tufin The Security Policy Company.

Define the Segments. Control the Network.

How to simplify segmentation in complex hybrid environments

Technology never really sits still. What starts out small has a way of growing faster than anyone expects. Think about the internet. At first, it was just a handful of addresses, simple enough to scribble on a napkin. Before long, it became the backbone of global commerce and communication.

Enterprise networks have followed the same pattern. Not long ago, a security engineer could trace most traffic through a familiar set of devices and environments. The topology might have been large, but it was still predictable.

Today, those same engineers are contending with sprawling, interconnected systems that span data centers, multiple clouds, and remote infrastructure. The attack surface has expanded, dependencies are layered across environments, and lateral movement is harder to contain. That reality makes consistent policy enforcement and granular segmentation essential.

For this reason, CISOs and their security teams are under increasing pressure to implement Zero Trust strategies. The success of that initiative hinges on one foundational challenge: making segmentation work across hybrid networks.

In the following sections, we will demonstrate how organizations can approach segmentation in complex hybrid environments with a practical, repeatable model that reduces risk, accelerates implementation, and ensures long-term manageability.

Zero Trust? Not without a solid segmentation strategy

Network segmentation forms the foundational principle of Zero Trust architecture by preventing lateral movement within compromised networks and minimizing the “blast radius” of successful breaches.

The data is sobering: over 70% of successful breaches leverage lateral movement techniques, including the spread of ransomware and malware across enterprise networks[1]. This reveals a critical gap in traditional security models. Once an attacker gains initial access, segmentation becomes the essential last line of defense.

Over 70% of successful breaches leverage lateral movement techniques, underscoring the need for deep segmentation controls.

The threat landscape reinforces this imperative. Networks themselves were a direct attack front in 58% of incidents responded to by Palo Alto Networks’ Unit 42[2], highlighting the vulnerability of internal infrastructure. While initial access vectors vary, with phishing accounting for approximately a quarter of incidents and software or API vulnerabilities favored by nation-state actors, attackers are highly successful at spreading once inside. Without strong internal controls, perimeter security offers little resistance.

Regulatory standards now make segmentation a requirement rather than a best practice. The Payment Card Industry Data Security Standard (PCI DSS) v4.0, effective April 1, 2024, explicitly requires organizations to "Install and maintain network security controls" (Requirement 1) and "Restrict access to system components and cardholder data by business need-to-know" (Requirement 7)[3].

Most significantly, PCI DSS 4.0 calls for zero trust and segmentation to prevent lateral movement and to control connections between trusted and untrusted networks[3]. Other global frameworks, including ISO 27001, NIST CSF, SOC 2, HIPAA, and NERC CIP, incorporate similar requirements that push organizations toward architectural changes instead of surface-level compliance.

> **"Segmentation is the control layer that makes Zero Trust operational."**

Human error remains the silent amplifier of these risks. Research indicates that 95% of all data breaches are caused by human error, with insider threats, credential misuse, and missteps accounting for the majority of incidents[4]. Analysis from the InfoSec Institute confirms that 74% of breaches involve a human element, from privilege misuse to social engineering[5]. Complex, manually managed security policies magnify these risks, increasing the chance of misinterpretation and misconfiguration.

Segmentation mitigates this by enforcing the principle of least privilege automatically, reducing opportunities for mistakes to become vulnerabilities. These pressures set the stage for the real obstacle, which is translating segmentation from policy into day-to-day operations across diverse, constantly changing environments.

Networks weren't really designed to be governed this way

It sounds straightforward. Shift to an approach that assumes breach and enforces continuous verification. But as anyone on the Zero Trust journey will tell you, implementing segmentation in complex hybrid networks is a multi-step process that cannot be done overnight.

In a greenfield environment, segmentation can be designed into the network from the start. In a brownfield or hybrid environment, it means working within a complex mix of legacy systems, overlapping vendor technologies, and shifting architectures that resist simple solutions. These conditions make segmentation one of the most difficult, but also most critical, parts of the Zero Trust journey.

The reality of implementing segmentation in hybrid environments presents overwhelming operational challenges. The complexity is compounded by significant security tool sprawl across enterprise environments. According to Secureframe's cybersecurity research, enterprises typically deploy an average of 53 security solutions, with other industry estimates placing this number between 70 and 90 products for the average enterprise[6].

The cybersecurity market itself is highly fragmented, comprising more than 3,700 vendors selling over 8,000 products, many of which are point solutions designed for specific vulnerabilities[6]. According to Silicon Angle's analysis, this proliferation of tools is expected to worsen, with cybersecurity tool sprawl becoming an increasingly critical challenge for enterprise security teams[7]. While each tool aims to address a specific security need, their sheer volume and lack of native integration create significant operational overhead, leading engineers to spend more time managing the tools themselves than on proactive security.

Case Study: Johnson Controls (2023–2024 Ransomware Incident)

In late 2023, Johnson Controls suffered a ransomware attack that resulted in the theft of 27 terabytes of data. The root cause, identified in post-breach analysis, was a lack of internal network segmentation. Once inside, attackers were able to move laterally across the network without encountering meaningful boundaries. With proper segmentation in place, the breach could have been contained to a single segment. Instead, the attackers traversed systems across the enterprise, turning a localized compromise into a full-scale data exfiltration event.

The most paralyzing aspect of hybrid segmentation is the protracted time required for network change implementation when changes are handled manually. This administrative burden creates a cascading effect where fear of outages due to potential misconfigurations leads to extensive manual validation processes and multiple approval layers. The concern about changes that could "damage your reputation and revenue" drives organizations toward increasingly conservative change management practices that prioritize safety over agility.

The longest lead times for network changes are often administrative and procedural, not technical, meaning that technical expertise is frequently stifled by organizational inertia. This environment breeds "policy drift," where security policies become inconsistent across the network due to manual processes and siloed teams. The constant concern about the impact of changes creates a cautious, often paralyzing approach to network evolution that inhibits both security improvements and business agility.

tufin

Case Study: Major AWS Breach (Cloud Misconfiguration, 2024)

In early 2024, a cloud misconfiguration involving exposed environment variables (.env files) led to a breach affecting more than 230 million records. The attackers exploited insufficient segmentation across AWS domains to move laterally, deploy malicious Lambda functions, and launch phishing campaigns. With stronger segmentation controls like VPC peering restrictions and well-defined security groups, the damage could have been isolated to a single service.

Reigning in complexity. Start with a shared visual.

Addressing these challenges requires a platform that gives teams visibility across all environments, the ability to define and enforce consistent policy, and the automation to keep pace with change. Tufin helps engineers move from manual, fragmented policy work to an organized, scalable segmentation strategy built for hybrid networks. Instead of configuring devices one at a time or relying on spreadsheets to track network boundaries, teams can visualize, design, and enforce segmentation using a centralized platform that supports both on-prem and cloud environments. Every part of the process, from asset classification to policy enforcement, is supported by tools that reduce manual work and help teams make confident decisions faster.

Tufin gives teams the ability to see their entire environment clearly, map segments accurately, and enforce policy based on how the network operates. This helps reduce exposure, clean up legacy access, and bring distributed networks under a single policy model.

Visualizing and organizing complex networks

Engineers start by defining segmentation segments using data pulled from routers, firewalls, and cloud platforms. These segments form the foundation for segmentation and allow teams to group infrastructure into meaningful categories such as business units, applications, or shared services. Segments can be assigned manually or automatically using integrations with IP address management (IPAM) systems.

Once segments are defined, teams can build hierarchies that reflect how the business is structured. For example, segments for site A and site B may be grouped under a broader application or region-level tier to simplify policy enforcement across multiple locations.

In cloud environments, tags help streamline segmentation at scale. Even when consistent tagging is not yet in place, strong naming conventions can be used to generate tags automatically. This enables security teams to apply consistent policy across cloud environments using metadata that already exists.

Ask Yourself: Are your segments built around how your network operates or how your documentation is organized?

Mapping assets to segments with accuracy

Many organizations enter their segmentation journey with a general understanding of what segments they want to build and how those segments should communicate. But the challenge begins when they try to assign actual subnets to those segments. Most organizations see the network association to the segments as the biggest challenge. Poor documentation, dynamic subnet growth due to M&A or scaling, and inconsistent IPAM data make it nearly impossible to maintain an accurate mapping manually. Without automation or real-time context, they either misclassify subnets or delay segmentation efforts altogether.

Defining and refining segmentation policy

Even when initial guidelines exist, execution can stall at the policy layer. Most of the real pain surfaces across three stages: creating segments, assigning subnets, and defining traffic rules. In practice, defining traffic between segments often comes down to trial-and-error without visibility into how assets are used. This makes it difficult to protect crown jewel systems without either over-permitting or disrupting legitimate flows.

After segments are defined, organizations need a structured way to describe how they can communicate. Tufin provides this through a Unified Security Policy (USP), typically represented as a matrix, where each row and column corresponds to a defined segment.

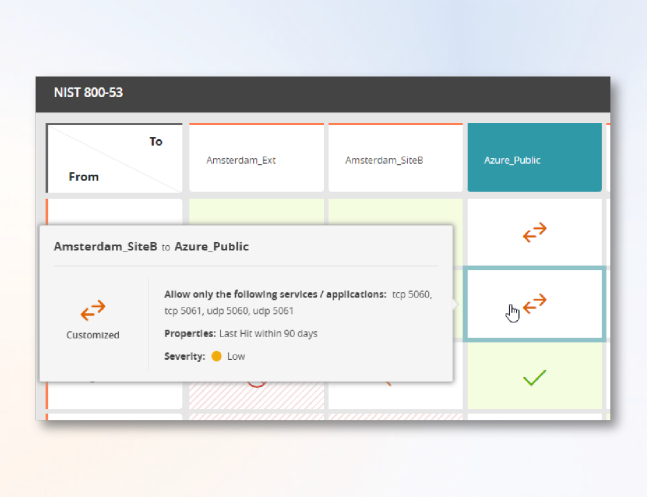

At the intersection, each cell captures decisions about which services are allowed or restricted between those segments. For example, a USP matrix may specify that Sales can access the Data Center over SSH and HTTPS. Beyond traffic rules, the USP can also enforce rule hygiene standards, such as requiring logging, comments, and avoiding the use of wildcards.

tufin

What makes segmentation work:

- You know what you're protecting
- You group systems by purpose, sensitivity, or usage
- You define which systems can talk and under what conditions
- You enforce those rules across firewalls, cloud, and infrastructure
- You adapt as things change

Tufin compares these policies to what exists in production to help engineers identify mismatches. If a policy allows HTTPS between two segments but only SSH is actually used, the system highlights the difference. If a firewall rule exists that isn't covered by any security policy, it becomes visible for investigation. These comparisons make it easier to clean up policies and align enforcement with intent.

Start with one use case

Your first use case becomes your blueprint. It shows you what good looks like.

1. Choose a system or application you understand deeply.
2. Document what should connect and what shouldn't.
3. Apply segmentation there first. Then scale it.

Alignment is easier when every team sees the same thing

Segmentation works best when all teams involved in network, cloud, and security operations are working from a shared model. In most cases, separate teams manage on-prem firewalls, public cloud controls, SD-WAN, and SASE platforms. And each of these platforms have their own tools and change processes. As environments grow more complex, coordination becomes harder. Segmentation policies may be defined in one system but enforced inconsistently across others, which increases the likelihood of overlap, drift, or gaps.

tufin

Establishing a Unified Security Policy makes it easier to coordinate across teams without disrupting existing responsibilities. Tufin supports this by providing a centralized control plane and a shared abstraction layer. Tufin's Unified Security Policy gives NetSec, CloudOps, and DevOps teams a common language for how segmentation should be defined, implemented, and maintained (regardless of where it's enforced).

This approach helps organizations maintain consistency even as the environment evolves. Each team can continue managing their own tools, but the underlying segmentation policy remains aligned across the infrastructure. By focusing on visibility, collaboration, and shared ownership, enterprises can scale segmentation in a sustainable and reliable way.

Fear not. AI-assisted segmentation is coming

AI may soon play a central role in reducing one of the most time-consuming aspects of segmentation. Mapping subnets to segments is a foundational step, but in large or hybrid environments, it can take months of manual work and coordination. An AI-assisted approach would reduce this workload, improve consistency, and help teams keep pace as networks evolve.

Machine learning models can be trained on existing subnet-to-segment relationships. By analyzing connectivity patterns and usage across security rules, the model can learn how different segments behave. It could then examine unassociated objects and recommend which segment they belong to. Each recommendation could be scored based on confidence, giving administrators clear guidance on whether to approve or escalate.

What's Next: AI will learn from network patterns to assign segments

Once trained, this model will support both initial setup and long-term maintenance. Teams would no longer need to research every object from scratch. Instead, they could validate AI-generated suggestions and focus their time on high-impact policy work. This approach would also support ongoing alignment by flagging gaps or inconsistencies as networks change.

AI can help identify poor segmentation practices, highlight ambiguous or overlapping segments, and uncover opportunities to create new segments based on observed behavior. It can also support ongoing accuracy by monitoring network changes and recommending updates to Unified Security Policies. Over time, this would help teams maintain clean segmentation without relying on ad hoc cleanup projects.

As networks continue to grow in size and complexity, AI will help teams make more confident decisions at scale. The result is a more sustainable segmentation strategy that can keep pace with operational change.

tufin

A clear path forward: visibility, policy alignment, and sustainable control

Segmentation has become essential to modern security operations. It supports the containment of threats, simplifies compliance, and provides the enforcement layer required to operationalize Zero Trust. When segmentation fails, it is usually because visibility is limited, policy ownership is unclear, or infrastructure is too fragmented to manage through a consistent model.

The research shows that success comes from applying a structured approach. Teams start by mapping their environments, grouping assets by logical segments, and applying broad controls to reduce risk quickly. Over time, they layer on more precision, introduce automation to maintain alignment, and adopt a shared policy model that allows different teams to work from the same foundation.

As the scale and complexity of networks continue to grow, organizations that combine clear segmentation strategy with intelligent automation will be best positioned to adapt quickly, contain threats, and maintain consistent control.

Learn more at tufin.com