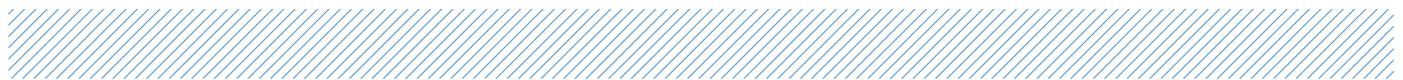


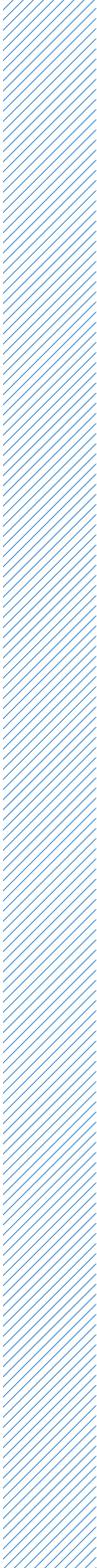
WHY IT'S TIME TO REPLACE YOUR VPN



Whitepaper

Contents

Overview	3
VPN Security Flaws	4
VPN Complexity Flaws	6
A Better Approach to Remote Network Security	7
How Software-Defined Perimeter Works	8
Evaluating a Software-Defined Perimeter	9



Overview

Today's global, distributed IT landscape makes it virtually impossible to secure infrastructure with perimeter-based security like the VPN, a solution that has not been updated in any meaningful sense in over 20 years.

VPNs were created in 1996 when Microsoft first developed the peer-to-peer tunneling protocol, around the same time that Blackberry was just launching two-way pagers and the term "cloud computing" was first coined.

Today's network landscape is one of incredible complexity with distributed applications, people, and data. Companies have taken the standard method of protection, the trusted private network, and applied hundreds or thousands of VPN and firewall rules with complex topologies to manage the chaos. Our expanding cloud and mobile ecosystems have rendered perimeter-based tools obsolete. In the meantime, networks are laden with unsanctioned, insecure devices. To complicate matters, in an increasingly distributed work environment, cyber threats are just as likely to come from inside an organization as they are from the outside.

VPNs were simply not designed for today's complex and changing IT infrastructures. We're a modern workforce, working anywhere and everywhere, and the days of a fixed and easily identifiable perimeter are long gone. Using VPN technology to secure how we work today simply defies progress.



“By 2021, 60% of enterprises will phase out network VPNs for digital business communications in favor of software-defined perimeters.”¹

Gartner



VPNs are not a security solution, but simply a means of connectivity that allows users access to environments from remote locations. A Software-Defined Perimeter is a modern security solution that overcomes the limitations of the VPN.

VPN Critical Flaws

- **VPNs authenticate to everything** because they trust blindly. Once a user's device is authenticated, he or she can typically gain complete access to an entire network – including unauthorized assets.
- **VPN access rules are too limited** and unable to keep up in complex environments. Rules based on IP address are either set to be too broad, allowing for wide-open access, or overly restrictive to the point of inhibiting work.
- **VPNs provide static, perimeter-based security.** This is ineffective when users are accessing data in multiple locations, public clouds, or SaaS applications that are hosted by third parties.
- **VPNs are a siloed solution** only intended for remote access by remote users. They do not help organizations secure on-premises users or networks.

¹ Gartner, "It's Time to Isolate Your Services From the Internet Cesspool," Steve Riley, Neil MacDonald, Greg Young. Refreshed: 17 November 2017. Published: 30 September 2016.

VPN Security Flaws

VPNs are widely used to provide remote employees, third parties, and contractors access to a corporate network. VPNs use simple, IP-based security to authenticate users to a network, and there are a number of reasons why VPNs are vulnerable.

VPN solutions often provide users access to far more resources than they actually need to do their jobs. As VPNs are difficult to administer, they're often configured to grant broad access to entire subnets where they rely on a flawed idea that's as old as basic TCP/IP networking: that a user should be allowed to connect to an entire network segment prior to authentication at an application level. Often times VPN users can see and potentially access much more than they really should, without any regard for their responsibilities or the context of their connection.

There are a number of ways malicious actors can use VPNs to get into networks. Here are four commonly exploited issues:

Issue 1: VPN Backdoor

Backdoors are a major problem with VPNs. A backdoor is a method to remotely access a computer system that bypasses customary security mechanisms. A simple web search for 'VPN backdoor' turns up millions of results involving just about every VPN vendor on the planet. These articles are all about flaws (either intentional or unintentional) in the technologies that these companies created and released to the market. There is also another kind of "backdoor" that Evan Gilman and Doug Barth talk about in their book [Zero Trust Networks: Building Secure Systems in Untrusted Networks](#):

"A VPN...It's the greatest backdoor that no one ever suspected."

Why? Because VPNs lack intra-zone traffic inspection, do not have enough flexibility in host placement, and create single points of failure.

Issue 2: Open Ports

Every VPN concentrator, without exception, is deployed in such a way that it has a presence on the internet with an open, continuously-listening port. A VPN concentrator is simply a networking device that enables thousands of remote users to establish a VPN connection. This means that all VPN concentrators are "out there on the public internet", accepting connection attempts all day, every day, regardless of who is on the other side of that connection. This inherent design flaw exposes an attack surface/vector for any nefarious person to exploit.

Issue 3: Authentication

Consider what happens once a user authenticates and establishes a VPN tunnel through the concentrator. That user is issued an IP address and effectively dropped onto the "trusted" internal network behind the concentrator. Without other tools (like ACL policies) or technologies (like NACs or Routers), that user is free to do whatever they desire on that network, and any endpoint is now an attack vector to move laterally inside the network. This is the other backdoor that Evan Gilman and Doug Barth speak of in their book. Unless other defenses or protections are in place, all the individual services running on those systems are open and listening for connection attempts.

Issue 4: Cookies

Once a user is authenticated, that authorization is written to the user's computer, usually in the form of an unencrypted cookie. If no special precautions are taken, that cookie will also exist in memory, which is also normally unencrypted. If that user's laptop is already under a threat actor's control, meaning an attacker has already deployed malware to it, then that authorization cookie is available to the bad actor and can be stolen, exfiltrated, and used in credential-based attacks. The open port issue above allows attackers to easily gain access to a network using a VPN once they've gained access to the unencrypted cookie. Again, once connected, the backdoor is wide open to the entire network.

There are architectural flaws with VPNs that won't go away by merely encrypting the authorization cookie:

- The concentrator will still be listening on an internet accessible port
- The user will still be connected to the "trusted" network
- The various data centers will still be connected by wide-open LAN-to-LAN connections, enabling cross-data center lateral movement
- The VPN application still won't support simultaneous connections to multiple destinations
- The concentrator will still lack integration with other important business systems

VPN Complexity Flaws

VPN policy management introduces too much complexity that requires significant resources to manage. VPN administrators that set policy start with an important choice: either create open policies that offer broad network access or create restrictive policies that offer limited network access.

The problem is that these policies are often set for broad access to the network because managing restrictive policies is complex, error prone, and difficult to manage. Strict access leads to a proliferation of rules to be managed, maintained, and audited. It's a challenge: open broad access and introduce significant security risks, or restrict access and make admins spend valuable time manually providing or "fixing" access.

Worse yet, as VPNs are used across on-premises and cloud environments, managing access based on static IP addresses doesn't work. New IP addresses are assigned dynamically by the cloud provider and can overlap with other IP addresses. VPN policies must be modified or new rules written so, once again, admins spend too much time applying and deciphering overbearing sets of VPN rules. Managing this access complexity becomes resource intensive and overbearing.

VPN clients are limited to one concurrent connection per device. As a result, VPN administrators with multiple data centers, office locations, or public cloud providers have two options:

1. Provide enterprise-wide connectivity between all data centers, offices, and cloud platforms for all users with one VPN connection. The security risk with this approach is that if one location is compromised, all other locations become vulnerable to compromise as well.
2. Limit each VPN client to connect to only one location. This requires users to know which VPN connection to use for the specific resource they are attempting to connect to. This is not an ideal user experience and can impact productivity. End users have to internalize which connections connect to which resources. It is not efficient for users to constantly disconnect and reconnect the VPN to access resources required to perform daily tasks.

Lastly, most organizations deploy firewall rules alongside the VPN. Firewalls are often configured and then left alone because any deviation usually results in a significant change ticket. This adds another layer of complexity for an insecure technology.

A Better Approach to Remote Network Security

A new approach to network security is needed to overcome the VPN's critical flaws and common exploits. The new approach must be rooted in Zero Trust, the notion that we must not trust unless we've extensively verified the identity of a user.

The Software-Defined Perimeter (SDP) is rooted in Zero Trust. It is a network security model that dynamically creates one-to-one network connections between users and the resources they access. SDP reduces the attack surface in real-time by creating a discrete, encrypted network segment of one, making everything else invisible and inaccessible. A network segment of one is an individualized, micro-segmented network tailored for each individual user, device, and session. Using Single-Packet Authorization (SPA) technology, the infrastructure is cloaked so that only verified users can communicate with the system. It's invisible to port scans and cryptographically hashed for additional defense. Further, this solution is holistic – it provides a single secure access control platform for both remote and on-premise users accessing remote and on-premise resources, as well as unprotected IoT devices.

Identity-Centric

A Software-Defined Perimeter is designed around identity, not the IP address. SDP seamlessly integrates with existing directory services and IAM solutions. It is able to build a multi-dimensional profile of a user and their device. Before allowing a connection to the network, it analyzes the context of the actual user such as:

- What are their roles and privileges?
- Where are they requesting access from?
- What time of day is it?
- Is their device secure or infected?

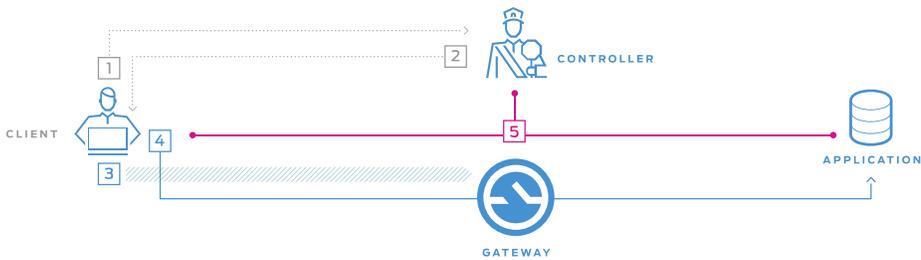
Adaptive and Extensible

Access privileges automatically adapt based on user context, device, and security conditions in real-time. It is fully programmable, scales and applies policies automatically to new cloud deployments, and integrates with operational systems for automated and time-based access.

Enables Zero Trust

SDP supports Zero Trust in two critical ways, starting even before a connection is ever made. Utilizing SPA to authenticate first, connect second, the external network is made invisible to attackers, reducing your external attack surface. After the successful authentication, SDP creates a discrete, encrypted segment of one, allowing access only to required resources, thereby reducing lateral movement attack surface.

How Software-Defined Perimeter Works



1. Using Single-Packet Authorization, the client makes an access request to a controller. The client device authenticates to the Controller, which evaluates credentials and applies access policies based on the person, environment, and infrastructure.
2. The controller checks the context and passes Live Entitlements to the client. The Controller returns a cryptographically signed token back to the Client, which contains the authorized set of network resources.
3. Using SPA, the client uploads these Live Entitlements, which the Gateway uses to discover applications matching the user's context. When the user attempts to access a resource – for example by opening a web page on a protected server – the network driver forwards the token to the appropriate cloaked Gateway. The Gateway applies additional policies in real-time to control access based on network location, device attributes, time of day, and more. The Gateway may permit access, deny access, or require additional action from the user, such as prompting for a one-time password (OTP).
4. A dynamic Segment of One network is built for this session. Once granted, all access to the resource travels from the Client, through the Gateway, and finally to the resource across a secure, encrypted network tunnel. Access is logged through the LogServer, ensuring that there is a permanent, auditable record of user access.
5. The Controller and Gateway continuously monitor for any context changes and will adjust the Segment of One accordingly.

Evaluating a Software-Defined Perimeter

Not all SDPs are created equal. Some include many restrictions: being limited to specific vendor clouds or even being limited to web-based applications. Look for an SDP that includes support for all protocols, including RDP and SSH. A true SDP will adhere to the Cloud Security Alliance SDP specifications and use Single Packet Authorization (SPA) to cloak all ports and provide DDoS protection. A comprehensive SDP will be able to use auto-resolvers that quickly define resources and auto-provision user access.

If you are seeking to implement a Zero Trust model, gain granular control of remote network access across offices, and support remote employees regardless of location, then Software-Defined Perimeter is the solution.

Accelerate Your Journey to Zero Trust With AppGate SDP

AppGate SDP is the industry's most comprehensive Software-Defined Perimeter. Get more information on how it allows enterprises to kill their VPN.

DOWNLOAD THE DEFINITIVE GUIDE TO A SOFTWARE-DEFINED PERIMETER

https://ww2.appgate.com/definitive_guide_to_sdp

VISIT THE APPGATE SDP WEBPAGE

<https://www.appgate.com/software-defined-perimeter>

TEST-DRIVE APPGATE SDP NOW

<https://testdrive.appgate-sdp.com/>