Zero trust microsegmentation

Visualizing security through a mission-first lens

Clarity around zero trust matters more than ever

The concept of zero trust security sounds relatively straightforward: Trust nothing and no one, continually verify the identity of every device, system, or person requesting access to a network or application resource, and ensure they're given least privileged access to it.

Despite these apparently unambiguous terms, there appears to be some confusion around what they really mean. Having a clear understanding of the key concepts behind zero trust is required to build consensus and adoption across any government organization. It is much more than a compliance effort. Adopting zero trust is essential to help ensure that government agencies—especially those with existential responsibilities—are able to consistently and reliably achieve their missions in a world where technology has become both an indispensable asset and a significant vulnerability.

One of those key concepts is **microsegmentation**. Put simply, it's how computer network resources—application servers, databases, and related supporting devices—are grouped or organized under zero trust. While the details may get a bit technical, the concept itself is actually based on looking at your network resources not through a technology-first lens but through a mission-first lens—the purpose they serve to achieve a specific outcome.

This requires a shift in thinking. But it's a shift that can help put agencies on the path to vastly improved security and far lower risk.

Why modern government is important

Government agencies in the U.S. must modernize in order to keep up with changing user needs, regulations, and health and public safety requirements. Leaders of modern governments rethink business processes and service delivery models to more effectively achieve their mission. This article is one of a series that features how modernizing affects the government workforce and the user experience, improves security and public trust, and accelerates the digital journey. KPMG team members offer insights intended to help guide governments in their modernization efforts to encompass all processes, technologies, policies, and the workforce so each works together to create connected, powered, and trusted organizations.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Visualizing the problem

While hackers may immediately come to mind as the "bad guys" when discussing zero trust, the real villain in this story is the complexity of the typical network environment. It doesn't take too many years for almost any network to become an extraordinarily complex bowl of spaghetti. Devices are added. Applications and services are developed. APIs are implemented. Ports are opened. Routes are added.

On the flip side, rarely are any of these undone when an application or service is changed or decommissioned. It's more the rule than the exception that unused ports on a device are left open or that a firewall or router will contain thousands of "allow" rules that no one remembers or understands why they were added. Arbitrarily closing or removing any of these can often lead to sudden failures of mission-critical processes or systems, and so an "if it ain't broke don't fix it" philosophy is generally applied. The spaghetti continues to grow over time, and along with it, the number of vulnerabilities that an intruder can exploit.

For many people, zero trust security equates to identity and access management (IAM)—passwords, certificates, multifactor authentication and so on. In their minds, it may primarily be about people, and so it doesn't matter how complex their network environment has become. They may believe that the zero trust mandate doesn't have to apply to literally *every* device or system as long as those devices or systems are inside a network ostensibly protected from the outside world by perimeter defenses such as firewalls. In this traditional model of network security, every device on a subnet or virtual local area network (VLAN) can implicitly trust every other device on the same subnet or VLAN.

Yet this is antithetical to zero trust's literal mandate of "trust nothing and no one." These implicitly trusted intra-VLAN connections are often the vectors by which intruders gain access to critical resources. An intruder will exploit a vulnerability to gain access to one device, and then once inside the perimeter defenses, can hop freely from device to device—sometimes across VLANs or subnets or even entire networks—until the pot of gold is discovered.

To help our clients visualize the magnitude of the problem, we work with one of the leading technology providers in the zero trust space: Illumio. **The Illumio solution begins by scanning every port on every device, detecting all traffic flows and protocols used where one device connects to others. It then produces an interactive visual map showing each of these connections.**

The map, as you might imagine, can actually look like a bowl of spaghetti, but thankfully it enables users to drill down to examine each connection in detail. You can see, for example, that "these two servers, which are used for this application, are talking to each other on this specific port using this specific protocol for this purpose." The servers may be in the same VLAN, in different VLANs within the same datacenter, across different datacenters, even across agencies.

The results are often eye opening for our clients. It's not uncommon to discover that servers believed to be tucked safely behind a firewall are actually exposed to the internet, or that multiple ports are open on a server for no apparent reason. In one case, an agency could see that thousands of devices were using Microsoft Active Directory servers that it was sure had been shut down more than a year earlier. Considering these servers were no longer being patched, what vulnerabilities might they have had?

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

A different approach to network segmentation

VLANs have historically been the way large networks are segmented into smaller, more manageable zones. Routers and switches are used to make the physical connections and define the logical groupings based on each device's IP address, even if they span different physical networks.

VLANs afford some protection from intrusion or attack by isolating one VLAN from another, but their primary purpose has always been to simplify network management and reduce traffic volume—not to enhance security. There can be a relatively large number of devices within a single VLAN, with devices used for different applications. It's also not at all unusual for a device in one VLAN to be connected to a device in another VLAN, or to a device in a different organization or to the entire internet, for that matter—hence the spaghetti that network mapping often reveals. The granularity of VLANs and thus the protection they afford are limited.

Zero trust takes a fundamentally different approach to segmenting a network: by workload—that is, by all of the computing resources required to run a particular application or service or to achieve a specific mission or "business" outcome, independent of whatever VLANs they may be part of. There are almost surely to be far fewer network resources involved in a single workload than would be found on a VLAN, and therefore these "microsegments" offer far more granular divisions, with each microsegment dedicated to a single purpose.

This is where having an in-depth understanding of the agency's operations is critical to help develop zero trust capabilities that support the agency's mission at the highest level. **How network resources are technically or physically connected doesn't matter as much under zero trust; what matters is the purpose they serve—why they are connected to achieve a specific outcome—and the broader agency objectives and priorities, compliance and regulatory mandates, data sensitivity, and so on tied to that purpose.**

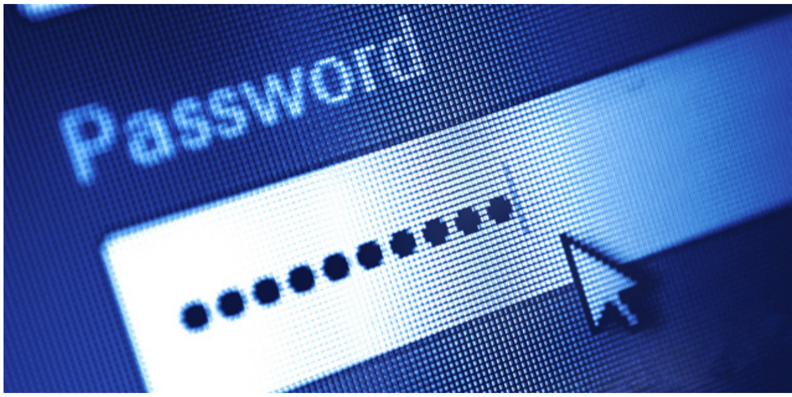

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

One step at a time

Reimagining your entire network as a series of microsegments can sound like—and actually be—a daunting task, especially after viewing one of the Illumio "spaghetti" maps. The good news is that it doesn't have to be an all-or-nothing proposition. You can score some quick wins simply by uncovering and patching critical vulnerabilities that may exist in your network just with an Illumio mapping scan. You can also start small by creating just a single microsegment, for example, perhaps to protect your most valuable application, and then grow from there.

The same Illumio solution that enables our clients to visualize the magnitude of their connection challenges also helps to define microsegments by identifying the data flows that a specific application or workload requires. It can harvest metadata from a configuration management database (CMDB) or similar source to create labels that better identify each connection and build a rich map of all upstream and downstream dependencies.

Armed with this knowledge, you can then create policies defining the microsegment "borders" that describe precisely what access is allowed from what sources on which ports using which protocols.

With the Illumio solution, microsegmentation is achieved by installing a small software agent into the operating system of networked devices. These agents serve as policy enforcement endpoints (PEPs), which achieve a least-privilege policy model by denying connections by default. At the macro level, automated monitoring and enforcement help prevent infiltrations by intruders or malware. Of course, policies must be tested and validated before they're enabled—something that here, too, the Illumio solution simplifies by providing a test mode where traffic is not blocked but where actual data flows are measured against those anticipated by policies to detect where policy violations are occurring.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Reliable and sustainable zero trust

One question that may come to mind is, "Isn't this just a different way to create another bowl of spaghetti that may have its own vulnerabilities, perhaps an even larger bowl given there will be now far more segments than under a VLAN/perimeter-based security approach?"

There are some things built into the zero trust model that can help. The nature of microsegmentation means that each segment is dedicated to a single purpose rather than having devices used for different purposes or workloads mixed together on the same VLAN. The least-privilege policy model and automated monitoring and enforcement help as well.

But even those things can't automate your way to a zero trust organization and avoid creating more unnecessary complexity and vulnerabilities in the future. Before zero trust practices can be reliably and sustainably implemented at the device and system level, the problem must first be understood in its entirety. It's more complicated than just the technology side—that is, untangling the connection "spaghetti" by understanding which ports are unnecessarily open and which firewall rules are obsolete.

Whether he actually said it or not, Albert Einstein is often quoted as saying, "We can't solve problems by using the same kind of thinking we used when we created them." Something must change in the organization and in the business processes that created the spaghetti in the first place. Otherwise, the bowl will simply refill again over time.

While technology is indispensable, zero trust isn't a technology problem and so it can't be solved by technology alone. Some define it as a security model or architecture. We call it a business philosophy and an organizational transformation. You can implement all the available tools and technologies with an ideal architecture and still not get it right.

In our experience, it's rarely the technology that's the complicated part, even in government environments, with their tangled web of aging legacy systems, cloud-based solutions, and cybersecurity and compliance challenges. It's almost always the organization that's the real challenge—the "business" side.

A deep understanding of your agency, its mission, and its security and regulatory mandates and constraints are table stakes. No tool, no matter how powerful or sophisticated, can understand these things. It can't understand agency objectives and priorities. It can't understand the criticality of a system or the sensitivity of the information within. Such understanding is a prerequisite for providing a roadmap and recommendations for how zero trust can benefit your organization and advance its mission—instead of becoming a hindrance to it or something that creates unexpected and unwelcome surprises.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

How KPMG can help

KPMG LLP (KPMG) has worked with federal, state, and local governments for more than a century, so we know how agencies work. Our team understands the unique issues, pressures, and challenges you encounter in the journey to modernize. Our alliances with leading technology providers such as Illumio help give us the breadth and depth of skills and experience required to provide a wide-ranging solution. We draw on our technical acumen and government operations knowledge to offer zero trust methodologies tailored to help you overcome challenges and work with you from beginning to end to deliver the results that matter. In addition, KPMG has significant experience implementing zero trust in both the private and public sectors and can bring those experiences to every government client.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

About KPMG

KPMG has worked with federal, state, and local governments for more than a century, so we know how agencies work. Our team understands the unique issues, pressures, and challenges you encounter in the journey to modernize. We draw on our government operations knowledge to offer methodologies tailored to help you overcome these challenges and work with you from beginning to end to deliver the results that matter.

The KPMG team starts with the business issue before we determine the solution because we understand the ultimate mission. When the way people work changes, our team brings the leading training practices to make sure your employees have the right knowledge and skills. We also help your people get value out of technology while also assisting with cloud, advanced analytics, intelligent automation, and cybersecurity. Our passion is to create value, inspire trust, and help government clients deliver better experiences to workers, citizens, and communities.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Contacts

Tony Hubbard
Principal, Government Cyber Security Leader
KPMG LLP
202-486-4945
thubbard@kpmg.com

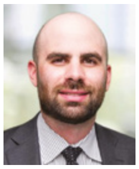

Tyler A. Carlin
Director, Advisory
KPMG LLP
571-243-5655
tcarlin@kpmg.com

Gary Barlet
Federal CTO
Illumio
571-969-4190
gary.barlet@illumio.com

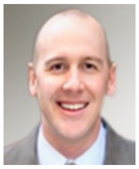

Nate Deshong
Director, Advisory
KPMG LLP
843-327-6641
ndeshong@kpmg.com

read.kpmg.us/modgov

Illumio is a cybersecurity software company using **zero trust microsegmentation** to stop the lateral movement of a cyberattack. Illumio provides real-time visibility, reduces the dynamic attack surface, and enables faster implementation using existing hardware and software infrastructure to improve cybersecurity.

kpmg.com/socialmedia

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.