ZIVARO

···· BRILLIANT IT ····

# Zero Trust Networking

Scott Hogg
Principal Architect – Network and Security

## INTRODUCTION:

The slow-but-steady transformation of the traditional isolated internal corporate WAN has been occurring over the past 20 years. SD-WAN and cloud-hosted applications have accelerated the disruption of the contiguously-connected corporate enterprise intranet. Soon, the use of 5G wireless technology and further expansion of cloud environments will move even more end-user traffic toward the Internet. The enterprise network is no longer completely trusted, and now, the corporate intranet has metamorphosized into the Internet backbone. With the COVID-19 pandemic and more people working remotely, the use of the Internet as the enterprise network is even more profound. Enterprises can no longer definitely say that their corporate networks are completely trusted and follow a least-privilege security policy.

This paper will discuss the concept of Zero Trust Networking (ZTN) and describe its applicability to the modern enterprise network architecture. There are many solutions that assist by providing contextual information about the end-users and devices accessing on-premise and cloud applications using a least-privileged model. These solutions lend visibility, put the application access into context, and give the enterprise policy-level control over their multi-cloud environments. Zivaro provides network architecture, design guidance, and technology roadmap development for enterprises seeking to gain the benefits of this new approach.

## PROBLEM STATEMENT:

Enterprises frequently lack situational awareness of the devices connected to their network and in cloud environments and they don't have confidence in their user authentication systems. Enterprises also often lack knowledge of who is connected to their networks. The lack of visibility to the devices being used on their internal networks also extends to their cloud environments, and outward onto the Internet (particularly with the remote workforce arrangements during a global pandemic). This absence of "context" makes it nearly impossible to secure the applications and control valuable and sensitive data.

Most corporate networks exhibit a "most privilege" model whereby all users, devices, and applications are implicitly trusted if they are connected to the internal network. To deal with this situation, enterprises have deployed complicated data center or core-network firewalls. These firewalls can be expensive bottle-necks with complex policies that can be operational risks to the organization. The proliferation of core and network border firewalls performing Network Address Translation (NAT) has

been used to handle the overlap of private IPv4 address space that results from mergers and acquisitions, expansion into cloud environments, and the demands of software container applications and Internet of Things (IoT) devices.

Access to applications is frequently controlled by usernames and passwords. Enforcing strong passwords can be difficult when the end-users chose their own password and frequently re-use passwords to compensate for the complexity requirements. Enterprises need to strike the right balance between ease of use and creating a seamless single-sign-on experience with the address protections of Multi-Factor Authentication (MFA).

## BACKGROUND:

The enterprise network has been slowly transforming over the past decade. The Internet has been come the former corporate intranet, and what was previously known as the extranet. User's computers have become more mobile with the proliferation of wireless access networks. These changes have applied pressure to the traditional perimeter-based approaches to corporate network security.

**Orbital Security Model:** For decades, enterprise networks have been created with an Internet security perimeter to prevent against external threats, however, the internal network remains a completely trusted realm. If a computer is connected to the internal enterprise network, it is explicitly trusted and can connect to any other internal system without any security measures. Attackers changed their strategy toward the development of malware that compromises the internal hosts and APT to retain remote access to those computers and control them. Very weak perimeter firewall security policies also allowed the attackers to create back-channels and exfiltrate data from the compromised hosts. Once the attacker has a foothold within a single computer in the enterprise, they can easily expand their attack of the internal systems.

Organizations realized that their network architecture must change. However, 15 years ago, security researchers realized this enterprise network design weakness and worked on an improved security model. The Jericho Forum (now The Open Group ) was a group of security architects that advocated for de-perimeterization because most corporate perimeters were so porous that they were ineffective. Instead, the Jericho Forum advocated for an orbital security model whereby outer rings were less trusted and low risk and the inner rings surrounded the more important data and contained the highest risk assets. There Jericho Forum documented their vision and created a set of commandments of this new model. They also created this concept of a Collaboration Oriented Architecture (COA) whereby systems would allow nodes to interact over the untrusted Internet but only based on the security policy. The re-architecture of the enterprise network will utilize these concepts and principles.

The modern design of the "orbital security model" has transformed into what is now known as the Zero Trust Network (ZTN) architecture. The concept of Zero Trust Networking is to assume that end-user devices are already compromised and place them into an isolated environment where they can't be used to attack other enterprise systems and applications. Users are required to prove their

authenticity and the security of their mobile devices prior to being granted access to sensitive applications. In this ZTN security model, there is no security difference between internal network and Internet, everything is untrusted until it is proven to be trustworthy. The ZTN model simplifies the network architecture, reduces complexity, and improves security for the current methods of attack.

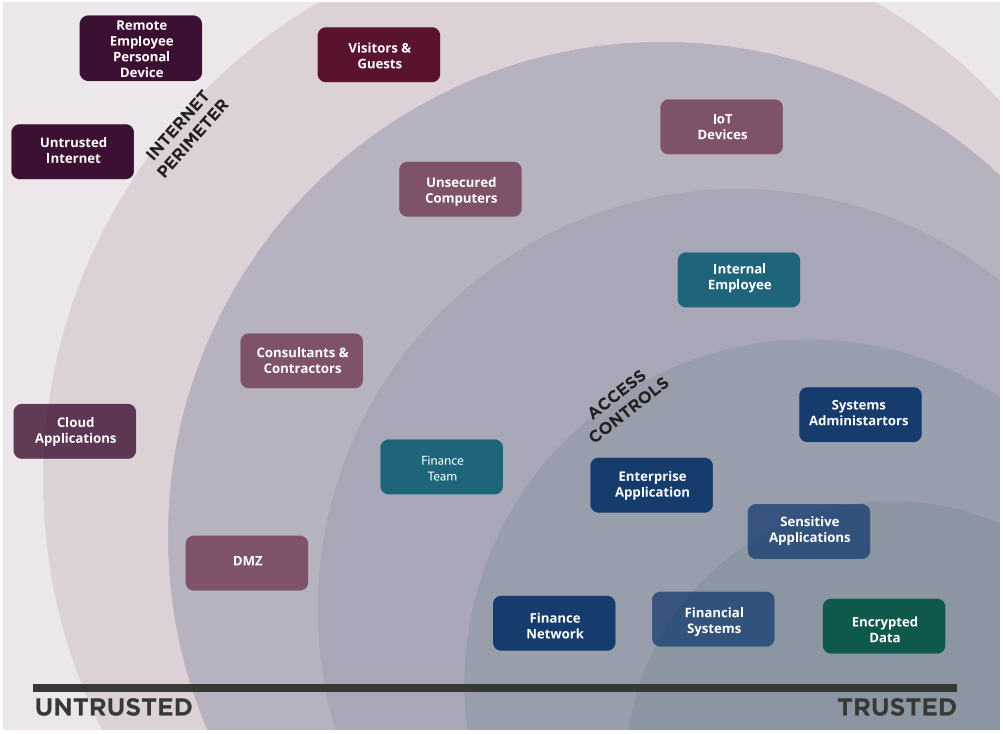


Users are already accustomed to having >100Mbps Internet access at their homes and they have the same expectation when in the office. However, many enterprises traditionally connected branch offices over costly and lower bandwidth Multiprotocol Label Switching (MPLS) service provider connectivity back to the headquarters with a smaller Internet connection. Once more enterprise applications became hosted on the Internet, this drove the changing branch office toward Software-Defined Wide Area Network (SD-WAN) with Direct Internet Access (DIA) connectivity to the branch. More corporate applications are moving to the cloud and being accessed over the Internet. Users will eventually just want fast Internet access, but still need Internet security such as protections against malware.

Today's attacks against end-user devices are becoming more sophisticated. Web-based and e-mail-based attacks are getting more creative and APTs are becoming difficult for security protections to detect. Attackers are using Domain Generation Algorithm (DGA) attacks using fast-flux for generating fake domain-names to avoid reputation filtering. The greater use of Carrier-Grade-NAT (CGN) and Large-Scale-NAT (LSN) by service providers due to the scarcity of IPv4 address provides Internet anonymity for attackers and makes reputation filtering ineffective. Furthermore, most reputation filters, web-security filters, and e-mail filters are sometimes not IPv6-capable. Therefore, end-user devices are under attack and organizations are now assuming that the end-user devices are already compromised and must be treated "at arm's length" from the internal corporate IT assets.

Unfortunately, the vast majority of the enterprise network are not utilizing the zero-trust least-privilege approach.

Here is a picture of the current set of corporate security "zones" and their level of trust plotted against the amount of authenticity of the connected nodes.

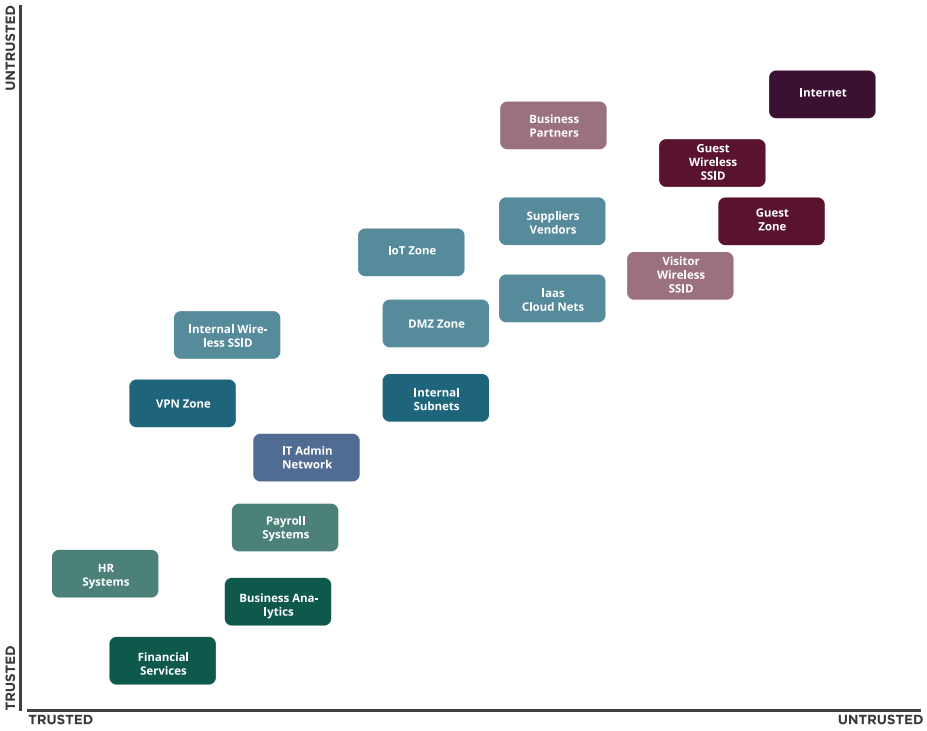


Zivaro performs assessments of enterprise networks and develops modern network architectures. We first start by assessing the current architecture and documenting findings and recommendations for improvements. Then we rank those findings using a formula that takes into account, the damage potential resulting from the issue, or the benefits resulting from the improvement, the CAPEX costs to make the improvement, and the OPEX time to correct the issue or migrate to the recommended design. Then we resolve order of the tasks, the project constraints and dependencies, and relationships between projects and create a short/near/long-term roadmap to realize the benefits. We document our findings and recommendations, then create the critical path and the roadmap (along with investment schedule) to achieve the target architecture.

**Context Aware Application Access:** Most enterprise have traditional networks that provides network access to those devices that are physically located within the on-premises facilities or connected by a remote access Virtual Private Network (VPN). Wired Ethernet access ports are often not restricted and can be connected to so long as the end-user has access to the building. Wireless access is provided freely to guests, visiting users, and only occasionally is the internal wireless access is controlled based on X.509 certificate and multi-factor user login credentials. Wireless access is often only controlled and provided to those devices within the radio frequency range of the 802.11 access point.

A modern context-aware network access model has been conceptualized. The idea of granting access based on an individual proving their identity based on strong authentication (using two or more factors), their level of authorization to perform some activity on an application, and the verification of the device they are using is now feasible. Once these characteristics are verified and access granted, this level of permission must be continually evaluated and if anything changes making the permit policy invalid, then access is revoked. Following is a diagram of this multi-factor network access control model.

This context-aware network access model relies heavily on the end-user providing information about their identity. Identity is often performed using strong multi-factor authentication (two-factor authentication) to prove the authenticity of the end-user who is using a device to connect to the

network. Depending on the user's identity, their role within the organization, their level of network access may vary.

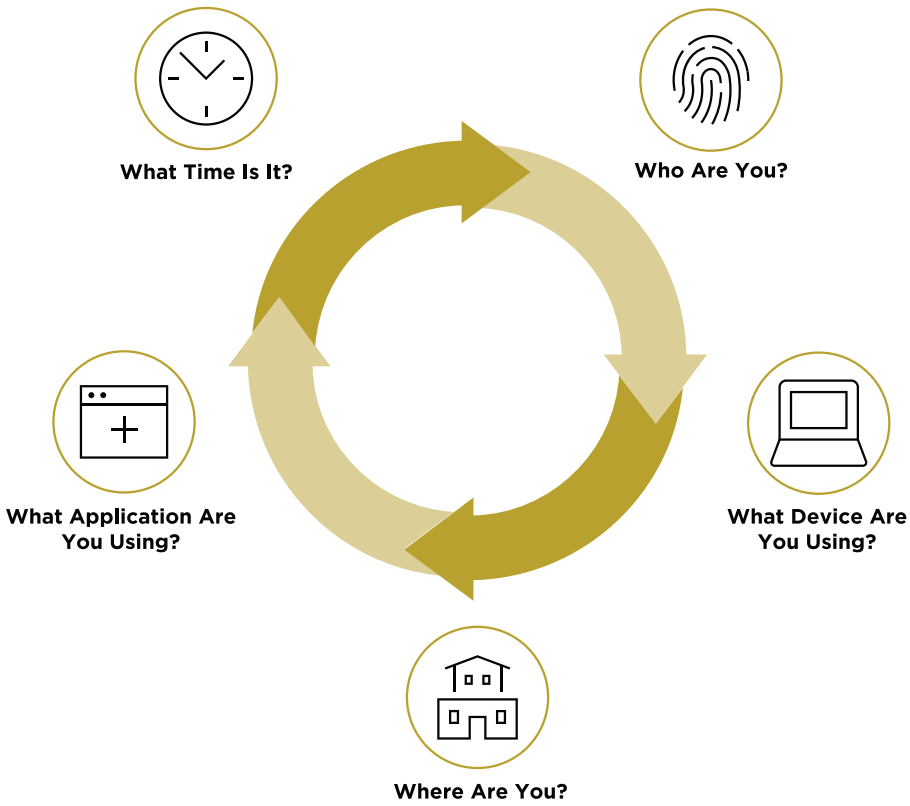


The end-user's device must also prove its characteristics of trustworthiness and ownership. The device posture could be an individual's own computer or a corporate laptop. Validating the device, its level of security, up-to-date operating system and application patches, ownership, and other characteristics need to be ascertained prior to granting network access.

Location is another important factor to consider prior to granting access. If the device is connected to an internal enterprise network then the network could ascertain its geography. However, on the Internet, there may be specific geo-location rules that may govern access. For example, a finance user accessing an internal finance system from their office may be accepted, but accessing from a coffee shop may not be permitted. Context-aware security systems may also track historical access information for the purposes of anomaly detection. For example, a user logging in from a foreign country may be denied if there is evidence that the same user just used their access badge to walk into the corporate headquarters an hour ago.

The nature of the application also plays a factor. Some applications are more sensitive than others and may require a higher level of access restrictions. For example, a financial system is deemed riskier than a web site showing open-to-the-public information about products for sale. Taking into account the type of application and the security sensitivity is important for a risk-based access system.

Time of day can also be used to provide valuable context to determine if network access should be permitted. A system could detect anomalies and prevent access that is well outside normal working hours or as in the previous example of access from a foreign country.
Most enterprises lack any type of context-aware system in place. Organizations must establish access policies, define end-user roles, and compile a device inventory. Few organizations have end-user behavior analysis (EUBA) system, or User and Entity Behavior Analytics (UEBA) that can determine if there is a compromised internal computer or a malicious insider. Enterprises often have a Security Information Event Management (SIEM) system that are able to perform root-cause analysis or

correlation or analysis of these types of access network anomalies. Enterprises will need to begin to establish this type of context-aware network access system as they move to a zero-trust network and to cloud-based applications.

## SOLUTION:

As this transformation of the corporate WAN toward the internet and the erosion of the traditional Internet perimeter has occurred, new solutions have emerged to help enterprises regain awareness and control.

Cisco's approach to Zero-Trust Networking is to address is by focusing on 3 pillars: Workforce, Workload, and Workplace . For the workforce, enterprises must to make sure the end-users are properly authenticated and are using legitimate and secured devices to access only the applications they require. For the workloads, enterprises must validate that applications are not threatened, are only accessed by authenticated users, and application visibility exists. For the workplace, enterprises must ensure the networks containing user devices (or IoT devices) are defended from unauthorized access and groups of systems are segmented and compromised devices are isolated. The goal is to create a consistent application of enterprise-wide security policies to allow valid users to use their protected devices to access the applications they need to perform their work.

Cisco has a wide variety of solutions that can be combined to create a Zero Trust architecture . Now we'll describe the products and solutions that fill each of these roles and how they create a comprehensive ZTN strategy.

**Compendium of Zero-Trust Network Segmentation Methods:** Thanks to the early work by the Jericho Forum, there are now many methods of breaking up a network into various segments based on authentication, trust, user role, and topology. There are now many different techniques, protocols, vendor products and services that can aid in implementing a network architecture that follows a "least-privilege" model. This model aims to minimize lateral movement from compromised hosts and limit the damage of such a security incident. Segmenting, separating, and isolating systems based on their context provides stronger enterprise security.

The following image shows a vast array of different types of methods that organizations could consider for improving the current network security architecture.

# Compendium of Zero-Trust Networking Methods

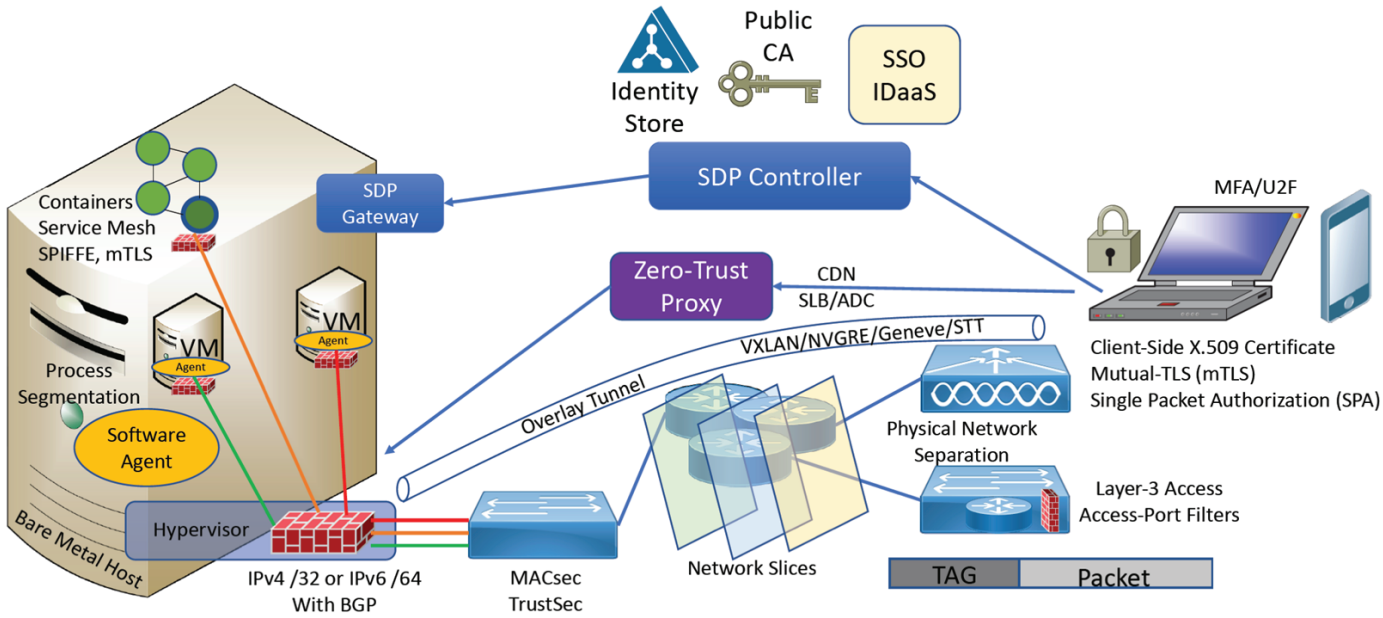


The National Institute of Standards and Technology (NIST) has created a Special Publication (SP) 800-207, Zero Trust Architecture (ZTA) . This document discussed the logical components of a ZTA and provides some various styles of how these can be implemented. The NIST document describes how a Policy Decision Point (PDP) and Policy Enforcement Point (PEP) are used to authenticate and validate the user and grant access to systems, resources, and applications. U.S. federal departments and agencies as well as corporate enterprises are encouraged to review this SP and consider the seven Tenets of Zero Trust that are discussed in the document.

These methods of micro-segmentation and host/application isolation can be divided into techniques used on the server/application, in the network, middleboxes, or in the end-user's node. There is a wide array of methods that can be implemented in the host operating system, in the software container virtual networks, in the hypervisor, in the network port, tagged packets (tunnels, overlays), in the firewall/proxy/SLB/CDN/CASB, and Software Defined Perimeter (SDP) or Identity-Aware Proxy (IAP).

There is a veritable cornucopia of micro-segmentation, host-isolation, and zero-trust networking techniques. However, not all these techniques be implemented in concert, but rather applied as-needed to particular points in the enterprise topology when applicable.

**Firewall Disaggregation:** Enterprise core firewalls must be highly-performant yet placing CPU-intensive filtering into the network slows down large amounts of streaming data or humongous data transfers. Increasing the amount of firewall logging slows down throughput even more. Therefore, the cost of large firewalls with 10Gbps interfaces and the cost to maintain the complex centralized policy can be substantial. When the rule-base gets large because the single HA pair of firewalls must support numerous applications, each with their own TCP/UDP service ports and communities of interest. In order to keep the administrative burden under control, its human nature to make the objects and rules coarser. Objects tend to define whole subnets and zones, rather than individual IP addresses of the server or client device. Groups of TCP/UDP service ports are often re-used among rules making policies less granular. All this can lead to coarse and overly-permissive policies that grant too much access and don't follow the least-privilege best practices.

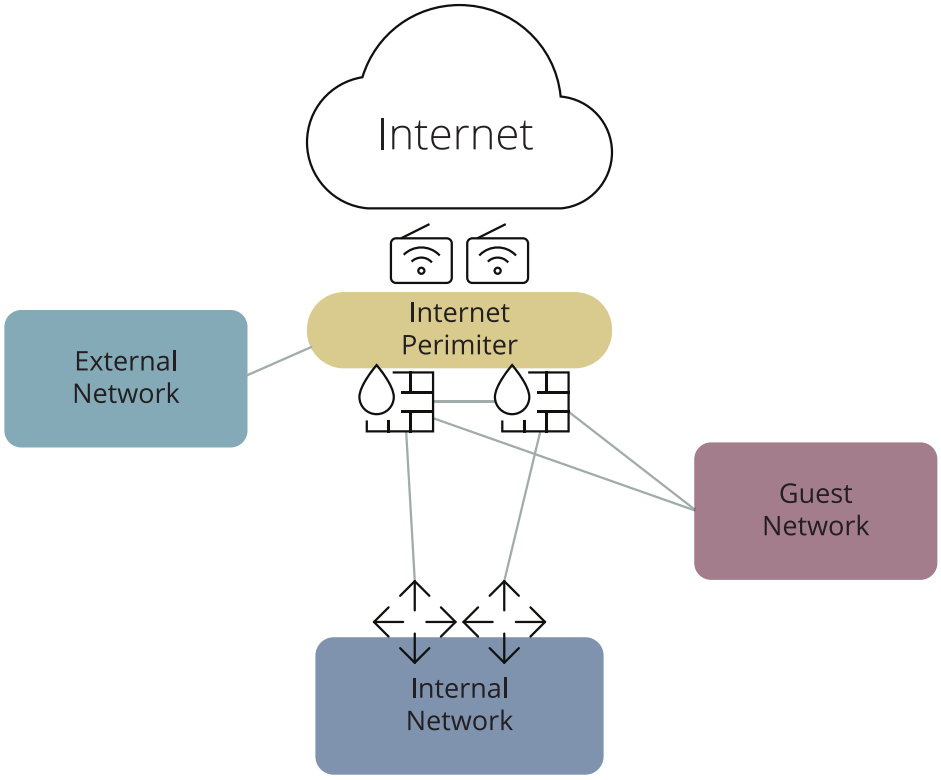


There is also increased risk of having an HA pair of firewalls placed at the core of the network that all traffic must pass through. A problem for that firewall pair could cause a complete service-affecting outage for the organization and all its applications. The complexity is higher with a large centralized policy with many rules and many different types of NAT taking place. This could make remediation take longer.

The additional problem with data center or perimeter firewalls is that they only observe TCP port 443 (HTTPS) encrypted packets traversing their interfaces. The core firewalls can only secure applications based on source/destination IP address (essentially reduced to being an ACL). The process of decrypting the SSL in the middle causes problems for applications validating the authenticity of certificates, require nodes to accept the firewalls cert, and consumes a tremendous amount of CPU resources (typically using specialized hardware) that can be costly. With firewalls close to the application edge, they can observe the un-encrypted communications at the server itself. The host-based firewall has full visibility to all the end-to-end application traffic is its raw unencrypted form.

An alternative design to using centralized "data center" firewalls placed at the core of the network is to distribute that functionality closer to the applications themselves. This design leverages the abundant CPU resources at the end-nodes to perform the complex filtering. This trend of moving the firewall function, and their stateful packet filtering, closer to the application is gaining popularity. This approach (e.g., cloud instance security groups) moves filtering closer to the server, allowing for smaller policies that can be more granular, permit lower-impact changes, and make operations and troubleshooting easier.

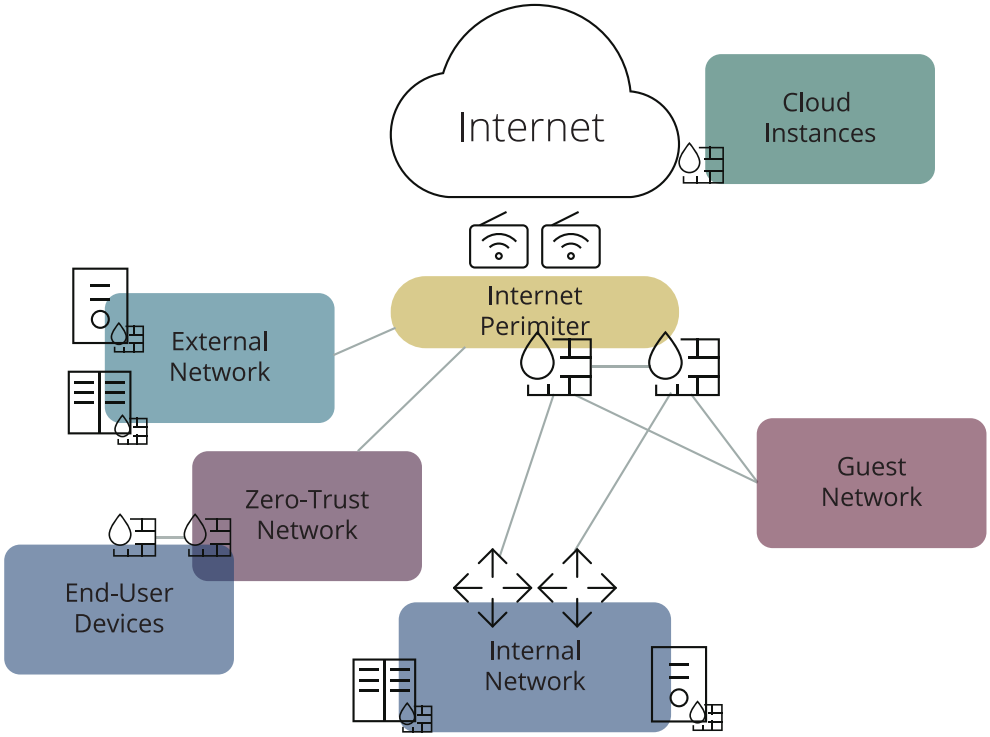


Instead of using one firewall for 10 domains, an organization could use 10 smaller firewalls (one for each domain). The maintenance of the single-domain firewalls is simpler and each has a smaller "blast radius"/"failure domain".

Firewall disaggregation is adjacent to the concept of micro-segmentation where we try to isolate each server with its own security policies and filtering. In the micro-segmentation architecture, there are lots of tiny firewalls, each with a much smaller policy. These smaller, less complex policies are easier to configure, maintain, troubleshoot, and provide simpler change management. The smaller rules can leverage software automation and rules can be generated with a least-privilege model. Many smaller policies are easier to maintain, automate, and troubleshoot resulting in a lower Mean-Time-To-Repair (MTTR).

An example of how firewalls can be virtualized and decentralized can be realized with the virtual edition of the Cisco Adaptive Security Appliance (ASA) firewall (ASAv) and the Cisco Firepower Next-Generation Firewall Virtual (NGFWv) . This software-based firewall can be run in hypervisor environments, in cloud networks, and in server environments using service chaining or Segment Routing (SR). These virtual firewalls can operate as a VPN concentrator for AnyConnect clients and it has a flexible and scalable license model.
This concept of moving firewalls closer to the servers is not new. The concept of Network Functions Virtualization (NFV) has been around for many years. The idea is that middleboxes such as firewalls, load balancers, Web Application Firewalls (WAFs), proxies could be placed closer to the servers where the CPU resources exist. Service chains can be configured to direct the traffic through these virtual security functions as needed based on the nature of the application.

**Identity and Authentication:** It has been a long-accepted fact that username and password are not completely secure methods of user authentication. Organizations now commonly embrace Multi-Factor Authentication (MFA) by using two-factor authentication (2FA) to validate the user and then validate their level of application access (e.g. One Time Password (OTP), software token, hard token, mobile app, text message). Additional methods of verifying identity as a component of a Zero Trust network include using an IDentity as a Service (IDaaS) system. This type of a system could also be used as the central point of authentication and Single Sign On (SSO) capabilities that could be leveraged in a Zero Trust Network.

The Cisco AnyConnect SSL VPN agent can run on the end-user's mobile devices when they are working remotely, as is often the case these days, and provide device security posture assessment, strong access authentication, and end-to-end encryption to on-premises and cloud environments. Cisco AMP for endpoints can provided added security for end-user's mobile devices and help defend them from malware compromise. Duo Access can also perform end-user device security checks and make sure the operating system and application patches are applied and confirm the device's location.

**Software Defined Partner:** The concept of a Software Defined Perimeter (SDP) was first defined by the Cloud Security Alliance (CSA) in 2013 . The concept is that there is a centralized controller that authenticates users and has knowledge of their role and privileges and access requirements. When the initiating user authenticates, the controller contacts the SDP gateway nearby, or within, the application server and unlocks access for that individual user. This provides granular control and closes off applications to any other host that hasn't first authenticated through the SDP controller. Google BeyondCorp helped to define this concept. The concept of an SDP gateway is similar to that of an Identity-Aware Proxy (IAP) in that authentication must take place first before application access is granted, but an IAP is more like a pass-through proxy. The benefit of using SDP/IAP for Zero Trust Access (ZTA) is that there is no back-haul of the traffic through a centralized VPN concentrator because the proxy (SDP Gateway) is located closer to the application server (on-premises or in the cloud). There are no MTU size issues due to encapsulation/tunneling. This solution leverages strong authentication, identity and role of the user and doesn't allow compromised client computers access to all the other enterprise networks.
Using Software Defined Perimeter (SDP) systems can also have this topology where the SDP controller is centralized, but the SDP gateway is close-to or inside-of the application server. The access controls and policies are close to the server and users are treated as "less-than-trusted". The traffic is not back-hauled through a central location because the filtering and policy enforcement is located at the server, application, or Service Mesh software container environment.
Identity is a key component of a zero-trust network design whereby Duo (now Cisco) tokens are issued to employees, visitors, and guests. Enterprises can easily expand their use of MFA with the Cisco Duo platform and begin to establish a proof-of-concept of Duo Beyond and Duo Access. These Duo applications and services provide the software that can form a Software-Defined Perimeter (SDP) environment, whereby users strongly prove their authenticity and their role defines their application access permissions. The access policies are defined in the Duo gateway software and these gateways validate the user's identity and their device characteristics. The Duo gateways act as the SDP gateway

that unlocks application access based on the user's MFA strong authentication. Furthermore, Duo Access can provide that Single Sign On (SSO) seamless experience for users that ties in with their Duo tokens for MFA.

**Service and Mesh Security:** One feasible option for improving security of microservices and container-to-container or service-to-service communications is to use a Service Mesh. Application segmentation can be performed using separate containers for each application component, decoupling an application, using containers, with a Service Mesh. A Service Mesh provides added security using sidecar proxies adjacent to application containers. This sidecar proxy performs the policy enforcement, encryption of communications, and communications visibility between containers. Encryption between containers uses Mutual Transport Layer Security (mTLS) with Secure Production Identity Framework For Everyone (SPIFFE) . Enterprise DevOps and application development teams should start to learn about the Service Mesh method (e.g. Istio Envoy, HashiCorp Consul) and test it in their container architectures. Products like Tigera Calico Enterprise can aid in the creation of access policies for Kubernetes environments. Calico policies define the roles, labels, permissions and policies for container applications which create a zero-trust environment within and external to the k8s pods.

The goal is to ascertain the authenticity of end-user computers, instances, and containers in the environment before application access is granted. Some enterprises provision X.509 certificates for many end-user computers, mobile devices and server instances in their environments. By placing an X.509 computer certificate onto each computer, these can be verified by a Certificate Authority (CA). These certificates can also be used for increased application security using mTLS. Istio Citadel can be used in a Service Mesh (described later) architecture to issue SPIFFE IDs for all workloads. Single Packet Authentication (SPA) is a method whereby the client browser/application sends a specific set of packets to the SDP controller or application server that identifies the user and their device and indicates that it is authorized to access the application. These are emerging models that leverage certificates and SDP-methods to provide Zero Trust application access.

**Network Segmentation:** Enterprise networks should be designed such that access is granted only to validated users with mobile devices that are secure and uncompromised. If devices are vulnerable or actively controlled by an attacker, then they should be segmented off from the rest of the corporate networks.
Cisco Software-Defined Access (SD-Access) is an enterprise overlay intent-based enterprise fabric networking solution that uses Cisco Digital Network Architecture (DNA) as the policy controller that defines which users can access which systems/applications. SD-Access performs identity-based micro-segmentation host isolation telemetry and analytics and software automation.

Cisco DNA Center (DNAC) can define the policies that allow access and control applications and coordinate with the wired and wireless network infrastructure components. IEEE 802.1x access control using dynamic ACLs to enforce traffic permissions. Security Group Tagging (SGT) policies can be created and deployed to TrustSec-capable network devices. This security concept involves tagging traffic and using the tag applied to packets to let the network infrastructure validate the connections

against the security policy. DNAC integrates with Cisco Identity Services Engine (ISE) as the mechanism to validate users, their devices, and the contextual characteristics. Cisco Meraki network devices can also perform this checking and validation of users. Policies and device management can be configured and verified using the Meraki Systems Manager.

In a data center environment, Cisco Application Centric Infrastructure (ACI) can have automated configuration of security policies, contracts and filters to protect the servers and applications. ACI can perform segmentation and host isolation as defined in the policies.

**Internet Security Based:** In Zivaro's previous article titled " Enterprise Network Security Architecture Transformation" the metamorphosis of the corporate intranet has now evolved into the Internet. As we anticipated, enterprises are now using more cloud-based security solutions to cover a larger footprint than the typical Internet perimeter. A clear example of this is displayed by how Cisco Umbrella can use DNS-based policies to protect on-premises as well as remote users accessing cloud applications. Another example is how Cloud Access Security Brokers (CASBs) provide the control and access policies for end users accessing cloud applications.

The Cisco ASAv or NGFWv can function as a VPN concentrator and can run in a private data center, near the corporate Internet edge, or in cloud environments. This firewall could be used as the "cloud hub" that connects and secures geographically diverse mobile user devices. The Cisco Cloud Services Router 1000v (CSR 1000v) virtual router can run in highly-available cloud environments to facilitate connectivity between clouds, connect hybrid topologies, and connect remote locations and user devices. SD-WAN, and Cisco Viptela has now transformed the traditional corporate WAN, to Direct Internet Access, and now to cloud connector.

This evolution has culminated into the Gartner-coined term Secure Access Service Edge (SASE) . Combining these traditionally on-premises (or Internet perimeter) security functions (DNS security, firewall, IDS/IPS, content filtering, VPN) into cloud environments. Nomadic end-user mobile devices can now be secured and a global policy of access controls can be administered and enforced across the widest topologies.
Once users are authenticated, secure connect to the SASE service, the corporate Zero Trust access policies can be applied based on the role of the user and the context of how they are connecting, from where, and the device they are using.

**Network and Application Visibility:** The full-stack cross-functional IT teams need visibility to the entire global environment in order to be able to proactively manage and reactively troubleshoot applications. Network, systems, security, and DevOps administrators need visibility to on-premises networks, private data centers, cloud networks, and everything in between. Without visibility into application traffic, enterprises would be "flying blind" and unable to perform capacity planning, perform forensics, or have any understanding of what is occurring with their IT assets.

Cisco Tetration has been that solution that enterprises have uses to gain deep insights into network

device communications (via sensors) and create an Application Dependency Mapping (ADM). Tetration is the system that performs the Cloud Workload Protection Platform (CWPP) function to gain visibility to network flows, regardless of where they occur in the global topology. Tetration Security Dashboard provides visibility to the communications and can validate that the Zero Trust Network is allowing the correct level of access. Tetration can be integrated, using pxGrid with the ISE platform to integrate with the current view of connected users and their devices.

Cisco AppDynamics also provides the proactive and reactive Application Performance Management (APM) function required to keep a global array of applications running smooth for the highest level of end-user experience and productivity.

Corporations also need to have visibility to all security events in their diverse environments. Security Information Event Management (SIEM) systems are what comes to mind first. Cisco's SecureX service is an example of how these same security events, and security configurations can be placed into a cloud-native service. The SecureX dashboard creates that simple single-pane-of-glass for security incidents for cross-functional teams all responsible for secure management and operations.
Many organizations use NetFlow (IPFIX) to gather data about application access. Traditionally, this information was used for Quality of Service (QoS) configuration validation or for proactive or reactive network performance management. This flow data can also be used for assessing security incidents. Cisco Secure Network Analytics (Stealthwatch) can observe the application traffic crossing the corporate network and identify indicators of compromise that triggers forensic investigations. This can be the early-warning-system if anything gets past the other layered diverse defensive mechanisms.

## CONCLUSION:

It is impossible to solve the enterprise security conundrum using just one of the techniques mentioned here. To holistically develop a Zero Trust Network Architecture, organizations must adopt a combination of techniques. Each of these methods adds to the comprehensive set of compensating controls. Should one protection measure fail, the diversity of defense allows other functions to back-up each other and reinforce the environment.
To create a layered Zero Trust Network Architecture, it requires first understanding your security zones, application accessibility, trust boundaries, and context of end-users and their mobile devices. Zero Trust Network Architectures are not solely deployed within the network engineering team or by the security team. They require a cooperative effort across the entire IT organization. Techniques in the end-user device, across the Internet, on the corporate networks, in the cloud environments, and in server instances are all required elements. Key pillars of a Zero Trust Network include identity and strong authentication, encryption using TLS and certificates, global application access policies based on role and context.

Zivaro performs Zero Trust Network architecture, design, and roadmap projects. Zivaro helps organizations gain visibility to their networks, authenticate and secure application access, adjust network architectures to have a least-privilege, segmented, isolation, and distributed security. Zivaro

can help your enterprise develop a Zero Trust network selecting from the wide variety of solutions tailored to your exact requirements, needs, and specifications.

## REFERENCE (END-NOTES)

**Wikipedia Definition of Zero Trust Networks (ZTN)**
https://en.wikipedia.org/wiki/Zero_Trust_Networks

**The Open Group**
https://www.opengroup.org/forum/security
https://collaboration.opengroup.org/jericho/vision_wp.pdf
https://collaboration.opengroup.org/jericho/commandments_v1.2.pdf
https://en.wikipedia.org/wiki/Collaboration-oriented_architecture

**Cisco Zero Trust Solution Overview**
https://www.cisco.com/c/en/us/products/collateral/security/solution-overview-c22-742591.html

**Cisco Zero Trust Products**
https://www.cisco.com/c/en/us/products/security/zero-trust.html

**NIST 800-207 - Zero Trust Architecture**
https://csrc.nist.gov/publications/detail/sp/800-207/final
https://blogs.cisco.com/security/an-overview-of-zero-trust-architecture-according-to-nist

**Cisco Adaptive Security Virtual Appliance (ASAv)**
https://www.cisco.com/c/en/us/products/collateral/security/-adaptive-security-virtual-appliance-asav/datasheet-c78-733399.html

**Cisco Firepower Next-Generation Firewall Virtual (NGFWv)**
https://www.cisco.com/c/en/us/products/collateral/security/-firepower-ngfw-virtual/datasheet-c78-742858.html

**Cisco AnyConnect Secure Mobility Client**
https://www.cisco.com/c/en/us/products/security/anyconnect-secure-mobility-client/index.html

**Cloud Security Alliance - Software Defined Perimeter**
https://cloudsecurityalliance.org/artifacts/software-defined-perimeter/

**Google Beyond Corp**
https://beyondcorp.com/

**Cisco Duo**
https://duo.com/use-cases/industry-solutions/zero-trust-security
https://duo.com/docs/beyond-overview

https://duo.com/product/single-sign-on-sso
https://duo.com/docs/access-overview

**Secure Production Identity Framework for Everyone (SPIFFE)**
https://spiffe.io/docs/latest/spiffe/overview/

**Service Mesh Security**
https://istio.io/
https://www.consul.io/segmentation.html
https://www.tigera.io/tigera-products/calico-enterprise/
https://www.projectcalico.org/

**Cisco Software-Defined Access (SD-Access)**
https://www.cisco.com/c/en/us/solutions/enterprise-networks/software-defined-access/index.html

**Cisco Digital Network Architecture (DNA) Center (DNAC)**
https://www.cisco.com/c/en/us/products/cloud-systems-management/dna-center/index.html

**Cisco Identity Services Engine (ISE)**
https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html

**Meraki Systems Manager**
https://meraki.cisco.com/products/systems-manager/

**Cisco Application Centric Infrastructure (ACI)**
https://www.cisco.com/c/dam/en/us/solutions/collateral/da-
ta-center-virtualization/application-centric-infrastructure/at-a-glance-c45-732546.pdf

**Enterprise Network Security Architecture Transformation, Dec 20, 2019**
https://blog.zivaro.com/enterprise-network-security-architecture-transformation

**Cisco Umbrella**
https://umbrella.cisco.com/trends-threats/secure-access-service-edge-sase

**Cisco Cloud Services Router 1000v**
https://www.cisco.com/c/en/us/products/routers/cloud-services-router-1000v-series/index.html

**Cisco SD-WAN (Viptela)**
https://www.cisco.com/c/en_id/solutions/enterprise-networks/sd-wan/index.html

**Secure Access Service Edge**
https://www.cisco.com/c/en/us/products/security/what-is-sase-secure-access-service-edge.html

**Cisco Tetration (Cisco Tetration Workload Security)**
https://www.cisco.com/c/en/us/products/collateral/data-center-analytics/tetration-analytics/datasheet-c78-737256.html
https://www.cisco.com/c/en/us/products/security/tetration/index.html

**Cisco AppDynamics**
https://www.appdynamics.com/

**Cisco SecureX**
https://www.cisco.com/c/en/us/products/security/securex/index.html

**Cisco Secure Network Analytics (Stealthwatch)**
https://www.cisco.com/c/en/us/products/security/stealthwatch/index.html