White Paper

Removing the Complexities from Network Segmentation

How Advances in Discovery and Automation Make Segmentation Easier to Implement

By John Grady, ESG Analyst; and Jon Oltsik, Senior Principal Analyst & ESG Fellow

October 2019

This ESG White Paper was commissioned by Cisco Systems and is distributed under license from ESG.

Contents

Executive Summary

Network segmentation is a well-established best practice, but one that has traditionally been exceedingly difficult to implement. Organizations understand the need to move beyond the binary approach of recognizing everything inside the perimeter as good and everything outside as bad, and work toward a zero-trust or least-privilege approach. However, the intensely manual processes of configuring VLANs, access control lists (ACLs), and firewalls to accomplish this has forced many organizations to abandon segmentation projects before ever completing them.

Vendors are introducing significant innovations that will simplify network segmentation by improving automation and centralizing policy management.

However, there are strong indications that these dynamics are changing. Vendors are introducing significant innovations that will simplify network segmentation by improving automation and centralizing policy management. With many of the complexities associated with network segmentation removed, organizations of all types should revisit network segmentation as a core tenet of their security strategy.

Network Segmentation Should Be a Component of Every Security Program

The Benefits of Network Segmentation

Organizations gain several advantages when network segmentation is employed. First, the attack surface is reduced. A flat network does nothing to prevent attackers from quickly exploiting their access and moving laterally once they have breached the perimeter.

Network segmentation prevents this lateral movement, limits insider threats, and can secure IoT devices that may not support on-device security capabilities by isolating them from the public internet. Ultimately, a layered network that has implemented segmentation represents a much more difficult target for attackers to navigate.

A layered network that has implemented segmentation represents a much more difficult target for attackers to navigate.

There are business benefits to employing network segmentation, as well. For example, compliance regulations such as PCI can be met at a lower cost by separating in-scope and out-of-scope systems and ensuring that customer cardholder data is isolated to a specific subnet. PCI-mandated controls would then only need to be applied to the specific network segments in question, and the auditing of data flows, vulnerability scans, penetration testing, and other requirements become much simpler by focusing on a limited subset of the environment.

Finally, performance and availability can be improved via network segmentation. Employing subnets results in fewer devices and resources on each segment. This makes it easier for network administrators to monitor and maintain quality of service for business-critical resources. From a security perspective, attacks often manifest themselves as operational problems, so reducing the noise to a limited segment of the network can enable network and security teams to quickly determine if an availability issue is malicious and react quickly if it is.

Zero-trust and least-privilege strategies have been gaining momentum of late and segmentation is one way to implement some of the core tenets of these approaches. They are not the same thing, but the implementation process and resulting network posture drive toward many of the same goals:

- Identifying and inventorying the most valuable or vulnerable corporate assets.
- Implementing context-based policy to limit resource access to only those entities that have a valid business need.
- Maintaining a consistent state of analysis and visibility to understand when new users, devices, or resources come online, regardless of location.

Heterogenous Networks Make Proper Segmentation Much More Difficult

Despite the litany of benefits and its recognition as a security best practice, network segmentation is not employed ubiquitously. Historically, there has been an inverse relationship between the need for segmentation and the ease with which it can be implemented. Organizational, business, and technological factors drive this complexity.

- Network segmentation requires cross-functional coordination between the network and security operations teams to determine project roles, ownership, and budget.
- Undertaking a strategic, preventative project can be difficult as ESG research found that 36% of organizations say that one of their biggest threat detection and response challenges is not having enough time for strategy and process improvements due to addressing high priority or emergency issues.[1]
- Data considered sensitive can include everything from employee information, customer PII and financial information, business strategy and intellectual property, to health and medical information.[2] Developing a granular understanding of all of the entities on the network, all of the organization's important data, where it resides, and who requires access makes it daunting for an enterprise to even know where to begin.

Internet of Things

Technology introduces additional complexities. IoT adoption is growing rapidly, with ESG research indicating that 67% of organizations either have IoT initiatives underway or are planning to launch them in the next 12 months.[3] In fact, the number of devices and subsequent traffic on the network are the two most cited reasons for network security complexity (see Figure 1).[4]

[1] Source: ESG Master Survey Results, *The Threat Detection and Response Landscape*, April 2019.
[2] Source: ESG Master Survey Results, *Trends in Cloud Data Security,* January 2019.
[3] Source: ESG Master Survey Results, *2019 Technology Spending Intentions Survey,* March 2019.
[4] Source: ESG Research Insights Report, *Navigating Network Security Complexity,* June 2019.

Figure 1. The Number of Devices and Amount of Traffic on the Network Have Caused Security Complexity

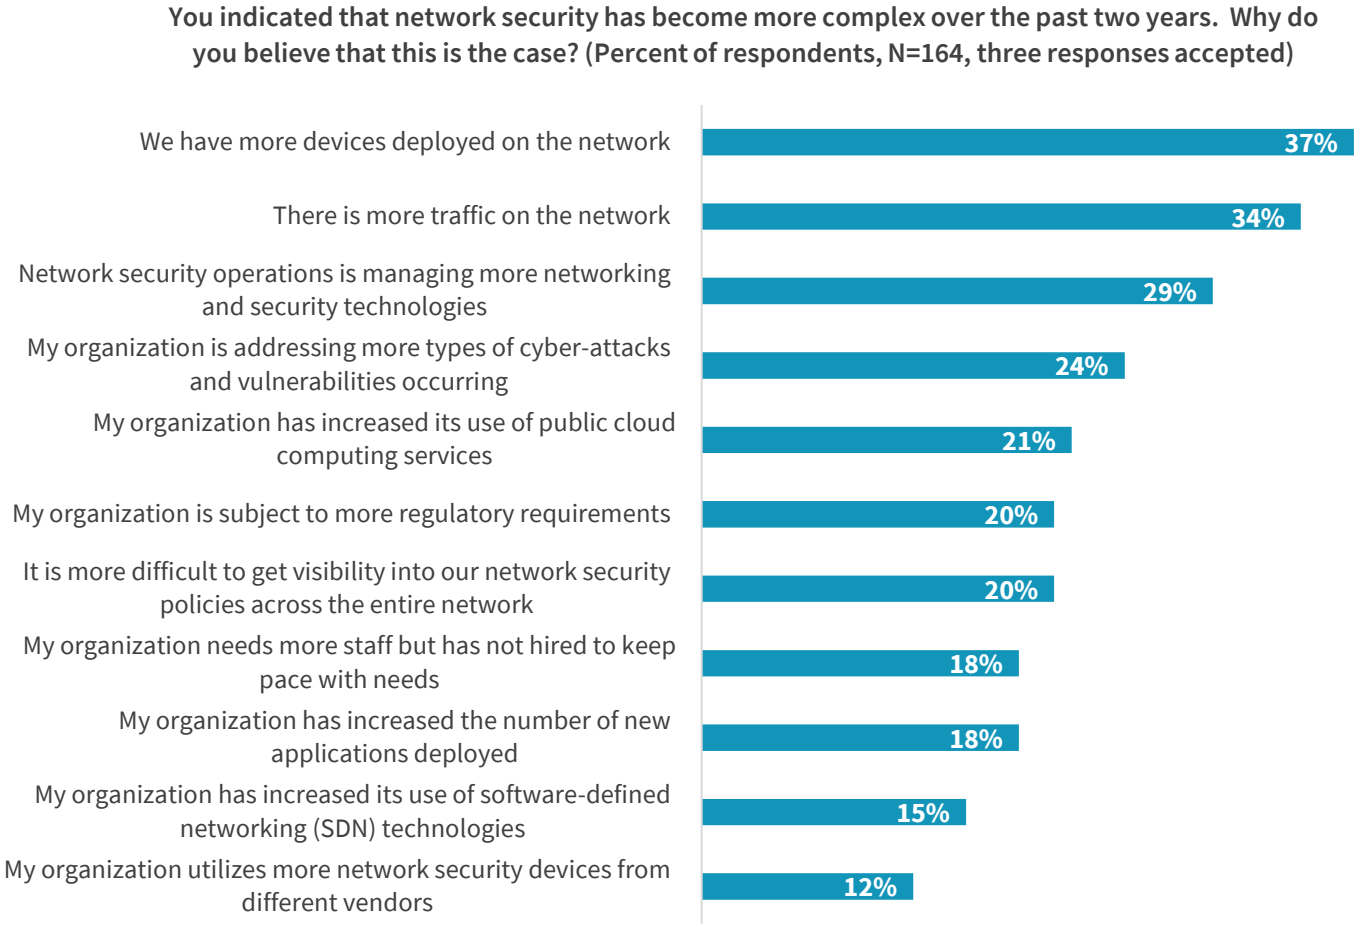

Source: Enterprise Strategy Group

Further, as both the volume and types of devices connecting to the network explodes, accurately identifying and classifying them becomes incredibly difficult. Without the accurate identification and classification of every device on the network, proper segmentation is not possible. Even when using traditional network access control solutions, organizations often find a large number of unknown devices connected to their networks due to the ever-increasing scope of IoT. Without the proper device context, it is impossible to intelligently create segmentation policy.

Cloud

Cloud adoption has complicated even the basic idea of securing the perimeter. It is well established that enterprise applications are moving to the cloud. In fact, ESG research has found that 39% of organizations now follow a "cloud-first" policy when deploying new applications, up from 29% just a year ago.[5] Yet many applications do remain on-premises due to underlying code, compliance, data residency, or any number of reasons. Even parts of the same application can reside in different locations. This becomes especially problematic relative to segmentation as on-premise segmentation technologies do not readily extend to the cloud. In a hybrid application environment, the application or web tier may be in one cloud environment while the data tier resides in a different cloud or even on-premises. In this scenario, segmenting the application and associated users and devices requires the VLAN to span multiple locations and policies to be applied consistently across different segmentation technologies and as many as three disparate environments.

[5] Source: ESG Master Survey Results, *2019 Technology Spending Intentions Survey*, March 2019.

Mobility and Remote/Branch Office Access

Furthering this issue is the fact that users are increasingly mobile and working away from the main campus. They may be at a branch or remote office location, or completely off-network. In any of these cases, this location information, as well as the type of device the employee is using, can provide important context in determining the security policy that should be applied. An employee attempting to access sensitive data from a corporate device while off-campus may be perfectly acceptable to an enterprise. However, if that employee tries to access the same data while off-network, from a personal device outside of business hours, the context may predicate blocking the action. Without granular detail about the attempted connection (employee, device, location, and application) the allow/deny decision must be made without all relevant information and risks limiting productivity or weakening security.

Even if an organization has this kind of visibility, it may not have centralized policy management. Remote offices may have an IT generalist, or no onsite support at all. Creating VLANs and firewall access rules via command-line interface (CLI) is an imperfect solution and is one example of where a network segmentation strategy can begin to break down. According to ESG research, managing security across application locations (on-premises, SaaS, or IaaS-hosted apps), visibility of user activities, lack of collaboration across network and security teams, and challenges enforcing consistent policies are the top cybersecurity challenges related to ROBOs and roaming users reported by respondents.[6] All would contribute to the complexity of a segmentation initiative using today's tools.

Manual Processes

Making things worse, manual processes make it difficult to keep pace with a dynamic network. When undertaking a network segmentation initiative, the first step organizations should take before any network rules are implemented is understanding everything that resides on the network, the relationships between users, the applications and devices, and where the greatest risk exists. Only at that point can enterprises begin the equally intensive process of configuring policies based on the business logic of the relationships and risk.

The traditional method of creating VLANs, creating ACLs, and configuring firewalls is time-intensive. ESG research has found that 80% of organizations indicate it takes hours or days to change network security policies, versus only 15% that report it takes minutes (see Figure 2).[7] Further, this approach buries the business logic into lines of code. As staff turns over and network changes occur over time, either organically or with new business units being added through acquisition, it becomes easy to lose visibility into the original purpose of the policy, making it more difficult to maintain a secure state.

[6] Source: ESG Research Insights Paper, *The Rise of Direct Internet Access (DIA): Securing Remote Users and Branch Offices*, May 2019.
[7] Source: ESG Research Insights Report, *Navigating Network Security Complexity,* June 2019.

Figure 2. Time Spent on Manual Processes Makes Segmentation Difficult

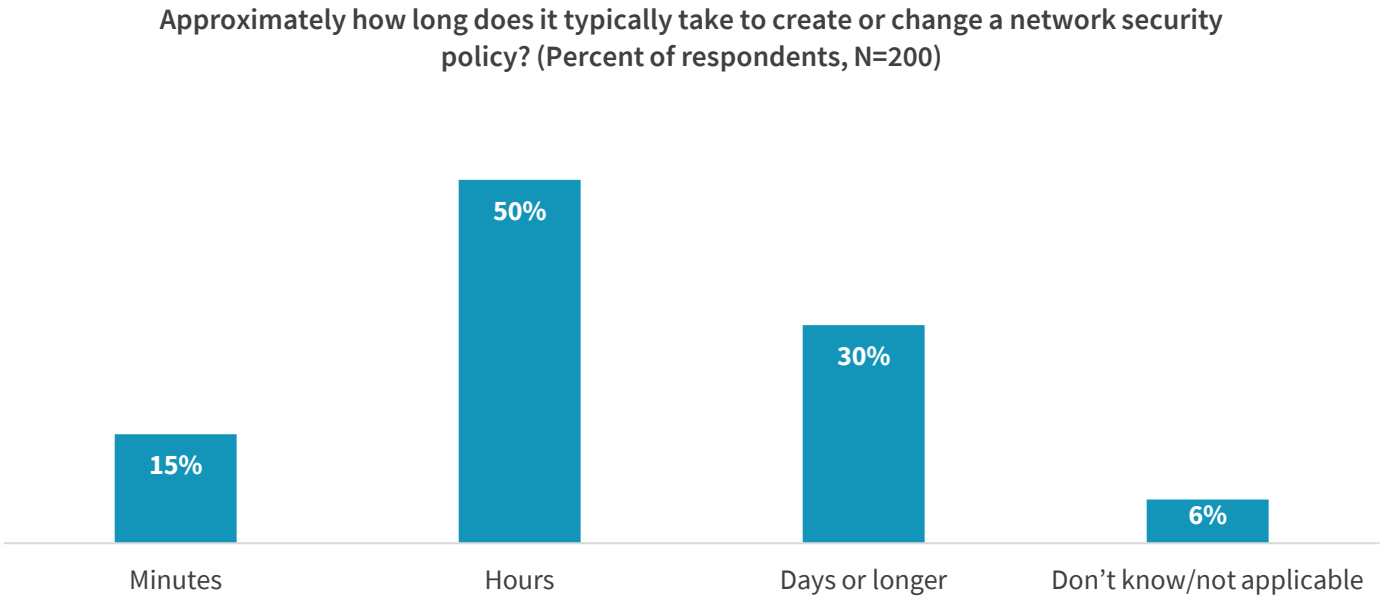

Source: Enterprise Strategy Group

As Innovations Simplify Network Segmentation, Organizations Should Reassess These Initiatives

Although the complexity issues previously described have held back segmentation initiatives, there are positive signs of change. Automation innovations that are already being incorporated into network and security technologies will make network segmentation simpler and more effective, allowing organizations that have considered segmentation strategies but not moved forward to revisit these projects.

For those that have implemented network segmentation to one extent or another, these enhancements will improve the efficiency and effectiveness of the projects. Ultimately, these innovations must support a merging of network, security, and data policy into a more holistic intent-based, business-centric approach to networking overall and segmentation specifically.

A More Automated Approach Through Intelligence and Analytics

Specifically, innovations focusing on improving visibility and automation through network intelligence and analytics will help organizations get the right things on the network, with the right privileges, quickly and efficiently. Automated device discovery and classification will identify all the entities on the network. Automated behavioral analysis will discern access requirements and recommend policies. Intuitive management GUIs will simplify confirming or adjusting policies based on business logic. Finally automated enforcement will abstract the underlying segmentation method for universal enforcement.

Profiling

Network access control (NAC) has moved well beyond simply granting guest access and running basic posture checks on devices attempting to connect to the network. Current solutions leverage improved profiling capabilities to identify not only laptops and mobile devices, but IP connected phones, printers, and security cameras. The ability to integrate with a configuration management database (CMDB) is increasingly important to help the profiler identify the device in question and maintain an up-to-date view of the environment by communicating back to the CMDB when new devices are identified.

With the number and types of IoT devices constantly increasing, though, it is difficult for operations teams to keep pace. Many solutions have templates to detect and classify IoT devices as well, but the sheer scope results in many remaining unknown. To that effect, the more advanced solutions leverage passive network monitoring and deep packet inspection (DPI) to gather telemetry and apply machine learning (ML) algorithms to identify previously unknown assets by understanding what they are accessing on the network. By using all these capabilities (standard profiling, CMDB querying, and telemetry analysis/ML), a holistic view of all devices is created.

Policy Recommendations

Even after all the devices on the network are accounted for, the process of connecting the dots between those devices, users, and applications to create a segmentation policy has typically been a significant undertaking. However, additional advances are removing many of the hurdles with this part of the process as well. Network management solutions have started to shift away from using IP addresses as the identifiers for network entities.

Once the identity of the users and applications are understood, it becomes much easier to create groups based on business logic and then apply policies specific to those groups. Further, once the data is normalized, the network management solution can incorporate analytics to apply real-world scenarios and recommend policies based on similar actions seen elsewhere.

Once the data is normalized, the network management solution can incorporate analytics to apply real-world scenarios and recommend policies based on similar actions seen elsewhere.

For example, if it is understood who the users are within the human resources organizations, it is fairly straightforward to map them to typical HR-specific applications such as Workday or Kronos and, with network analytics then being used to fill in the gaps based on traffic flows that show what else users from that group may be accessing on a regular basis, recommend a VLAN policy based on those factors. Similarly, with the advanced profiling features being introduced, IoT medical devices could be automatically identified as such with a segmentation policy recommended that segregates the devices from financial systems. Even after the initial rollout, continuous monitoring is becoming a prerequisite to ensure that, as the network changes, policies can be updated as required. If an application shifts to the cloud, an existing policy may become problematic and a recommendation made for how it could be adjusted.

Simplified Management

This centralized, group-based policy becomes even more important as manual modifications are needed. Rather than cumbersome CLI-based management tools, more intuitive GUI-based solutions are being introduced. This provides a more graphical representation of the network to enable administrators to easily see, understand, and incorporate business context into policy creation. Rather than wasting cycles on IP address mapping and writing multiple iterations of the same policy for different users, operations teams can drag and drop a policy to a group or vice versa, saving time and maintaining consistency. This approach allows organizations to focus on the higher-level strategy of how things should be grouped based on business criteria, rather than getting into the weeds of where things are, what IP addresses should talk to each other, and so on.

Heterogeneous Technology Support and Distributed Enforcement

The final component that ties the three previous innovations together and delivers the automated capabilities that ultimately enable simpler network segmentation is universal enforcement. Specifically, the ability of the management layer to abstract the underlying business policy logic and translate it across all the disparate technologies required for segmentation across the campus, data center, and cloud. This includes all networking and security capabilities employed (switches, routers, firewalls, SDN, microsegmentation, and NAC) across all form factors (physical, virtualized, and cloud).

> The final component that ties the three previous innovations together and delivers the automated capabilities that ultimately enable simpler network segmentation is universal enforcement.

Without this write-once, enforce-everywhere capability, the model quickly breaks down.

This is especially important relative to the cloud and mobility issues described earlier. Previously, if something moved once the initial policy was set, manual updates were required. Now, users can be tied to devices, applications to devices, and with location and other context, policies can be preset based on those criteria. Because group-based policy is not based on IP address, the policy is mobile and remains consistent. For example, an on-premises application may be tagged as part of the finance group and the segmentation policy set accordingly. If that application then moves to the cloud, the policy follows, even though the location has changed.

The Bigger Truth

There are many reasons people may procrastinate or decide to put off an undesirable task. When something is difficult, boring, frustrating, or ambiguous, it becomes that much easier to set it aside for another day. Unfortunately, network segmentation has historically fallen into more than one of these categories—in fact, maybe all of them except boring. And it has only gotten harder. Network technology has become software-defined and cloud, mobility, and IoT have driven fundamental changes in the way networks are architected. Yet the process and tools for applying granular segmentation to isolate portions of the network to enhance security and improve quality of service have remained largely the same.

However, this is finally changing. Through automation and analytics, improved device identification and streamlined policy creation will greatly simplify network segmentation, even in the most dynamic environments. Regardless of whether an organization has fully implemented, started and abandoned, or never considered a network segmentation project, all would benefit from these recent innovations. Reducing the attack surface and improving security posture are two goals all businesses should have, and properly segmenting the network is a good way to achieve both.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.

 www.esg-global.com

 contact@esg-global.com

 P. 508.482.0188