WHITE PAPER

Secure Segmentation Prevents Flat Networks from Failing When Attacked

Building Effective Enterprise Security Requires Network and Business Leaders to Think Differently

Executive Summary

Hybrid IT and adopting work-from-anywhere (WFA) strategies have led to the exponential expansion of new network edges. And for many organizations, this has resulted in an expanded and fragmented attack surface that has become a perfect opportunity for bad actors to launch cybersecurity attacks from new attack vectors, undermining the ability of network and security leaders to maintain business operations.

Traditional flat networks, even those using network-based segmentation or microsegmentation techniques, cannot detect or prevent many of today's more sophisticated attacks. Part of the problem is that many of these networks still provide single-time authenticated users and devices with unfettered access to virtually any application. Such implicit trust policies provide free rein across permitted segments while reducing visibility across the network, especially into encrypted paths. And the lack of integration between security and network elements constrains their ability to perform essential firewall functions—let alone advanced security inspection—at the growing number of dynamic network edges and junctions, rendering them unable to contain cyberattacks.

78%

Fortinet survey finds 78% of organizations felt prepared for ransomware attacks, yet half still fell victim.

The Challenge of Securing Disparate Networks

Organizations are deploying hybrid IT architectures comprising campuses, data centers, interconnecting branches, home offices, mobile workers, and multi-clouds to accelerate digital innovation and optimize and develop new products. And nearly all these networks are being enhanced with 5G, which adds hyper-performance to an already complex network environment.

The recent transition to a new hybrid workforce approach has compounded this challenge. Many employees work at least part-time from home, with their devices following them everywhere they go. And applications continue to migrate to one or more clouds, including data center and private clouds, as well as Software-as-a-Service (SaaS)-based solutions. Looking at the mobility of users and the disparate locations of applications, the question facing many IT teams is: How do we deliver consistent security everywhere? And how can users safely consume applications from any location, on any device, at any time?

These new hybrid worker and IT paradigms have led to an exponential expansion of the network edge, resulting in an expanded attack surface and fragmented visibility and control. The result is a perfect platform from which bad actors can successfully launch cyberattacks and undermine business continuity. And while some of this network transformation is the result of intentional digital acceleration, some of it is also happening organically. For example, mergers and acquisitions activity often results in a diverse and fractured infrastructure with limited coordination or visibility between different parts of the organization.

One challenge arising from these expanding and fragmenting attack surfaces is that they create numerous new paths through which criminals can attack, along with new devices and interconnected applications and network environments for them to target. The need for new devices and software to support digital acceleration efforts has contributed to the growing volume of vulnerabilities targeted by new or improved cyberthreats. Common Vulnerabilities and Exposures (CVEs), the list of publicly disclosed computer security flaws, reached an all-time high in 2022, with critical vulnerabilities up 59% over 2021.[1] That list is only expected to grow. This has caused many IT teams to struggle with keeping their distributed devices and applications patched, especially as home networks leverage personal technologies to access business applications deployed in hybrid cloud and on-premises environments—a fact that cybercriminals have been all too eager to exploit.

And at the same time, threats are increasingly sophisticated, automatically seeking and exploiting vulnerabilities with advanced malware, making security a reactive exercise in many organizations. Increasingly sophisticated threats, many enhanced with automation and artificial intelligence (AI), regularly target high-priority sectors such as critical infrastructure, healthcare, information technology, financial services, and energy. Ransomware, in particular, has become a significant concern for most organizations. Although 78% of organizations felt prepared for ransomware attacks, half still fell victim to attacks.[2] Organizations are more concerned about ransomware than any other cyberthreat.

Difficulty Managing Disparate Networks: Is Segmentation the Answer?

Network engineering and operations leaders have responded to these challenges for years by building strong perimeter defenses to prevent attacks and segmenting their networks internally for operational controls.

Traditional network segmentation techniques based on IP addresses have primarily been augmented with VLANs, with VXLAN-based segmentation techniques supporting large-scale virtualization deployments and enabling more granular controls. Other methods include VMware NSX segmentation for virtualized workloads and Cisco ACI Application segmentation using physical switches. And there is a plethora of host-based segmentation techniques that leverage agents running on hosts that need to be segmented.

These microsegmentation techniques enable access control policies to be defined by workloads, applications, or architectural attributes such as the virtual machines (VM) on which the applications, data, and operating systems reside.

However, such segmentation and microsegmentation approaches are not the panacea they are sometimes hailed to be. Segmented and microsegmented networks must still perform advanced security inspection at each segmentation edge and juncture. Otherwise, they cannot prevent intrusions from moving laterally across the devices and applications that connect to and traverse the resulting flat network, whether within a single segment or for the many applications and workflows moving across multiple segments.

Why Traditional Segmentation Fails to Protect the Enterprise

Access control for internal network segments tends to be designed from the architecture up. As a result, security is neither intrinsically nor deeply integrated into networking. Instead, it is applied as an overlay, which may be fine for static networks and largely predictable workflows and transactions. But such tactical approaches mean that security policies, inspection, and enforcement cannot quickly adapt to evolving business needs or dynamic networks, and such changes leave security gaps targeted by cybercriminals.

There are three critical reasons why segmentation alone will not protect today's dynamic hybrid networks.

1. The trust valuations on which access policies are based tend to be static, implicit, and unrestricted. The inability to continually verify users and devices creates compliance and control challenges, especially when a user or device becomes compromised.
2. Access control policies cannot be effectively enforced due to a lack of advanced (Layer 7) security detection and inspection across the hybrid IT. Isolated legacy security solutions cannot see and control these components efficiently or adapt in real time to changes in the network.
3. These problems often stem from network engineering and operations staff planning their segmentation architecture without adequate attention to identity, visibility, and security. Understanding these issues and their aggregate impact can lead to a more risk-aware and responsive approach to segmentation.

Why a Traditional Bolted-on Network Security Approach Is Ineffective

Organizational needs usually dictate corporate network design, with the rules governing who and what can access which network resources being determined by business policies, industry standards, and government regulations. The network operations team then uses these rules to configure the access control settings in their routers and switches that permit users, devices, and applications to access specific network resources. While this approach may seem straightforward, network engineering and operations leaders should immediately recognize some critical downsides.

First, the needs of today's organizations are evolving, and the growing demand for flexibility and agility is impacting corporate network design. As a result, business processes, compliance requirements, and network access requirements have become vastly more complex than the network structure. Consequently, it is ineffective to use the network architecture to define and secure network segments for those resources that must be simultaneously accessible to all authorized users and applications (and utterly inaccessible to all others).

To effectively manage security risks, network engineering and operations leaders must instead rely on current and accurate information on the trustworthiness of users, applications, and network assets at all times. Unfortunately, traditional network connectivity—including intelligent application-driven solutions such as SD-WAN deployed in hybrid IT architectures—does not include seamless security integration. Other issues, such as the proliferation of unknown Internet-of-Things (IoT) devices and ongoing OT and IT integration, create additional challenges around visibility and security.

Trust Valuations Based on Statistics and Implicit Access Allow Breaches

Many of today's most damaging security breaches are due to compromised user accounts and passwords, and users with inappropriate access levels exacerbate this problem. To effectively manage security risks, network engineering and operations leaders must always have current and accurate information on the trustworthiness of users, applications, and network assets. As a result, internal firewalls and other access control mechanisms that manage internal traffic flows between network segments must constantly identify, verify, and monitor users, devices, and applications. If trust assessments are out of date, segmentation technologies become useless at preventing threats from moving laterally across the network.

Some organizations have responded to these dangers by employing a zero-trust network access (ZTNA) strategy, which controls access to applications to verify users and devices before every application session. Zero-trust network access confirms that they meet the organization's policy to access that application, grants access to specific applications per user, and then monitors those connections to detect threats and maintain compliance.

Security Requires End-to-End Visibility. Without It, Security Controls Mean Little

Most traditional approaches to segmentation assume that all necessary network security components are in place and ready to execute whatever access control policies the IT team defines. Unfortunately, this is usually an unsafe assumption.

First and foremost, the rising volume of encrypted web traffic has reached 94%.[3] While this is great news for organizations looking to provide secure, encrypted access to applications, it also allows bad actors to hide their activities in secure channels. Making things worse, many network teams intentionally turn off SSL/transport layer (including TLS 1.3) inspection in their next-generation firewalls (NGFWs) to optimize network performance because they fear the impact on performance. The inability of nearly all legacy firewalls to inspect encrypted traffic at digital speeds means that criminals can find their way in and out of an enterprise network undetected to launch ransomware attacks and exfiltrate data.

Second, due to budgetary constraints or because deployment and management require too many resources, many network engineering and operations teams hesitate to deploy advanced network security and other solutions everywhere they are needed—within the enterprise, in every cloud, and on every endpoint and IoT device. And the ones they do deploy tend to operate in isolation. Unfortunately, point security solutions cannot easily share threat intelligence on known, emerging, or zero-day threats or easily participate in a coordinated response.

Acting promptly is essential to disrupting an attack sequence, as outlined by MITRE.[4] However, the overall effectiveness of security components is severely compromised when they are not tightly integrated. For example, when an isolated firewall detects a suspicious packet, it may take hours (or longer) for the information to be seen by a security administrator and disseminated to the rest of the network.

Third, organizations cannot respond effectively to mitigate the impact of breaches without dealing with malicious websites, known malware, and unknown attacks. This requires integrating extended detection and response (XDR), an intrusion prevention system (IPS), and sandboxing technologies to automatically quarantine and test suspicious packets. Conversely, the lack of integration between security elements and between security and the network makes orchestration and automation across hybrid networks impossible. And the subsequent reliance on manual operations invariably leads to breaches, as they are far too slow and error-prone.

Segmentation with Network and Security Convergence Becomes Ineffective

What's required is an integrated, coordinated approach to security. A fully integrated and unified security solution is the only way to ensure consistent, adaptable threat detection and response across today's segmented hybrid IT architectures. That's where the hybrid mesh firewall (HMF) approach comes into play. A hybrid mesh firewall secures the convergence of on-premises and cloud-native domains with consistent policy enforcement and unified management. This unified security platform approach provides coordinated protection across every area of enterprise IT, including corporate sites, branches, campuses, data centers, public and private clouds, and remote workers. It's even better when unified management and analytics span the entire secure networking framework. This single-pane-of-glass strategy results in complete visibility and protection against security threats. The hybrid mesh firewall approach simplifies operations, ensures compliance, and reduces complexity with automation to increase operational efficiency, whether you have all on-premises firewalls, all cloud firewalls, or a hybrid mix of both.

Seeing Is Understanding

Without centralized management, orchestration, and control, network engineering and operations leaders who believe their segmented network is well-protected are likely working under a false sense of security. But without real-time data, it is impossible to know. And the best way to determine whether the security strategy being used to protect a dynamic, hybrid network is effective is to run ongoing end-to-end security assessments. However, without the end-to-end visibility provided by a fully integrated security platform—a security fabric able to touch and adapt to every edge of the network—a reliable assessment is impossible, preventing IT leaders from accurately reporting on and optimizing their company's security posture.

It is up to network engineering and operations leaders to ensure that the access control policies applied to internal network segments can withstand today's perpetually expanding and fragmenting attack surfaces. Addressing this challenge starts by converging network and security into the hybrid network architecture. Only with careful attention to segmentation design can a company be confident in its ability to thwart attackers looking to sow destruction by moving laterally across the network.

[1] "We analysed 90,000+ software vulnerabilities: Here's what we learned," The Stack, January 9, 2022.

[2] "The 2023 Global Ransomware Report," Fortinet, April 2023.

[3] "HTTPS encryption on the web," Google, accessed May 17, 2022.

[4] ATT&CK Matrix for Enterprise," MITRE ATT&CK, accessed January 31, 2022.

FORTINET

www.fortinet.com

June 22, 2023 12:23 PM

1440498-A-0-EN