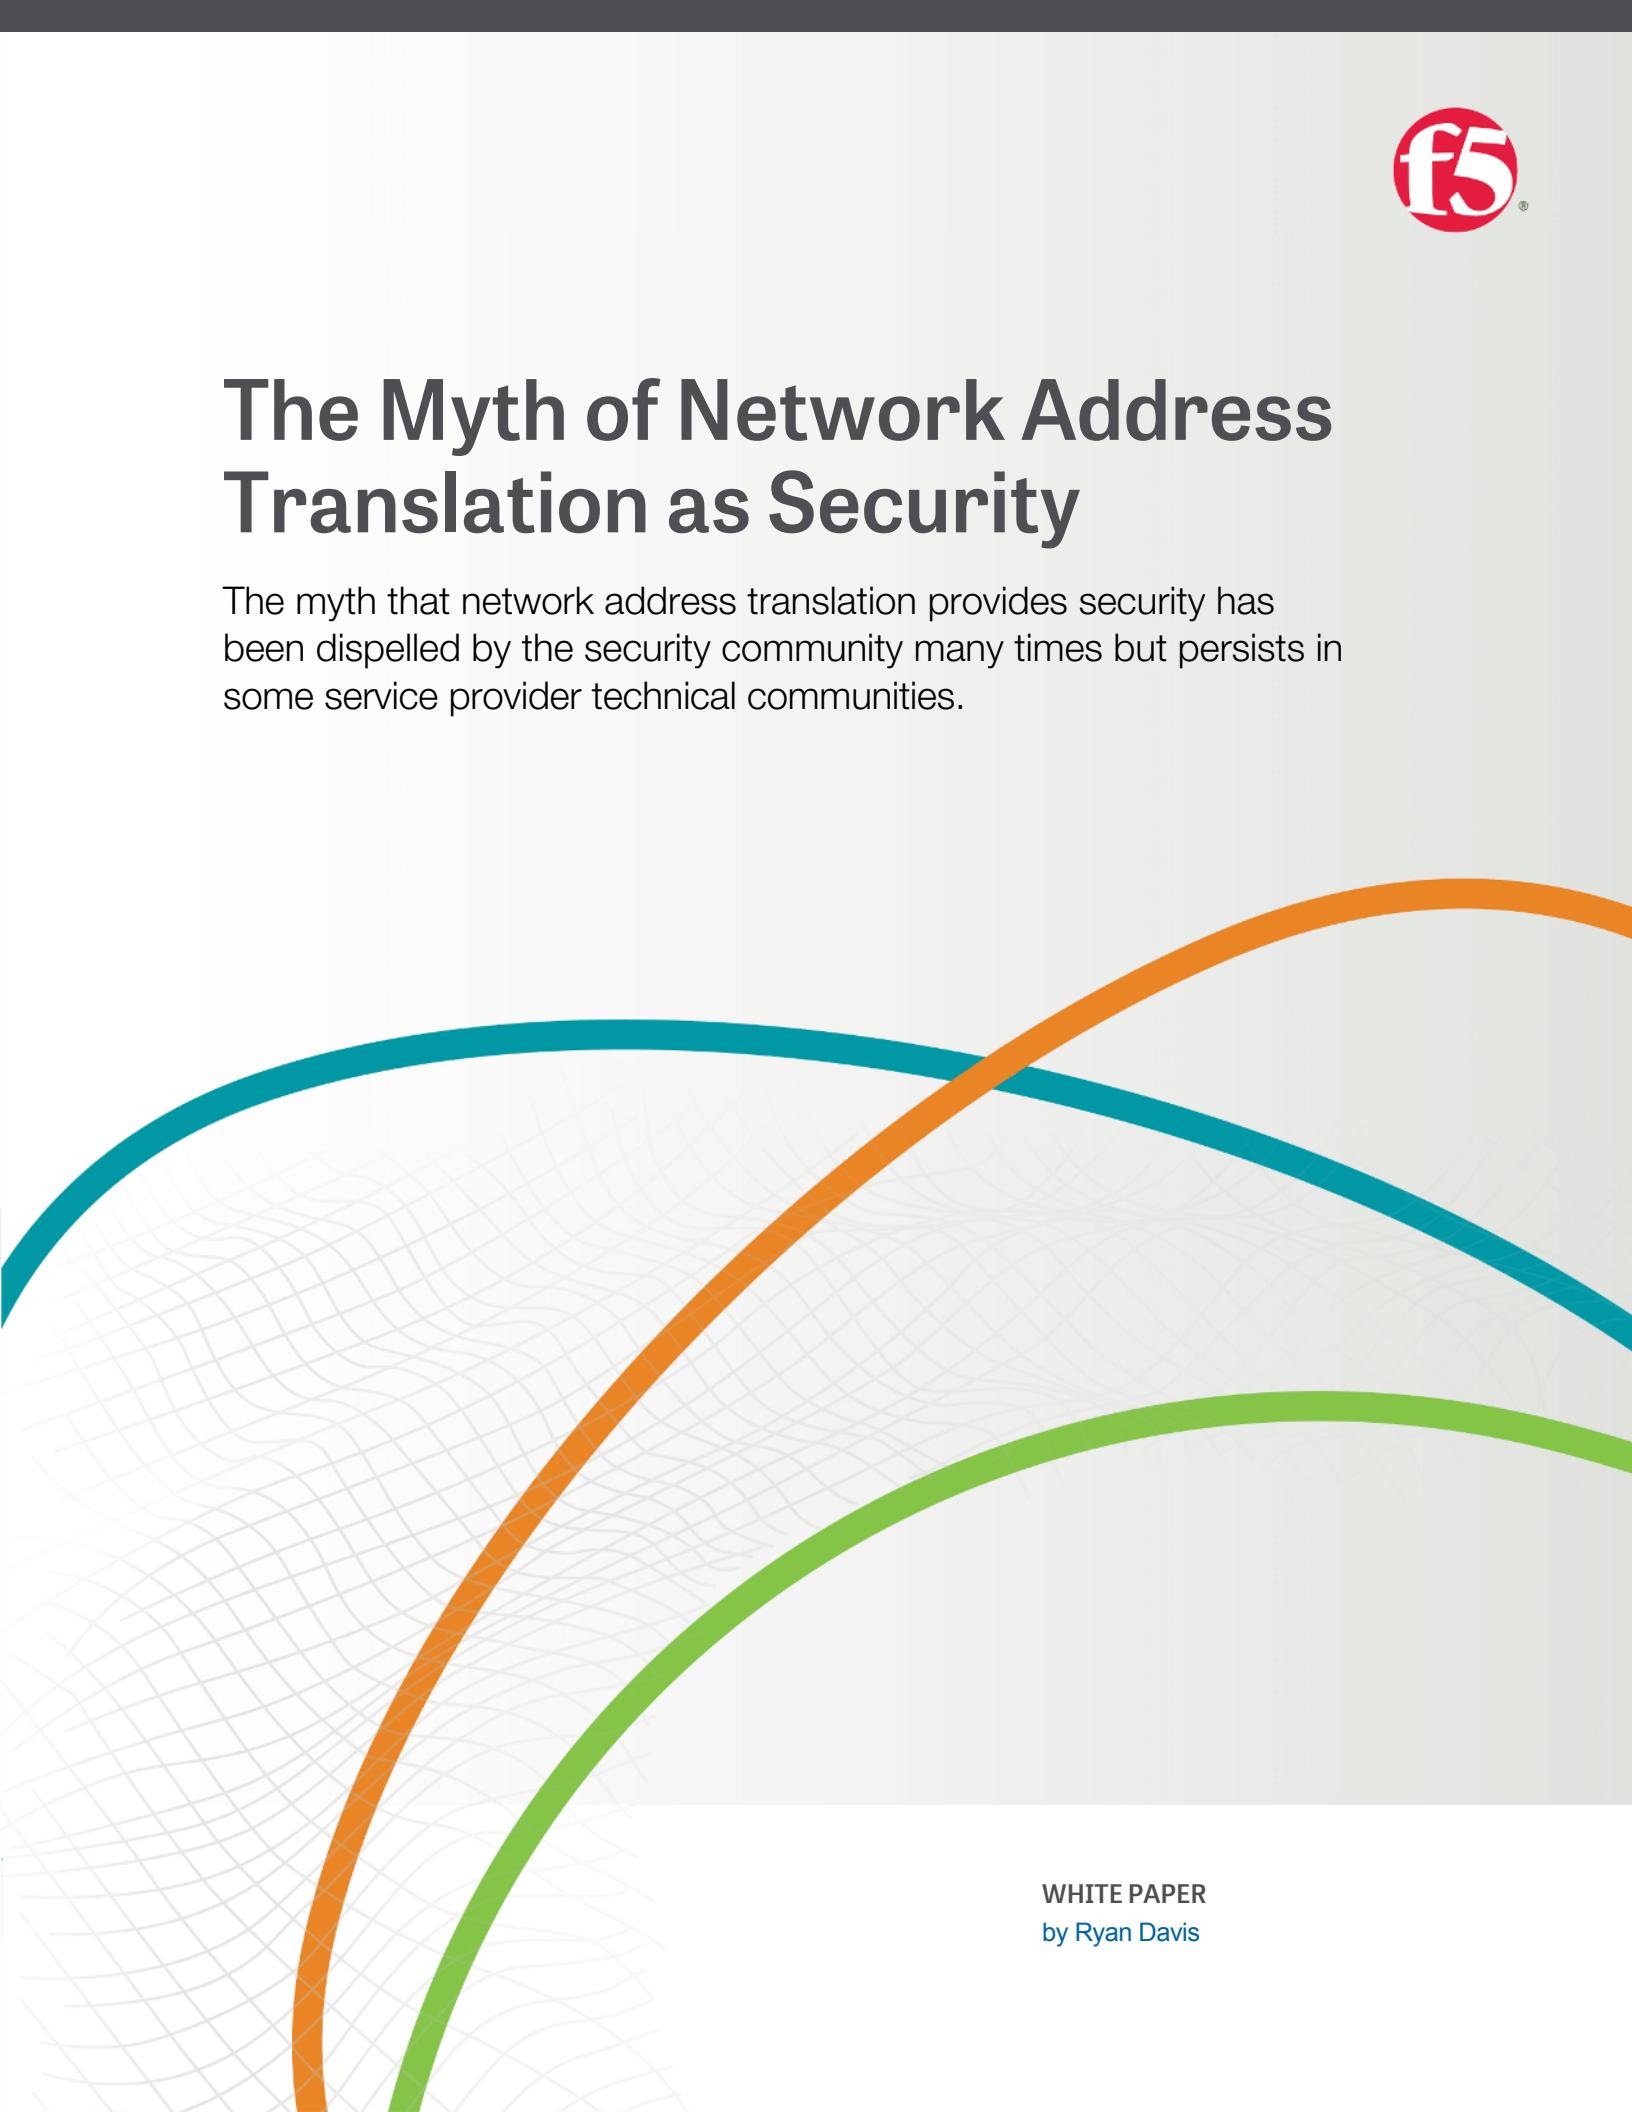

The Myth of Network Address Translation as Security

The myth that network address translation provides security has been dispelled by the security community many times but persists in some service provider technical communities.

WHITE PAPER

by Ryan Davis

Introduction

In service provider networks, the largest use of network address translation (NAT) tends to be at the point of the subscriber Internet edge, but unfortunately this point is also the largest attack surface, carrying the greatest threats, within the service provider network. In mobile networks, this footprint is called by various terms, including Gi LAN, SGi LAN, and mobile edge. In more general terms, it is the location where pure Internet connectivity meets a gateway that manages a specific access technology (such as wireless, cable, or fiber). Gateways are excellent at managing subscriber connectivity on a specific access network, but they are not well suited for applying security controls or address translation due to limited security functionality or excessive cost.

Service providers in the past have labored under the mistaken assumption that NAT can provide both address translation and security at the subscriber Internet edge. The security community has tried to dispel that myth, but it persists, and mobile and fixed service providers in today's environment of escalating attacks need to understand why NAT is insufficient. This understanding starts with the recognition that there are many different types of security services at the point of subscriber aggregation in addition to denial of inbound traffic.

Subscriber aggregation security

A common architectural solution consists of a stateful system between the gateway and the Internet that can provide various services to subscribers and the network. Common types of security services include:

- DDoS protection. Typically protecting against high-scale attacks from the Internet, distributed denial-of-service (DDoS) attack mitigation consists of eliminating packets that serve no legitimate purpose and limiting packets that deviate from typical behavioral patterns.
- Port and protocol limiting. Over a decade ago, service providers would provide wide-open access to and from all subscribers. While this did facilitate connectivity for many applications, it also provided a wonderful environment for the spread of viruses and worms, which would excessively consume network resources. Today, most service providers do prevent some protocols on their networks, as well as limiting use of others in an effort to reduce the risk of enabling malware to flourish.
- Infrastructure protection. The beauty of IP connectivity is that it enables anyone with an IP address to potentially communicate with any other IP address on the planet. However, this also enables subscribers to potentially connect to systems on the service provider's network without any legitimate reason. Implementing controls to protect infrastructure at the subscriber aggregation point is quite important.
- Botnet mitigation. At any point in time, many subscribers have systems infected with viruses and malware that are communicating to command and control (C&C) systems somewhere on the Internet. Restricting this communication provides benefits to the subscriber as well as to the network by reducing the use of subscriber devices in botnet attacks.
- Types of service enforcement. Many service providers have various service offerings that enable subscribers to perform different types of connectivity. For example, enterprise customers may require static IP addresses with unlimited protocol usage and the ability to host servers. However, a lower tier of consumer service may restrict the running of servers on mobile devices. Or perhaps a specialized type of IP address class could be used for administrative connectivity to update firmware and software on the CPE. When mobile service providers deploy voice over LTE (VoLTE), they will use a dedicated APN that must have limited connectivity—or the risk of revenue leakage increases significantly.

These five examples of security controls are not exhaustive, but they do represent a base set of functionality that is required on modern networks at the subscriber aggregation point. With this in mind, NAT solutions do not provide sufficient functionality to secure and protect modern networks. (Note that if a company offers a carrier-grade NAT (CGNAT) solution that performs any of the above functions, it is an admission that NAT alone is insufficient for protecting a service provider's most valuable product—the ability to connect people and things.) The proper solution for securing the subscriber Internet edge through services such as those above is not NAT but a carrier-class network firewall such as F5® BIG-IP® Advanced Firewall Manager™ (AFM).

Technical Arguments

Having established that there are many other security threats that NAT can do nothing to protect against, let's dispel the technical and business myths of using NAT as a security service, even for minimal denial of unsolicited traffic.

Stateful and stateless ingress and egress: IPv6

While created to facilitate communication between hosts in overlapping private address space, NAT has seen wide deployment for the main purpose of providing connectivity to more hosts in a network than individual IP addresses are available. Security issues associated with NAT have long been documented. (See the Security Considerations sections of the Internet Engineering Task Force's (IETF's) RFC2663 and RFC2993). These concerns were a large driver for the creation of IPv6. Now that the world has been moving to IPv6 for some time, NAT will no longer be required for hosts that use IPv6. Therefore it will be unavailable to provide any protection to IPv6 hosts.

An argument could be made that a stateful NAT46 or NAT64 gateway may provide some security. However, this security would only be partial and short lived. CGNAT is mainly designed to translate traffic between hosts when one side has not converted to IPv6, which means that the ultimate goal is to remove the CGNAT device and allow native IPv6 traffic without translation. Without a carrier-class network firewall, the NAT device becomes the obstacle to full IPv6 migration for the operator.

Further, many service provider networks are now providing IPv4/IPv6 dual-stack configurations. In this case, even if the service provider has NAT44 enabled, the IPv6 interface does not, meaning that the organization has decided that NAT suffices for security on the IPv4 interfaces, but the same host remains completely open on the IPv6 interface—which is logically incongruent. Why would a network architect believe that protection is required in IPv4 and not IPv6 when the exact same threats apply to both?

Carrier-class network firewalls must be positioned as solutions today, as operators will not accept lack of security as the obstacle to IPv6 conversion. Deploying a CGNAT point product today will force the operator in the future to repurchase equipment and software as well as to reengineer the network, leading to much higher costs, as change in a service provider network is a difficult process. By contrast, implementing CGNAT functionality on a carrier-class network firewall will ensure a smooth transition to IPv6 in terms of security feature parity and capacity.

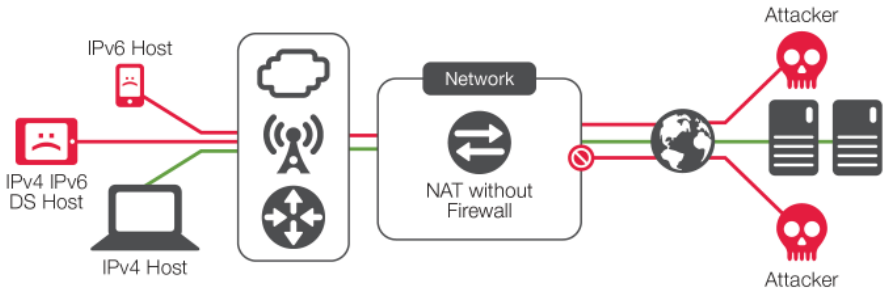

Figure 1: A CGNAT point product may serve to provide limited security for IPv4 hosts, but leaves IPv6 hosts completely unprotected.

Stateless ingress and egress: IPv4

Stateless NAT in general is a rare use case, but it may be employed in some unique situations. Stateless devices allow all traffic, regardless of the request from the subscriber device, to ports that are allowed for reverse NAT. Since they are stateless, they do not know what traffic was sent by the host that required NAT; therefore they must allow all return traffic. It should be apparent that stateless NAT may not be suitable even for most NAT use cases, and is completely unsuitable for any security use cases.

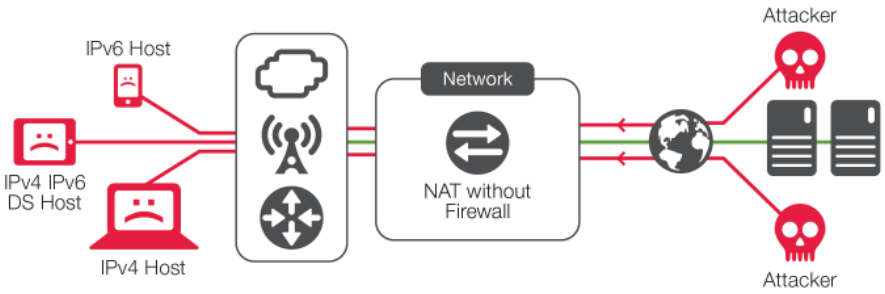

Figure 2: A stateless NAT solution, which must allow all return traffic, is completely unsuitable for providing security.

Stateful egress: IPv4

Stateful egress IPv4 NAT gives an operator the ability to partially protect internal hosts from externally initiated traffic in most cases. However, it does not provide protection for internal hosts, nor does it enable the possibility of response to attacks from those internal hosts to other network resources connected to the NAT device. In a large number of cases, many interior hosts that will be compromised may access the Internet for botnet command and control. In a number of rare but most severe cases, these compromised hosts are used as launching pads for advanced attacks on the internal service provider network or externally directed attacks toward higher-level corporate or government targets that will draw unwanted publicity. The costs associated with cleaning up after these types of incidents far outweigh the cost of the solution.

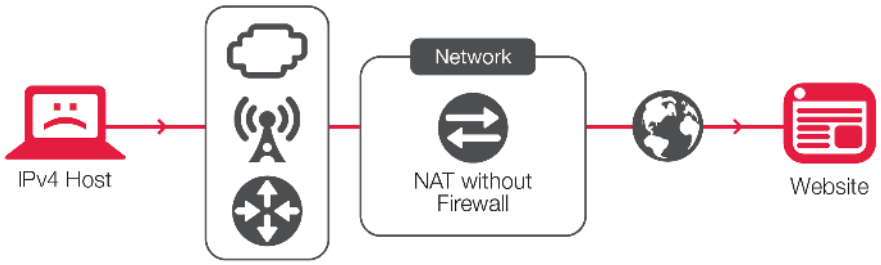

Figure 3: Stateful NAT provides no protection for willing or unwilling internal hosts, which may by used by botnets to attack the network or external targets.

Stateful ingress: IPv4

Stateful ingress IPv4 traffic is the only traffic type where anyone could reasonably argue that security protection is provided by NAT, and then only under certain conditions. This is a myth that persists, but as the above explanations make clear, that IPv4 traffic is only a fraction of the attack surface presented by networked hosts. Furthermore, stateful NAT does not provide much protection even for IPv4 ingress, given that modern attack techniques assume there will be a NAT device in the path, one that must be subverted. Static and destination translation of hosts provides no security.

Some of NAT's supposed security relies on obfuscation, which is not considered by the security community to be a real solution. Obfuscation only makes it more challenging to find information that can be gained in other ways, so it prevents nothing. The other component of the stateful ingress NAT security myth is that it is thought to provide a "one-way street"—however, it really does not. While it is true that stateful ingress IPv4 NAT will reject externally initiated TCP traffic, that does not mean that an external host cannot in certain situations send traffic to internal hosts or use other methods to circumvent the NAT. In fact, most network-based attacks assume this as a requirement of the compromise.

There are several ways to accomplish this circumvention, all of which can be prevented by a firewall. First, an attacker can either use a targeted or a sweep attack to send traffic to ports that are open in the NAT device's state table. The purpose of this attack could be to create a denial-of-service (DoS) by invalidating an existing session on the host or NAT state table, to footprint an internal network, or to inject a malware payload into a third party's existing session in an effort to compromise the internal host. Serious implications are seen in UDP traffic that is by design stateless; however, the same could be accomplished (given host susceptibility) in TCP or other protocols. In addition, NAT may not provide protocol conformance, sequence number checking, or any other layer 2 or layer 3 DoS security measures that firewalls or advanced security devices inherently provide. NAT also provides no tools to respond should security breaches occur.

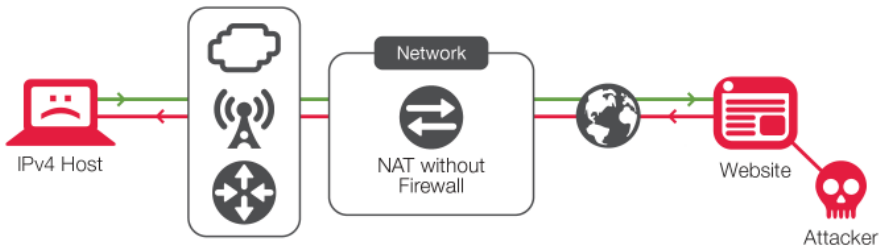

Figure 4: Even for IPv4 hosts, modern attacks frequently can compromise a NAT device in the path.

Business arguments

The business arguments for positioning a carrier-class network firewall are simple: First, no service provider can afford today's damaging and sometimes high profile security compromises. That's why firewall and network address translation services frequently come bundled together. The monetary damage that can occur on a modern mobile network can easily exceed millions of dollars, with some attacks potentially crippling entire brands.

Second, in the case where only a NAT device is present, the service provider has no tools to respond to the attack and must helplessly endure it until ad-hoc solutions are found.

Finally, savvy service providers using an advanced firewall device can add additional security services to their customer offerings. Traditionally such additions attract enterprise business customers who have a clear case for protecting their business assets. Without a firewall device with these capabilities, a service provider will never have the opportunity to earn that revenue. In fact, by not showcasing a combined NAT and firewall solution with advanced security features, the service provider enables its customers to assume there is a gap in the service provider's product line or expertise.

Conclusion

The myth that NAT provides any significant security in light of today's sophisticated attacks needs to be put to rest. From a technical viewpoint, in fact, NAT provides:

- No security to IPv6 hosts, as NAT is unnecessary for them.
- No security for stateless NAT hosts.
- No security for stateful NAT host outbound attacks.
- Minimal protection for stateful NAT host ingress attacks, since modern attacks assume the presence of a NAT device and readily compromise or circumvent those devices.
- No tools for responding to security attacks that routinely occur.

In business terms, neglecting to deploy a carrier-class network firewall such as BIG-IP AFM as part of an edge NAT and security solution risks both severe and pernicious revenue leakage and shows a lack of innovation by a service provider.

By contrast, service providers who deploy an appropriate and feature-rich carrier-grade firewall like those available from F5 gain realistic confidence in their network security, mitigate the associated financial and reputational risks of attack, and can take advantage of the opportunity to offer their customers cutting edge and added-value security services for increased revenues.

F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com