Tempered

Five common micro-segmentation mistakes and how to avoid them

You don't need to compromise security or budget to protect critical infrastructure

This guide is designed to help you avoid the most common pitfalls organizations make when implementing micro-segmentation.

Introduction

The benefits of micro-segmentation have already been established by industry experts, government frameworks, and regulatory agencies like PCI DSS and HIPAA. The symptoms of failed micro-segmentation initiatives are nearly identical in every case. Regardless of the type of business or industry, here's what we've heard:

- ✓ The approach is too costly
- ✓ The configurations are too complex to maintain
- ✓ Wholesale changes of existing infrastructure are required
- ✓ Significant security and compliance gaps still exist

While the concepts behind micro-segmentation are straightforward, determining the best approach for achieving zero trust micro-segmentation has been elusive, complex, and confusing.

Until recently, restricting access could only be accomplished by enforcing some combination of IP addresses, ports, VLANs, tags, protocols, and certifications. Nearly any combination of these has proven to be complex and unsustainable at scale.

With the objective of locking down access to protect critical assets, we see companies inadvertently preventing connectivity and mobility of devices, workloads, and data. Employees need access to the tools to execute, so this can create frustration and interrupt business as usual.

IT are often told they must do some combination of the following:

- ✓ Reassign IP addresses to resources
- ✓ Modify routing Access Control Lists (ACLs)
- ✓ Provision private Access Point Names (APNs)
- ✓ Update, create, and maintain new nested firewall ACLs
- ✓ Tightly couple groups to extend Virtual Local Area Networks (VLANs)
- ✓ Distribute, manage, and revoke security certificates
- ✓ Manage hundreds, if not thousands, of IPsec tunnels

No combination of these actions creates a simple micro-segmented path for authenticated, authorized, and encrypted network connections between distributed devices. Not only are these methods complex, its interlocking dependencies are also slow to implement and unmanageable at scale.

Here are common mistakes made when attempting to micro-segment network resources:

mistake #1

Not understanding the difference between micro-segmentation and segmentation

Organizations often assume that micro-segmentation is possible with traditional network segmentation techniques. There is in fact a significant difference between the two approaches in isolating network resources, and they are not interchangeable.

Network segmentation is the practice of creating sub-networks or zones within the overall network. This includes isolating a particular type of data or traffic, such as payment data as prescribed by PCI DSS, preventing an attacker from moving laterally once inside the network perimeter or increasing system performance by efficiently routing data.

The new zones are often created based on geographic region or existing network tiers – such as data, applications, or network. Network segmentation is widely considered a north-south network traffic control. Once inside a designated zone of the network all communication, software applications, and users are trusted.

Typically, organizations build network segments via VLANs or firewalls, using up to a thousand course-grained policies to control each segment. Managing security based on network characteristics is an ineffective approach to isolation given today's extensive use of public cloud and container environments.

The rise of Software-Defined Networks (SDN) and network virtualization have paved the way for micro-segmentation. This security technique enables fine-grained policies to be assigned to isolated workloads within a data, center, or edge deployment.

Micro-segmentation is often called a "zero-trust model" of virtualized security, meaning that only necessary actions and connections are specifically enabled in a workload or application and everything else is blocked. It reduces the network attack surface by limiting east-west communication through the application of very granular security controls. Furthermore, it creates a Software-Defined Perimeter (SDP), regardless of whether it involves a virtual machine (VM), container, or function.

mistake #2

Using VLANs for micro-segmentation

VLANs serve an important role in networking, breaking up a flat network into smaller segments that can be managed and secured more easily. However, attempting to use VLANs to create a zero-trust security environment is too complex. There are VLAN vulnerabilities to overcome, as configurations are a challenge to maintain and manage.

Creating several dozen VLANs for isolation or segmentation purposes on a few switches are a standard practice, but creating and maintaining VLANs across hundreds or thousands of distributed physical and virtual switches are not. Even the promise of SDN hasn't solved the VLAN management challenge. One customer said it best after an abandoned SDN project: "All we did was transfer the VLAN problem from one domain to another. It was still too complex and prohibitive."

This trade-off pattern often repeats itself among traditional segmentation strategies. Organizations are forced to trade simplicity for security or vice versa. Either way, it can come at tremendous cost or major risk to an organization.

mistake #3

Using ACLs for micro-segmentation

Like VLANs, Access Control Lists (ACLs) are not a path to simple or effective micro-segmentation. Applying ACLs consistently and effectively is even more complex and fragile than VLANs because of so many cascading dependencies.

To properly apply ACLs, security experts recommend a network engineer:

- ✓ Understands all the data flows
- ✓ Knows the appropriate placement of ACLs within a rule
- ✓ Ensures they aren't creating redundant rules or objects that accidently create downtime or open security holes in the network
- ✓ Adheres to a common object naming convention
- ✓ Knows how to apply ACLs, depending on whether it's a stateful device like a a stateful device or stateless device

Not knowing how to properly apply ACLs or doing any of this incorrectly can open significant security holes. It's not uncommon for a large enterprise to have more than 5,000 firewalls and two times the number of switches and routers with tens of thousands of rules. This ends up becoming an impractical and unsustainable burden for any IT team.

The biggest mistake in equating ACLs with micro-segmentation is that the base enforcement attribute is dependent on using non-verifiable IP addresses and macro objects as the deciding factor in allowing a network connection. The lack of verifiable device identity creates complexity that ACLs do not solve.

mistake #4

Overlooking the depth and breadth of your micro-segmentation needs

Micro-segmentation is not just about servers, virtual machines, or the data center. A successful micro-segmentation architecture should encompass all the connected 'things' on your network. If not, attack vectors will be left exposed and this creates cyber risk.

Most CIOs and CISOs want a unified micro-segmentation architecture that can span all potential environments, devices, and transport options. Supporting transport options like cellular, Wi-Fi, radio, and Ethernet is extremely important to ensure organizations can micro-segment anything — from anywhere.

mistake #5

Miscalculating the degree of the human factor

There is a reason “fat fingering” is a term. Human error is perhaps the biggest problem when it comes to effective network security and micro-segmentation. Complex solutions with many dependencies create too many opportunities for inadvertent mistakes. IT teams are frequently overworked and understaffed, especially when it comes to network security. Additional staffing or expertise is not the answer. The ratio of open IT positions to qualified network and security candidates has been estimated at three to one, making new headcount not only very difficult to find, but also expensive.

Organizations can’t hire their way out of this problem, yet that’s the expectation many vendors have of their customers. Our experience tells us that any micro-segmentation strategy relying on humans and too many dependencies on the underlying network and security infrastructure is a recipe for disaster. Be wary of any solution that requires an army of consultants, specialists, or personnel to deploy and maintain.

Micro-segmentation at scale

Effective micro-segmentation can be easy to implement and cost-effective. Make your organization highly resilient and improve security posture. Learn how you can build a less complex, more secure network today. **#protectyourprivates** and keep your private assets, well, private.

Schedule a call with our experts to learn more.

experts@tempered.io | +1 206.452.5500