Beyond the VPN:

Zero Trust Access in a Federal Hybrid Work Environment

zscaler™

How can agencies secure the new hybrid workforce within the parameters of federal requirements?

The answer: a cloud-enabled zero trust architecture that enables secure, least-privileged access to private applications by establishing dynamic connectivity from authorized users to specific applications based on identity and context.

The virtual private network (VPN) has been a powerful tool in the network security administrator's toolbox for decades, because it has provided a means for remote computers to communicate securely across an untrusted network such as the internet. Whether for branch offices communicating with headquarters (site-to-site), or an employee working from home (remote access), the VPN provided a secure point-to-point tunnel back to resources on protected networks.

VPNs were an effective tool for federal agency security when the majority of network traffic was on-premises. When VPNs were introduced in the late '90s, proprietary applications were in the data center. Cloud hosting of apps didn't exist. Users going off-premises was the exception rather than the rule.

But times—and the ways we work—have changed.

According to a recent survey, the majority of federal employees expect to continue teleworking post-pandemic. Respondents said that before COVID-19, they telecommuted an average of two days per week. During the lockdown, that jumped to nearly the entire work week. It's clear that for federal employees, remote and hybrid work is the new normal, with those survey respondents not expecting to go back to the office full time, citing that they are likely to work remotely at least three and a half days per week.

One challenge with VPN in these new circumstances is that VPN connectivity comes with a performance tradeoff: Routing traffic through narrow, dedicated, persistent point-to-point tunnels via indirect security gateways—secure as it may be—adds considerable latency as compared to direct on-the-network data travel. And the full-tunnel VPN approach to protecting outbound traffic—carrying it inbound first, for inspection by an on-premises security stack—only exacerbates the problem. When those inconveniences affected only that small minority of remote-work network traffic, the pros may have outweighed the cons. While the federal government had been slowly expanding remote work, the pandemic response radically sped the process, bringing the drawbacks of VPN into the spotlight: the more remote users, the more constrained the bandwidth. And the worse the user experience.

What is the security implication of this new hybrid workforce?

The rapid push to telework has stressed our network-security-based controls to the breaking point.

In practice, the rapid evolution of the federal workspace to a hybrid on-prem/off-prem model means we must now guard against a new, wider range of attacks. Now that the majority of network traffic is remote, agencies must move beyond VPN to secure network traffic and optimize connectivity performance. A cloud-delivered Zero Trust Architecture (ZTA) is the optimal path to secure provisioning for both remote access and on-premise users.

What about complying with guidance?

Fortunately, oversight bodies recognize the need to evolve. The National Institute of Standards and Technology (NIST) has started updating guidance to match the new hybrid-workforce reality. In December 2020, NIST made a minor revision to SO 800-53 Revision 5 (Security and Privacy Controls for Information Systems and Organizations) that indicates agencies have more flexibility in applying controls. The section reads:

> "(7) BOUNDARY PROTECTION | SPLIT TUNNELING FOR REMOTE DEVICES Prevent split tunneling for remote devices connecting to organizational systems unless the split tunnel is securely provisioned using [Assignment: organization-defined safeguards]. Discussion: Split tunneling is the process of allowing a remote user or device to establish a non-remote connection with a system and simultaneously communicate via some other connection to a resource in an external network. This method of network access enables a user to access remote devices and simultaneously, access uncontrolled networks. Split tunneling might be desirable by remote users to communicate with local system resources, such as printers or file servers. However, split tunneling can facilitate unauthorized external connections, making the system vulnerable to attack and to exfiltration of organizational information. Split tunneling can be prevented by disabling configuration settings that allow such capability in remote devices and by preventing those configuration settings from being configurable by users. Prevention can also be achieved by the detection of split tunneling (or of configuration settings that allow split tunneling) in the remote device, and by prohibiting the connection if the remote device is using split tunneling. A virtual private network (VPN) can be used to securely provision a split tunnel. A securely provisioned VPN includes locking connectivity to exclusive, managed, and named environments, or to a specific set of preapproved addresses, without user control." (**Quoted from SC-7 Boundary Protection**)

Secure provisioning by organization-defined safeguards empowers an organization to optimize its network traffic based on the traffic's destination, as long as all flows are controlled and secured. Cloud-enabled ZTA delivers these security requirements.

As TIC 3.0, SP 800-171 (Protecting Controlled Unclassified Information in Nonfederal Information Systems and Organizations), and the Cybersecurity Maturity Model Certification (CMMC) take their historical foundation from SO 800-53, it is likely this guidance will also evolve.

VPN yesterday and today

Thirty years ago, the enterprise network was relatively simple: security consisted of protecting network-connected applications and systems by building a secure perimeter around them. The original full-tunnel VPN (and subsequent split-tunnel option) provided the security controls needed at the time for those looking to connect to the network from somewhere outside that network's physical location. While the user experience was not good and bandwidth consumption was high, the security benefits outweighed the drawbacks for remote users, at least for IT leaders tasked with establishing secure connectivity.

And then work started moving outside the network perimeter. Applications moved to the cloud, work required internet access, and employees grew to rely on systems and resources outside of the enterprise's control. The user experience issue became more visible, as more users sought to work off-network and from any device, anywhere—but couldn't. Concurrently, the operational decision for split tunneling became exponentially riskier, as well.

The future of work is in the cloud, remote-enabled, and device-agnostic. Remote access VPNs worked for a network-centric world. But they cannot scale to secure or support work in this new age. Poor user experience drives users to bypass security; hub-and-spoke legacy architectures cannot optimize data flows in a cloud and mobile environment; and network-centric appliances are insufficiently flexible for today's demands.

How can agencies secure this new way of work within the parameters of federal requirements? A cloud-enabled zero trust architecture that enables secure access to private applications by establishing connectivity from authorized users to specific applications on a dynamic identity- and context-aware basis.

New connectivity for a new way of work: cloud-enabled zero trust security

Full-tunnel and split-tunnel VPNs are predicated on the idea of an encrypted tunnel connecting a remote endpoint to a protected network, resulting in large, complex, network-based access configurations across multiple network devices. That network-centric security control no longer fits with the reality of cloud, mobility, and the hybrid workforce.

Zero trust access built on a flexible, reliable, scalable cloud platform redefines how user traffic reaches target applications. Whether from user to app, user to internet, or app to app, connectivity is direct, secure, and ephemeral. Security is delivered by a close-to-every-user edge-computing proxy, served from the cloud, and scalable to include full inspection of all data traffic, even if it is encrypted.

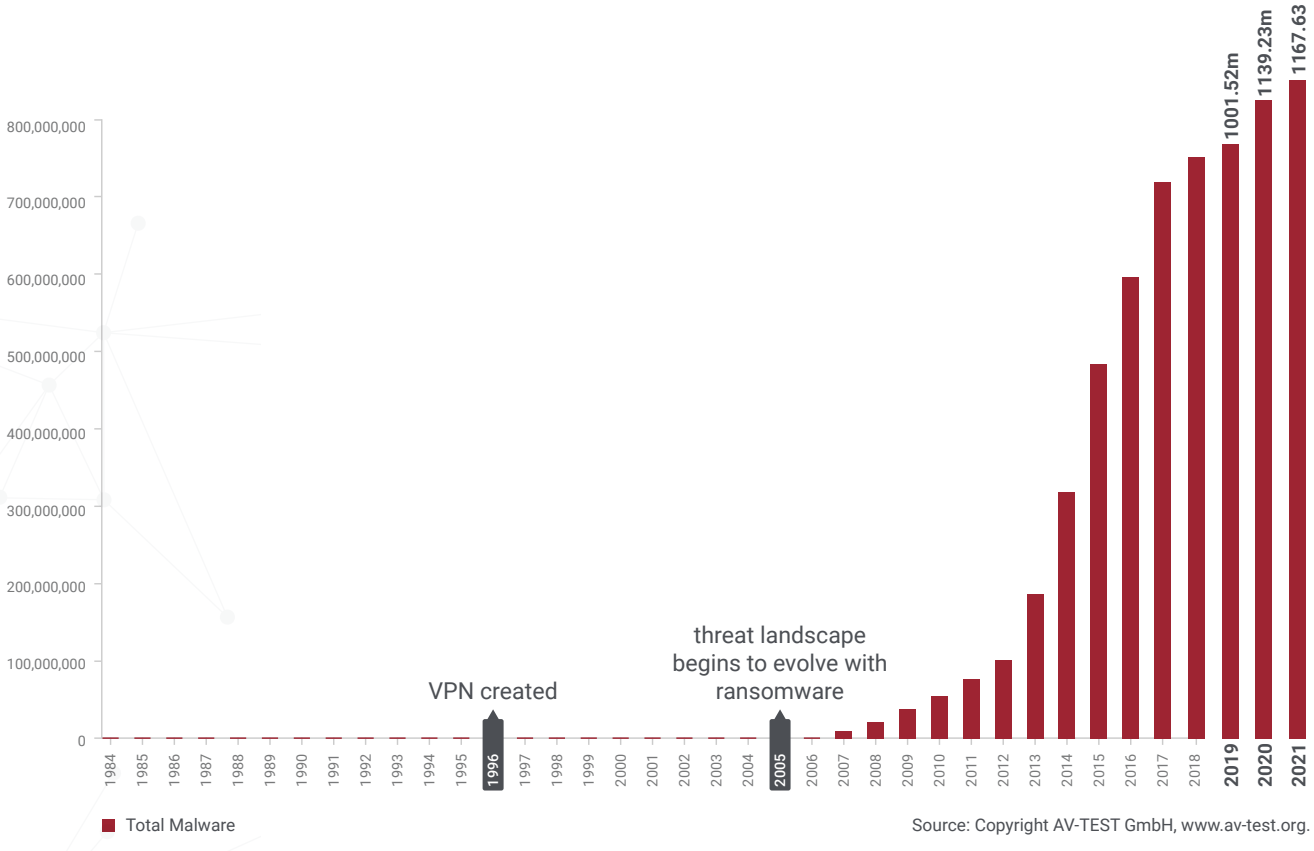

Figure 1: Threat landscape over time

Network architecture technology and its accompanying perimeter-based security model remain stubbornly entrenched in organizational IT environments. As the technology has aged (in 2021, the VPN turned 25!), threat risk has risen dramatically. Hackers have had decades to discover and exploit network security and VPN vulnerabilities. Figure 1 above shows the total number of reported malware attacks since 1984.

Cloud-enabled ZTA supplants the need for a VPN, enabling users to connect to anything, from anywhere, securely, with no compromise to performance. It's a move from a full-tunnel or split-tunnel model to an encrypted application microtunnel. Endpoints no longer connect to a network; instead, there is a protected connection securing authorized traffic to and from private resources.

Zscaler™ delivers this ZTA model via the Zscaler Zero Trust Exchange™—a highly-distributed, cloud-native platform with security served from more than 150 data centers around the world for general-purpose traffic, plus a FedRAMP-authorized GovCloud that can be extended to meet agency needs. Protection is inline, proxied, and immediate: data traffic moves directly to the cloud, the internet, or the (FedRAMP-authorized private-edge) data center without lag-inducing backhauling trips over costly MPLS infrastructure. Data is inspected, even if encrypted: Nothing bad comes in, nothing good goes out.

The combination of both internet-bound and private traffic management through the cloud-based Zscaler Zero Trust Exchange delivers on the ideals of a full-tunnel VPN: providing full visibility and control of user traffic, preventing U-turn attacks, protecting outbound traffic, and allowing only authorized inbound traffic at a very granular level. At the same time, Zscaler offers scalability, reduced overhead, and (greatly) reduced risk, supplanting open inbound listeners, complex network-based controls, lateral-movement risk, and other remnants of the legacy network-centric model.

Ignoring its performance and scalability limitations, a VPN can offer some semblance of a zero trust security posture. But a VPN is by default an open Layer 3 connection. Therefore, security admins must apply access controls on the tunnel to control what a user can send into it, and may need to combine that with network-segmentation policies inside the protected network. By contrast, in the Zero Trust Exchange, application micro-tunnel access is authorized or rejected immediately based on user identity, so only authorized traffic can get into the tunnel in the first place.

The Zero Trust Exchange enables agencies to deliver true least-privileged access. If a user is not authorized to access an application, that user cannot see (let alone connect to) the application.

VPN vs. cloud-enabled security: a quick comparison

A VPN and the Zero Trust Exchange differ in three primary areas: security, user experience, and complexity.

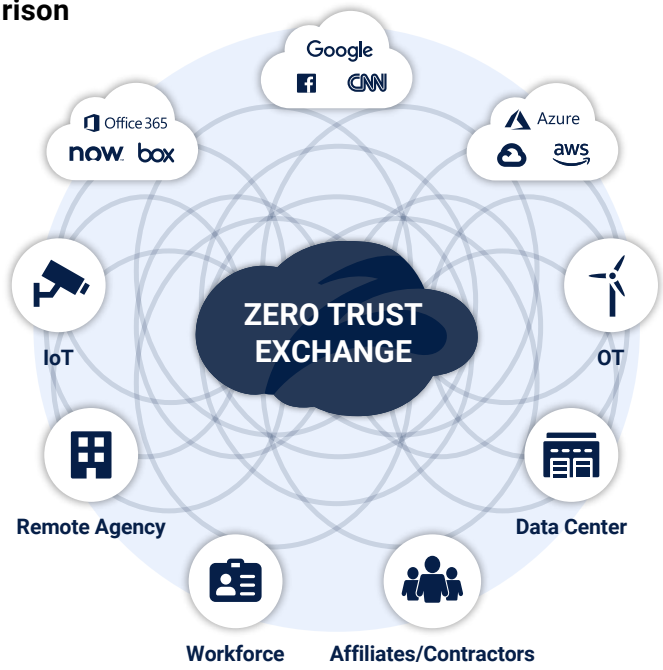

Security

VPN: Providing application access requires placing users on the network, while also presenting an exposed external attack surface as the VPN concentrators listen for inbound connections.

Zero Trust Exchange: Providing access to private apps does not require network access. Direct, outbound-only connections reduce external-facing attack surface, and application micro-tunnels ensure least-privilege access.

User experience

VPN: All traffic is backhauled to the data center—and potentially further backhauled from there to other application-hosting environments—slowing connectivity for the user. As more users employ the VPN for remote access, performance degrades further for all. Access is inconsistent between on-premise and remote users.

Zero Trust Exchange: Enforcement is inline to the user, ensuring fast, seamless access to applications across data center and cloud environments. User experience is consistent regardless of device, location, or application.

Complexity

VPN: Expensive inbound security stacks are replicated across multiple data center locations, each stack requiring tedious manual management and configuration of ACL and FW policies. Troubleshooting across multiple hops and administrative interfaces is tedious.

Zero Trust Exchange: Cloud-delivered security controls offer centralized configuration and management for distributed access, with flexibility and scalability for rapid deployment and simplified operations. Policy follows the user and can be as granular, or as general, as necessary.

Conclusion: VPNs cannot secure the new way of work. The Zero Trust Exchange can.

The threat space has evolved since VPNs were introduced. The threat landscape has migrated upward from the network layer to the application layer, and protections must migrate upward too.

IT leaders' fundamental goal is to prevent unauthorized connections that make the system vulnerable to attack and exfiltration. A cloud-enabled zero trust architecture like the Zscaler Zero Trust Exchange achieves that objective, and is an equivalent control—a modern alternative to a VPN.

Some counsel to agency IT leaders evaluating network security:

- Review NIST SP 800-207 "Zero Trust Architecture." The document establishes a new standard for defining a ZTA.
- Learn about the cloud-enabled technologies that are superseding legacy network-based approaches and identify use cases that could be better served by a modern Zero Trust solution.

It is possible to have the best of both worlds when it comes to network security: the user experience and performance benefits of routing traffic directly to the desired resources AND the full security granularity and visibility of connecting only authorized users to permitted applications. It's possible with the Zscaler Zero Trust Exchange.

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so that customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SASE-based Zero Trust Exchange is the world's largest inline cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

Zscaler, Inc.
120 Holger Way
San Jose, CA 95134
+1 408.533.0288
www.zscaler.com