White Paper: **January 2022**

Zero Trust in U.S. Federal Agencies' Private Cloud

Implications of the president's executive order

vmware®

Table of contents

NIST on Zero Trust

The National Institute of Standards and Technologies (NIST) has published an abstract definition of Zero Trust architecture and a selection of generalized deployment models. According to NIST's definition, Zero Trust:

- Assumes that no implicit trust is granted to assets, services, or user accounts based solely on their physical or network location
- Focuses on protecting resources (including assets, services, workflows, and accounts) rather than network segments

NIST's logical components of a Zero Trust Architecture include network micro-segmentation (segmenting the internal network at a fine-grained level).

Introduction

2021 began with a wave of high-profile cyberattacks targeting the private sector and U.S. federal agencies. In early May of that year, a ransomware attack forced Colonial Pipeline, one of the largest pipeline operators in the U.S., to suspend operations and shut down its IT systems. This brought fuel supplies to a temporary halt within pipelines that regularly transport more than 100 million gallons of fuel daily across an area extending from Texas to New York.[1] This attack followed closely on the heels of a mass compromise that took advantage of a software update from a network management tool provider to deliver malicious code to an estimated 18,000 victims, including several U.S. federal agencies.[2]

Incidents like these dramatically highlighted the scope and severity of the threat that such highly sophisticated and malicious cyber threats pose to national security and individual organizations. In their aftermath, U.S. President Biden released an Executive Order on Improving the Nation's Cybersecurity on May 12, 2021.[3] The Executive Order was intended to galvanize public and private sector efforts to deter, detect and defend against such threats.

The Executive Order calls for establishing private-public partnerships to foster bold and significant improvements in federal agencies' ability to prevent, detect and remediate cyber incidents. It also stipulates that the U.S. government should lead by example by building information systems that meet or exceed current cybersecurity standards. In particular, the Executive Order requires federal agencies to modernize technology infrastructures to increase cyber resilience, stating that this should be achieved by implementing Zero Trust Architectures.

Within the Executive Order, Zero Trust Architecture is explicitly defined as "a security model, a set of system design principles, and a coordinated cybersecurity and system management strategy based on an acknowledgement that threats exist both inside and outside traditional network boundaries." (This model) eliminates implicit trust in any one element, node or service and instead requires continuous verification of the operational picture via real-time information from multiple sources."[3]

1. "Colonial Pipeline attack: Everything you need to know," ZDNet, May 2021.
2. "A 'Worst Nightmare' Cyberattack: The Untold Story of the SolarWinds Hack," NPR, April 2021.
3. "Executive Order on Improving the Nation's Cybersecurity," The White House Briefing Room, May 2021.

It's important to note that the Executive Order calls for this core Zero Trust principle—the elimination of implicit trust—to be applied not just when granting users access to applications and resources but also to all services within the environment, including those communicating with one another. In other words, to become fully compliant with the new Executive Order, federal agencies must apply the Zero Trust model in ways that go beyond merely securing user account access to the organization's network. **In particular, they must also develop strategies to secure communications between the workloads that make up modern applications.**

This new focus on inter-workload communications is becoming especially important as federal agencies move toward microservices-based architectures. In turn, today's attackers are increasingly attempting to exploit the security control and policy gaps that opened when intra-application traffic patterns shifted with the adoption of microservices architectures. A focus on inter-workload communications is essential for countering today's advanced persistent threats (APTs), who seek to dwell within private cloud environments for long periods, concealing their activities by hiding them within normal internal traffic patterns. And federal agencies have long been a primary target for APTs.

However, leveraging traditional network security tools to enforce a Zero Trust-based approach can be challenging. Most network security tools don't offer a clean and scalable method for securing inter-workload communications over existing network deployments.[4]

The VMware NSX Distributed Firewall does precisely this. It provides an efficient means for federal agencies to secure workload-workload communications no matter which underlying networking hardware and solutions they have deployed.

Micro-segmentation to Secure East-West Traffic

Most federal agencies already have a network solution within their private cloud. These agencies have also virtualized the majority of workloads in their private cloud via VMware vSphere. In most cases, though, the existing networking solutions were not designed to protect virtualized workloads. Nor were they built to secure complex and distributed IT infrastructure. As a result, micro-segmentation is challenging to achieve with these networking solutions.[4]

However, micro-segmentation is an essential component in Zero Trust network Architectures. It's also necessary for preventing the lateral movement of attackers. To enforce micro-segmentation for its internal traffic, an organization must have visibility into the traffic flows between workloads within applications so that they can allow only those necessary for the application to function.

4. "Internal Firewalls for Dummies," VMware, October 2020.

Protecting Physical Workloads in the Private Cloud

In cases where federal agencies have some workloads that run on physical servers rather than virtual machines, micro-segmentation can be achieved for these workloads as well. There are several strategies to consider:

- Installing an NSX agent on the physical workload
- Configuring NSX Distributed Firewall policies for communication between physical workloads and virtual workloads at the virtual machine. It's worth noting that this approach can only be used to protect physical workloads that communicate with virtualized workloads. It cannot be used for physical-server-to-physical-server traffic.
- Using the NSX Distributed Firewall in conjunction with the NSX Gateway Firewall, a sister product designed to work with the NSX Distributed Firewall. With the NSX Gateway Firewall, VMware makes it possible to extend the same consistent access control and threat prevention capabilities offered by the NSX Distributed Firewall across all workloads in the private cloud.

Some security solutions rely on agents (installed on the physical server or virtual machine) to enforce micro-segmentation. Such solutions are generally designed to provide a mechanism for security policy enforcement using IP addresses, port numbers, and the protocol identifiers found in IP packet headers. This is a limited subset of capabilities that's not flexible enough to address the full array of micro-segmentation use-cases within federal agencies. In addition, such solutions typically leverage the host operating system's firewall capabilities to enforce security policies within the operating system kernel. These capabilities are also limited in nature.

With this model, it's not possible to enforce application-based or user-based policies. Nor can such solutions provide advanced threat prevention capabilities such as Intrusion Detection/Prevention Systems (IDS/IPS), network sandboxing, network traffic analysis (NTA), and network detection and response (NDR). However, such capabilities are necessary for adequately securing east-west traffic within the private cloud in the face of today's sophisticated threats.

Micro-segmentation with Advanced Threat Prevention

As mentioned, most federal agencies already have a network solution in place. This may be a software-defined network (SDN) or a traditional network solution (non-SDN). However, these solutions typically don't have the capabilities to secure application communication and workload access within the private cloud.

Effectively protecting east-west traffic requires two distinct sets of capabilities. These are:

- **access control,** which necessitates deploying internal firewalls that can enforce fine-grained policies at the workload level
- **advanced threat prevention**, which involves the use of technologies such as IDS/IPS to analyze live traffic, contain suspicious files and objects, and detect anomalous behavior that may represent malicious activity.

The NSX Distributed Firewall incorporates the full suite of required capabilities within a software-only implementation of a distributed internal firewall.

The NSX Distributed Firewall can be deployed independently from NSX's network virtualization capabilities. It can work with any underlying network solution. This makes it possible for federal agencies to implement the NSX Distributed Firewall to achieve Zero Trust for workload-workload communications within the private cloud irrespective of which network solutions they already have in place and regardless of whether those are hardware- or software-based.

The NSX Distributed Firewall can be turned on at each workload to achieve macro-or micro-segmentation. This makes it possible to define security controls and deliver services for each individual workload. When fine-grained security policies are tied to individual workloads, it's difficult for attackers to move laterally across a private cloud.

The NSX Distributed Firewall's advanced threat prevention (ATP) capabilities can also be turned on at the individual workload level. These capabilities include IDS/IPS, network sandboxing, and NTA/NDR. This means that the NSX Distributed Firewall brings capabilities traditionally available only in an edge firewall solution to inter-application communications.

NSX Firewall: Modern Network Defense

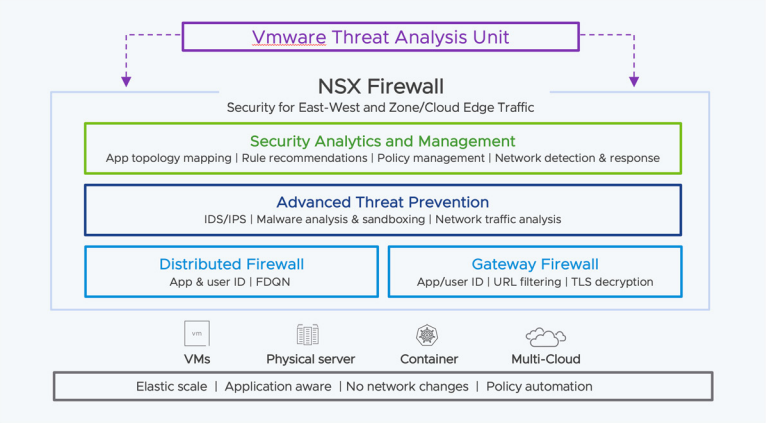

Conclusion

While implementing Zero Trust Architectures has long been desirable for organizations moving towards cloud-hosted and microservices-based architectures, it's now required for U.S. federal agencies. As a result, federal agencies must deploy solutions that will enable them to enforce access control policies and advanced threat prevention at the individual workload level.

The VMware NSX Distributed Firewall incorporates all required capabilities in a software-only implementation of a distributed firewall. Because it's a separate product that can be deployed independently from the NSX network virtualization solution, federal agencies can use it to achieve Zero Trust for their workload-workload communications no matter the networking solution they have in place.

The recent Executive Order may be a harbinger of things to come in the international regulatory landscape. Although it's only immediately applicable to federal agencies within the U.S., it's likely to serve as a model that will be emulated by other regulatory bodies around the world. It's also likely to become a standard that private sector organizations will voluntarily adopt and implement. In today's world, we cannot ignore the fact that general principles of Zero Trust are increasingly relevant for all organizations across sectors as computing environments grow in complexity and cyberthreats become more sophisticated.

Item No: vmw-wp-zero-trust-in-fed-private-cloud-v2 1/22