



ZERO TRUST ENTERPRISE STRATEGY



Stephen S. Choolfaian

Table of Contents

Overview	3
Summary	3
Why Adopt Zero Trust?	3
The Risks of Zero Trust	4
PISTOL Methodology	5
Goals of this Document	5
Principles	6
Ideals	7
Strategies	8
Overview	8
Sub Strategies	11
Identity and Access Management	12
Identity Stores (Directories)	12
<i>Overview</i>	12
Identity Lifecycle	14
<i>Overview</i>	14
Access Management	15
<i>Overview</i>	15
Analysis	16
<i>Gaps</i>	16
Privileged Access Management	17
Password Vaulting	18
Secrets Management	18
Privileged Session Management	18
Analysis	19
Network Infrastructure	19
Network Firewalls	20
Domain Name System	21
Wide Area Network	22
Load Balancers, Application Delivery Controllers, and API Gateways	22
Virtual Private Networks	23
Micro Segmentation	25

Analysis	26
Intrusion Detection and Prevention Systems	27
Host Based	27
Network Based	28
Analysis	28
Data Protection	29
Data Types and Classification.....	29
Data Lifecycle	30
Data Security.....	31
Analysis	32
Infrastructure and Platform as a Service	32
Cloud Services	33
Service Meshes	34
Analysis	34
Software as a Service	35
Federated Identity Management.....	35
Secure Web Gateways	36
Cloud Access Security Brokers	36
Analysis	37
Tactics, Operations and Logistics	37
Next Steps	38

Overview

Summary

Zero Trust is a true security philosophy and set of principles, which taken together represent a significant shift in how enterprise IT and security should be approached. The results can be enormously beneficial for lowering enterprise risk and ensuring compliance. The strategy is broad in scope and a true implementation of the strategy will require implementation of many processes and tools, not least of which is a long organizational commitment to continuous improvement of the enterprise's security posture. It can be thought of by the financial minded as zero-based budgeting for security: all assumptions must be discarded, and all prior decisions need to be re-examined.

A Zero Trust system is an integrated security platform that uses contextual information from identity, security and IT Infrastructure, and risk and analytics tools to inform and enable the dynamic enforcements of security policies uniformly across the enterprise. Zero Trust shifts security from an ineffective perimeter centric model to a resource and identity-centric model. As a result, organizations can continuously adapt access controls to a changing environment, obtaining improved security, reduced risk, simplified and resilient operations, and increased business agility.

Why Adopt Zero Trust?

A safe assumption is that any enterprise with a competent security team is secure; at least it is secure in the sense that it uses the best of breed security tools to protect and scan the environment, implements the tightest security settings, and actively addresses vulnerabilities that appear in the environment. A fair question then is why adopt the Zero Trust strategy? 'It's better and cheaper' are reasonable but unsatisfying answers.

The best way to understand the value of the Zero Trust strategy is to examine it through the prism of how businesses operate. For a business to compete and thrive it must accomplish two primary goals: (a) be more efficient and productive than its competitor and/or (b) have capability, capacity, or advantage that its competitor does not have.

There are operational and logistical challenges to meeting these goals. Does the business have the right tools? Does the business have the right people? Is the business able to get the right tools to the right people, keep them operational and the input/output organized?

Information technology spent the back half of the 20th century dealing primarily with logistical challenges. Today, there are many powerful, customizable information tools that

service almost all facets of industry. Companies have workers located across the world and specialization is not limited to regions. There are a lot of options to get the right tool, the right person to use the tool and there are a lot of ways to keep the process organized.

While logistical challenges were alleviated, they were replaced by operational concerns, particularly regarding data security and privacy. The more decentralized a company became, the more it became susceptible to malicious or reckless behavior by criminals (e.g., hackers), employees, vendors, and partners. This caused many IT departments, rightly concerned about risk, to develop a fortress mentality, where access past the network perimeter was exceptionally difficult.

The tension is that while companies are driven by global competition to decentralize, IT departments beholden to the classic approach must build bigger and bigger fortresses which become exceedingly more expensive and difficult to manage. Instead of providing efficiency and competitive advantage information technology becomes onerous!

The Zero Trust strategy looks to address this issue by dismantling the fortress by embedding security right into the system. Users have simplified access to the networks, applications and data needed for their jobs, they, the network, and data are constantly monitored for inappropriate activity. External threats do not have the obstacle of a single, strong perimeter to breach, but dozens of controls operating behind the scenes.

The intended result is that the enterprise's security posture has greatly increased, but the controls are more transparent to the end users. The company can continue to drive efficiency and advantage with information technology and not be hampered by security.

The Risks of Zero Trust

Zero Trust strategy does not posit the substitution of any proven security control or technology with something unproven; in fact, it encourages and extends tried and true tools and tactics. It can be declared with reasonable certainty that successful implementation of a Zero Trust strategy will be MORE secure than what is yielded by classic security methodologies. However, that does not mean the strategy is without risk. Failure can usually be manifested in three scenarios:

First, Zero Trust is an absolutist strategy. The benefits cannot be realized unless policies AND policy enforcement have been implemented. Deploying tools without configuring them with well thought policies will only provide a false sense of security.

Second, Zero Trust is a strategy not a project; it requires organizational continuity. To claim to have 'implemented' Zero Trust would be a misnomer; Zero Trust requires constant re-

evaluation, implementation, and adjustment. If the organization does not change the way it operates, the issues that Zero Trust is attempting to solve will remain.

Lastly, there is a danger of 'cherry picking' pieces of Zero Trust strategy as way to decrease security risk without attempting to realize business benefit. In this case, perimeter controls remain intact and Zero Trust strategy is overlaid over classic controls. In this scenario, security becomes even more complicated and onerous and provides exceedingly diminished returns. The 'fortress' becomes a 'prison'!

PISTOL Methodology

Given the challenges of extensive, complicated, continuously evolving strategy e.g. Zero Trust, this document uses the PISTOL Methodology as a way to organize assessments of current state and recommendations. PISTOL is broken up into six categories:

- Principles – A core assumption or goal from which everything else by proximity can derive business value e.g. "Provide excellent IT end user experience to all employees irrespective of their role or location" or "Data security is our first and greatest priority."
- Ideals – A standard of perfection that defines whether we have been successful in implementing our goal e.g. "All users have the ability to work remotely at the same level of productivity as they would if they were in the office."
- Strategies - Plan of actions or policies designed to achieve a major or overall aim, goal or ideal e.g. Flexible Workplace, Digital Democracy or Zero Trust.
- Tactics -Specific actions or plans used to implement the initiatives outlined in a strategy e.g. VDI, MFA and BYOD being tactics that could be used to support a Flexible Workplace strategy.
- Operations – Organized and concerted activities required to implement and support a strategy e.g. a technology rollout project or support team.
- Logistics - The required coordination of people, facilities, or supplies to support an operation e.g. a project plan, organization structure or committee.

Goals of this Document

A Zero Trust IT strategy touches the entire enterprise from badge readers all the way to the most business-critical data storage systems. An in-depth analysis of existing and recommended strategies would be hundreds of pages and required contribution from dozens of IT groups. Instead, this document strives to provide high level analysis and recommendation from Enterprise Architecture group's perspective based on interpretation of discussions with other technology groups coupled with industry standards that define the Zero Trust strategy. In doing so, this document looks to provide five things:

1. The philosophical principles that should govern the [REDACTED] implementation of the Zero Trust methodology.
2. A set of high-level ideals or requirements that a successful Zero Trust methodology would meet.
3. Exploration of the strategies that are used by Zero Trust implementations, how [REDACTED] is presently implementing them and gap analysis.
4. Commentary on the tactics, operations and logistics required to implement Zero Trust in the [REDACTED] enterprise.
5. Next steps on the Zero Trust Journey.

Principles

There are three core security principles that are generally accepted across the industry as being essential to any Zero Trust implementation that were initially defined in the “No More Chewy Centers” paper published by Forrester¹ In addition to these principles, we have added several other core concepts described within the NIST Zero Trust Architecture document². These governing principles are:

- **Ensure all resources are accessed securely, regardless of location.** Zero Trust should secure access by all identities (human and machine), to all resources (data, applications, servers)—regardless of the location of the identity, and regardless of the location or technology of the resource being accessed. The traditional corporate perimeter should be dissolved and replaced with an alternative security paradigm.
- **Adopt a least privilege strategy and strictly enforce access control.** Least privilege accessed must be consistently assessed and managed across locations and resource types, and at both the network and application layer, using security and identity context. Users must authenticate and authorize at the network level as well as at the application and data layers; broad access to networks must be avoided (e.g. large VLANs or VPN connections with wide access.)
- **Inspect and log all traffic.** Networks are how distributed components connect and communicate with one another, providing excellent opportunities to detect threats. Zero Trust systems should broadly examine and log network traffic metadata but be judicious in the inspection of network traffic content due to processing and storage costs.
- **Ensure all components support APIs for event and data exchange.** Components (e.g. third party systems, home grown applications or data connections) should use access

¹ <https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf>

² <https://csrc.nist.gov/publications/detail/sp/800-207/final>

points that support standardized API protocols, so that solutions can fit seamlessly into a Zero Trust architecture.

- **Automate actions across environments and systems, driven by context and events.** Zero Trust is based on a set of dynamic access control rules, which change in response to identity, device, network, and system context. These rules are centralized in Policy Decision Points (PDP) connected with a distributed set of Policy Enforcement Points (PEPs) that automate changes to the enforced policies via integration/APIs. This is critical for a Zero Trust system to work since access changes based on context. (For example, an administrator's super user rights are disabled until they are needed for a task and then they are temporarily enabled until the task is completed.)
- **Deliver tactical and strategic values tied to business value.** Given the significant impact security changes have on infrastructures, teams, operations, and end-user experience it is important to tie changes to delivered business value such as the reduction of risk, technical debt or increase of competitive capability (e.g., a flexible workplace.)

Ideals

Zero Trust ideally generates sets of requirements that stem from the principles described above. While many of these requirements (especially APIs and Integration) are specifically tied to IT and security functions there are some general requirements that a Zero Trust implementation should meet. They are:


- Data plane communications must be encrypted. Any exceptions must be deliberate (e.g., DNS).
- System must be able to enforce access controls for all types of resources.
- Access control mechanisms must be driven by identity-centric and contextual policies.
- Data resource protections should be able to use identity and contextual policies to control access.
- System and policy model must support securing all users in all locations.
- Policy model and controls must be consistent for remote and on-premises users.
- Devices must be able to be inspected for their security posture and configuration prior to being granted access, and periodically thereafter.
- It must be possible to distinguish BYOD from corporate-managed devices and control the level of access accordingly.
- Access to any network resource must be explicitly granted by policy.
- No user or device should inherently have broad network access.
- Access controls must be able to distinguish between different services on the same network resource. (For example, access to HTTPS must be granted separately from access to SSH.)

- Access to specific data elements contained within applications or containers that have different classifications must be enforced based on business policy.
- Network traffic metadata must be logged and enriched with identity context.
- Network traffic must be able to be examined for security and data loss purposes.
- Workloads transferred into the cloud should include the same access control policies as defined by on-premises solutions.
- Automation must include identity-centric details to provide efficient and effective incident response.
- Logs must be included in analytics tools for effective and dynamic enforcement of policies.


Strategies

Overview


Zero Trust strategy is at its heart an expansion, not a replacement for the classic security strategy. Each architectural layer is built on the previous foundation and refined. The first level of any architecture is basic and unsecured and purely functional. A client traverses a network and gets access to a resource or some type of workload.




The first Zero Trust principle states **‘Ensure all resources are accessed securely, regardless of location.’** At least the first half of this axiom is consistent with the classic security strategy. It requires hardening, protecting and encrypting all transactions. The client is protected through anti-virus (or more sophisticated forms of threat protection), the client, network and resources are encrypted, and the resource has some form of data protection, whether through permissions or more role based and complex policies. (It is important to note, the architecture only half complies with the zero-trust principle because though ‘hardened’ it relies on its location inside the corporate perimeter, as foundational to its security profile.)



The second principle '**Adopt a least privilege strategy and strictly enforce access control**' builds on this architecture in three ways. First, it requires strict identity access management since identity is key to establishing privilege. Second, it adds the need for strict network access control, since Zero Trust does not assume any privilege based on location. Third, it adds privileged access management, to ensure sure clients receive the least amount of privilege required. This expanded model, begins to demonstrate potential interdependencies of security systems:



The third principle is '**Inspect and log all traffic.**' Logging is not just for historical purposes (e.g., behavior modeling) but for active monitoring of bad behavior, either by external or internal actors. Strategically this means the implementation of effective, rigorous, and timely Intrusion Detection and Prevention capability across all components.




The fourth principle **‘Ensure all components support APIs for event and data exchange’** is perhaps the most straightforward. In other words, all components in a system should be designed using the core security standards that support Zero Trust. While this is a strategic decision, it is tactical in implementation; something that must be done consistently but on an individual basis.

The fifth principle is **‘Automate actions across environments and systems, driven by context and events.’** This is key for Zero Trust, since it is a strategy designed to create an adaptive, reflexive, agile security environment. To achieve this principle, Zero Trust creates two conceptual categories to make the environment dynamic:

- Policy Decision Point (PDP) - A mechanism that evaluates access requests to resources against the authorization policies that apply to all requests for accessing that resource to determine whether specific access should be granted to the client that issued the request.
- Policy Enforcement Point (PEP) - A system responsible for enabling, monitoring, and terminating connections between a subject and an enterprise resource based on decisions made by a PDP.

To conceptualize this, PEPs feed data to PDPs from points in the topology. PDPs use data and policy to make decisions. These decisions can range from allowing access to a particular resource, to in the case of a detected breach, locking a client completely out of all systems. Ergo, a Zero Trust architecture becomes organized something like this:



Lastly, the final principle is **‘Deliver tactical and strategic values tied to business value.’** This is not an architectural strategy but a business and political one. We have, through application of the strategy built a security environment that in layman terms is:

- Hardened and well protected.
- Locked down to the least privilege necessary for a client to do its job.
- Monitoring itself for bad actors
- Built to open standards.
- Dynamically reacting to events e.g., breaches or questionable activity

We can therefore consider the zero-trust environment extremely secure based on reasonable standards and can expect to provide and should communicate the following business value:

- Lower risk for data loss, better compliance with audits and therefore increased corporate reputation.
- The removal of the requirement of a hard corporate perimeter as a constraint on facility decisions or employee location
- The ability to integrate with zero trust capable partners and vendors using streamlined, efficient, and fast security processes.

Sub Strategies


As Zero Trust encompasses a massive number of overlapping technologies and processes, it would be impossible to address them all in a single strategy. However, there are opportunities to assess larger areas, offer enterprise assessment, gap analysis and recommendations. The sub strategic areas that will be examined are:

- Identity and Access Management
- Privileged Access Management
- Network Infrastructure
- Intrusion Detection and Prevention Systems
- Data Protection
- Infrastructure and Platform as a Service
- Software as a Service

Identity and Access Management

Identity and Access Management (IAM) comprises both the technical and business process aspects of controlling access by providing the right access at the right time to the right person. A reasonably well-run identity management program is the key to success for a Zero Trust program since it is an identity-centric approach to security. Identity management systems should serve as the authoritative source of information about identities (person and non-person entities) and be used as the “keystone” system for many technical integrations as well as business processes.

IAM is broken down into these basic functional areas:



3

Identity Stores (Directories)

Overview

The core element of any identity management system is its identity store, often referred to as a Directory (or Directory Service.) This is where the logical authoritative information about entities, their frequently used properties and metadata (e.g. home office or telephone number) are stored. Directories often include user credentials, making directories a centrally managed source against which users are authenticated. Identity Stores can be generalized into three types:

³ Garbis, Jason; Chapman, Jerry W.. Zero Trust Security (p. 117). Apress.

- **Databases** - Databases provide a raw centralized identity store which can be accessed across a network. Standards-based directories rely on an underlying database however it is considered bad design to give remote applications database access to user information (especially credentials) even if it is read-only.
- **LDAP** - Lightweight Directory Access Protocol (LDAP) is a protocol specification, defining a set of messages (effectively an API) for interacting with directory services across a network. LDAP is a well-established standard, outlined in a series of RFCs from the Internet Engineering Task Force (IETF).³ LDAP v3, initially published in 1997, has been a very successful standard, in the sense that many directory providers (both open source and commercial) support the protocol, and components from different vendors can successfully interoperate. LDAP provides a straightforward API for operations against the set of entities in the directory and is also very commonly used to authenticate users' credentials (passwords). LDAP is very widely deployed and supported today, with broad support across identity, security, application, and infrastructure vendors.
- **IDaaS** - Identity-as-a-Service are vendor managed cloud-based directories, which relieve organizations from having to operate on-premises directory servers, provide a modern web-based UI, and deliver end-user-friendly features such as Single Sign On (SSO) and, increasingly, passwordless authentication. These services typically provide both newer APIs such as SAML and older APIs such as LDAP and RADIUS.

Assessment of Current Situation

- [REDACTED]
- [REDACTED]
- [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]

- [REDACTED]
 - [REDACTED]
- [REDACTED]

Identity Lifecycle

Overview

Every identity has a lifecycle, they exist for a period, they potentially have different roles over time, and eventually they are disabled and then destroyed. Organizations use technical tools and business/IT processes to manage and control identity lifecycles. These areas within IAM are important to a successful Zero Trust initiative and include:

- **Lifecycle Management** – The core IAM policies that include initial provisioning of accounts and privileges, assigning rights by role or business need, removal of unnecessary rights by request or automated assessment and lastly the retirement of identities when users leave an organization.
- **Identity Governance** - The policies or rules that determine “who should get access to what” as part of the identity lifecycle. Governance is commonly driven by both compliance and security requirements, often with a heavier compliance driver.

Assessment of Current Situation

- [REDACTED]
- [REDACTED]
 - [REDACTED]
 - [REDACTED]
- [REDACTED]

- [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
- [REDACTED]
- [REDACTED]

Access Management

Overview

Access management is the core of identity management: how entities prove that they are who they claim to be and the model for defining and expressing the set of actions that a given entity is permitted to perform. The components include:

- **Authentication** – The common authentication protocols, mechanisms, standards, and trends used to check if a user is who they claim to be. These generally include simple authentication (i.e. something that you know e.g. a password), multi factor authentication (the use of more than one authentication factor as part of an authentication process e.g. a physical token, smartphone app, or biometric mechanism to validate something you have, or something you are) or step up authentication (the process of prompting a user for an additional form of authentication, after some event or trigger has taken place e.g. an already-authenticated user attempting to access a high-value resource.)
- **Authorization** - Authorization is the end goal of access management, which after all is ultimately responsible for mapping from some policy model (technical or business policy) to enforcement points. Authorization can be granular based (i.e. assigned to a specific resource), group based, role based or context based (e.g. based on an identity's location)

Assessment of Current Situation

- [REDACTED]
- [REDACTED]

- [Redacted]
- [Redacted]
- [Redacted]

[Redacted]

- [Redacted]
- [Redacted]
- [Redacted]

[Redacted]

Analysis

Gaps

[Redacted]

- [Redacted]
- [Redacted]
- [Redacted]
- [Redacted]
- [Redacted]

- [REDACTED]

Recommendations

As Identity Access Management is a cornerstone of determining privilege, it is a critical sub strategy to get right for Zero Trust. This is an area where [REDACTED] should invest significant intellectual and technical capital. We recommend:

- [REDACTED]
- [REDACTED]
- [REDACTED]

Privileged Access Management

Privileged access management (PAM) consists of strategies and technologies for exerting control over the elevated (“privileged”) access and permissions for users, accounts, processes, and systems across an IT environment. By dialing in the appropriate level of privileged access controls, PAM helps organizations condense their organization’s attack surface, and prevent, or at least mitigate, the damage arising from external attacks as well as from insider malfeasance or negligence.

While privilege management encompasses many strategies, a central goal is the enforcement of least privilege, defined as the restriction of access rights and permissions for users, accounts, applications, systems, devices (such as IoT) and computing processes to the absolute minimum necessary to perform routine, authorized activities. From a Zero Trust perspective, least privilege is an ideal state.

PAM consists of three general strategies/technologies: Password Vaulting, Secrets Management and Privileged Session Management.

Password Vaulting

Overview

Password vaults store passwords for privileged accounts in a secure repository instead of relying on admin users individually maintaining them. The vault provides secure storage and access management for passwords, and automates their lifecycle, including expiration and rotation. These systems implement the required business processes, including access request and approval processes for “checking out” passwords for use (i.e., they act like a library for passwords, and users “check out” passwords in the same way that patrons check a book out of a library.)

Assessment of Current Situation

[REDACTED]

Secrets Management

Overview

Secrets are not limited to passwords, but they can include any type of information necessary to directly or indirectly secure systems. The following are examples of elements other than passwords that may be stored in a secret sharing solution:

- Hashes
- Certificates
- Cloud tenant information
- API keys
- Database storage information
- Personal Information
- SSH connection information

Applications can be granted access to Secret Management services so that keys and certificates can be managed centrally.

Assessment of Current Situation

[REDACTED]

Privileged Session Management

Overview

PSM solutions provide three main functions for enterprises:

- Keylogging or session recording of admin access, ensuring that all such activities are logged for audit, compliance, and forensic purposes.
- “Supervised” admin access, where a second individual can view an admin’s session in real time, to ensure oversight of high-risk activities.

- Enforcing role-based access on a privileged system, by providing users with just enough permissions to support necessary tasks. This can take the form of restricting permissions granted to the admin’s account, as well as blocking certain commands from being executed on the target device.

Assessment of Current Situation

[REDACTED]

- [REDACTED]
- [REDACTED]
- [REDACTED]

Analysis

Gaps

[REDACTED]

[REDACTED]

Recommendations

This is an area of growth for the organization, not only in the IT sphere. We recommend:

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

Network Infrastructure

Modern network infrastructure is a mix of hardware devices, software applications, and network services, generally broken into three classifications:

- Hardware Infrastructure - typically includes routers, switches, hubs, repeaters, gateways, bridges, and modems.
- Software Infrastructure - includes monitoring and management tools and operating systems.
- Network Services – Virtual representation of hardware or software infrastructure, usually hosted by a vendor.




Network Firewalls


Overview

A Firewall is a network security device that monitors and filters incoming and outgoing network traffic based on an organization's previously established security policies. At its most basic, a firewall is essentially the barrier that sits between a private internal network and the public Internet.

The traditional firewall uses on a complex but static rule base directly assigning IP addresses and ports to control network access:



Next generation firewalls (or NGFWs) implement the capability to build dynamic rulesets that enforce access policy based on user or role:



Assessment of Current Situation

- [Redacted]
- [Redacted]
- [Redacted]
- [Redacted]

Domain Name System

Overview

DNS is the system that translates domain names and hostnames to IP addresses. Standard DNS does not provide any encryption or privacy, and has some interesting complexities associated with it when users are remote. The system generally encompasses two services: Public DNS where devices query a DNS server on their local area network that acts as a recursive server which relays uncached queries to a set of external recursive, root, top level domain, and authoritative servers or Private DNS where the contents are intended to be only available to a limited audience e.g., internal resource addressing.

Obviously, in a zero-trust world where device location is not supposed to be a factor, differentiating public vs private addressing becomes more complex, since those categories are based on network location.

Assessment of Current Situation

[Redacted]

Wide Area Network

Overview

Wide Area Networks (WANs) connect geographically distributed enterprise sites and networks. WANS are generally privately routed over carriers or network providers and may or may not use additional encryption.

Assessment of Current Situation

- [REDACTED]
- [REDACTED]

Load Balancers, Application Delivery Controllers, and API Gateways

Overview

Load Balancers, Application Delivery Controllers (ADCs), and API Gateways are components of networking and IT infrastructure that provide a layer of abstraction between the provider and consumer of a service used to provide better performance, higher scalability and resiliency for applications and services. They can be complex, and often can achieve their goals through several technical approaches. For example, even simple load balancers may use one or more techniques (such as round-robin, random, or load-based) to assign workloads to servers. ADCs and API Gateways reduce server workload by performing certain network, content optimization, and API consolidation functions in front of their back-end servers. This can include capabilities such as SSL termination, content caching, connection multiplexing, traffic shaping, and microservices abstraction or consolidation. In general, these systems provide network and application functions and are not typically considered as security appliances.

The functions provided by these systems are valuable, and will continue to exist in Zero Trust systems, largely unchanged. However, consideration is the potential impact of network topology changes and new use of tunneled encrypted traffic within the Zero Trust system. This traffic will likely become opaque to intermediary components such as these—it all depends on where the PEPs are, and how the PEPs enforce policies.

Assessment of Current Situation

[REDACTED]

Web Application Firewalls


Overview

Web Application Firewalls (WAFs) are security components which sit in front of web servers, protecting them by parsing, monitoring, and securing HTTP traffic. The term WAF is

perhaps a bit misleading, as they are less of a network firewall and more of a reverse proxy, examining incoming HTTP traffic to detect and prevent attacks such as SQL injection and Cross-Site Scripting. WAFs are frequently deployed to protect public-facing web servers, which are probed, scanned, and attacked on a regular basis. WAFs are also deployed to protect internal applications which are only accessible to internal users. In this case, they're protecting against malicious insiders and compromised devices on the internal network.

From a Zero Trust perspective, additional security even for internal applications is recommended, even though it presumes compromise of the environment. I.e., if a Zero Trust system limits access to only those users with legitimate business needs, if 10% of the user population uses this application, the Zero Trust system will eliminate the ability of the remaining 90% of devices to even attempt to attack it. However, the remaining users may well be hosting malicious software which can attempt to attack the application.

Assessment of Current Situation




Virtual Private Networks

Overview

Virtual Private Networks (or VPN) use either hardware, software, or a combination to establish an encrypted network tunnel between remote nodes, permitting application traffic to be securely transmitted through this tunnel, across an untrusted intermediary network. There are three types of VPNs:


- Consumer VPNs are designed to shield end-user Internet-bound traffic from intermediaries, for privacy and security and to bypass regulations.
- Enterprise VPNs connect remote users to an enterprise network through a secure tunnel.
- Site to Site VPNs connect two physical locations through a secure tunnel and are one of the ways enterprises construct Wide Area Networks.

From the perspective of Zero Trust, the concern is primarily with Enterprise VPN as it is a way for network ingress, primarily for user remote access. The client connects to a VPN Server via the Internet, credentials are checked against an IAM (usually the Enterprise IAM) and if successful, access is provided to the internal WAN. The VPN Client may allow split tunneling to Internet addresses or it may force all traffic internally through the tunnel.




Enterprise VPN

Site to Site VPN peers two VPN servers to be able to traverse the Internet and form a seamless contiguous network, either via hardware or software (i.e. SD-Wan). It is generally considered more of an infrastructure than a security solution. However, it is essentially creating a trust relationship between two locations which is antithetical to the Zero Trust strategy.



Site-to-Site VPN



Assessment of Current Situation

[REDACTED]


[REDACTED]

Micro Segmentation

Overview

Micro-segmentation is a network security technique that enables security architects to logically divide the data center into distinct security segments down to the individual workload level, and then define security controls and deliver services for each unique segment. Micro-segmentation enables IT to deploy flexible security policies deep inside a data center using network virtualization technology instead of installing multiple physical firewalls. Also, micro-segmentation can be used to protect every virtual machine (VM) in an enterprise network with policy-driven, application-level security controls. Because security policies are applied to separate workloads, micro-segmentation software can significantly bolster a company's resistance to attack.

In classic perimeter topologies, resources are segmented based on physical or organizational parameters. The concern is more based on addressing and performance than on providing capability for security control. In Zero Trust, network access control is critical to fulfilling the least privilege principle of the strategy and isolating resources from attack. To facilitate this control, it makes sense to segment resources based on workload, at least to the level where policy would be shared (e.g., a segment for all web servers for Application X, etc.)



Assessment of Current Situation

[REDACTED]

[REDACTED]

Analysis

Gaps

[REDACTED]

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

Recommendations

Zero Trust began as a network strategy, so to say that network policy is critical to successful implementation may be an understatement. The Infrastructure and Security teams should initiate a comprehensive review and redesign of the network using next generation concepts. Some recommendations:

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

- [REDACTED]

Intrusion Detection and Prevention Systems

Intrusion Detection Systems (IDS) provide the ability to detect, log, and alert on suspicious activity, and Intrusion Prevention Systems (IPS) add the ability to respond, by blocking or terminating the activity in some way. Intrusion Detection and Prevention Systems (IDPS) typically rely on signatures (pattern matching) and/or anomaly-detection mechanisms (often using statistical analysis or machine learning) to identify unwanted activity. They also often integrate with threat intelligence systems, to obtain updated data to inform their algorithms. These solutions are commonly available as both standalone solutions, as well as within many Next-Generation Firewalls. They can broadly be categorized into host based and network based.

IDPS are most effective when they can be placed into network "chokepoints," and obtain visibility into traffic, and perform deep packet inspection. In a Zero Trust architecture, PEPs are a natural place for these functions to be performed and most modern IDPS, when integrated into a Zero Trust system, effectively can become a PEP by consuming and enforcing Zero Trust policies, and act as a source of events to the PDP, which will in turn trigger other PEPs to act (e.g., initiate changes in user risk levels or access.)

Host Based

Overview

Host-based intrusion detection and prevention systems utilize a software agent running on devices. Host-based systems have the advantage of being able to deeply examine everything occurring in the OS, including process and network activity, and to take local actions. Local presence on the host is advantageous in many Zero Trust deployments, which typically utilize encrypted traffic tunnels across networks.

The feature summary of host based IDPS:

Detection	Prevention
• File integrity monitoring • Process behavior analysis • Network metadata analysis • Local log and event analysis • Log and event forwarding • Device or user behavior monitoring • Software installation or download monitoring.	• Process whitelisting • Process termination • Prevention of software download or installation. • Network connection termination

• Privilege escalation or rootkit detection	
---	--

Assessment of Current Situation

[REDACTED]

[REDACTED]

Network Based

Overview

Network-based IDS and IPS are deployed into an organization’s network, where they can monitor and potentially modify network traffic. The scope and capability of any network-based IDS/IPS system is dependent on the network segments where the system nodes are deployed. For example, an IDS deployed within an employee LAN subnet can potentially examine traffic among devices on that one subnet, or traffic between those devices and remote resources. An IDS on a WAN link connecting distributed data centers may be able to examine the inter-data center traffic but cannot see any local LAN traffic within a given data center or network.

The feature summary of network based IDPS:

Detection	Prevention
• DNS monitoring • Network metadata analysis • Network traffic inspection (deep packet inspection)	• DNS filtering • Network content blocking • Network connection termination • Sandbox “detonation” of suspicious content

Assessment of Current Situation

[REDACTED]

Analysis

Gaps

[REDACTED]

- [REDACTED]
- [REDACTED]

- [REDACTED]

Recommendations

[REDACTED]

- [REDACTED]
- [REDACTED]
- [REDACTED]

Data Protection

Data protection is the process of safeguarding important information from corruption, compromise, or loss. From the Zero Trust perspective, data, which is often the primary target of attackers, is a key enterprise resource that must be secured by PEPs through integration with a PDP, via an identity and metadata-centric policy model.

Organizations are used to looking at data from business value, performance, or compliance perspectives. Zero trust requires looking at data with the additional perspectives of classification, lifecycle, and security.

Data Types and Classification

Overview

Both structured and unstructured data require a level of classification to inform a data security system of how it should be treated. Classification is the process of identifying the level of risk associated with the data, based on its potential impact to the organization. One model for classifying risk is level is presented in the “Standards for Security Categorization of Federal Information and Information Systems” which defines the following levels:

- Low: Loss of confidentiality, integrity, and availability with limited adverse effects on business functions (e.g., marketing or website content)
- Moderate: Loss of confidentiality, integrity, and availability with serious adverse effects on business functions (e.g., customer information, price lists, business plans, or strategy documents)
- High: Loss of confidentiality, integrity, and availability with severe or catastrophic effects on business functions (e.g., source code, banking credentials, or signing credentials)

It is important to note that these classifications do not replace standard permissions, they enhance them by providing an additional level of policy driven protection. A ‘general’ may have permission to both the lunch menu and the launch codes. However, nothing should stop

the general from mailing the lunch menu to a lower-level subordinate, whereas the launch codes should be handled differently.

Assessment of Current Situation


[REDACTED]

Data Lifecycle

Overview


Much like an identity, data has a concrete lifecycle. Data's lifecycle begins at data creation, continues through data usage, and finally ends with data destruction. Each of these stages requires different security methods and approaches:

- Creation - Data can be created by a user or application in a variety of ways either in a structured or unstructured format.




- Usage - There are multiple stages to consider in data usage—data-at-rest, data-in-motion, and data-in-use. These stages all provide both challenges and benefits in data

management and security.



- Destruction - The final stage of the data lifecycle is data destruction. Organizations, especially those in regulated industries or which manage sensitive data, need to define, and enforce data retention policies, which determine how long data should be stored and remain accessible before it's destroyed.

Assessment of Current Situation



Data Security

Overview

Data security is achieved differently across different stages of the data lifecycle. Data can be easily secured for data-at-rest (via full drive encryption) and for data-in-motion (via encrypted transmissions) however for the phase of data-in-use there are Data Loss Prevention (DLP). DLP is comprised of:

- Device control: The means to define how data can be utilized at the device level (e.g., preventing the ability to print or copy-and-paste, or whether USB ports can be utilized on a device).
- Content-aware control: Enforcing and adjusting security controls for data based on the content of the data. This may include data obfuscation.
- Enforced encryption: Ensuring that data-at-rest is encrypted at the drive or physical storage level. Its purpose is to ensure the stored data remains inaccessible even if the device is lost or stolen.
- Data discovery: One of the most important aspects of data security, discovery solutions provide organizations the means to not only find unknown sensitive data but also automate its classification.

DLP solutions provide the capability to enforce access control policies and are a relevant component to Zero Trust.

An additional subset of technologies are Digital Rights Management, which encrypts an object so that accessing the requires a check with a PEP, therefore making sure policy is followed even if the object is outside of the control of the Zero Trust environment.

Assessment of Current Situation

[REDACTED]

Analysis

Gaps

[REDACTED]

Recommendations

[REDACTED]

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

Infrastructure and Platform as a Service

Infrastructure as a Service (IaaS) is straightforward: dynamic provisioning of a full operating system, deployed in a Cloud Service Provider (CSP) environment. Customers are responsible for configuring and maintaining the full OS and its surrounding network, as well as deploying any desired software onto the virtual server. Effectively, the infrastructure that

enterprises are using as a service is a virtual “bare metal” machine onto which they deploy and configure an OS image of their choosing.

Platform as a Service (PaaS) also is hosted by a CSP but encompasses a wide variety of functions, models and services. The term serverless computing is also frequently used when discussing PaaS, referring to the ability to deploy custom code without the need to deploy a full server operating system. Serverless functions are deployed into a PaaS environment, which provides the surrounding infrastructure for accessing, managing, and launching.

IaaS and PaaS are highly relevant to Zero Trust, as these platforms represent a large proportion of how applications are built and deployed today. While these platforms do have sophisticated and powerful access control models, they’re largely designed to protect services based within their cloud environments, and not serve as broad enterprise security solutions for all users across multiple heterogeneous environments. However, given the wide capability of these services and that they were built with Zero Trust in mind, it is important to consider how they fit into the strategic landscape.

There are three key areas to consider when looking at IaaS and PaaS Zero Trust strategy:

- Cloud Services
- Service Meshes


Cloud Services

Overview

Cloud services refer to any IT services that are provisioned and accessed from a cloud computing provider. This is a broad definition for services that range from IaaS hosting of a VM, to a PaaS serverless databases, all the way to simple Time web service that provides the current time of day via a cloud base API.

The challenge in implementing Zero Trust with Cloud Services is deployment of a PEP on resources which is foundational to the strategy. This is not difficult IaaS since the PEP can be placed on the resource itself. However, PaaS by its nature does not allow modification. Similarly, PaaS services usually have their own security control system, built by the vendor.

The work around for this difficulty is to rely on the vendor tools for policy enforcement or if not possible, place a PEP on the access gateway between the resource.



Assessment of Current Situation

[REDACTED]

Service Meshes

Overview

Service meshes are platforms for operating, controlling, and managing large-scale containerized (microservices) workloads, with an emphasis on managing the communications between microservices. If configured, they provide automated baseline traffic resilience, service metrics collection, distributed tracing, traffic encryption, protocol upgrades, and advanced routing functionality for all service-to-service communication.

While they are not Zero Trust out of the box, the granularity of control lends itself to a service mesh to be configured to support PEPs, usually through the implementation of internal traffic proxies. (Essentially internally based PaaS micro segmentation)

Assessment of Current Situation


[REDACTED]

Analysis

Gaps

[REDACTED]

Recommendations



We recommend that system development with these technologies consider the necessity of a PEP or PDP integration with the holistic zero trust environment.

Software as a Service

SaaS are publicly accessible web applications in which the service provider (vendor) hosts, manages, and maintains the infrastructure and with which the subscriber performs (and manages) the designated application function across the Internet. SaaS applications are typically multi-tenant for efficiency, with each subscriber restricted to accessing only their private data.

From a Zero Trust perspective there does not seem to be much to do regarding architectural security, at least in comparison to IaaS and PaaS. A SaaS system is publicly accessible by design, with any user on the Internet able to access them over an encrypted (HTTPS) connection. As a result, there will be little or no opportunity to secure access on an infrastructure level via PEPs.

However, this does not mean that there are not Zero Trust friendly controls that can be applied, and many SaaS systems have extensive Zero Trust controls on the application layer. That said, this strategy examines infrastructure styled controls that can be implemented:

- Federated Authentication
- Secure Web Gateways
- Cloud Access Security Brokers

Federated Identity Management

Overview

Federated identity management is where the SaaS application leverages an external identity provider for user authentication, through standardized mechanisms such as SAML and OpenID Connect. Effectively, this is another way to enforce identity-centric aspects of Zero Trust access controls, which is independent of network-level security.

From a technical perspective, users cannot authenticate directly into the SaaS application; the SaaS app either retrieves a current authentication token from the user's browser or redirects the browser to the identity provider for authentication. This then uses whatever authentication factors and contextual controls are configured within the identity provider.

Is important to note that this is authentication only, most SaaS applications handle authorization internally and do not have mechanisms to consume external contextual information to make authorization decisions.

Assessment of Current Situation

[REDACTED]

Secure Web Gateways

Overview

A secure web gateway is an on-premises or cloud-delivered network security service. Sitting between users and the Internet, secure web gateways provide advanced network protection by inspecting web requests against company policy to ensure malicious applications and websites are blocked and inaccessible.

Generally secure web gateways are designed to prevent access to inappropriate sites and scan traffic for data that could contain malicious code, otherwise the gateway does not involve itself on what happens during application access.

From a Zero Trust perspective, there is some value in using an Enterprise Secure Web Gateway as an additional layer of authentication and authorization for SaaS services, particularly if the SaaS service can be configured to only accept traffic from the Secure Web Gateway. In other words, the gateway becomes a public PEP. (It should be noted that things like latency should be considered in repurposing a Secure Web Gateway as a PEP, since it the role could add significant transaction overhead.)

Assessment of Current Situation

[REDACTED]

Cloud Access Security Brokers

Overview

A cloud access security broker (or CASB) is designed to be a PEP between enterprise users and cloud service providers. CASBs can combine multiple different security policies, from authentication and credential mapping to encryption, malware detection. Unlike a Secure Web Gateway, CASBs are designed to actively scan for SaaS use and offer controls on API or application use.

From a Zero Trust perspective, it would generally make sense to use internal SaaS controls as PEPs for defined policies. However, a CASB's role is not to serve as a PEP for a properly secured SaaS system but for those that are being used by "Shadow IT" or groups that sit outside of IT control.

Assessment of Current Situation

[REDACTED]

Analysis

Gaps

- [REDACTED]
- [REDACTED]
 - [REDACTED]
 - [REDACTED]

Recommendations

We recommend that the following actions be taken with SaaS systems to meet Zero Trust strategy criteria:

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

Tactics, Operations and Logistics

As this is a strategic document, it does not attempt to solve tactical, operational, or logistical issues. However, some general commentary can be made on some of the adjustments that will need to be made in these areas for a Zero Trust strategy to succeed.

Tactically, [REDACTED] has a diverse and powerful set of security tools, many of which are designed to support Zero Trust. Although some capability gaps exist and will need to be addressed and some tools should be replaced as better fit the strategy, there is no reason that a Zero Trust initiative could not begin immediately. What will need to change is that tactics will need to be adjusted to view tactical decisions through the interconnected prism of the Zero

Trust strategy versus siloed assessment (i.e., this tool provides better security for an individual task.)

Operationally, while [REDACTED] has excellent operational and administrative Security staff, focus is needed to transition on making the processes they manage automated and driven by well written, comprehensive policy. This is not to drive efficiency or any commentary on staff quality but instead, needed to drive the requirements of the Zero Trust strategy for the security environment to be dynamically adaptive, flexible, and reactive to threats.

Logistically, management should consider this as a 2–3-year transformational process, that will require steady investment and organizational attention. Zero Trust is a strategic initiative, not a project and will require both technical and educational efforts to affect a long-term culture change.

Next Steps

Recommendations have been presented in depth in the individual strategy sections. Here are the five steps our organization should take to begin the Zero Trust transformation. In order:

1. The management towers come to final consensus on Zero Trust answer these questions:
 - a. How will Zero Trust provide value to [REDACTED]?
 - b. What are the agreed principles, ideals(requirements) and strategy?
 - c. What does IT and the business community need to do to transform our thinking about security?
2. Provide workshops on Zero Trust value, concept and strategy to ALL IT employees and [REDACTED] liaisons including the topics of:
 - a. Strategy
 - b. Zero Trust Networks
 - c. Policy Enforcement Points
 - d. Policy Decision Points
 - e. Monitoring
 - f. Data Security and Classification
3. Security and the [REDACTED] work together to develop secure but reflective of business needs documented policies for:
 - a. Identity management for External Users (e.g., partners, vendors, customers, etc.)
 - b. File Sharing with External Participants
 - c. Next generation authentication (e.g., Windows Hello, etc.)
 - d. Bring Your Own Device (a.k.a. BYOD)
 - e. Data classification and retention

4. Implement a cloud network segment that implements these Zero Trust strategies:
 - a. No perimeter
 - b. Dynamic network access control based on user credentials.
 - c. Micro segmentation
 - d. IDPS monitoring and automated response.
5. Move several low sensitivity systems into this network as part of a pilot phase. Out of this pilot we will determine:
 - a. If additional tooling is required.
 - b. If the existing organization can support a Zero Trust infrastructure or if changes to structure or personnel are required
 - c. How to move systems from classis to modern methodologies
 - d. Metrics for success (e.g., a scorecard for Zero Trust implementation)

Figure credits: Several diagrams in this document are adapted from Jason Garbis and Jerry W. Chapman, *Zero Trust Security: An Enterprise Guide* (Apress, 2021).