IMPLEMENTING ZERO TRUST WITHIN VIDEO DISTRIBUTION SOLUTIONS

ABSTRACT

The United States Government and private industry face increasingly sophisticated and persistent threat campaigns that target its technology infrastructure. Video, Audio and KVM Distribution solutions need to shift towards a Zero Trust Architecture. VDS can no longer be a "stovepipe" physically secured architecture, but must embrace zero trust tenets delivering a distributed information and situational awareness capability. It is part of an integrated network.

Learn about the process, and the importance of beginning your implementation planning now. Discover the steps in creating and executing your Zero Trust VDS roadmap.

Howard Sutton, P.Eng MBA CFA

[January 2022]

Implementing Zero Trust within Video Distribution System (VDS) Solutions

"The United States Government and private industry face increasingly sophisticated and persistent threat campaigns that target its technology infrastructure, threatening public safety, privacy, corporate data, damaging the American economy, and weakening trust in Government."[1]

While Zero Trust primarily targets cybersecurity networks, many of the same risks and responses apply in closed and segregated VDS networks. Legacy Perimeter-based (location centric) VDS security platforms and practices fail to protect against both internal and network connected external threats. Live video and data distributed within physically secure environments, such as a Sensitive Compartmented Information Facility (SCIF) or briefing centers, are not safe and secure according to Zero Trust Tenets. As a result, any lack of a VDS Zero Trust strategy places current installations at risk and will inhibit any future transformation to a secure, distributed, anytime, anywhere architecture (data/video centric). As the Federal Government shifts towards cloud-based services which include VDS, Zero Trust practices need to be designed into all VDS solutions today. Agencies must plan and adapt to VDS de-perimeterization. The Time is Now.

What is Zero Trust and the Zero Trust Maturity Model?

In the current threat environment, the Federal Government and industry can no longer depend on perimeter-based defenses to protect critical video systems at rest and in transit. Meeting this challenge will require a major shift in how agencies approach their VDS solutions.

As described in the OMB Department of Defense Zero Trust Reference Architecture, "the foundational tenet of the Zero Trust Model is that no actor, system, network, or service operating outside or within the security perimeter is trusted. Instead, we must verify anything and everything attempting to establish access. It is a dramatic paradigm shift in philosophy of how we secure our infrastructure, networks, and data, from verify once at the perimeter to continual verification of each user, device, application, and transaction."[2]

This strategy envisions a Federal zero trust architecture that:

- Bolsters strong identity practices across Federal agencies;
- Relies on encryption and application testing instead of perimeter security;
- Recognizes every device and resource the Government has;
- Supports intelligent automation of security actions; and
- Enables safe and robust use of cloud services.

---

[1] Federal Zero Trust Strategy OMB

[2] DoD Zero Trust Architecture

Zero Trust is a security concept anchored on the principle that organizations need to proactively secure all access to data and resources to reduce security risks. It is an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. Zero Trust assumes no implicit trust granted to assets or users based solely on their network or physical location. The use of authentication and authorization for both device and subject are discrete functions before a resource is established.

Its goal is to ensure the trustworthiness of the user, device or service requesting access to an agency resource at any time. The ZT infrastructure should also allow for continuous assessment and authorization based on various conditions – such as location, device or time of day and any others while monitoring threats, vulnerabilities, risk, behaviour and other relevant information. If changes in circumstance or environment are detected, earlier permissions may be restricted or revoked.

The Zero Trust (ZT) Foundation:

- ZT provides a consistent security strategy of users accessing data that resides anywhere, from anywhere in any way;
- ZT assumes a "'never trust and always verify" stance when accessing services and/or data;
- ZT requires continuous validation and authorization based on agency criteria; and
- ZT increases visibility and trust in decisions.

Zero Trust Assertions:

- The network is always assumed to be hostile;
- External and internal threats exist on the network at all times;
- Network locality is not sufficient for deciding trust in a network;
- Every device, user and network flow is authenticated and authorized; and
- Policies must be dynamic and calculated from as many sources of data as possible.

Source: ACT-IAC American Council of Technology – Industry Advisory Council

Executive Order 14028 directs agencies to focus on meeting key baseline security measures across the government, such as universal logging, multi-factor authentication (MFA), reliable asset inventories, and ubiquitous use of encryption, and to adopt a zero trust architecture.

This memorandum requires agencies to achieve specific zero trust security goals by the end of Fiscal Year (FY) 2024. Grouped using the five pillars that underpin the zero trust maturity model of the Cybersecurity and Infrastructure Security Agency (CISA), those goals include:

1. **Identity:** Agency staff use an enterprise-wide identity to access the applications they use in their work.
2. **Devices:** The Federal Government has a complete inventory of every device it operates and authorizes for Government use and can detect and respond to incidents on those devices.
3. **Networks:** Agencies encrypt all traffic within their environment and begin segmenting networks around their applications.

4. **Applications**: Agencies treat all applications as internet-connected, routinely subject their applications to rigorous testing, and welcome external vulnerability reports.
5. **Data/Video**: Agencies are on a clear, shared path to deploy protections that make use of thorough data/video categorization. Agencies are taking advantage of IP (cloud security services) to monitor access to their sensitive data and have implemented enterprise-wide logging and information sharing.

Source: OMB Federal Zero Trust Strategy (Sept 2021) adapted to VDS by the author

VDS Architectures

VDS platforms are driven by operational functional requirements that include, amongst others, connecting to a variety of sources, delivering ultra high resolution (4k), low latency video and using KVM. The technology to deliver this capability was, and to some degree, still is, baseband architectures. Baseband Routers have dominated the market share because they have historically been the only platform available that could deliver a reliable, high resolution, low latency (1ms) capability These uncompressed signals are driven in many formats including DVI, HDMI, SDI and KVM baseband switches. However, this architecture was developed decades ago for simple plug and play use within stovepipes and has critical security issues. There is NO encryption capability nor use of Keys nor Certificates for authorization. All signals are openly viewable. As such, for current certifications, including Secure KVM Switching, the Security is based on physical security and use only within 100% TRUSTED end user environments. Baseband routers provide security thru routing tables; the mapping of inputs to outputs. Very primitive, but widely implemented. There is no device authentication let alone the concept of mutual authentication. This architecture cannot support Zero Trust implementation and must, over time, be replaced with a ZT capable architecture.

WHAT IS A COMMON VDS ARCHITECTURE TODAY?

VDS = Video Distribution System – It distributes information from multiple sources to multiple destinations

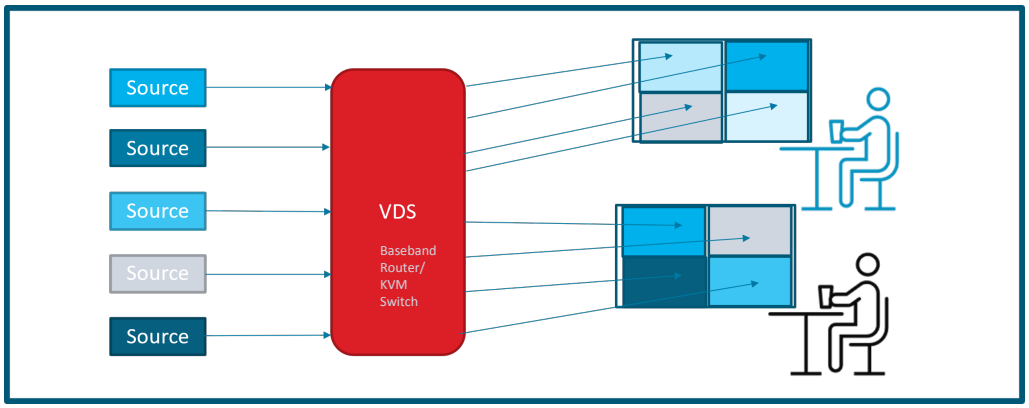

Capabilities: Reliable, low latency, mostly 720, 1080, some 4Kp30, Locally attach to router (includes Secure KVM), Closed Room/Campus Setups , Small to hundreds endpoints, “plug and play” openly viewable video flows, Requires physical security, 100% trusted users

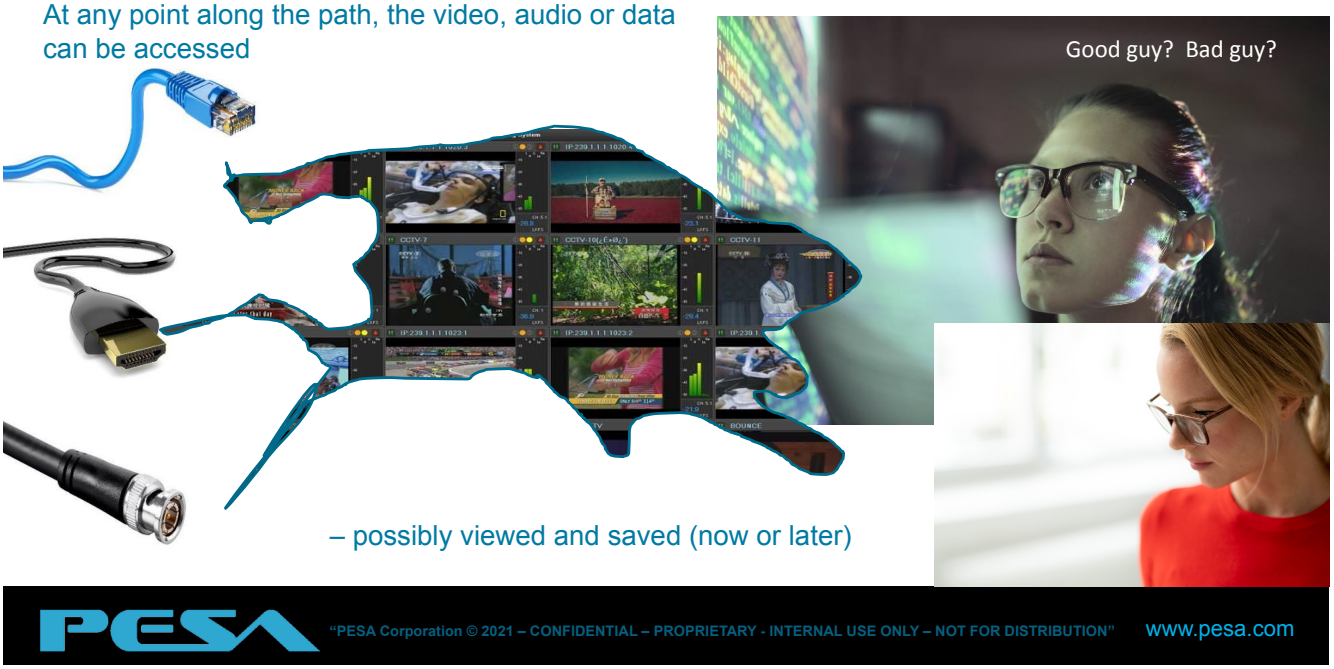

VDS legacy platforms do not provide the architecture to build a strategic ZT VDS. So what does?

Secure IP VDS Zero Trust Design Considerations

The Secure IP VDS breakthrough for live environments including command and control centers, conference rooms, sensitive compartmented information facilities (SCIFs), operations centers, amongst others, has primarily resulted from the tremendous improvement in quality of visually lossless, low latency video encoding/decoding (codecs) along with the significant drop in total cost of ownership. Because IP is somewhat sensitive to payload, IP codecs utilize lower bandwidth in transmission versus baseband systems that require uncompressed signals. For example, at 4K-60 frames per second (fps) mezzanine IP codecs require approximately 850Mb/s of bandwidth per video flow versus 12Gb/s per video signal in baseband (93% less bandwidth). Using an MPEG or motion codec, 4K-60 can be transmitted using only 20MB or less (98% less bandwidth). Combining this efficiency with the drop in price of 10Gb/s and100Gb/s IP switches, the functionality of a secure IP design cannot only match a baseband solution, but provides significant feature capabilities beyond legacy systems including a scalable, distributed, mobile, cloud and application interoperability. Most importantly, IP VDS can be architected to support Zero Trust Models.

VDS Baseband vs IP Zero Trust Report Card:

Let's evaluate Baseband vs IP VDS against the strategy for Federal zero trust architecture:

- Bolsters strong identity practices across Federal agencies;

  Baseband Fail: Baseband VDS systems do not use authentication, certificates or keys. There is no device authentication. Limits identity practice.

  IP Pass: IP VDS system uses authentication, certificates and keys in every connection. Includes all device authentication. Supports strong identity practice

- Relies on encryption and application testing instead of perimeter security;

  Baseband Fail: Baseband VDS does not support encryption. All video, audio, USB and KVM are viewable. Security relies on physical "Castle and Moat".

  IP Pass: IP VDS Networks encrypt every flow (audio, video, USB, KVM).

- Recognizes every device and resource the Government has;

  Baseband Pass: Baseband VDS keeps track of connected devices.

  IP Pass: IP VDS keeps track of connected devices.

- Supports intelligent automation of security actions;

  Baseband Fail: Baseband VDS can use automation, but there are few security parameters to monitor other than routing tables.

  IP Pass: IP VDS tracks every port, device, location, individual, geography.

- Enables safe and robust use of cloud services.

  Baseband Fail: Baseband VDS utilizes protocols that are not supported on the cloud nor are they networkable.

  IP Pass: IP VDS architecture supports cloud services.

Elements for Zero Trust VDS Implementation

Requirements for ZT VDS include integrating a secured content flow while utilizing a control system plane to provide an identity practice with comprehensive monitoring, action, reporting and auditing.

A typical base set of capabilities is illustrated in Figs. A, B, C, below. Each and every content and control flow must be secured utilizing FIPS 140-2 (Moving to FIPS 140-3). The control system must deliver all elements and inputs required to drive a policy engine from source to destination, including the network switches. This may be coupled with automation, AI, and behavior identity analysis applications, amongst others. With each and every flow mutually authenticated and encrypted utilizing dynamic rotating keys, it is possible to implement Multiple Independent Levels of Security (MILS). MILS implementation should have the capability to be authorized based on user level, device level, group level and area level. Each site will have it's own policy whether to use air gaps or not for MILS. While certifications are evolving to include cloud-based practices, many dated VDS profiles still rely on physical security for stovepipe implementations. IP VDS certification requirements should include a Network Protection Profile for switches, DoD APL and NIST FIPS140-2 (minimum) crypto module libraries implemented within the application. Ensure the crypto libraries are fully implemented in all the applications, otherwise, the integrity of platform is compromised.

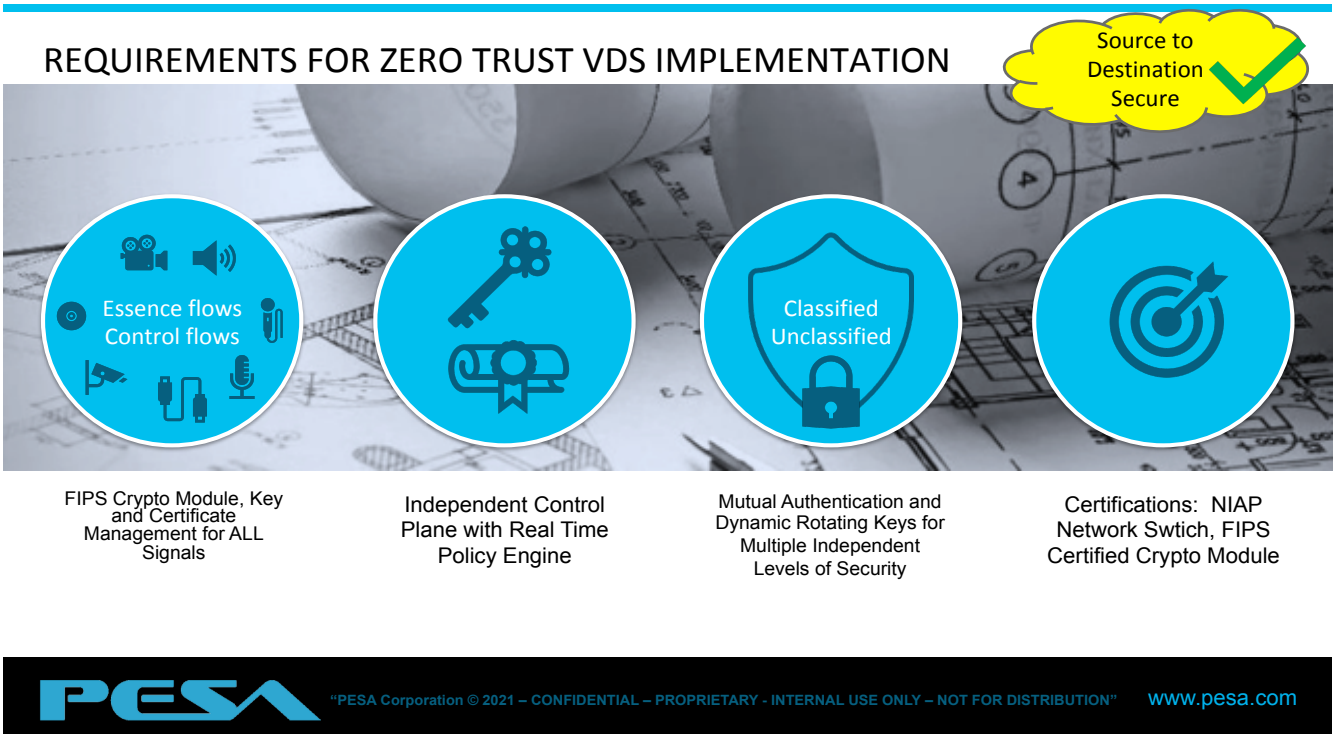

Fig A

FIPS 140-2 ENCRYPTION

- FIPS 140-2 Encryption
- DoDIN mandates ALL IP use FIPS 140-2/3 certified modules
- PESA Crypto Module (CMVP #4021)
  - ✓ FIPS 140 = Cryptography
  - ✓ It is NOT…
    - Certificates
    - Public Keys
    - Tokens
    - Others

ALL IP communications = FIPS!

 www.pesa.com

Fig B

AUTHORIZATIONS FOR USERS, DEVICES, GROUPS AND AREAS

CROSS DOMAINS – (MILS)

Secure
Network MILS

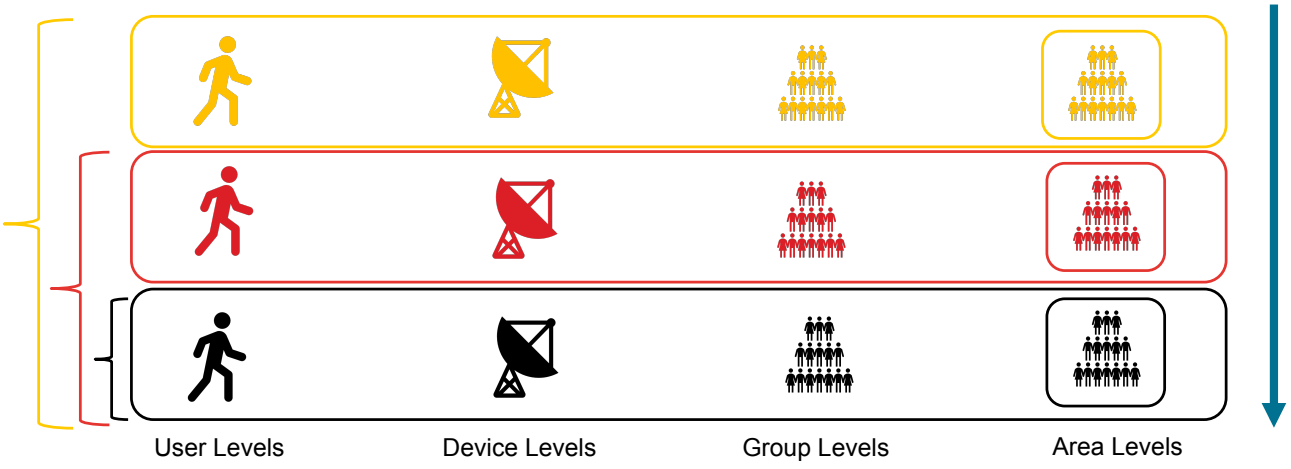

 www.pesa.com

Fig C

VDS Network/Environment Components

In a VDS ZT environment, there should be a separation (logical or possibly physical) of the communication flows used to control and configure the network and application/service communication flows used to perform the actual work of the organization. This is often broken down to a *control plane* for network control communication and a *data plane* for application/service communication flows [Gilman].

The control plane is used by various infrastructure components (both enterprise-owned and from service providers) to maintain and configure assets; judge, grant, or deny access to resources; and perform any necessary operations to set up communication paths between resources. The Data/Video plane is used for actual communication between software components. This communication channel may not be possible before the path has been established via the control plane. For example, the control plane could be used by the Policy Administrator (PA) and Policy Enforcement Point (PEP) to set up the communication path between the subject and the enterprise resource. The application/service workload would then use the data plane path that was established.

Network Requirements to Support ZTA

The NIST 800-207 Zero Trust Architecture document applies to VDS networks. Below is a summary of important implementation capabilities from this document.

1. Enterprise assets have basic network connectivity. The local area network (LAN), enterprise controlled or not, provides basic routing and infrastructure. The remote enterprise asset may not necessarily use all infrastructure services.
2. The enterprise must be able to distinguish between what assets are owned or managed by the enterprise and the devices' current security posture. This is determined by enterprise-issued credentials and not using information that cannot be authenticated information (e.g., network MAC addresses that can be spoofed).
3. The enterprise can observe all network traffic. The enterprise records packets seen on the data plane, even if it is not able to perform application layer inspection (i.e., OSI layer 7) on all packets. The enterprise filters out metadata about the connection (e.g., destination, time, device identity) to dynamically update policies and inform the PE as it evaluates access requests.
4. Enterprise resources should not be reachable without accessing a PEP (policy enforcement point). Resources accept custom-configured connections only after a client has been authenticated and authorized. These communication paths are set up by the PEP.
5. The data plane and control plane are logically separate. The policy engine, policy administrator, and PEPs communicate on a network that is logically separate and not directly accessible by enterprise assets and resources. The data plane is used for application/service data traffic. The policy engine, policy administrator, and PEPs use the control plane to communicate and manage communication paths between assets. The PEPs must be able to send and receive messages from both the data and control planes.

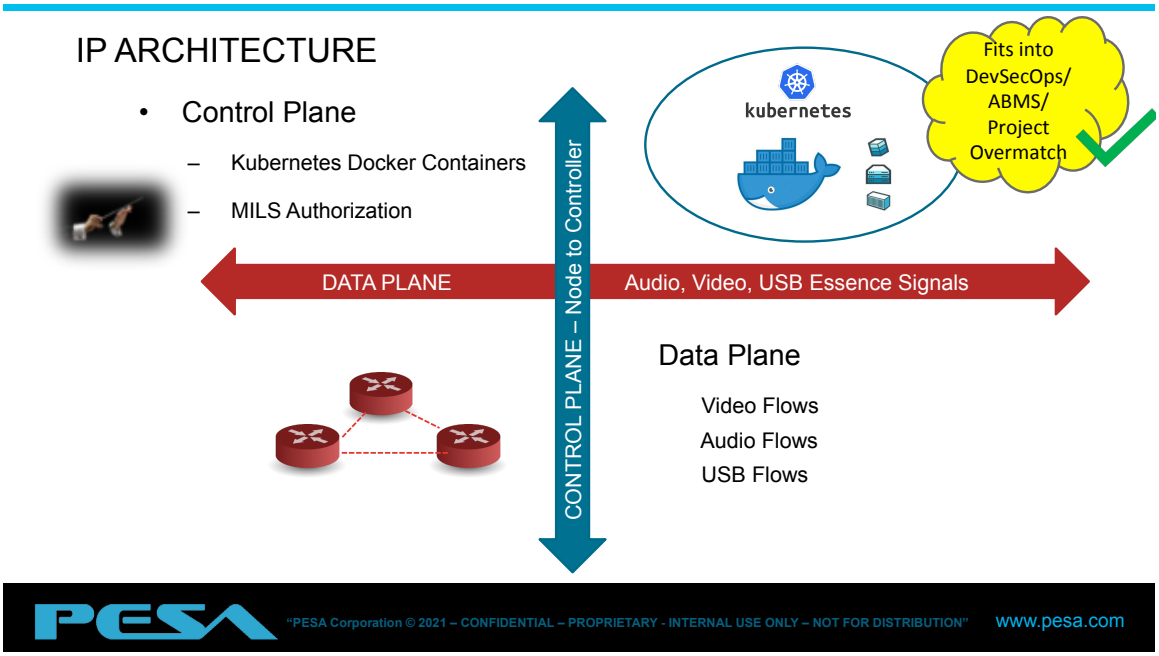

Fig D

6. Enterprise assets can reach the PEP component. Enterprise subjects must be able to access the PEP component to gain access to resources.
7. The PEP is the only component that accesses the policy administrator as part of a business flow. Each PEP operating on the enterprise network has a connection to the policy administrator to establish communication paths from clients to resources. All enterprise business process traffic passes through one or more PEPs.[3]

Implementing ZT VDS

In August 2020, NIST published Zero Trust Architecture (Pub. 800-207). Section 7.3 defines steps to "Introduce ZTA to a Perimeter-Based Architected Network". VDS environments today are traditionally Perimeter-Based Networks using only physical security to protect openly viewable content from baseband switches. The "castle and moat" approach towards VDS is common practice and assumes everyone in the "castle" is trusted. As such, these steps are a very good introductory guideline for integrators, designers and departments to follow when moving beyond risky Stovepipes towards IP ZT VDS. In this section, we take the roots of Pub. 800-207 and apply it to VDS. All quoted text below is from NIST Publication 800-207.

"Before undertaking an effort to bring ZTA to an enterprise, there should be a survey of assets, subjects, data flows, and workflows. This awareness forms the foundational state that must be reached before a ZTA deployment is possible. An enterprise cannot determine what new processes or systems need to be in place if there is no knowledge of the current state of operations. These surveys can be conducted in parallel, but both are tied to examination of the business processes of the organization. These steps can be mapped to the steps in the RMF [SP800-37] as any adoption of a ZTA is a process to reduce risk to an agency's business functions." The pathway to implementing a ZTA can be visualized in Figure E, below.

[3] NIST 800-207 Zero Trust Architecture

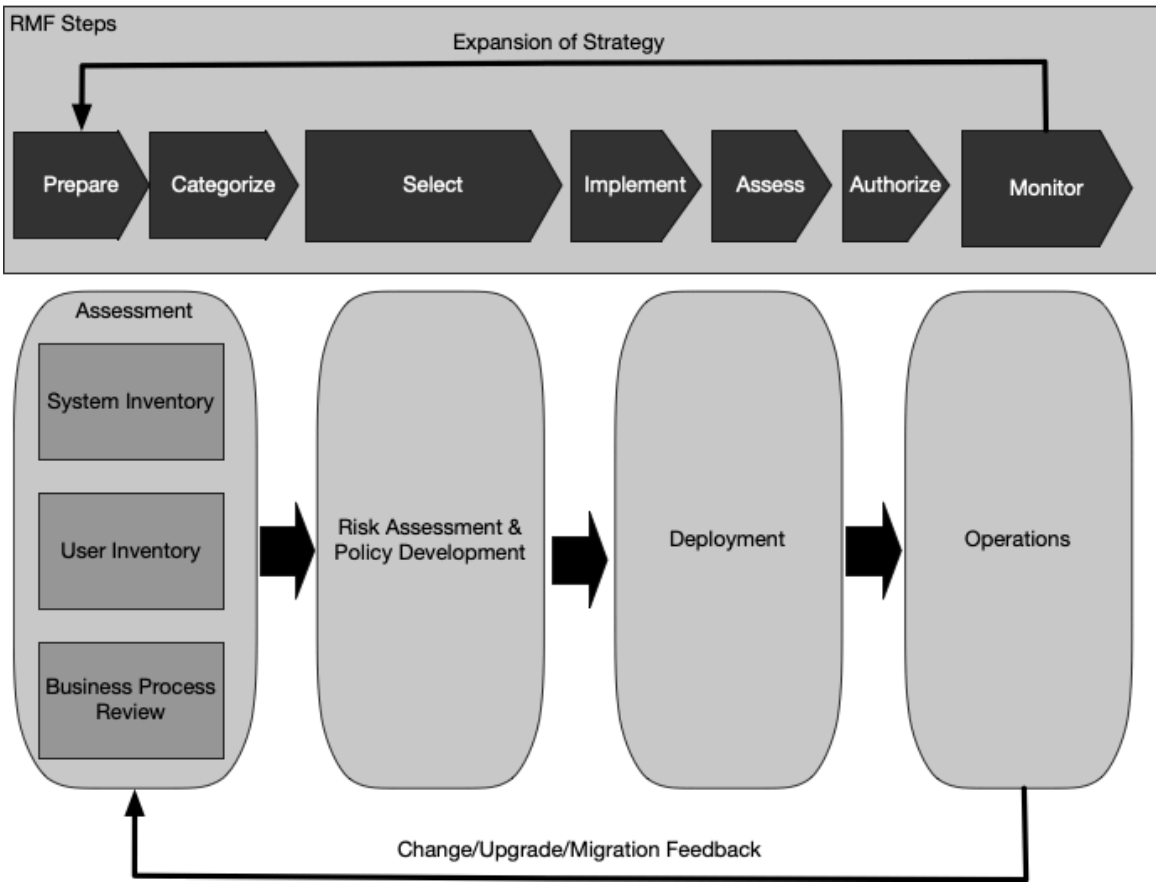

Figure E: ZTA Deployment Cycle (Source: NIST Pub 800-207)

Identify Actors on the Enterprise

“For a zero trust enterprise to operate, the PE (Policy Engine) must have knowledge of enterprise subjects. Subjects could encompass both human and possible service accounts that interact with resources.”

Identify Assets Owned by the Enterprise

One of the key requirements of ZTA is the ability to identify and manage devices. ZTA also requires the ability to identify and monitor non enterprise-owned devices (may not apply to specific sites) that may be on enterprise-owned network infrastructure or that access enterprise resources. The ability to manage enterprise assets is key to the successful deployment of ZTA. This includes hardware components (e.g., video transmitters, video receivers, application servers, network switches, cameras, DVR, IoT devices etc.) and digital artifacts (e.g., user accounts, applications, digital certificates). It may not be possible to conduct a complete census on all enterprise-owned assets, so an enterprise should consider building the capability to quickly identify, categorize, and assess newly discovered assets that are on enterprise-owned infrastructure.

This goes beyond simply a catalog and maintaining a database of enterprise assets. This also includes configuration management and monitoring. The ability to observe the current state of an asset is part of the process of evaluating access requests. This means that the enterprise must be able to configure, survey, and update enterprise assets, such as virtual assets and containers. This also includes both its physical (as best estimated) and network location. This information should inform the PE when making resource access decisions. For a VDS implementation, this requires the Control System to collect this relevant information and be applied and acted upon within the Policy Engine.

Identify Key Processes and Evaluate Risks Associated with Executing Process

"The third inventory that an agency should undertake is to identify and rank the business processes, data flows, and their relation in the missions of the agency. Business processes should inform the circumstances under which resource access requests are granted and denied. An enterprise may wish to start with a low-risk business process for the first transition to ZTA as disruptions will likely not negatively impact the entire organization. Once enough experience is gained, more critical business processes can become candidates."

Formulating Policies for the ZTA Candidate

"The process of identifying a candidate service or business workflow depends on several factors: the importance of the process to the organization, the group of subjects affected, and the current state of resources used for the workflow. The value of the asset or workflow based on risk to the asset or workflow can be evaluated using the NIST Risk Management Framework [SP800-37]."

Identifying Candidate Solutions

"Once a list of candidate business processes has been developed, enterprise architects can compose a list of candidate solutions. Some deployment models are better suited to particular workflows and current enterprise ecosystems. Likewise, some vendor solutions are better suited to some use cases than others."

Enclave-Based Deployment

In this model, the resource components may not reside on assets in front of individual resources but instead reside at the boundary of a resource enclave (e.g., on-location data center) as shown in Figure F. Usually, these resources serve a single business function (C2, Conference Rooms etc.) and may or may not be able to communicate directly to a network. This deployment model may also be useful for enterprises that use a private (or secure cloud) network to communicate between enclaves for business processes (e.g., Distributed C2). In this model, the entire VDS network (possible Leaf/Spine switch architecture) is integrated source to glass and controlled thru a single Control Plane. Of vital importance is the implementation of a SDN (Software Defined Network) to enable not only Quality of Service Video capability, but enhanced Control Plane functionality delivering real time data to the PE.

Figure F: Enclave Model

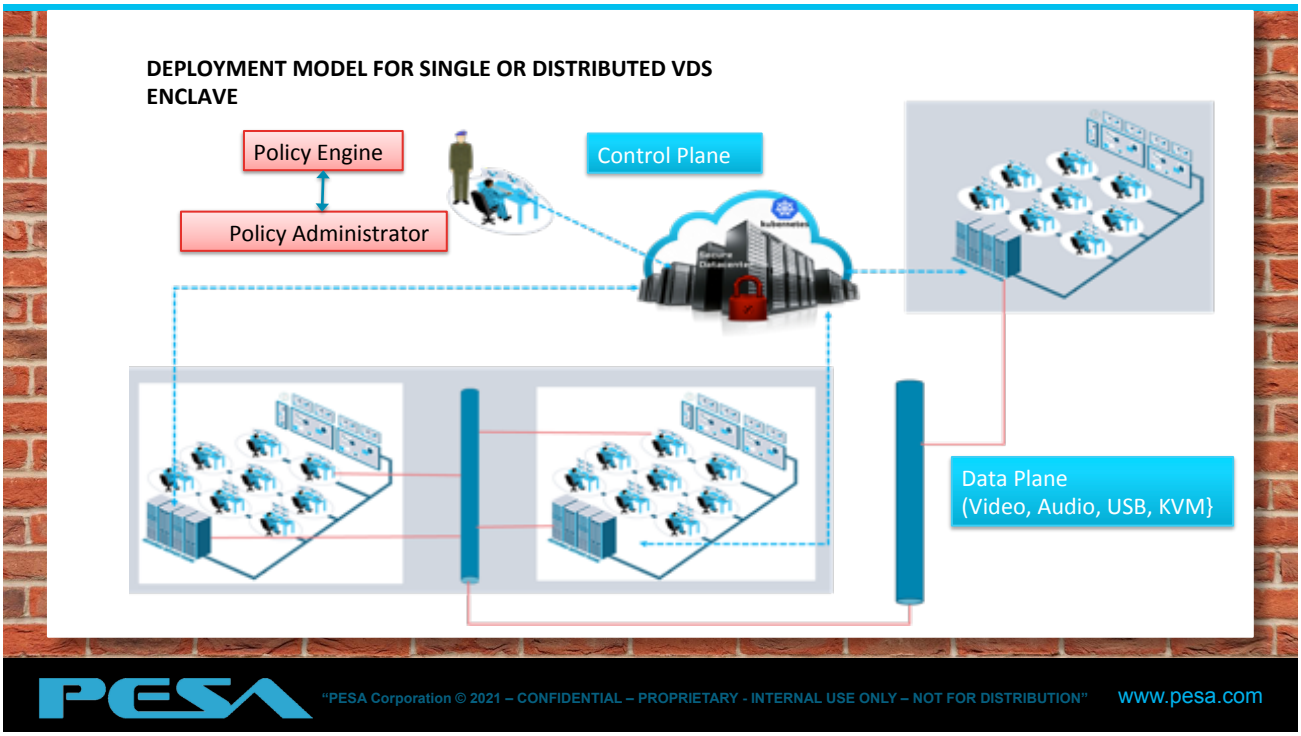

The enclave model is useful for enterprises that have single or multiple control rooms that require secure connectivity. The environment needs a robust asset and configuration management program in place to install/configure the device agents along with real time information from every device and network connection. Each individual resource must be protected.

Initial Deployment and Monitoring

"Once the candidate workflow and ZTA components are chosen, the initial deployment can start. Enterprise administrators must implement the developed policies by using the selected components but may wish to operate in an observation and monitoring mode at first. Few enterprise policy sets are complete in their first iterations: important user accounts (e.g., administrator accounts) may be denied access to resources they need or may not need all the access privileges they have been assigned".

"The new ZT business workflow could be operated in reporting-only mode for some time to make sure the policies are effective and workable. This also allows the enterprise to gain an understanding of baseline asset and resource access requests, behavior, and communication patterns. Reporting-only means that access should be granted for most requests, and logs and traces of connections should be compared with the initial developed policy. Basic policies such as denying requests that fail MFA or appear from known, attacker controlled or subverted IP addresses should be enforced and logged, but after initial deployment, access polices should be more lenient to collect data from

actual interactions of the ZT workflow. Once the baseline activity patterns for the workflow have been established, anomalous behavior can be more easily identified. If it is not possible to operate in a more lenient nature, enterprise network operators should monitor logs closely and be prepared to modify access policies based on operational experience." [4]

Expanding the ZTA

When enough confidence is gained and the workflow policy set is refined, the enterprise enters the steady operational phase. The network and assets are still monitored, and traffic is logged, but responses and policy modifications are done at a lower tempo as they should not be severe. The subjects and stakeholders of the resources and processes involved should also provide feedback to improve operations. At this stage, the enterprise administrators can begin planning the next phase of ZT deployment. Like the previous rollout, a candidate workflow and solution set need to be identified and initial policies developed.

However, if a change occurs to the workflow, the operating ZT architecture needs to be re evaluated. Significant changes to the system—such as new devices, major updates to software (especially ZT logical components), and shifts in organizational structure—may result in changes to the workflow or policies. In effect, the entire process should be reconsidered with the assumption that some of the work has already been done. For example, new devices have been purchased, but no new user accounts have been created, so only the device inventory needs to be updated.

Conclusions

The road to Zero Trust is a journey rather than an entire replacement of infrastructure and processes. As Federal agencies will be reporting on zero trust adoption and automation efforts in the annual cybersecurity reports to the Office of Management and Budget, organizations should look to incrementally implement zero trust principles and policies as well as VDS technology solutions that comply and protect high value video assets. The shift to ZT will occur most likely in a hybrid fashion with baseband legacy systems gradually shifting towards ZT IP VDS platforms as IT modernization initiatives are implemented. Investment in IT modernization should include moving to an architecture based on ZT principles.

Additional References

The Federal Government has been preparing for the transition to a zero trust architecture. Several agencies have published architectural models that can be helpful to other agencies:

- CISA's Zero Trust Maturity Model is a high-level overview of zero trust "pillars" that shows how agencies may progress to "Advanced" and "Optimal" states and describes how CISA service-offerings align to these pillars.

[4] NIST 800-207 Zero Trust Architecture

- CISA's Cloud Security Technical Reference Architecture, co-authored with the United States Digital Service and FedRAMP, provides a more granular reference for secure cloud architectures and migration strategies.
- NIST's SP 800-207, Zero Trust Architecture provides a consensus definition and framework for the key tenets of zero trust architecture, while describing several different approaches to zero trust architecture that organizations with different risk postures and skillsets can adopt.
- The NIST National Cybersecurity Center of Excellence (NCCoE) has initiated "Implementing a Zero Trust Architecture," a collaboration with industry partners to apply the concepts in NIST SP 800-207 to a conventional enterprise architecture.
- The Department of Defense's Zero Trust Reference Architecture comprehensively describes potential security features and architectural controls that the Department plans to execute across its systems.