signatures and other maintenance, as well as a direct reduction in successful attacks. The other returns on investment would be the ability to use information gathered by the firewall to develop better security practices and the closing of vulnerabilities that are related.

Test Execution and Results

NGFWs are especially challenging to test because they are content aware and responsive. Each network environment is unique, making any general datasheet parameter irrelevant to the target deployment.

The following baseline tests drive the overall test plan:

- **Performance.** During the requirements-generation process, an organization must determine what types of applications need to be supported and their respective throughput requirements. This information should drive the performance section of the testing. Creating realism in the simulated traffic to match the intended protocols and applications that the NGFW will experience is important and will return the most accurate information. Some of the basics in this area are:
 - **Simple performance test.** This includes basic testing and analysis of the device to confirm that it works. During this phase, bandwidth is also tested to fully understand the performance impact a device will have in the current environment.
 - **Application protocol test.** Networks no longer just move packets; they are connection points for applications and business process. All applications need to be tested through the device, making sure that all applications that are supposed to work are properly functioning at expected average expected full load.
 - **Baseline attack test.** While the performance and protocol tests focus on testing the positive, this phase focuses on testing the negative, making sure that attacks and malicious activity are properly detected and blocked. Although zero-day attacks won't have a signature, sandboxing technology can detect their use.
 - **Infinite loop test.** Knowing the limitations of a device and when it is no longer going to work in a specific environment is critical for proactive planning. This final performance test is used to push the limits of the device and determine when it no longer works given a certain set of traffic.
- **Functionality.** Organizations should have a list of the required and optional functions the device needs to perform based on the organization's risk analysis and gap assessment. Each organization should develop a rating scheme to prioritize those items, but no organization should purchase a device unless it meets all of the mandatory requirements. The following are some of the key types of functionality testing:
 - **Positive testing versus negative testing.** Positive testing confirms that applications are working as intended; for example, testing email and web surfing to confirm that they are still functioning correctly. Conversely, negative testing verifies that the forbidden activity is actually blocked; for example, setting up a C&C channel and establishing whether the device actually blocks it.
 - **Understanding types of traffic.** The ability to distinguish among the range of applications in the enterprise network and to flag traffic that doesn't meet the expected type definition—including protocol port, source/destination and packet size—is paramount.
 - **Understanding attack vectors.** This implies the ability to incorporate newly discovered methods of attack into the device's inspection efforts in addition to those that are defined out-of-the-box.

- **Integration.** No single device is going to make an organization secure. Defense in depth is a critical component of effective security—any device an organization purchases needs to work with other devices in the environment. The key questions that need to be answered are:
 - Does the device work with other products from the vendor?
 - Does the device work with products from other vendors?
 - Does the device integrate into current system management schemes?
 - Do all the organization's monitoring tools support the devices?

Guidelines to Follow

Testing guidelines reflect the vulnerabilities or high priority risks that the organization is trying to reduce to an acceptable level. The testing guidelines will ultimately serve as the decision point for the purchase. Professionals should use the requirements document as the foundation for the testing guidelines, which should be granular enough for appropriate testing. These guidelines typically include:

- **Functional testing**

- Ease of installation
- Ease of use
- Usability of management console or interface
- Functionality of logging and reporting
- Ability to detect attacks
- Ability to respond to attacks
- High availability, throughput and accuracy
- Integration and operation with other products
- Feature and functional parity for both IPv4 and IPv6 traffic

- **Informational evaluation**

- Company stability
- Price
- Documentation and training
- Support

Questions to Ask

The testing guidelines can also be in question form. These sample questions gather the required information for devices under review:

1. Does it inspect application and data or only ports and protocols?
2. Does it analyze, detect and prevent encrypted traffic?
3. Does it decode and analyze obfuscated traffic?
4. Can it perform code emulation to detect unknown threats?
5. Can it do these types of analyses without impairing good traffic?
6. Does it integrate and work with devices from other vendors?
7. Does it detect and/or prevent the leakage of sensitive data?
8. Does it detect or analyze compliance-related information?
9. Does it accommodate customer-defined rules and signatures?
10. Does it protect the application infrastructure and users?
11. Is it effective against zero-day attacks or other attacks lacking a signature?
12. Does it adapt policies for dynamic applications, users and devices?
13. Does it include user-based enforcement, as well as device and application enforcement?

The remaining steps of narrowing the field, picking the winner and trying before buying are generic enough to require no further discussion.

Summary

In today's enterprises, multiple vulnerability monitoring and detection, application monitoring and firewall tools are used to manage network threats. Unfortunately, attackers are still succeeding and maintaining beachheads in organizations that don't even know they've been infiltrated until a third party tells them so. This indicates that today's mixed bag of tools for detection, monitoring and firewalling are not enough.

NGFWs grew out of necessity, combining features of multiple tools to give them better visibility and accuracy for detection and prevention of malware and attacks. Many attacks succeed in the absence of detailed traffic analysis; however, by combining such analysis with Layer 7 inspection, firewalls become much more effective at protection by thoroughly examining inbound and outbound traffic and performing decryption functions when required.

Although NGFWs can increase an organization's network security and decrease the associated risks, all devices are not up to the task, and proper testing is required to ensure the appropriateness of a particular device. This starts with knowing your systems, their usage and the risks associated with your environment. Armed with this information, organizations need to compare their baselines against the tools they already have in place to identify weaknesses and create a next-generation system that can meet today's demands for more integrated and comprehensive network protections.

About the Author

Dr. Eric Cole is a SANS Faculty Fellow and an industry-recognized security expert with over 20 years of hands-on experience in information technology, focusing on helping customers build dynamic defense solutions that protect organizations from advanced threats. Dr. Cole has a master's degree in computer science from the New York Institute of Technology and a doctorate from Pace University with a concentration in information security. Dr. Cole is the author of several books, including *Advanced Persistent Threat*, *Hackers Beware*, *Hiding in Plain Site*, *Network Security Bible*, and *Insider Threat*. He is credited on more than 20 patents; is a researcher, writer and speaker; and is also a member of the President's Commission on Cyber Security and several executive advisory boards. Dr. Cole is founder of Secure Anchor Consulting and has served as CTO of McAfee and chief scientist for Lockheed Martin. He is actively involved with the SANS Technology Institute (STI) and SANS working with students, teaching, and maintaining and developing courseware.

SANS would like to thank its sponsors:





Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS Chicago 2017	Chicago, ILUS	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VAUS	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS San Francisco Fall 2017	San Francisco, CAUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Tampa - Clearwater 2017	Clearwater, FLUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Network Security 2017	Las Vegas, NVUS	Sep 10, 2017 - Sep 17, 2017	Live Event
SANS Dublin 2017	Dublin, IE	Sep 11, 2017 - Sep 16, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MDUS	Sep 25, 2017 - Sep 30, 2017	Live Event
Data Breach Summit & Training	Chicago, ILUS	Sep 25, 2017 - Oct 02, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, DK	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS London September 2017	London, GB	Sep 25, 2017 - Sep 30, 2017	Live Event
Rocky Mountain Fall 2017	Denver, COUS	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS SEC504 at Cyber Security Week 2017	The Hague, NL	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS DFIR Prague 2017	Prague, CZ	Oct 02, 2017 - Oct 08, 2017	Live Event
SANS Oslo Autumn 2017	Oslo, NO	Oct 02, 2017 - Oct 07, 2017	Live Event
SANS October Singapore 2017	Singapore, SG	Oct 09, 2017 - Oct 28, 2017	Live Event
SANS AUD507 (GSNA) @ Canberra 2017	Canberra, AU	Oct 09, 2017 - Oct 14, 2017	Live Event
SANS Phoenix-Mesa 2017	Mesa, AZUS	Oct 09, 2017 - Oct 14, 2017	Live Event
Secure DevOps Summit & Training	Denver, COUS	Oct 10, 2017 - Oct 17, 2017	Live Event
SANS Tysons Corner Fall 2017	McLean, VAUS	Oct 14, 2017 - Oct 21, 2017	Live Event
SANS Brussels Autumn 2017	Brussels, BE	Oct 16, 2017 - Oct 21, 2017	Live Event
SANS Tokyo Autumn 2017	Tokyo, JP	Oct 16, 2017 - Oct 28, 2017	Live Event
SANS Berlin 2017	Berlin, DE	Oct 23, 2017 - Oct 28, 2017	Live Event
SANS Seattle 2017	Seattle, WAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS San Diego 2017	San Diego, CAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Gulf Region 2017	Dubai, AE	Nov 04, 2017 - Nov 16, 2017	Live Event
SANS Miami 2017	Miami, FLUS	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Amsterdam 2017	Amsterdam, NL	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Milan November 2017	Milan, IT	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Sydney 2017	Sydney, AU	Nov 13, 2017 - Nov 25, 2017	Live Event
Pen Test Hackfest Summit & Training 2017	Bethesda, MDUS	Nov 13, 2017 - Nov 20, 2017	Live Event
SANS Paris November 2017	Paris, FR	Nov 13, 2017 - Nov 18, 2017	Live Event
SANS Adelaide 2017	OnlineAU	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced