

WHITEPAPER

MANAGING THE RISKS OF THIRD-PARTY ACCESS

Why Old Model Security Solutions Aren't Stopping Data Breaches

EXECUTIVE SUMMARY

Appgate SDP provides secure user access to network resources. Recently some of the most catastrophic data breaches in history made US companies answerable to millions of customers whose data was compromised, as well as costing those companies millions in remediation and lost revenues. In the most notable cases, the theft of usernames and passwords from third-party vendors was to blame and was the entry point.

Many of today's security solutions, even when used in combination, simply aren't designed to mitigate the risks associated with third-party access. To address this, organizations must adopt a user-centric context-aware model that is built on the principle of least privilege and leverages the Software-Defined Perimeter model.

THIRD-PARTY CREDENTIALS: THE EASY WAY INTO YOUR NETWORK

Attackers have discovered a ripe opportunity to attack some of the biggest companies in the world, and it might take nothing more than a user name and password pair from third-parties with access to systems.

As recently as a few years ago, third-party credential theft was practically unheard of as a means to gain access to the most sensitive areas of a business network. Now, though, it's one of the biggest threats out there. If you're a large enterprise like Target or Home Depot, the easiest way for attackers to get into your network is likely to go through your third-party vendors.

Each third-party vendor could potentially offer attackers a direct route into your most sensitive network segments. Think of how a third-party data breach is typically carried out:

1. The attackers identify their target's vendors.
2. They use spear-phishing techniques to acquire those vendors' credentials for access to the target company's network.

3. Once inside, the attackers can look for ways to widen their foothold in their target company's systems by moving laterally across VLANs. If access is provisioned via VPNs, they may have direct access to the underlying network infrastructure and be able to start scanning for open ports and unsecured devices in seconds.
4. The attackers might then spend weeks or months preparing to strike, studying the network's weaknesses and installing sophisticated malware that could take just as long again to detect.

Looking at this attack pattern, it's easy to understand how tens of millions of records are compromised in a single incident. How, then, can you mitigate the impact of third-party credential theft? Working alongside each of your business partners to strengthen their individual security profiles is a noble goal, but not a practical one. What's needed is a better way to manage the risks of third-party access to your networks and applications. Unfortunately, most organizations are relying on old technologies.

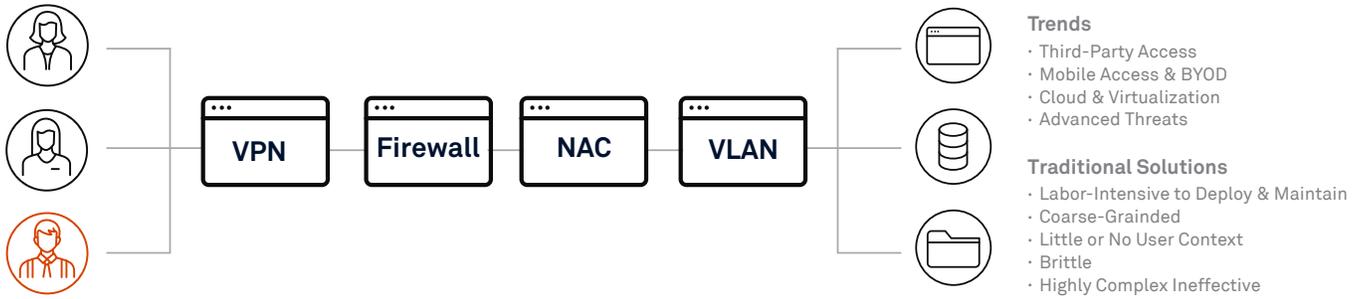
THE PROBLEMS WITH OLD MODEL SECURITY SOLUTIONS

Current best practices recommend a laundry list of security technologies: VPNs, VLANs, NAC, Next Generation Firewalls, Privileged Access Management (PAM) solutions, and so on.

But too much technology results in 'spend in depth', and not necessarily improved security. And if you're still using the same principles you were using ten or twenty years ago, you might have the strongest network perimeter in the world, but no ability to respond to internal threats.



Traditional Security Architectures



Result Users Have Too Much Network Access

NAC

Network access control (NAC) is a method of bolstering network security by restricting the availability of network resources to endpoint devices that comply with a defined security policy. A traditional network access control server performs authentication and authorization functions for potential users by verifying client device profiles (such as the presence of antivirus software and spyware-detection programs) before permitting access to the network.

Through a combination of client agents and network server components, NAC systems enforce policies about which network segments users can access. NAC (which often follows the 802.1X protocol), uses client profile and authentication information to make these policy decisions. Based on these policy decisions, NAC permits access to network segments or VLANs. NAC systems may also require or perform remedy actions on non-compliant devices (such as enabling a client firewall).

NACs do incorporate some (limited) client profile information to make network access decisions, and can (in some ways) remediate non-compliant clients. And they integrate into existing network infrastructure components such as VLANs. Ultimately though, NAC solutions fall short for several reasons:

- Most importantly, they cannot provide fine-grained control of which network resource users can access. They rely on existing (and separately managed) network segments, firewalls, or VLANs.
- Due to the management issues around adding devices and firewall rules, enterprises have expressed doubt about the practicality of NAC deployment in networks with large numbers of diverse users and devices, the nature of which constantly change.
- They typically have limited ability to make access decisions based on user context.
- NACs do not provide secure, encrypted communications between clients and services.
- NAC customers must use another solution (such as a VPN), which adds more cost, complexity, and management effort.

VPNS

VPNs are a common way for organizations to set up secure, encrypted tunnels across untrusted networks. They provide a secure way to allow remote users to access the trusted corporate network across the untrusted internet. They work by using client software to create and manage a secure, encrypted network tunnel from the client device to the VPN server. Once authorized, users are effectively inside the corporate firewall, and have access to all corporate resources. VPNs are inexpensive, commodity software. Users are accustomed to using them, and they integrate into many multi-factor authentication platforms (e.g. RSA SecurID).

However, they lack the ability to prevent third-party breaches in that:

- They are only effective in environments that have a well-defined perimeter around on-premises software.
- They don't help control access to cloud-based solutions (SaaS or IaaS).
- They provide coarse-grained "all-or-nothing" access to the network—made worse by the fact that most organizations typically only separate services into two or three segments (such as "guest", "employee", and "admin") each with many services.
- They cannot easily adapt to the user's situation—for example adjusting authentication strength or access levels up or down based on access location, device or user context.

VLANS

VLANS are virtual (software-defined) networks running on top of a (differently configured) physical network infrastructure. A VLAN allows a network of computers and users to communicate as if they exist in a single LAN and are sharing a single broadcast and multicast domain. This gives organizations more flexibility to group hosts together independent of their physical network setup.

Like physical LANs, VLANs ensure that users (or hosts) on a VLAN have network access to all other resources on that VLAN. Since they can quickly adapt to changes in network requirements and relocation of workstations and server nodes, VLANs are implemented to achieve scalability, security and ease of network management. Traffic patterns can also easily be controlled by using VLANs, and they can reduce network latency / improve network performance.

On the downside from a security perspective, while users with access to one VLAN cannot see systems on other VLANS, users can access all systems on the VLAN they're on. And because VLAN access rules are labor-intensive to administer, most organizations only have a small number of VLANs (typically five or fewer). This means that each VLAN has dozens or hundreds of hosts on it, which represents a large security risk—an attacker with a foothold on one client device can immediately attempt lateral movement, and probe hundreds or thousands of services for weak points. VLANs also cannot easily adjust to what users have access to, based on user context, and configuring

VLANS often requires a dizzying number of firewall rule sets, making it complex and hard to manage for IT Admins.

PRIVILEGED ACCESS MANAGEMENT (PAM)

PAM solutions help organizations better manage and track how privileged users (system admins) access systems. With PAM, enterprises control who accesses key systems, when they access them, and what actions they perform on those systems.

PAM solutions provide the following:

- Admin account credentials vault—for user checkout/checkin of logins to privileged accounts.
- Ability to control, filter, and log commands performed by admins—via proxied sessions of Secure
- Shell (SSH) and Remote Desktop Protocol (RDP) (and others).
- Application-to-application credential management.

Because they strengthen the security and better control admin access to key systems, PAM solutions have obtained wide enterprise adoption. They do a good job at obscuring admin credentials and forcing users to gain access via the PAM solution (often combined with multifactor authentication).

They also do a good job at proxying access to systems, and controlling and logging user admin activities, with some access controls based on user context. And, because they control passwords, they can be used to manage access to cloud-based admin apps, such as the AWS console.

PAM solutions fall short from a security perspective because:

- Although they do strengthen security and auditability for privileged users and key systems,
- they do so primarily by just strengthening the “front door” of authentication.
- They do not protect or prevent users from accessing unauthorized resources at a network level, and cannot prevent attackers from exploiting unpatched server vulnerabilities.
- They're incomplete, as they don't provide any remote access capabilities.
- They're known to be expensive and only used to cover a few key systems.

NEXT GENERATION FIREWALLS (NGFW)

NGFW is a hardware or software-based network security system that is able to detect and block sophisticated attacks by enforcing security policies at the application level, as well as at the port and protocol level. NGFWs bring additional context to the firewall's decision-making process by providing it with the ability to understand the details of the web application traffic passing through it and taking action to block traffic that might exploit vulnerabilities. They build on the strengths of traditional firewall technologies, with the addition of deep protocol knowledge to better enforce security.

NGFWs work by analyzing packet flows. Their security downfall is that not every single packet can be analyzed. NGFWs examine on the initialization of the flow, not every single packet of the flow. Further, NGFWs can identify a user and an application, but they can't analyze the context of the user and/or application. Did the user log in with only a password or did he use two-factor authentication? NGFWs generally don't obtain or leverage device context to make access decisions, instead they use the source IP to identify traffic and try to build some context around that IP.

IDS, IPS AND SIEM

Most organizations do, of course, have technologies at their disposal to detect when their perimeter defenses have been breached. These include intrusion detection systems (IDS), intrusion prevention systems (IPS) and security information and event management (SIEM). These are mature and effective solutions that can quickly identify malicious activity within a network.

They're also prone to picking up false positives, which can make it nearly impossible for organizations to pinpoint real threats before their data is compromised. A large enterprise might encounter hundreds of false positives per day, with otherwise harmless applications like instant messaging clients causing suspicious-looking traffic on sensitive network segments. Attackers, meanwhile, aren't known for broadcasting their activities in a way that makes them easy to catch. They might

spend weeks or months hidden inside a network before they strike, at which point they know exactly how to exfiltrate data as efficiently as possible.

Because of these circumstances, IDS, IPS and SIEM can be resource-intensive and reactive rather than proactive, hindering their ability to stop hackers in their tracks.

MITIGATING THE DAMAGE POTENTIAL OF THIRD-PARTY RELATED BREACHES

Organizations need to better manage the risks of third-party access to decrease the chances that attackers can penetrate through each of the defense layers. An effective security solution should be able to tell if the context of a remote connection is suspicious, such as if it originates from an unusual location or time of day, or from a device with no antivirus software installed. And it should be able to ask for additional authentication steps like one-time passwords (OTP), adjust user permissions on the fly, and ultimately block access according to the level of risk.

By using a solution that leverages the Software-Defined Perimeter security framework, organizations can ensure that all endpoints attempting to access a given infrastructure are authenticated and authorized prior to accessing any resources on the network. All unauthorized network resource are made inaccessible. This not only applies the principle of least privilege to the network, it also reduces the attack surface area by hiding network resources from unauthorized or unauthenticated users. A Software-Defined Perimeter overcomes the constraints of traditional tools by effectively creating a dynamic, individualized perimeter for each user—a network 'segment of one'.

APPGATE SDP

Appgate SDP, an adaptive, identity-centric, security solution built for today's hybrid enterprise. Appgate SDP protects critical data from internal and external threats, while significantly lowering costs.

Unlike a traditional network that connects various roles or groups to a network segment and then relies on application level permissions for authorization, Appgate SDP creates individualized perimeters for each user, allowing for much more

fine-grained access control and giving individual users access to only what they need to do their jobs. Appgate SDP provides this access control with a real-time understanding of policy.

Appgate SDP ensures that all endpoints attempting to access a given infrastructure are authenticated and authorized prior to being able to access any resources.

Once the user initiates a session with an authorized resource, Appgate SDP creates an encrypted tunnel, allowing traffic to flow only from the user device to the protected resource. We call this a 'Segment of One' and we make the rest of the network completely invisible to the user.

Including the system itself. Meaning all resources, including Appgate SDP are completely dark to all unauthorized users. Gateways and controllers are completely cloaked so they cannot be probed, scanned, or attacked. So, a port scan of the system would show no open ports, reducing the network attack surface by preventing network reconnaissance and limiting lateral movement on the network.

Even while the session is open, Appgate SDP can detect changes in the posture of the user, his or her environment and infrastructure, including changes in the cloud, and automatically adjust access privileges. Appgate SDP may then force a step-up authentication or terminate the session completely based on this newly detected change in posture or context.

CONCLUSION

Chances are we've yet to hear the last of breaches tied to credential theft. Whether it's third-party or employee credentials, organizations need to change their security practices to not only better secure access, but also limit damage if bad actors find their way into your networks. And remember—even if your vendors are to blame, it's your customers whose data is being compromised, and your reputation that will suffer.