

Traditional Segmentation Fails in the Face of Today's Expanding Attack Surface

**Why Network Engineering and Operations
Leaders Should Be Concerned**

Table of Contents

Executive Overview	3
Introduction: Difficulty Managing Disparate Networks: Is Segmentation the Answer?	4
3 Reasons Why the Segmentation Status Quo Heightens Risk	6
A Bottom-up, Tactical Approach to Access Control	8
Trust Valuations Tend to Be Static	9
Access Control Means Little Without Enforcement	10
Conclusion: Key Segmentation Concerns	13

Executive Overview

An expanding and fragmenting attack surface—which results from mobility and multi-cloud adoption—is undermining the ability of network engineering and operations leaders to maintain network performance, security, reliability, and availability. Traditional network-based segmentation, in addition to even more recent microsegmentation techniques, are insufficient. Constrained by the network architecture, they are tactical rather than strategic and focused on business logic. They are also typically static, allowing once-trusted users, devices, and applications free rein in their permitted segments. Finally, they lack comprehensive security visibility, across the network and into encrypted flows, which is essential to effective risk management.

Introduction: Difficulty Managing Disparate Networks: Is Segmentation the Answer?

The user base on the typical corporate network is increasingly geographically dispersed, as are the devices and applications connecting to corporate IT resources. As corporate networks have accommodated mobile and Internet-of-Things (IoT) technologies and adopted Software-as-a-Service (SaaS) applications in multiple public clouds, their attack surfaces have become increasingly difficult to protect, even with strong perimeter security.

One challenge with these expanding, fragmenting attack surfaces is that they create an array of new paths through which criminals can attack. Another problem is that threats are increasingly more sophisticated, automatically seeking and taking advantage of any vulnerabilities. Further complicating the situation is the fact that merger and acquisition (M&A) activity may result in a diverse infrastructure with limited coordination or visibility between different parts of the organization. In many organizations, security has become a reactive exercise, because IT is unable to prevent lateral movement of intrusions across the devices and applications connected to and traversing the network.

For years, network engineering and operations leaders have responded to these challenges by segmenting their networks. Traditional segmentation techniques based on IP addresses have been augmented with VLAN segmentation and VMware NSX segmentation for virtualized workloads. Networks based on Cisco gear rely on Cisco ACI segmentation using physical switches and VXLANs. These microsegmentation techniques enable access control policies to be defined by workloads, applications, or by architectural attributes such as the virtual machines (VM) on which the applications, data, and operating systems reside.

In these segmentation approaches, firewalls are used to separate the network resources for each group. This prohibits any traffic that is not authorized from moving between segments. Thus, when an attack breaches network security in one area, this approach should prevent the spread of attacks laterally to other areas of the network—in theory.

Unfortunately, microsegmentation is not the panacea that it is sometimes hailed to be. The underlying ideas make sense, but if a network infrastructure that entails microsegmentation is not designed properly, it might actually hinder security. Dividing a complex corporate network into a large number of small segments may limit visibility into threats and attack mitigation activities across the network.¹

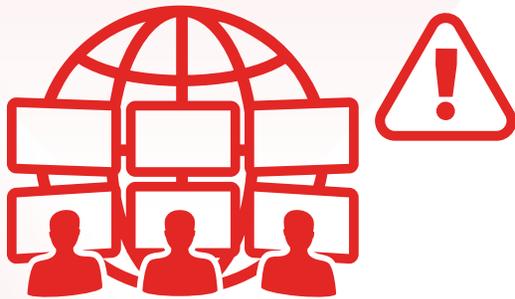
Dividing a complex corporate network into a large number of small segments may limit visibility into threats and attack mitigation activities across the network.²

3 Reasons Why the Segmentation Status Quo Heightens Risk

There are three primary problems with current segmentation techniques:

1. Access control for internal network segments is designed from the architecture up, a tactical approach that cannot easily adapt to changing business needs.
2. The trust valuations on which access policies are based tend to be static and become quickly outdated.
3. Access control policies cannot be effectively enforced due to a lack of advanced (Layer 7) security components from the data center to the network edge and moreover are unable to see and control these components efficiently.

These problems often stem from the fact that network engineering and operations staff plan the segmentation architecture without adequate attention to security. Understanding each of these issues and their aggregate impact can lead to a more risk-wise approach to segmentation.



“All too often, the network is designed without considering security design and its operations. IT teams fail to make the security blueprint part of the network blueprint when, in fact, the two go hand-in-hand. As a result, the two function as leader and follower rather than as equal partners in the IT helix. This misalignment becomes multiplied in highly segmented, complex network environments.”³

A Bottom-up, Tactical Approach to Access Control

Presumably, the design of the corporate network is dictated by the needs of the organization as it evolves. The rules governing who and what can access which network resources are determined by business policies, industry standards, and government regulations. Following these rules, the network operations team configures the access control settings in the routers and switches, which permit users, devices, or applications to access specific network resources.

Network engineering and operations leaders will immediately recognize two downsides to this approach. First, the business processes, compliance requirements, and network access needs of an organization are vastly more complex than the structure of its network. Consequently, it is very difficult to use the network architecture to define secure segments for network resources that will be simultaneously accessible to all authorized users and applications and completely inaccessible to all others. In practice, there will be security gaps—access scenarios that the network architects did not envision—which bad actors can take advantage of. With advanced, sophisticated malware, they are doing so already.

Second, any process, regulation, or organizational structure is liable to change. So, even if the optimally secure network design were achieved, it would have to be amended. Once again, there are numerous opportunities for security gaps, not to mention the time and cost involved in the reconfiguration, which few networking teams can afford.

It is very difficult to use the network architecture to define secure segments. In practice, there will be security gaps—access scenarios that the network architects did not envision—which bad actors can take advantage of.

Trust Valuations Tend to Be Static

To effectively manage risk, network engineering and operations leaders need to have current and accurate information on the trustworthiness of users, applications, and network assets. Their internal firewalls or other access control mechanisms that enable or prohibit traffic flow between network segments must always be working off of up-to-date trust data. If trust assessments are out of date, the segmentation technologies become useless at preventing potential threats from moving laterally through the network.

The quality of trust data is becoming a pressing issue in network segmentation security because the actual trustworthiness of network resources can change unexpectedly. Indeed, numerous organizations have been surprised by attacks from within the ranks of their trusted employees and contractors. More than one-third of reported breaches involve internal users, and 29% involve stolen credentials.⁴

Some organizations have responded to these dangers by practically locking down their networks, trusting no user or application and creating layers of verification before permitting access. Network engineering and operations leaders must protect sensitive assets, but without causing unnecessary burdens for those who legitimately require access to those assets.

“Trust is not absolute, binary, or static. It is an indication of the relative level of strength of the assurance of the belief. Further, the level of trust is dynamic and changes over time. Thus, access to the capabilities should be adapted.”⁵

Access Control Means Little Without Enforcement

Access control policies cannot work as expected if the network is missing key elements of an effective security infrastructure. Traditional approaches to segmentation assume that all the necessary network security components are in place to execute whatever access control policies the IT team defines. However, this assumption may not hold, for several reasons.

Total cost of ownership (TCO) is a major reason why organizations may not have ubiquitous advanced security.

For example, a network engineering and operations team driving segmentation may decide that some network segments with smaller attack surfaces are adequately protected without Level 7 Advanced Security enforcement. Due to budgetary reasons or simply because deployment and management requires too many resources, network engineering and operations teams may hesitate to deploy next-generation firewalls (NGFWs) and other advanced threat-protection solutions everywhere they are needed—within the enterprise, in every cloud in which they operate, and at every endpoint and IoT device.

The security components that are in place may not be fully functional. Some network teams might intentionally turn off secure sockets layer (SSL)/transport layer security (TLS) inspection in their NGFWs to optimize network performance. Handicapping security solutions in this way may help legitimate traffic move between network segments more quickly, but it opens the door to illegitimate traffic at the same time. And with 72% of network traffic now encrypted and cyber criminals leveraging it to infiltrate networks and exfiltrate data, this is a serious concern.⁶

The overall effectiveness of the security components is reduced if they are not tightly integrated. Lack of integration has several implications. First, when one firewall detects a suspicious packet, it may take several hours or longer until the information is picked up by the security administrator and disseminated to the rest of the network.

Second, disparate security solutions cannot easily share threat intelligence, neither globally acquired intelligence on known and emerging threats nor zero-day threat intelligence regarding newly discovered threats. It may be one reason why the mean time to identify a breach remains high, at 197 days.⁷

Third, organizations cannot respond effectively to mitigate the impact of breaches that are detected. Without integrated sandboxing technology to automatically quarantine and test all suspicious packets, extensive damage can occur by the time the security team deals with the threat manually.

Under these conditions, network engineering and operations leaders who believe their segmented network is well-protected may be working under a false sense of security. An ongoing end-to-end security assessment would tell them how their security platform is performing and whether their access control policies are achieving their business intent. Unfortunately, without breadth of security and end-to-end visibility, a reliable assessment is not possible, preventing many network engineering and operations leaders from reporting accurately on their company's security posture.



Disparate security solutions cannot easily share threat intelligence on known or newly discovered threats. It may be one reason why the mean time to identify a breach remains high, at 197 days.⁸

Conclusion: Key Segmentation Concerns

Network segmentation is necessary, but the status quo is insufficient. Companies that do not incorporate dynamic evaluations of trust into access control between segments leave their users and assets vulnerable. Networks in which the segmentation architecture constrains business intent do not support progress toward organizational goals. At the same time, if performance priorities supersede security concerns, segmentation may result in reactive and ineffective threat mitigation. And networks lacking adequate visibility into the security posture may not incorporate the Layer 7 security that is crucial to preventing advanced threats.

It is up to network engineering and operations leaders to ensure that access control policies for internal network segments are adequate in this era of perpetually expanding and fragmenting attack surfaces. Only with careful attention to segmentation design can a company be confident in its ability to thwart attackers looking to move laterally throughout the network.

¹ Keith Townsend, "[Get a Quick Primer on How Microsegmentation Can Improve Network Security](#)," BizTech, May 26, 2017.

² Ibid.

³ "[Friction in the IT Helix: How to Create Harmony between Network Design and Security](#)," Masergy, August 8, 2018.

⁴ "[2019 Data Breach Investigations Report](#)," Verizon, accessed July 8, 2019.

⁵ Neil MacDonald, "[Zero Trust Is an Initial Step on the Roadmap to CARTA](#)," Gartner, December 10, 2018.

⁶ John Maddison, "[More Encrypted Traffic Than Ever](#)," Fortinet, December 10, 2018.

⁷ "[2018 Cost of a Data Breach Study](#)," Ponemon, July 2018.

⁸ Ibid.



www.fortinet.com

Copyright © 2019 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.