Zero Trust Architecture and Multi-cloud environments.

Joseph Costantini, PhD.

"As agencies continue to use cloud technology, they shall do so in a coordinated, deliberate way that allows the Federal Government to prevent, detect, assess, and remediate cyber incidents. To facilitate this approach, the migration to cloud technology shall adopt zero trust architecture, as practicable. The CISA shall modernize its current cybersecurity programs, services, and capabilities to be fully functional with cloud-computing environments with zero trust architecture. The Secretary of Homeland Security acting through the Director of CISA, in consultation with the Administrator of General Services acting through the FedRAMP within the General Services Administration, shall develop security principles governing Cloud Service Providers (CSPs) for incorporation into agency modernization efforts." (1)

Introduction.

According to NIST, CISA and FEDRAMP cloud services can be offered in four distinct ways. The security model or architecture needed to secure the cloud depends directly on the components that are used to build the multi-cloud ecosystem. The basic security issues involved in using a cloud service model are establishing and maintaining "**visibility**" and "**control**" over critical assets including data, communications, and operations such as update, patching, audit, maintenance, and backup, and recovery. (3, 4)

The Executive Order moves the Federal government toward secure cloud services, zero-trust architecture, and mandates deployment of multifactor authentication and encryption within a specific time. The Order establishes baseline security standards for development of software sold to the government, a Cybersecurity Safety Review Board, a standard playbook and set of definitions for cyber incident response and enables government-wide endpoint detection and response system for improved information sharing within the Federal government. It also creates a cybersecurity event log, and requires amendments to the FAR to align with requirements of the Executive Order. (1, 4)

The Multi-Cloud Environment.

The components of a multi-cloud environment may include several different approaches which themselves must be secured (6). According to NIST these approaches are:

> *Private: The cloud infrastructure is provisioned for exclusive use of an organization comprised of multiple customers (e.g., an agency with multiple business units). It may be owned, managed,*

and operated by the organization, an authorized third party, or combinations of them. The infrastructure may exist on-premises with the organization or off-premises with a cloud provider.

Community: The cloud infrastructure is provisioned to a specific community of consumers that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more organizations, an authorized third party, or some combination of these entities. The infrastructure may exist on or off premises.

Public: The cloud infrastructure is provisioned for use by the general public. It may be owned, managed, and operated by one or more organizations, an authorized third party, or some combination of these entities. The infrastructure exists off-premises.

Hybrid: The cloud infrastructure is a composition of two or more of the above deployment models (i.e., Private, Community, or Public). In this instance, multiple deployment models are connected through a standardized or proprietary technology offered by the provider to maintain compatibility of data and applications. (4)

Cloud Services/Components.

According to CISA, a multi-cloud environment can include any or all the service/component types listed below; each with differing visibility and control issues, including:

- ***Infrastructure as a Service (IaaS):*** *Provides only a base infrastructure (Virtual machine, Software Define Network, Storage attached). End users have to configure and manage platform and environment, and deploy applications on it.*
- ***Software as a Service (SaaS****): Software "on-demand." Typically accessed by users using a thin client via a web browser. In SaaS everything can be managed by vendors: applications, runtime, data, middleware, operating systems, virtualization, servers, storage and networking components.*
- ***Platform as a Service (PaaS):*** *Provides a platform to allow end users to develop, run, and manage applications without the complexity of building and maintaining the infrastructure.*
- ***Container as a Service (CaaS):*** *Container engines, orchestration and the underlying compute resources are delivered to users as a service.*
- ***Function as a Service (FaaS):*** *A platform allowing customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure.* (2)

Zero Trust.

CISA describes an architecture for the Zero Trust Security Model: "*The Zero Trust security model eliminates implicit trust in any one element, node, or service and instead requires continuous verification of the operational picture via real-time information from multiple sources to determine access and other system responses. In essence, a Zero Trust Architecture allows users full access but only to the bare minimum they need to perform their jobs. If a device is compromised, zero trust can ensure that the damage is contained. The Zero Trust Architecture security model assumes that a breach is inevitable or has likely already occurred, so it constantly limits access to only what is needed and looks for anomalous or malicious activity. Zero Trust Architecture embeds comprehensive security monitoring; granular risk-based access controls; and system security automation in a coordinated manner throughout all aspects* *of the infrastructure to focus on protecting data in real-time within a dynamic threat environment. This data-centric security model allows the concept of least-privileged access to be applied for every access decision, where the answers to the questions of who, what, when, where, and how are critical for appropriately allowing or denying access to resources based on the combination of sever."* (2)

Zero-trust is a response to modern enterprise network trends that include remote users, bring your own device (BYOD), and cloud- based assets that are not located within an enterprise-owned network "boundary or perimeter." Zero trust focus is on protecting resources (assets, services, workflows, network accounts, etc.) rather than network segments, applications, or operating systems, as the location is no longer seen as the prime component to the security posture of the resource.

Agencies migrating to cloud deployments are to apply zero trust principles and transition their environment commensurate with their risk tolerance. To achieve this, agencies should focus on strengthening fundamental areas of cybersecurity including identity management, asset management, network security, data protection, application security, and visibility integrated across on-premise and cloud environments. That is, they should manage the "security posture" of the overall system in a consistent and coherent manner, rather than as simply a mass of independently operating components.

To begin this process, agencies must identify their "security posture" which includes all critical assets and processes that apply to their system, assign a "risk rating for each asset," and then manage the security of those assets according to a consistent policy. However, those assets are expected to change whenever the business environment changes, so this identification of should be seen as part of a continuous risk assessment process. To accomplish this goal, the security process must be automated and controlled. We call this "managing the security posture of the system". (5)

Security Posture Management (SPM).

Security Posture Management (SPM) tools take advantage of automation capabilities to identify assets, and to correct issues without human intervention, including:

- Identification of the environment footprint and monitoring creation of new instances including storage resources (often called buckets); provide policy visibility to ensure consistent enforcement across all providers in the multi-cloud environment; Scanning for misconfigurations and improper settings that are vulnerable to exploitation; Scanning storage buckets for misconfigurations that could make data accessible to the public; Audit for adherence to regulatory compliance mandates such as HIPAA, PCI DSS, and GDPR; Perform risk assessments against frameworks and external standards such as those put forth by the National Institute of Standards and Technology (NIST), DoD, the DNI, or FEDRAMP; Verify that operational activities are performed as expected; and promote automated remediation.
- SPM tools and capabilities can be used to support agencies' transition to a zero-trust approach to security. Any data that is acquired, moved, stored or used must be protected: 1. at rest, 2. in transit, and 3. in use. SPM tools provide continuous monitoring and alerting on anomalous activity in access logs and help to identify and prevent misconfigurations that may lead to data leakage and data loss. (5)

Data Security.

Of particular interest is the protection of data. Data must be protected: 1. at rest, 2. in transit, and 3. while in use. Generally, cryptography, continuous monitoring, and access controls are used to protect data at rest and data in transit; but only access control is available when data is in use*. Regardless, access to sensitive data should never be dependent on "where the data is currently located," but rather its "defined sensitivity". In a sense, protection must accompany data wherever it resides and however it is used. (2, 3, 4)

**Homomorphic cryptography is a research effort that may eventually allow 'data in use' to remain encrypted during use, but this technology is currently ineffective.* (6)

The Reference Monitor/Validation Mechanism. Zero trust could be summarized as a set of principles (or tenets) used to plan and implement an IT architecture as defined in NIST SP 800-207 but grouped as tenets relating to network identity, device health, or data flows. To enforcing these tenets, we choose to re-define the old concept of a reference monitor and reference validation mechanism.

The function of the reference monitor is to validate all references to programs, data, peripherals, and other assets made by an operating system for programs in execution against those authorized for the subject (user, process, etc.) The Reference Monitor not only is responsible to assure that the references are authorized to shared resource objects, but also to assure that the reference is of the right kind (i. e. read, or read and write, execute, move, delete, etc.)

An implementation of the reference monitor concept is traditionally called the Reference Validation Mechanism (RVM) - the RVM is a combination of hard-ware, software, and data that implements the reference monitor concept. The mechanism should be tamperproof, always invoked, and simple enough to be subject to analysis and test to assure its correctness. The reference validation mechanism for a time-sharing operating system (TSOS) was given the name "Security Kernel" by Mitre Corporation (Ames, Gasser, and Schell) in the early years of computer security and applied well to an independent entity (generally a single operating system) that was bounded and closed within that boundary. The idea of a reference monitor, must be extended to handle modern information technology environments such as the multi-cloud environment, which is may well be an open, or partially open system; that is where defined boundaries no longer have meaning – that is, there is simply not inside or outside, there is just 'the system.' (7, 8)

Tenets of Zero Trust Architecture.

The foregoing discussion leads us to the tenets of a zero-trust architecture, which according to NIST, are:

Policy. A typical enterprise has a wide collection of network identities: end users, service accounts, applications, data stores, mobile devices, data stores, compute resources (real and virtual), remote sensors/actuators, etc. Some end users may have multiple network identities, and some identities may only be used by hardware/software components.

The enterprise needs a "governance policy and structure" that defines "authorized operations," so that only authorized operations are performed, and only when the entity for which the operation is performed has been properly authenticated and authorized. The enterprise needs to consider whether their current identity governance policies are mature enough and where and how authentication and authorization checks are currently performed.

All resource authentication and authorization are dynamic and strictly enforced before access is allowed; and all data sources and computing services are considered resources. But it is the case that some components (e.g., IoT sensors) may not be able to support some solutions such as configuration agents, app sandboxing, etc. so alternatives that use the underlying network infrastructure may be needed. If the resource lacks security capabilities, then the enterprise may need to add a component to provide that functionality; but only, when necessary, as this will increase the threat surface of the overall system. The enterprise monitors and measures the integrity and security posture of all owned and associated assets for all aspects of cyber hygiene: configuration management, patching, application loading, etc.

Continuous activity monitoring. The state of resources should be monitored, and appropriate action taken when new information such as a new vulnerability, new threat, or an attack is reported or observed. The confidentiality and integrity of data on the resource should be protected. This requires enterprise admins to know how resources are configured, maintained, and monitored.

- **Data Flows.** All communications are secured regardless of network location, with the assumption that an attacker is present on the network and could observe/modify communications and data. Appropriate safeguards should be in place to protect the confidentiality and integrity of data in transit. If the resources cannot provide this functionality natively, a separate component may be necessary.
- **Access.** Access to individual enterprise resources is granted on a per-session basis. In an ideal zero trust architecture, every unique operation would undergo authentication and authorization before it is performed. This is not always possible and other mitigating solutions such as logging, and backups may be needed to detect and recover from unauthorized operations.
- **Enforcement.** Enterprise administrators will need to learn how to enforce fine grain access policies on individual resources. If the current set of tools do not allow this, other solutions such as logging, versioning tools, or backups may help mitigate risk. Access to resources is determined by dynamic, perhaps rapidly changing, policy including the observable state of client identity, application/service, and the requesting asset and may include other behavioral and environmental attributes.
- **Allow List.** In zero-trust, the default behavior for all resources is to deny all action with an “allow list” or other method. The members of this “allow list” must authenticate themselves and prove they meet the enterprise policy to be granted access. This may include meeting requirements such as client software versions, patch level, geolocation, historical request patterns, etc. It may not be possible to perform all checks immediately prior to the access request. The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications, and uses it to maintain or improve its security posture.
- **Audit and Monitoring.** System logs and threat intelligence are used to refine or change policy in response to new information. For example, whenever a new vulnerability in a software component is identified, a zero-trust enterprise would move quickly to quarantine the affected resources until they can be patched or modified. Enterprise admins will need to set up and maintain a comprehensive monitoring and patching program for the enterprise and should consider how automated tools could assist in responding to newly discovered threats and vulnerabilities. (6)

Security Strategy. A zero-trust security strategy authenticates and authorizes every device, network flow, and connection based on dynamic policies, using context from as many data sources as possible.

To expand, the zero-trust security model ensures data and resources are inaccessible by default. Users can only access them on a limited basis under the right circumstances, this is traditionally called "a least-privilege policy." A zero-trust security model verifies and authorizes every connection, such as when a user connects to an application or software to a data set via an application programming interface (API). It ensures the interaction meets the conditional requirements of the organization's security policies.

To successfully implement a zero-trust architecture, organizations need to connect information from across each security domain. Security teams across the company must agree on priorities and align access policies. They must secure all connections across the business, from data to users and devices to applications, workloads, and networks. This architecture requires a well-planned strategy and roadmap to implement and integrate security tools to achieve specific business-focused outcomes. To make a zero-trust model work, adopters should:

- Make an organization-wide commitment.
- Catalog all IT and data assets and assign access policy or rights based upon roles.
- Lock down some common vulnerabilities.
- Classify data according to its sensitivity.
- Segment networks to prevent lateral movement when data is breached.
- Isolate and protect workloads during virtual machine and cloud server cross-movement. (10)

Based on the principle of verified trust— "*to trust, you must first verify*"—Zero-trust eliminates the inherent trust that is assumed in the traditional corporate network. Zero-trust reduces risk across all environments by establishing strong identity verification, validating device compliance prior to granting access, and ensuring least privilege access to explicitly authorized resources.

Zero-Trust requires that every transaction between systems (user identity, device, network, and applications) is validated and proven trustworthy before the transaction can occur. In an ideal Zero-trust environment, the following behaviors are required:

- **Identities are validated and secured with multifactor authentication everywhere.** Using multifactor authentication to eliminate password expiration issues and may use biometrics ensures strong authentication for user-backed identities.
- **Devices are managed and validated as healthy.** Device health validation is required. All device types and operating systems must meet a required minimum health state as a condition of access to any resource.
- **Telemetry is pervasive.** Pervasive data and telemetry are used to understand the current security state, identify gaps in coverage, validate the impact of new controls, and correlate data across all applications and services in the environment. Robust and standardized auditing, monitoring, and telemetry capabilities are core requirements across users, devices, applications, services, and access patterns.
- **Least privilege access is enforced.** Limit access to only the applications, services, and infrastructure required to perform the job function. Access solutions that provide broad access to networks without segmentation or are scoped to specific resources, such as broad access VPN, must be eliminated.

Verify identity – multi-factor authentication is used to include all users accessing resources from outside the corporate network. The massive increase in mobile devices connecting to corporate resources. Microsoft, for example, uses the Azure Authenticator application, but other methods may also be useful.

Verify device - Devices can be enrolled using a device-management system. Microsoft uses Windows Autopilot for device provisioning, which ensures that all new Windows devices delivered to employees are already enrolled in our modern device management system. Devices accessing the corporate wireless network must also be enrolled in the device-management system. This includes both company-owned devices and personal BYOD devices. If employees want to use their personal devices, the devices must be enrolled and adhere to the same device-health policies that govern corporate-owned devices. For devices where enrollment in device management isn't an option, Microsoft uses Windows Virtual Desktop that creates a session with a virtual machine that meets the device-management requirements. This allows individuals using unmanaged devices to securely access select company resources.

Verify access – Users and devices are segmented across purpose-built networks, migrating all employees to use the internet as the default network, and automatically routing users and devices to appropriate network segments. A device registration portal is used for wireless network rollout. This portal allows users to self-identify, register, or modify devices to ensure that the devices connect to the appropriate network segment. Through the portal, users can register guest devices, user devices, and IoT devices.

Specialized segments support various IoT devices and scenarios used throughout the organization.

Verify services - All applications and services are verified. All legacy applications or implementing solutions for applications and services that can't natively support conditional access systems must be "modernized," to eliminate dependency on VPN and the corporate network. Auto-VPN automatically routes users through the appropriate connection. The goal is to eliminate the need for VPN and create a seamless experience for accessing corporate resources from the internet. (6, 11)

Use cases for Zero trust The zero-trust security model can be predicated on the principle of identity-based access. In this model for any machine or user to do anything, they must authenticate who or what they are, and then their identity and associated policies define what they're allowed to do. For this use case, the HashiCorp has defined its "reference monitor" as three system wide components called the vault, consul, and boundary, as follows:

- HashiCorp Vault - Machine Authentication & Authorization: Enterprises need to centrally store, access, and distribute dynamic secrets like tokens, passwords, certificates, and encryption keys across any public or private cloud environment.
- HashiCorp Consul - Machine-to-machine Access: Enterprises need to enable machine-to-machine access by enforcing authentication between applications and ensuring only the right machines are talking to each other.
- HashiCorp Boundary - Human Access and Authorization: Enterprises need to secure access to applications and critical systems with fine-grained authorizations without managing credentials or exposing a network. (11)

Issues and misconceptions with ZTA Implementations

We find many advisories and models for "attaining zero trust" in the sales and marketing literature of various organizations with a variety of different approaches advised; be careful to recognize that much of this marketing and sales "hype." Dan Lohrmann notices that "Wherever you turn, experts and thought leaders are singing its praises. An online search can easily find thousands of articles, speeches and presentations on why zero trust is the must-have paradigm for all things cybersecurity moving forward".

Consequently, some misconceptions have been identified by security practitioners based on experience in implementing a zero-trust environment; and some of them work and some do not. Jeff Capone advises "that some enterprises have implemented zero-trust regimes, followed shortly by a significant breach of security;" he suggests:

- **Microsegment Resources.** Very granular controls should be used including device, device posture, location, time and other elements as authentication factors. He also advises that data security should have similar granular controls. Finally, he suggests that endpoint applications, networks, and devices should have such controls. Finally, he suggests that established policies should be persistently enforced everywhere
- **Establish and Enforce Policies Everywhere.** Access policies should be established and enforced for users, data, applications, networks, and system tools such as clipboards. No data should flow between secured and unsecured applications. Data security has previously focused largely on data access. But once data is accessed, the user typically has broad rights to use and transfer the data without additional security control, policy needs to apply to data, not just files. As users create new content, compare that content to previously secured content; if the content is similar, automatically secure the new content with the same permissions as the previously secured data. Make sure to monitor small data segments as they move from file to file and apply permissions accordingly.
- **Introduce Visibility and Automation.** Visibility and automation are two of the cross-functional principles of zero trust. Granular logging and reporting should enable orchestration tools to look for anomalies and suspicious behavior. Log all data access attempts, regardless of whether you allow or deny the action. Your log should include user, application, device, location, time and other metadata. Proper logging will allow orchestration tools to detect potential malware and suspicious user behavior while also creating audit and compliance reports. (13)

Issues with Zero-trust.

Dan Lohrmann advises us that it's important to keep the approach in check – that is, the approach needs to be practical and limited to the issue. Some definitions that go beyond questioning trust in digital assets; delving into trusting people and facilities is impractical and inadvisable – we might better handle human resource issues with personnel security and physical access with physical security. He goes farther in advising: "The fatal flaw was anthropomorphizing the network and moving over concepts like trust have no business belonging in digital environments." (13)

NIST tells us to stick to the principle "never trust, always verify." Palo Alto Networks tells us, "Zero trust is not about making a system trusted, but instead about eliminating trust." (15)

Mary Pratt advises us that “Zero trust is a security concept centered on the belief that organizations should not automatically trust anything inside or outside its perimeters and instead must verify anything and everything trying to connect to its systems before granting access.” (13)

John Kindervag, the originator of the zero-trust idea (16), calls for simplicity, suggesting that: Online Trust is a vulnerability and is real big problem in the digital world; but People are not packets. “People aren’t the issue; packets are the issue.” (15)

Dan Lohrmann further suggests that: “Digital trust and human trust are two separate things; Zero trust only applies to digital systems. People are not necessarily untrustworthy, but at the same time they are not packets. Zero-trust only applies to the zeros and ones that traverse our various digital systems. Malcolm Gladwell calls human beings trust engines (17); Morton Deutsch talks about how trust is the willingness of one individual to be vulnerable to another individual; and applies this to business management - not digital systems.” (18)

References:

1. Executive Order 14028. https://www.cisa.gov/executive-order-improving-nations-cybersecurity
2. CISA Cloud Security Technical Reference Architecture
3. NIST Zero Trust (NIST SP 800-207 [1]). https://csrc.nist.gov/publications/detail/sp/800-207/final
4. FEDRAMP. https://www.gsa.gov/technology/government-it-initiatives/fedramp
5. Security Posture Management. https://cloudsecurityalliance.org/blog/2019/10/01/cloud-security-posture-management-why-you-need-it-now/
6. Homomorphic Cryptography. https://www.forbes.com/sites/bernardmarr/2019/11/15/what-is-homomorphic-encryption-and-why-is-it-so-transformative/?sh=4b71bc297e93
7. The security kernel (Ames, Schell and Gasser). https://people.cs.ksu.edu/~danielwang/Investigation/System_Security/01654439.pdf
8. Computer Security Planning Study (Anderson). https://csrc.nist.gov/csrc/media/publications/conference-paper/1998/10/08/proceedings-of-the-21st-nissc-1998/documents/early-cs-papers/ande72a.pdf
9. IBM Zero-trust. https://www.ibm.com/topics/zero-trust
10. Microsoft Zero-trust. https://www.microsoft.com/en-us/insidetrack/implementing-a-zero-trust-security-model-at-microsoft
11. HashiCorp Zero-trust. - cybercybersecurityHashiCorpHashiCorp BoundaryHashiCorp ConsulHashiCorp FederalHashiCorp VaultIC InsiderIC Insidersmeteredzero trust
12. Palo Alto Networks Zero-trust. - https://www.paloaltonetworks.com/cyberpedia/what-is-a-zero-trust-architecture
13. Are We Taking Zero Trust Too Far in Cybersecurity (Lohrmann). https://www.govtech.com/opinion/are-we-taking-zero-trust-too-far-in-cybersecurity
14. Mary Pratt - *CSO Online*
15. John Kindervag - https://deloitte.wsj.com/articles/john-kindervag-the-hallmark-of-zero-trust-is-simplicity-01618513330
16. John Kindervag - https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf

17. Malcom Gladwell - https://qideas.org/qmoments/human-beings-are-designed-for-trust/
18. Morton Deustch - https://www.beyondintractability.org/artsum/deutsch-cooperation