



Department of Defense (DoD) Zero Trust Reference Architecture

Version 2.0

July 2022

**Prepared by the Defense Information Systems Agency
(DISA) and National Security Agency (NSA) Zero Trust
Engineering Team**

July 2022

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

Document Prepared By	Date
Name: Robert Freter DISA Zero Trust Program Lead (ID2)	June 2022

- 1 PURPOSE AND STRATEGIC GOALS 9
 - 1.1 Introduction 9
 - 1.2 Purpose 9
 - 1.3 Scope 10
 - 1.3.1 Stakeholders 10
 - 1.3.2 Organization of the Reference Architecture 10
 - 1.3.3 Timeframe 12
 - 1.4 Vision and Goals (CV-1)..... 13
 - 1.4.1 Vision and High-Level Goals (CV-1) 14
 - 1.4.2 Zero Trust Strategy 15
 - 1.5 Cybersecurity (Transition) Problem Statement (OV-1) 16
 - 1.6 Overall Target Environment (OV-1) 18
 - 1.7 Assumptions 19
 - 1.8 Constraints 20
- 2 PILLARS AND PRINCIPLES 20
 - 2.1 Overview 20
 - 2.2 Concept and Tenets of Zero Trust 20
 - 2.3 Pillars 21
 - 2.4 Reference Architecture Principles (OV-6a) 23
- 3 CAPABILITIES 25
 - 3.1 Capabilities Taxonomy (CV-2)..... 25
 - 3.2 FFP: Pillars, Resources & Capability Mapping 31
- 4 USE CASES 35
 - 4.1 Data Centric Security Protections (OV-1)..... 35
 - 4.2 Data-Centric Security Protections (OV-2) 37
 - 4.3 Data Encryption Protections (OV-2) 39
 - 4.4 Coordinating Policy for Data-Centric Security Protections (OV-2)..... 41

- 4.5 Data Analytics & AI (OV-1)..... 42
- 4.6 Data Analytics & AI (SV-1)..... 44
- 4.7 Centralized Orchestration & Policy Management (OV-1)..... 45
- 4.8 Centralized Orchestration & Policy Management (OV-2)..... 46
- 4.9 Dynamic, Adaptive Policy Feedback Loop (OV-1)..... 47
- 4.10 VPN-Less Implementation (OV-1) 48
- 4.11 East-West Segmentation (OV-1) 49
- 4.12 Global Uniform Device Hygiene (OV-1) 50
- 4.13 Global Uniform Device Hygiene (OV-2) 52
- 4.14 Dynamic, Continuous Authentication (OV-1)..... 54
- 4.15 Dynamic, Continuous Authentication (OV-2)..... 56
- 4.16 Conditional Authorization (OV-1)..... 60
- 4.17 Conditional Authorization (OV-2)..... 62
- 5 TECHNICAL POSITIONS 63
 - 5.1 Emerging Technologies..... 63
 - 5.2 Standards, Associated Architectures and Guides 64
 - 5.3 Linkages to Other Architectures..... 65
 - 5.3.1 DoD Cybersecurity Reference Architecture (CS RA) Integration 65
 - 5.3.2 DoD ICAM Reference Design (RD)..... 66
 - 5.3.3 NIST Special Publication 800-207 Zero Trust Architecture 67
- 6 SECURITY ASSESSMENT 68
 - 6.1 Governance 68
 - 6.2 Data Governance (OV-2)..... 68
 - 6.3 Securing Supply Chain (OV-2) 70
- 7 ARCHITECTURE PATTERNS..... 71
 - 7.1 Architecture Patterns (CV-4) 71
 - 7.1.1 Domain Policy Enforcement for Resource Access (SV-1)..... 72
 - 7.1.2 Software Defined Perimeter (OV-2) 73

- 7.1.3 ZT Broker Integration (SV-1)..... 74
- 7.1.4 Micro Segmentation (SV-1)..... 74
- 7.1.5 Macro Segmentation (SV-1)..... 78
- 7.2 External Services 78**
 - 7.2.1 SvcV-1: External Services(SvcV-1)..... 79
 - 7.2.2 SvcV-2: Enterprise Federated Identity Service (SvcV-2) 80
- 8 TRANSITION ARCHITECTURE PLANNING (FFP)..... 81**
 - 8.1 Maturity Model (FFP)..... 81**
 - 8.2 Baseline (OV-1)..... 82**
 - 8.3 Transition (OV-1) 83**
- 9 APPENDIX (AV-2) 84**
 - 9.1 Systems 85**
 - 9.2 Services 90**
 - 9.3 General Terms 92**
 - 9.4 DIV-1 93**
 - 9.5 StdV-1-2 References 96**
 - 9.6 Capability Table..... 97**
- 10 REFERENCES 104**

July 2022

LIST OF TABLES

Table 1 Reference Architecture Principles (OV-6A).....	24
Table 2 Design Pattern Table (CV-4).....	71

July 2022

LIST OF FIGURES

Figure 1 Legend for Performers.....	12
Figure 2 Zero Trust Vision (CV-1).....	13
Figure 3 Cybersecurity Problem Statement (OV-1)	16
Figure 4 Target Environment (OV-1).....	18
Figure 5 Zero Trust Pillars.....	22
Figure 6 Capability to Pillars Mapping (FFP)	26
Figure 7 Zero Trust Authentication and Authorization Capability Taxonomy (CV-2).....	27
Figure 8 Zero Trust Infrastructure, Workload and Data Capability Taxonomy (CV-2)	28
Figure 9 Zero Trust Analytics and Orchestration Capabilities Taxonomy (CV-2)	29
Figure 10 Zero Trust Enabling Capabilities Taxonomy (CV-2).....	30
Figure 11 FFP: Pillars, Resources & Capability Mapping (CV-7).....	31
Figure 12 Data Centric Security Protections (OV-1).....	35
Figure 13 Data-Centric Security Protections (OV-2)	37
Figure 14 Data Encryption Protections (OV-2).....	39
Figure 15 Coordinating Policy for Data-Centric Security Protections (OV-2)	41
Figure 16 Big Data Analytics & AI (OV-1)	42
Figure 17 Data Analytics & AI (SV-1).....	44
Figure 18 Centralized Orchestration & Policy Management (OV-1).....	45
Figure 19 Centralized Orchestration & Policy Management (OV-2).....	46
Figure 20 Dynamic, Adaptive Policy Feedback Loop (OV-1).....	47
Figure 21 VPN-Less Implementation (OV-1)	48
Figure 22 East-West Segmentation (OV-1).....	49
Figure 23 Global Uniform Device Hygiene (OV-1).....	50
Figure 24 Global Uniform Device Hygiene (OV-2).....	52
Figure 25 Dynamic, Continuous Authentication (OV-1)	54
Figure 26 Dynamic, Continuous Authentication (OV-2)	56
Figure 27 Performers Requiring Authentication.....	58
Figure 28 Conditional Authorization (OV-1)	60
Figure 29 Conditional Authorization (OV-2)	62
Figure 31 Standards Profile for DoD Zero Trust Architectures	64
Figure 32 Securing the Supply Chain (OV-2)	70

Figure 33 Domain Policy Enforcement for Resource Access (SV-1) 72

Figure 34 Design Pattern: Software Defined Perimeter (OV-2)..... 73

Figure 35 SoS Design Pattern: Zero Trust Broker Integration (SV-1) 74

Figure 36 SoS Micro Segmentation (SV-1)..... 75

Figure 37 SoS Micro Segmentation (SV-1)..... 76

Figure 38 SoS Micro Segmentation (SV-1)..... 77

Figure 39 Design Patterns: SoS Macro Segmentation (SV-1)..... 78

Figure 40 External Services (SvcV-1) 79

Figure 41 Enterprise Federated Identity Service (SvcV-2)..... 80

Figure 42 ICAM Service (SvcV-2)..... 80

Figure 43 Maturity Model (FFP) 81

Figure 44 Transition Architecture Baseline (OV-1) 82

Figure 45 Transition Architecture Transition (OV-1) 83

1 PURPOSE AND STRATEGIC GOALS

1.1 Introduction

“Zero Trust is the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. Zero Trust assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the Internet) or based on asset ownership (enterprise or personally owned).”¹ Zero Trust (ZT) requires designing a consolidated and more secure architecture without impeding operations or compromising security. The classic perimeter/defense-in-depth cybersecurity strategy repeatedly shows to have limited value against well-resourced adversaries and is an ineffective approach to address insider threats.

The DoD Cybersecurity Reference Architecture (CS RA) documents the Department’s approach to cybersecurity and is being updated to become data centric and infuse ZT principles.

ZT supports the 2018 DoD Cyber Strategy, the 2019 DoD Digital Modernization Strategy, the 2021 Executive Order on Improving the Nation’s Cybersecurity, and the DoD Chief Information Officer’s (CIO) vision for creating “a more secure, coordinated, seamless, transparent, and cost-effective architecture that transforms data into actionable information and ensures dependable mission execution in the face of a persistent cyber threat.”² ZT should be used to re-prioritize and integrate existing DoD capabilities and resources, while maintaining availability and minimizing temporal delays in authentication mechanisms, to address the DoD CIO’s vision.

1.2 Purpose

An architecture is built for a defined purpose and should answer a specific set of questions to enabling data-driven, informed decisions. The Reference Architecture (RA) establishes a framework that provides guidance via architectural Pillars and Principles. It identifies which of the overall strategic needs (goals and objectives) are the focus of the RA. The RA is a conceptual, capability-centric description of the architecture and primarily supports capability planning, portfolio management, and Information Technology (IT) investment decisions. It establishes high-level service and operation concepts, architectural questions of importance, and technology opportunities and constraints that shape the domain of an approach. The RA also includes a synopsis of current industry and DoD approaches and identifies key determining standards that together describe constraints and opportunities.

¹ NIST SP 800-207 Zero Trust Architecture, August 2020

² DoD Digital Modernization Strategy, June 2019.

1.3 Scope

The DoD Zero Trust Engineering Team developed this Zero Trust Reference Architecture (ZT RA) to align with the DoD definition: “Reference Architecture is an authoritative source of information about a specific subject area that guides and constrains the instantiations of multiple architectures and solutions.”³

This Reference Architecture describes Enterprise standards and capabilities. Single products/suites can be adopted to address multiple capabilities. Integrated vendor suites of products rather than individual components will assist in reducing cost and risk to the government. This document will evolve as requirements, technology, and best practices change and mature. ZT promotes an individual journey to a collaborative goal of continuous enhancements, while also incorporating best practices, tools, and methodologies of industry.

1.3.1 Stakeholders

The DoD ZT RA will be used by DoD Mission Owners (MOs) to guide and constrain the evolution of existing DoD IT and Enterprise Environments. MOs are individuals/organizations responsible for the overall mission environment, ensuring that the functional and cybersecurity requirements of the system are being met.

The ZT RA provides an end-state vision, strategy, and framework for MOs across the DoD to utilize in order to strengthen cybersecurity and guide the evolution of existing capabilities to focus on a data centric strategy.

ZT embeds security principles throughout the architecture for the purpose of protecting data and service operations, preventing, detecting, responding, and recovering from malicious cyber activities. The perspective of the ZT RA is to guide the developer, operator, manager, and user of ZT in the development of solutions to implement a ZT framework within an existing environment.

This ZT RA’s intent is to:

- Provide stakeholders with operational context needed to better understand principles and rules when applying a ZTA.
- Define capabilities required to enable a ZTA.
- Provide baseline description of ZT for use in managing change and risk associated with evolving operational needs.
- Define the importance of ZT by showcasing how the model constantly limits access when required, continuously monitors, and identifies anomalies or malicious acts.

1.3.2 Organization of the Reference Architecture

³ DoD Reference Architecture Description – June 2010

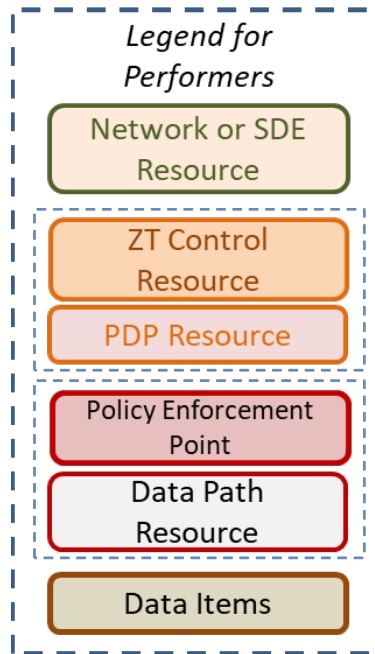
July 2022

This RA contains the following sections:

- Strategy and Vision (with broad Operational Views)
- Pillars and Principles
- Conceptual Capability Architecture (capabilities organized into a functional taxonomy, here associated with the Pillars)
- Use Cases and associated requirements
- Technical environment describing emerging technology, common industry approaches and key standards
- Security Assessment
- Architecture patterns (The scope of alternate ways to realize a conformant design and the refining of Performers into Systems and Services)
- Example, Transition Architecture direction meeting the above constraints and being pursued at the time of the RA (Maturity Model, baseline, transition, target, phases)

Following DoD standards, the artifacts in this RA are from the Department of Defense Architectural Framework (DoDAF). Because of the broad audience that needs to understand and adapt ZT, an informal style is used for the artifacts. Informal drawings are easier to understand by a wide audience, not all of whom are familiar with Unified Profile for DoDAF/MODAF (UPDM) model representations. These drawings should allow a common representative of the target stakeholder to grasp the meaning of the artifact. With the RA, it is the content that is important. However, this is still a digital architectural model and includes artifacts with descriptions, lists of definitions, and tables of interaction. Entities (the nouns of DoDAF) are defined and used in the artifact drawings which tell a story of function and entity relationships. The All View Integrated Dictionary (AV-2) is organized by type of entity and most of these tables are in the appendix. From this RA, Reference Designs (RD) can be created that capture a ZT logical architecture for specific environments and functional needs.

The conceptual capability architecture predominately is captured in several Operational Views [OV-1: High-Level Operational Concept Graphic, OV-2: Operational Resource Flow Description] and Capability Views [CV-1: Vision, CV-2: Capability Taxonomy]. Strategies are captured in a CV-1. Here, OV-1s describe the problem and the opportunities for a specific functional environment. Then capabilities are explained in relation to the OV-1 opportunities. The (entity type) capabilities appear in the drawings with a thin line. These are captured in a capability taxonomy (CV-2) organized by their associations with Pillars and resources. The other main view type is the OV-2: Operational Resource Flow Description. This captures specific resources and how they interact in a specific use case or architectural pattern (with some conceptual SV-1: Systems Interface Description & SvcV-1: Services Context Description).



Artifacts generally follow this presentation format except when a local legend is on the artifact:

- **Network & SDE Resource:** Resources controlling or providing network services are outlined in green.
- **Zero Trust enabled Resource:** (orange outlined in orange) There are no specific Zero Trust elements or services; however, a resource can act to provide control, usually by acting as an orchestrator or a controller with a NIST Policy Decision Point function (pink outlined in orange).
- **Data Path Resource** (grey outlined in red) may or may not also provide a NIST Policy Enforcement Point function (red outlined in red). Sometimes the PEP is in a service.
- **Data** when a data resource acted on by ZT, it is outlined in brown.

Figure 1 Legend for Performers

1.3.3 Timeframe

These are the general timelines associated with the development of the ZT RA.

- 30 September 2020: Initial ZT RA v0.9 submitted for review by DISA, NSA, DoD CIO, and United States Cyber Command
- 04 November 2020: ZT RA v0.9 submitted to Enterprise Architecture Engineering Panel (EAEP) for feedback
- 04 December 2020: Zero Trust Joint Engineering Team received feedback and began adjudication
- 24 December 2020: Submission of ZT RA v0.95 submitted to EAEP
- 04 January 2021: EAEP members voted on ZT RA release
- 11 February 2021: Digital Modernization Infrastructure Infrastructure Executive Committee approval of ZT RA v1.0
- 13 May 2021: ZT RA v1.0 published on DoD CIO Library
- 30 September 2021: ZT RA v2.0 draft development complete
- 21 November 2021: DCIO CS Chief Architect directed ZT RA 2.0 to be staffed through CS RA Steering Group on its way to EAEP and/or DMI EXCOM
- 7 February 2022: CS RA Steering Group - Joint O-6/GS-15 CATMS review of draft ZT RA v2.0 completed
- 24 May 2022: EAEP completed assessment
- 1 June 2022: Briefed the EAEP results of assessment with complete concurrence of the panel members

1.4 Vision and Goals (CV-1)

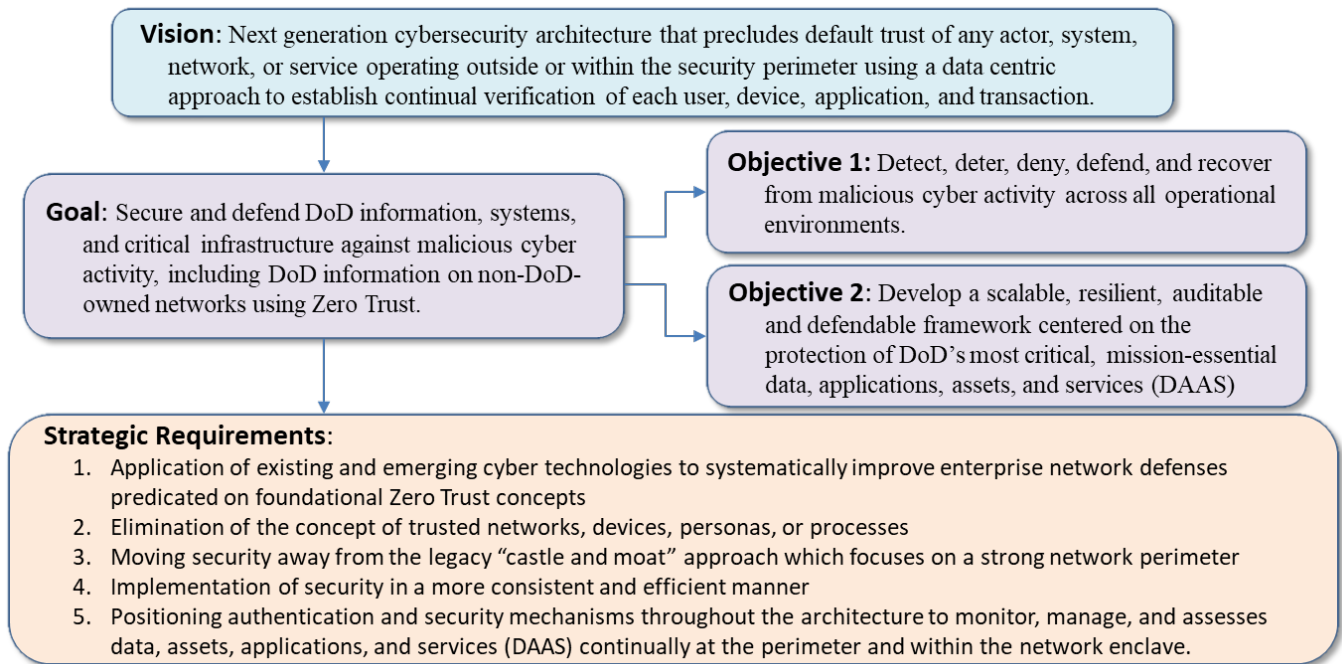


Figure 2 Zero Trust Vision (CV-1)⁴

By reconfiguring, reprioritizing, and augmenting existing DoD capabilities, the DoD will be able to evolve towards a next-generation security architecture, ZT. With these augmented capabilities, the agency will be able to secure and defend DoD information, systems, and critical infrastructure against malicious cyber activity, including DoD information on the non-DoD-owned environments. The ability to detect, deter, deny, defend, and recover from malicious cyber activities and develop a scalable, resilient, auditable, and defensible framework will require several different ways to strategically protect DoD environments. The concept of trusted networks, devices and endpoints geared towards perimeter based defenses will shift toward a never trust, always verify approach. Moving security away from the perimeter and towards an integrated security architecture focusing on protecting data, applications, and servers will be critical to achieving the ZT vision. As cyber threats evolve and become more and more sophisticated, ZT implementors will need to stay current on existing and emerging cyber technologies to systematically improve enterprise environment defenses that are in line with ZT concepts. These new strategic goals enable the implementation of security in a more consistent and efficient manner.

⁴ 2018 DoD Cyber Strategy

July 2022

1.4.1 Vision and High-Level Goals (CV-1)

Vulnerabilities exposed by data breaches inside and outside DoD demonstrate the need for a new and more robust cybersecurity model that facilitates mission enabling decisions that are risk aware. ZT is a cybersecurity strategy and framework that embeds security principles throughout the Information Enterprise (IE) to prevent, detect, respond, and recover from malicious cyber activities. This security model eliminates the idea of trusted or untrusted networks, devices, personas, or processes, and shifts to multi-attribute-based confidence levels that enable authentication and authorization policies based on the concept of least privileged access. Implementing ZT requires designing a consolidated and more efficient architecture without impeding operations to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services viewed as compromised.

ZT focuses on protecting critical data and resources, not just the traditional network or perimeter security. ZT implements continuous multi-factor authentication, micro-segmentation, encryption, endpoint security, automation, analytics, and robust auditing to Data, Applications, Assets, Services (DAAS). As the Department evolves to become a more agile, more mobile, cloud-instantiated workforce, collaborating with multiple federal and non-governmental organizations (NGO) entities for a variety of missions, a hardened perimeter defense can no longer suffice as an effective means of enterprise security. In a world of increasingly sophisticated threats, a ZT framework reduces the attack surface, reduces risk, and ensures that if a device, network, or user/credential is compromised, the damage is quickly contained and remediated.

State-funded hackers are well trained, well-resourced, and persistent. The use of new tactics, techniques, and procedures combined with more invasive malware can enable motivated malicious personas to move with previously unseen speed and accuracy. Any new security capability must be resilient to evolving threats and effectively reduce threat vectors, internal and external.

ZT end-user capabilities improve visibility, control, and risk analysis of infrastructure, application and data usage. This provides a secure environment for mission execution. Enabling ZT capabilities address the following issues and high-level goals:

- **Modernize Information Enterprise to Address Gaps and Seams.** Over time, DoD environments have been decentralized. Usability and security challenges stem from years of building infrastructure along organizational, operational and doctrinal boundaries, with multiple security and support tiers, enclaves and networks. Capabilities developed in silos have inevitably resulted in disconnects and gaps in the command structure and processes that preclude establishing a comprehensive, dynamic, and near-real time common operating picture. Adversaries have exploited these logical, technological, and organizational gaps and seams.
- **Simplify Security Architecture.** A fragmented approach to information technology and cybersecurity has led to excessive technical complexity, creating vulnerabilities in enterprise hygiene, inadequately addressing threats and results in high levels of latency. Complex security techniques render the user experience unresponsive and ineffective.

This is a factor that drives the use of unapproved or unsecure technologies as users look to complete their mission.

- **Produce Consistent Policy.** This is a critical lesson-learned from industry that automated cybersecurity policies must be consistently applied across environments for maximum effectiveness. System owners have a responsibility to define governance practices. This enforces reliability and consistency aligning with policy and requirements.
- **Optimize Data Management Operations.** The success of DoD missions, ranging from payroll to missile defense, are increasingly dependent on structured tagged data within and external to originating systems. Advanced analytics also depend on these dependencies. While data standards and policy exist, they are disparate and inconsistently implemented. This results in:
 - Interoperability challenges between applications, organizations, and with external partners
 - System inefficiencies and vulnerabilities
 - Poor user experience
 - Inability to fully leverage cloud computing, data analytics, machine learning, and artificial intelligence
- **Provide Dynamic Credentialing and Authorization.** Persona based identities, credentials, and attributes are not dynamic or context aware and come from disparate sources. Two factor authentication, in the form of the Public Key Infrastructure (PKI) Common Access Card (CAC), while secure, has not kept pace with more user-friendly multi-factor authentication advances in industry. In industry, Non-person entities (NPE) are not widely addressed beyond basic service accounts, nor are entities for bots or the Internet of Things (IoT). The DoD Identify Credential and Access Management (ICAM) Reference Design provides further guidance on identity, credentialing, and access management implementations.

1.4.2 Zero Trust Strategy

ZT principles, Pillars and culture will guide mission owners in their efforts to reconfigure, re-prioritize and augment existing DoD capabilities to evolve portfolios and resources towards a revised, data centric DoD Cybersecurity Reference Architecture (CS RA). It instantiates tenets of the 2019 DoD Digital Modernization Strategy, the 2018 DoD Cyber Strategy Lines of Effort, and the 2019 Cybersecurity Risk Reduction Strategy. It supports the DoD vision of “a more secure, coordinated, seamless, transparent, and cost-effective IT architecture... that ensures dependable mission execution in the face of a persistent cyber threat.”

ZT supports an incremental migration approach to cybersecurity with an end state of an interoperable, fully functioned, optimized cybersecurity architecture that secures our critical

assets and data from malicious threats and unintentional incidents. ⁵The desired outcome is the roll out of an employable set of enterprise ZT capabilities each consisting of standards, devices, and processes that are measurable, repeatable, supportable, and extensible, to any organization on the DoDIN, and federated across the DoDIN. The DoD CIO Zero Trust Strategy outlines the vision, approach, principles, goals and objectives, and roadmap for the DoD’s migration to ZT.

1.5 Cybersecurity (Transition) Problem Statement (OV-1)

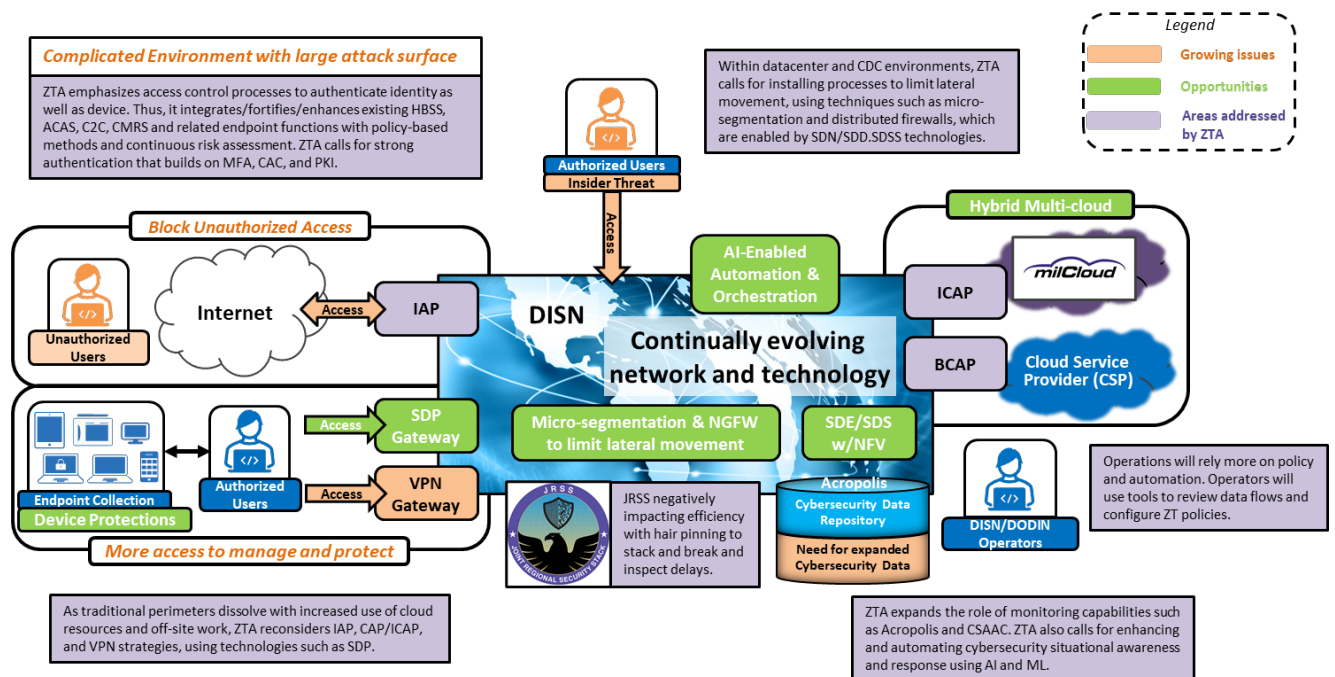


Figure 3 Cybersecurity Problem Statement (OV-1)

Traditional approaches to cybersecurity architectures, such as defense in depth, have resulted in complicated and redundant capabilities across the DoDIN. The prevalence of teleworking and adoption of cloud computing have caused a crossing of DoD data into new platforms; often hosted in industry and user environments. This change in the digital experience introduces new security challenges but also opportunities for leveraging important technology evolutions and ZT principals to revolutionize cyber defense.

The growing issues in security protections align with the evolution of endpoints to multiple platforms, inclusive both traditional and non-traditional devices such as IoT (Internet of Things), Supervisory Control and Data Acquisition (SCADA) and OT(Operational Technology), challenges in managing numerous security stacks, threats from privileged users and controlling access to cloud environments. Authentication and authorization of endpoints to the environment

July 2022

is inconsistently implemented and does not always consider device hygiene. Security stacks such as Joint Regional Security Stack (JRSS,) while providing numerous security capabilities, are challenging to manage and lack integration with data center and identity capabilities capable of providing a holistic security solution. Insider threats continue to compromise data and systems through unintegrated security policy and lateral movement. The adoption of the cloud computing has shifted data off-premises and has generated the need for a more robust security model beyond a perimeter defense. ZT concepts and principles impact each of these scenarios and provide centralized, standardized and integrated controls to mitigate these security challenges.

Beyond the issues identified above, there are existing capabilities within the enterprise which are key enablers for operationalizing ZT. Software defined enterprise (SDE) capabilities, such as automation, orchestration, and network function virtualization (NFV), combined with analytics for our defensive cyber operators enables the scale and confidence scoring necessary to achieve ZT security policies. Many of these capabilities have already been implemented throughout DoD and provide a great baseline for ZT when tuned with the proper configuration. NFV capabilities can also be rapidly deployed to protect legacy applications and enable resource authorization decisions for an enhanced security posture.

1.6 Overall Target Environment (OV-1)

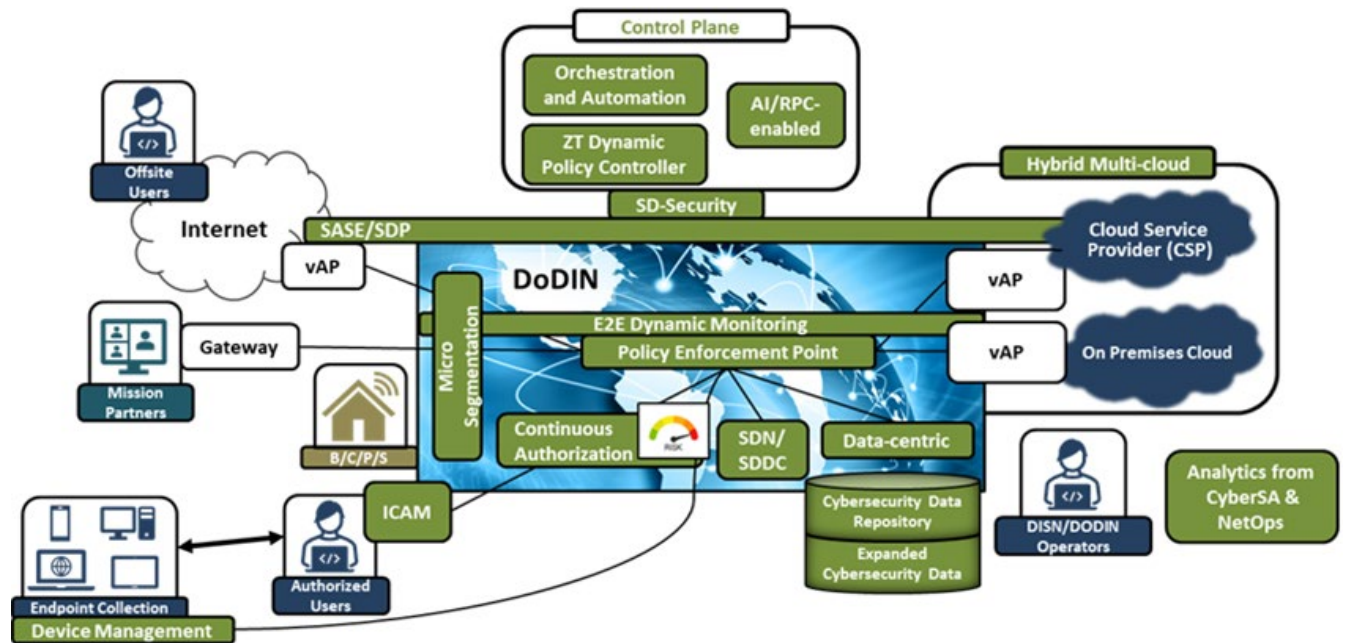


Figure 4 Target Environment (OV-1)

ZT cybersecurity architecture introduces new security concepts such as data centrality and conditional access to achieve the core concept of never trusting a request for data, applications, or resources. Beyond the notion of never trusting and verifying explicitly, assuming a breach in the environment brings new levels of granularity to the security policies implemented within these capabilities.

A data-centric security architecture starts with identifying sensitive data and critical applications for introducing ZT. This discovery process will include identification of the users and flows for development of the security policy. The control plane consisting of the ZT policy controller and automation and orchestration capabilities will be an insertion point for new conditional access policies. The integration between these technologies will be achieved via APIs. Evolution of artificial intelligence (AI) and robotic process automation (RPA) will modernize and enrich the policy deployed from the control plane.

The ZT security policy is executed on numerous policy enforcement points throughout the architecture. The first steps in a flow from user to data are authenticating and authorizing a user which requires integration with an enterprise ICAM solution, global device management and continuous vetting of identity and attributes. The attributes required for authorization will be specific to the user's level of access, hygiene of the device, and activities performed in the environment. The combination of these elements develops into a confidence score which dynamically changes based on conditions and telemetry.

July 2022

Virtual access points and gateways are the next phase in authorization. A Policy Decision Point provides a confidence score of the user or endpoint. A Policy Enforcement Point then enforces segmentation policy and connects the user or endpoint to the requested resource. Depending on the implementation, certain cloud access points should align to the DoD Cloud Native Access Point Reference Design (CNAP RD). Within these access points are numerous security capabilities to include firewall and inspection technologies. Software defined perimeter capabilities also align to the requirements for implementing virtual access points.

The deployment environment requires the adoption of software defined data center technologies such as software defined networking (SDN) to truly enable ZT controls. SDN technologies integrate at the host level to provide micro segmentation which is a key control to lateral movement. Beyond the traditional implementation of segmentation which focuses on port and protocol, processes should be evaluated along with identity to ensure east-west network flows within application components are not threats.

The data itself is protected through a combination of data loss prevention (DLP) and data rights management (DRM) to control data exfiltration. DRM will tie encryption to relevant security policies and attributes to protect access to the file. This will enable data-in-use protections to provide additional controls around how data can be manipulated and extracted from files.

Throughout each of these transactions data is logged, filtered, and analyzed. Unified analytics enrich confidence levels used in authorization decisions to provide relevant data beyond user attributes and device hygiene. User and entity behavior analysis (UEBA) will baseline normal activity and provide indicators of threats and additional risks to limit authorization transactions.

1.7 Assumptions

The following core assumptions drive planning, risks assessment, and implementation considerations for deploying a ZT architecture:

- The CS RA remains the authoritative cybersecurity reference architecture for the DoD. The ZT RA does not replace the CS RA but augments the CS RA. The CS RA and ZT RA will continue to converge over time whereby principles and Pillars of the ZT RA will be infused into ongoing agile updates to the CS RA.
- Technologies will exist, will be mature, and available/implementable to achieve a DoD ZT migration across the information enterprise.
- ZT assumes continued and mandated use of communication encryption to the greatest extent possible.
- Multiple decentralized Service pilots and proof-of-concepts will require integration and synchronization for a common ZT end-state.
- No single device or capability produces a ZT framework. ZT is a holistic approach to security that leverages several different technologies to enable a ZT end-state.
- Security policies will be universally and consistently automated and orchestrated at the macro level for the DoD enterprise. Granular security policies and access controls will be automated and orchestrated at the micro level by mission owners.
- Interoperability standards must emerge to enhance data security protections.

July 2022

1.8 Constraints

The following core constraints drive planning, risks assessment, and implementation considerations for ZT.

- Limited testing due to current environmental constraints has been completed on the capabilities that support the ZT RA version 2.0. Additional development and refinement stages should be completed to support design documentation.
- Coarse/fine grain policy is a design decision driven by constraints of DoD structure and policy management.

2 PILLARS AND PRINCIPLES

2.1 Overview

ZT Security “is an emerging initiative that DoD CIO is exploring in concert with DISA, United States Cyber Command (USCYBERCOM), and the National Security Agency (NSA). ZT is a cybersecurity strategy developing an architecture that requires authentication or verification before granting access to sensitive data or protected resources at a financial cost by reducing data loss and preventing data breaches. This security model helps transition and upgrade over time from trusted networks, devices, personas, or processes, and shifts to multiple attributes and multi-checkpoint-based confidence levels that enable authentication and authorization policies under the concept of least privileged access. Implementing ZT requires rethinking how we utilize existing infrastructure to implement security by design in a simpler and more efficient way while enabling unimpeded operations.”⁶

While straightforward in principle, the actual implementation and operationalization of ZT incorporates several areas which need to be smartly integrated and that include software defined networking, data tagging, behavioral analytics, access control, policy orchestration, encryption, automation, as well as end-to-end ICAM. Enterprise level considerations include identifying which data, applications, assets, and services to protect, and mapping transaction flows, policy decisions, and locations of policy enforcement. Apart from the advantages to securing our architecture in general, there are additional cross-functional benefits of ZT regarding cloud deployments, Security Orchestration and Automation (SOAR), cryptographic modernization and cybersecurity analytics.

2.2 Concept and Tenets of Zero Trust

The ZT security model re-thinks how to implement security access to resources and is determined by dynamic policy, including the observable state of user and endpoint identity, application/service, and the requesting asset and may include other behavioral and environmental attribute. Confidence levels are built from multiple attributes of the subject being

⁶ DoD Digital Modernization Strategy, June 2019

July 2022

authenticated (identity, location, time, device security posture) and allow a much more thorough evaluation of access requests beyond credential verification.

ZT has five major tenets. These tenets represent the foundational elements and influence all aspects within ZT.

- **Assume a Hostile Environment.** There are malicious personas both inside and outside the environment. All users, devices, applications, environments, and all other NPEs are treated as untrusted.
- **Presume Breach.** There are hundreds of thousands of attempted cybersecurity attacks against DoD environments every day. Consciously operate and defend resources with the assumption that an adversary has presence within your environment. Enhanced scrutiny of access and authorization decisions to improve response outcomes.
- **Never Trust, Always Verify.** Deny access by default. Every device, user, application/workload, and data flow are authenticated and explicitly authorized using least privilege, multiple attributes, and dynamic cybersecurity policies.
- **Scrutinize Explicitly.** All resources are consistently accessed in a secure manner using multiple attributes (dynamic and static) to derive confidence levels for contextual access to resources. Access to resources is conditional and access can dynamically change based on action and confidence levels resulting from those actions.
- **Apply Unified Analytics.** Apply unified analytics for Data, Applications, Assets, Services (DAAS) to include behavioristics, and log each transaction

The use of mutual authentication of users with PKI-based client-authentication or mutual authentication certificates to web applications has long been the effective standard. The DoD is making strides to improve access to data by approving multiple authenticators and authorization schemes to better improve usability and access while maintaining security and visibility.

2.3 Pillars

Zero Trust Pillars are identified in the ZT Strategy and are in alignment with the common industry identification of ZT Pillars. A Pillar is a key focus area for implementation of Zero Trust controls. ZT is depicted as interlocking puzzle pieces (Figure 5) that symbolize a data Pillar surrounded by Pillars of protection. All protection Pillars work together to effectively secure the Data Pillar.

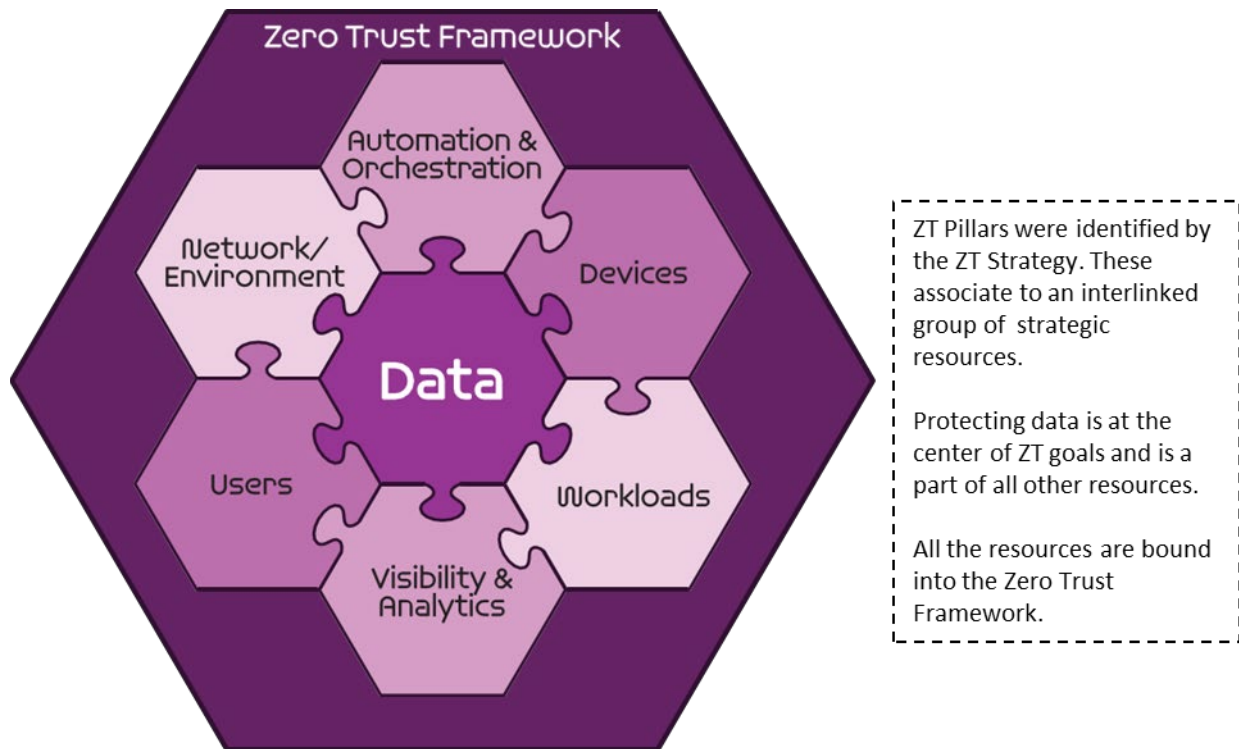


Figure 5 Zero Trust Pillars

The seven Pillars in the DoD ZT Architecture include:

User: Securing, limiting, and enforcing person and non-person entities' access to DAAS encompasses the use of identity capabilities such as multi-factor authentication (MFA) and Privileged Access Management (PAM) for privileged functions. Organizations need the ability to continuously authenticate, authorize, and monitor activity patterns to govern users' access and privileges while protecting and securing all interactions.

Device: Continuous real-time authentication, inspection, assessment, and patching of devices in an enterprise are critical functions. Solutions such as Mobile Device Managers, Comply to Connect programs, or Trusted Platform Modules (TPM) provide data that can be useful for device confidence assessments, authorization determination, and limiting access. Other assessments should be conducted for every access request (e.g. examinations of compromise state, software versions, protection status, encryption enablement, and proper configuration, etc.). Having the ability to identify, authenticate, inventory, authorize, isolate, secure, remediate, and control all devices is essential in a ZT approach.

Network/Environment: Segment (both logically and physically), isolate, and control the network/environment (on-premises and off-premises) with granular access and policy restrictions. As the perimeter becomes more granular through macro-segmentation, micro-segmentation provides greater protections and controls over DAAS. It is critical to, control privileged access, manage internal and external data flows, and prevent lateral movement.

July 2022

Applications and Workload: Applications and workloads include tasks on systems or services on-premises, as well as applications or services running in a cloud environment. ZT workloads span the complete application stack from application layer to hypervisor. Securing and properly managing the application layer as well as compute containers and virtual machines is central to ZT adoption. Application delivery methods such as proxy technologies, enable additional protections to include ZT decision and enforcement points. Developed Source Code and common libraries are vetted through DevSecOps development practices to secure applications from inception.

Data: A clear understanding of an organization's DAAS is critical for a successful implementation of a ZT architecture. Organizations need to categorize their DAAS in terms of mission criticality and use this information to develop a comprehensive data management strategy as part of their overall ZT approach. This can be achieved through the ingestion of consistent valid data, categorization of data, developing schemas, and encrypting data at rest and in transit. Solutions such as DRM, DLP, Software Defined Environments, and granular data-tagging support the protecting of critical DAAS.

Visibility and Analytics: Contextual details provide greater understanding of performance, behavior and activity baseline across other ZT Pillars. This visibility improves detection of anomalous behavior and provides the ability to make dynamic changes to security policy and real-time access decisions. Additionally, other monitoring systems, such as sensor data in addition to telemetry will be used, will help fill out the picture of what is happening with the environment and will aid in the triggering of alerts use for response. A ZT enterprise will capture and inspect traffic, looking beyond network telemetry and into the packets themselves to accurately discover traffic on the network and observe threats that are present and orient defenses more intelligently.

Automation and Orchestration: Automate manual security processes to take policy-based actions across the enterprise with speed and at scale. SOAR improves security and decreases response times. Security orchestration integrates Security Information and Event Management (SIEM) and other automated security tools and assists in managing disparate security systems. Automated security response requires defined processes and consistent security policy enforcement across all environments in a ZT enterprise to provide proactive command and control.

2.4 Reference Architecture Principles (OV-6a)

The ZT framework is an approach to security that utilizes a series of guiding principles in the creation of the RA and other future documents.

- Principle #1: Assume no implicit or explicit trusted zone in networks.
- Principle #2: Identity-based authentication and authorization are strictly enforced for all connections and access to infrastructure, data, and services.

- Principle #3: Machine to machine (M2M) authentication and authorization are strictly enforced for communication between servers and the applications.
- Principle #4: Risk profiles, generated in near-real-time from monitoring and assessment of both user and devices behaviors, are used in authorizing users and devices to resources.
- Principle #5: All sensitive data is encrypted both in transit and at rest.
- Principle #6: All events are to be continuously monitored, collected, stored, and analyzed to assess compliance with security policies.
- Principle #7: Policy management and distribution is centralized.

Table 1 Reference Architecture Principles (OV-6A)

July 2022

3 CAPABILITIES

3.1 Capabilities Taxonomy (CV-2)

The Seven Zero Trust Pillars assist with the categorization of capabilities and technologies that can perform Zero Trust functions in an environment. Capabilities are the ability to achieve a desired effect under specified (performance) standards and conditions through combinations of ways and means (activities and resources) to perform a set of activities. Pillars align with capabilities such as identity authentication and Software Defined Enterprise. Sub-Capabilities such as enterprise identity provider or Just-In-Time analytics support capabilities. Capabilities and sub-capabilities as defined reflect the current technologies that are applicable in ZT and are subject to change in future iterations of the ZT RA. This layered approach allows for flexibility in implementing ZT controls. Overarching governance will be required to achieve proper integration across Pillars and capabilities. The Pillar and capabilities enable maximum visibility and protection of data, which are the key focuses of any implementation of ZT.

With Figure 6 Capability to Pillars Mapping (FFP), the white arrows show which ZT aggregated capability acts on what ZT Pillars. A branch of capabilities may point to more than one Pillar and some point to the entire ZT framework. This section provides an overarching description of the ZT capabilities and is intended to provide capabilities that meet a ZT architecture end state rather than provide exact implementations. Certain capabilities do require enterprise scale enablers to include an enterprise federated identity service, enterprise analytics and enterprise orchestration. Proper attributes and labeling of data during the discovery process must also be implemented for a ZT architecture. Common to all Pillars is the implementation of continuous authentication and validating the identity of entities during all access transactions. This validation is based upon current identifying standards enhanced with behavioral metrics and additional identifying factors. Further mapping of the parent capabilities and mappings to service relationships are provided below.

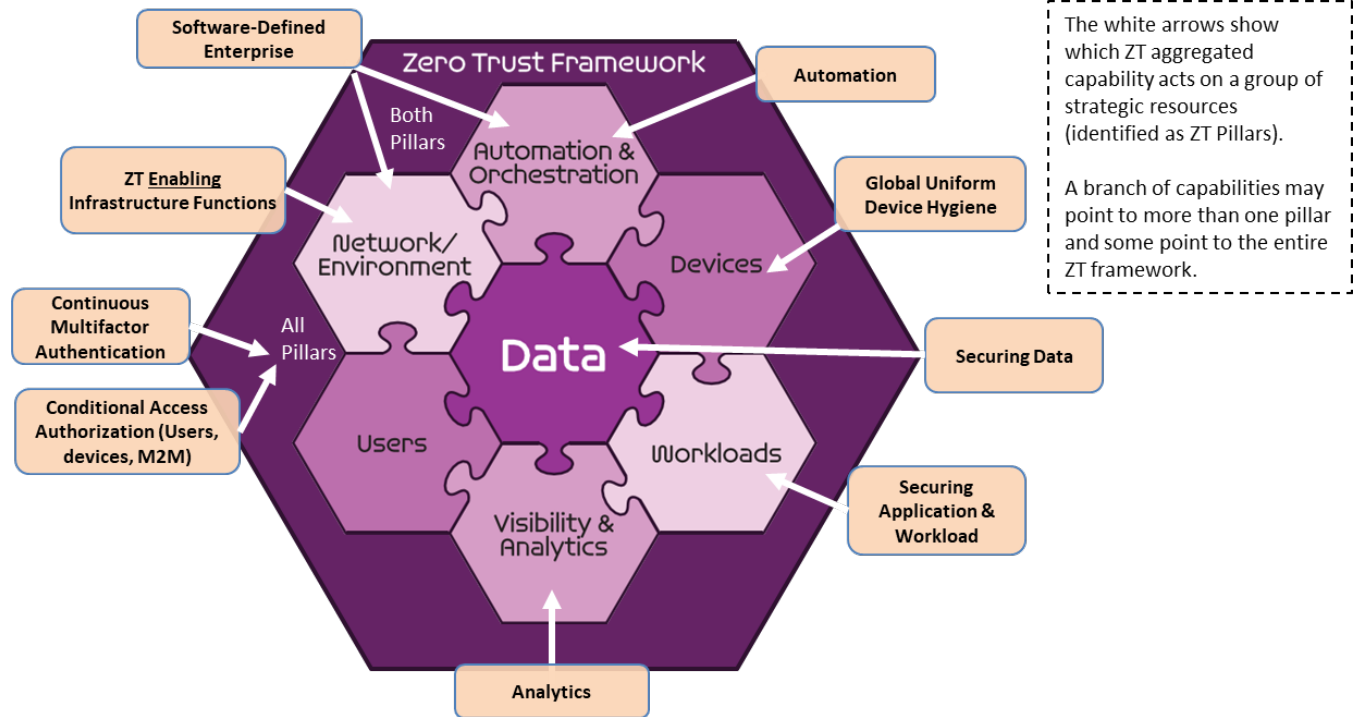


Figure 6 Capability to Pillars Mapping (FFP)

The ZT capability taxonomy is shown across multiple figures in this section. A table with definitions is located in Appendix (AV-2). The aggregated capabilities [main taxonomy branches] consist of continuous authentication, conditional authorization across, enabling infrastructure, securing application & workload, securing data, automation, analytics and orchestration. ZT has an interdependence with Data Governance, Risk Management, and the Software-define Enterprise. The full taxonomy CV 2 is provided below in several figures.

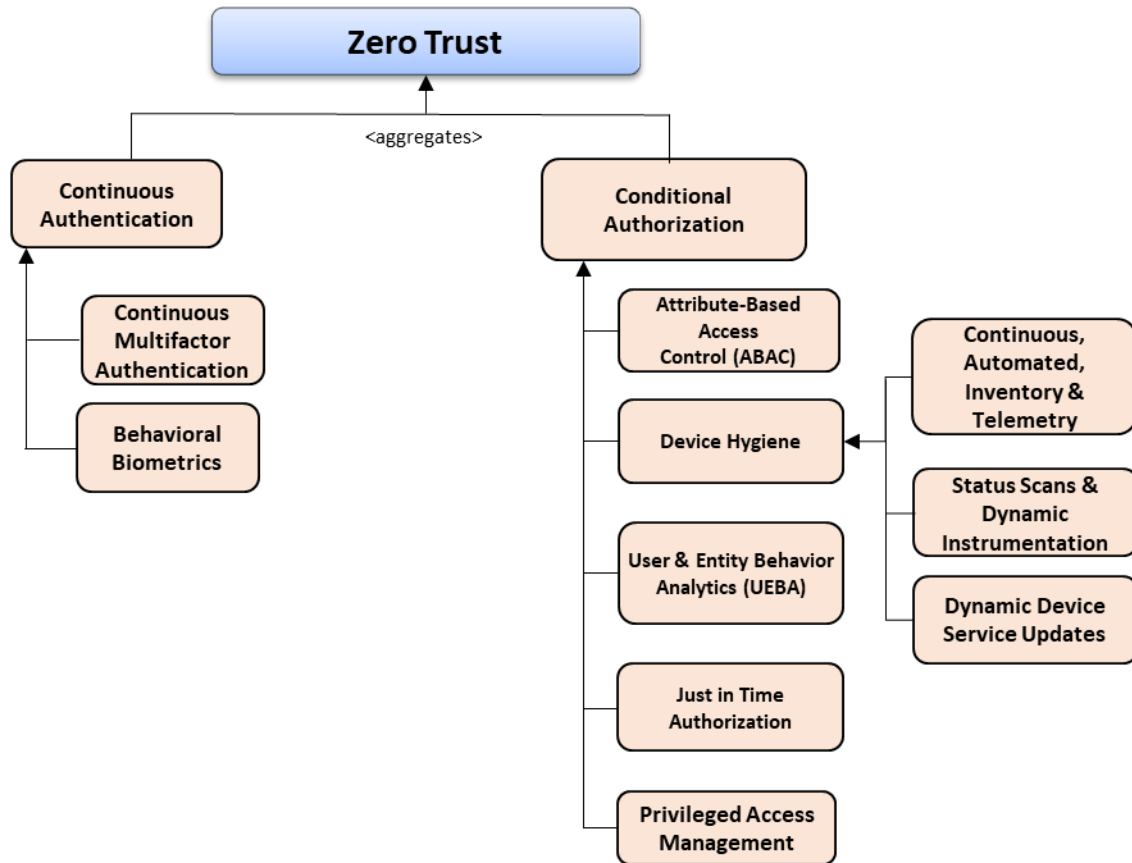


Figure 7 Zero Trust Authentication and Authorization Capability Taxonomy (CV-2)

When applied to the user’s Pillar, Conditional Authorization capabilities would focus on any object that would be considered a person-entity or non-person entity. The authorization to systems and resources would be conditional not only to standard roles but also attribute status, analytics of that entity, the requirement at a specific time and justification to access resources and data. When applied to the devices Pillar, Conditional Authorization capabilities center around systems and the enforcement of acceptable baselines and device state. Systems will be continually assessed for the current status of their inventories and telemetry data. Further information will be available through status scans and logging. Systems will be able to be updated on the fly or at the request of orchestration or other remediation methods. The degree of scrutiny and requirements for the systems accessing data will be relevant to the security level of the data that is trying to be accessed.

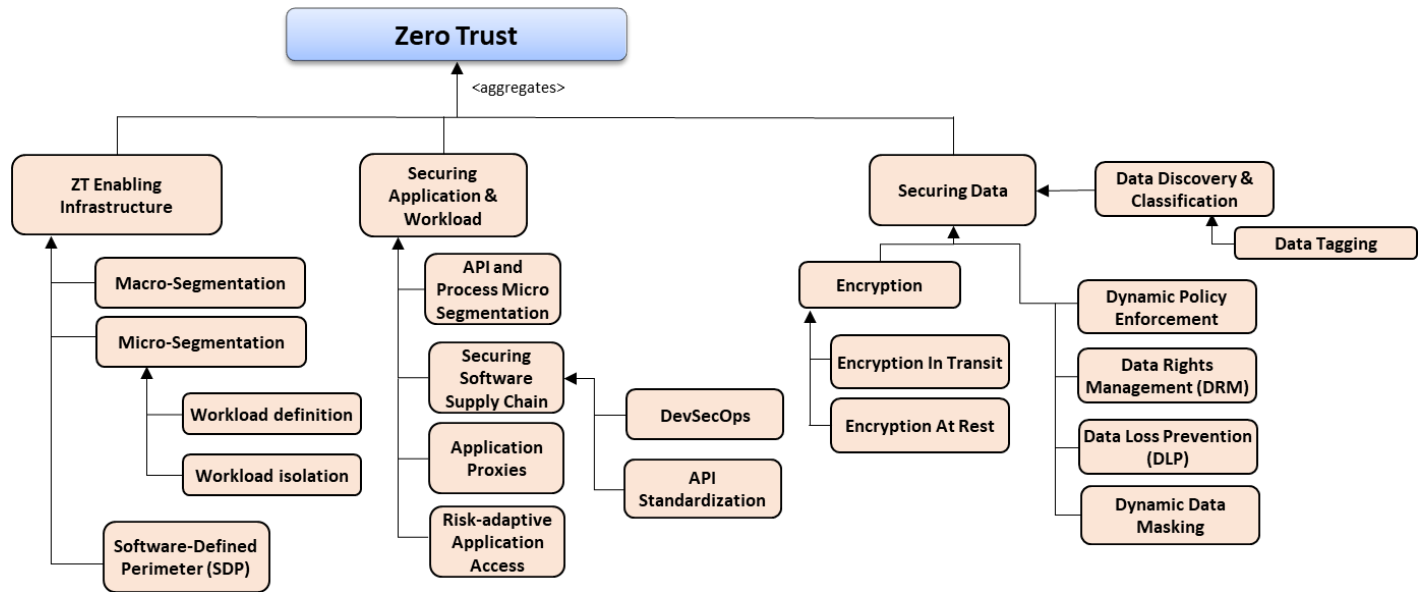


Figure 8 Zero Trust Infrastructure, Workload and Data Capability Taxonomy (CV-2)

The ZT-enabling Infrastructure aggregate capability includes all the capabilities that impact the Network and Environments resources Pillar. This will affect endpoints in the environment and nodes of travel among the enterprise. This can include not only on premises infrastructure but also cloud resources. Controls built around this Pillar relate to any ZT enabling infrastructure capabilities. A macro and micro segmentation policy can be designed around segmenting and isolating specific workloads as long as the workloads are strictly defined and validated. This allows for not only interconnection between required nodes and only those nodes but also the requirements of connection for Software Defined Perimeters.

The Securing Application and Workload aggregate capability includes all the capabilities that surround the Workloads Pillar. These capabilities will protect the application and devices serving data to end users. These capabilities aim to prevent lateral movement, validate good software practices, and segment the application into discrete highly contained secured areas. Connections into this zone are highly scrutinized and brokering between internal and external requests. The convergence to a standardization of application calls will aid in the proper implementation policy changes and updates.

The Securing Data aggregate capability include all the capabilities that surround the Data Pillar. These capabilities are the closest to the data that is to be protected. Its function is all about securing data whether it be tagged data, identified sensitive data, exfiltration protections or encrypting of sensitive data. Securing Data will constitute proper protections around sensitive information regardless of the effectiveness of the supporting Pillars and their capabilities.

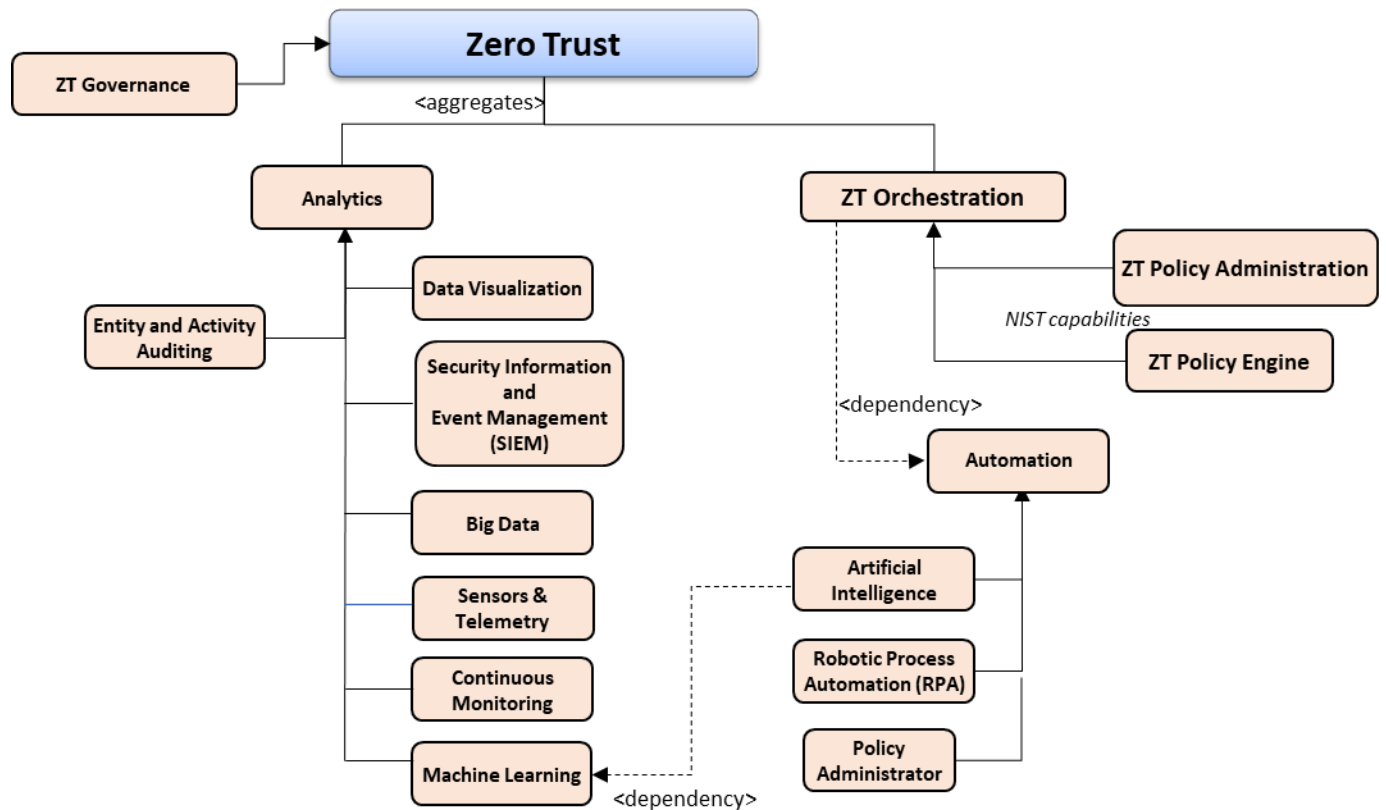


Figure 9 Zero Trust Analytics and Orchestration Capabilities Taxonomy (CV-2)

The Analytics aggregate capability includes all the capabilities that surround the Visibility & Analytics Pillar. The capabilities under this Pillar are an amalgamation of continuous entity monitoring, sensors, logging, event driven analytical tools, and machine learning. ZT will utilize Machine Learning to baseline environmental data and analytics. Machine Learning algorithms provide baseline data sets to enable ZT policy enforcement through Artificial Intelligence within ZT Orchestration.

The ZT Orchestration aggregate capability includes all the capabilities that surround the Automation & Orchestration Pillar. Its focus will be to provide automation for the deployment of policy changes in which to secure the enterprise and controlled around sensitive data. The automation and orchestration Pillar also be able to account for the ingestion of desired target state data from the Software-Defined Enterprise. While early capabilities will be centered around policy deployment, future iterations will augment the capabilities with artificial intelligence and robotic process automation into its core capabilities as the technologies evolve.

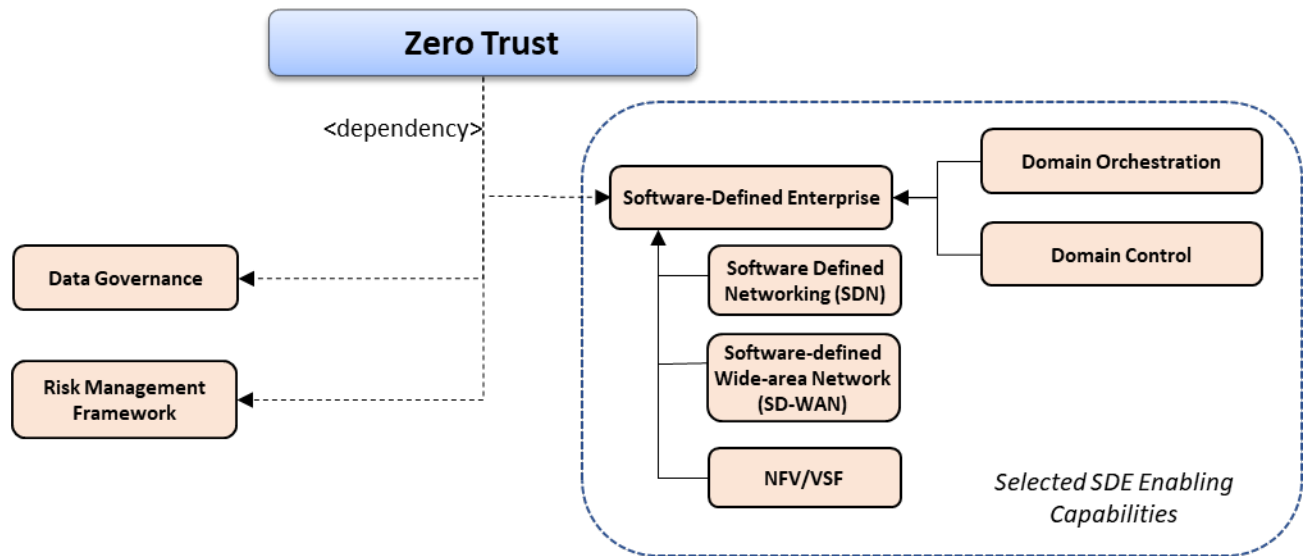


Figure 10 Zero Trust Enabling Capabilities Taxonomy (CV-2)

Data governance is a key element in the successful application of ZT security policy and provides the processes, tools, and framework for managing data from creation to disposition.

ZT and Risk Management are interdependent capabilities. ZT provides new discovery content to feed the Risk Management Framework (RMF). As a result of ZTA, the processes within the Risk Management Framework (RMF) will inform new discovery content for as well ZT, while adapting to modern application development practices such as DevSecOps. The major impacts will occur to the prepare, assess, and monitor steps. Prepare will require significant discovery, especially of the data flows to document and develop segmentation policies. Assessment will change as snapshots of systems become stale with DevSecOps capabilities rapidly modifying applications. ZT will require significant monitoring activities which improve feedback to the RMF process and incident response reactions.

Software Defined Enterprise is a key enabler to the breadth and depth of a ZT architecture implementation. As compute, network and storage infrastructure is virtualized and software defined this enables the ability to isolate data and applications at scale. Domain orchestration and control provide an enterprise control plan to push configuration and policy aligned with ZT controls.

3.2 FFP: Pillars, Resources & Capability Mapping

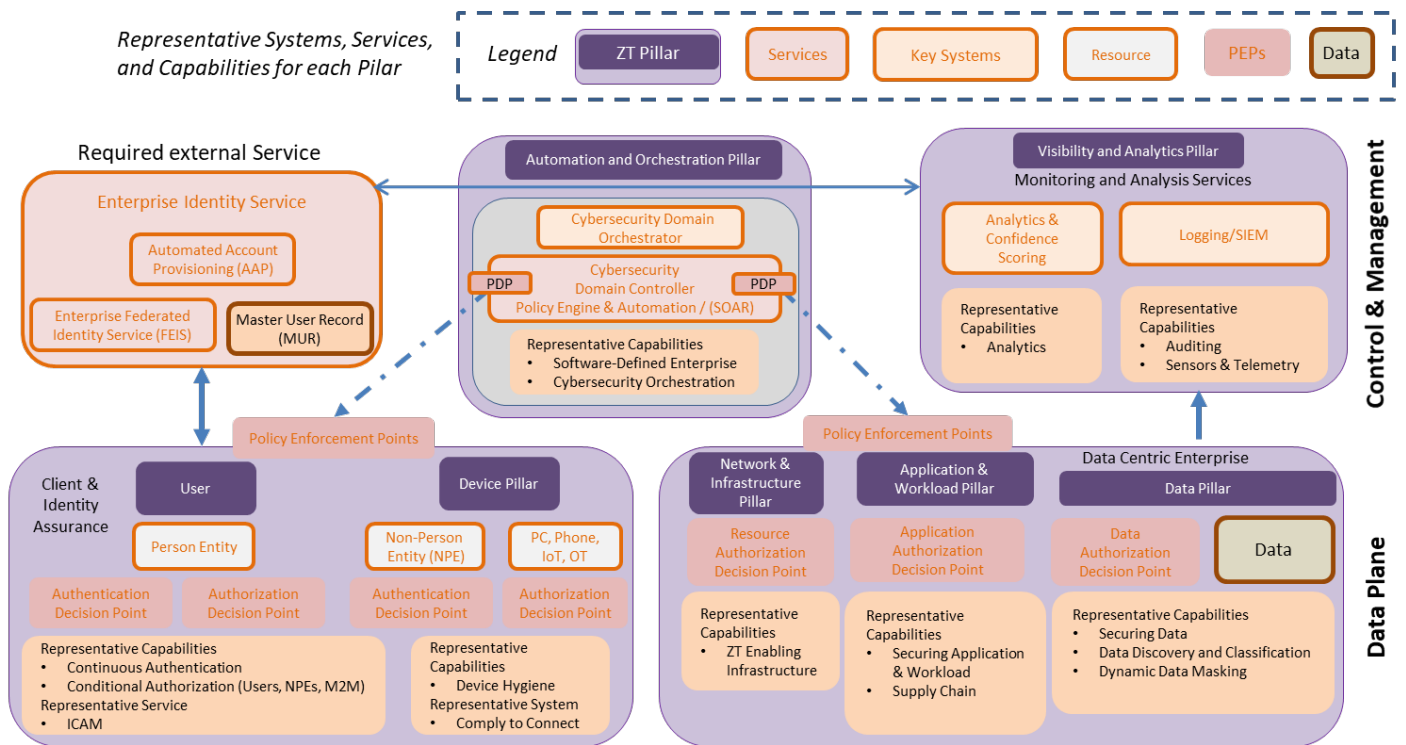


Figure 11 FFP: Pillars, Resources & Capability Mapping (CV-7)

The ZT Pillars, Resources & Capability Mapping concept provides an operational view on how security measures would be implemented within the architecture. NPE identity and person identity are tracked independently allowing for separate paths of validating confidence levels across enforcement points. Authentication and authorization activities will occur at numerous but focused points throughout the enterprise to include users and endpoints, proxies, applications and data. At each enforcement point, logs are sent to the SIEM and analytics are performed to develop a confidence level. Confidence levels of the device and user are independently developed and then aggregated where appropriate for policy enforcement. If the non-person entity or person entity has a confidence score above a measured threshold, then they are authorized to view the requested data. Data is protected along the way by Data Loss Prevention (DLP) which also feeds the SIEM to ensure the data is being used properly.

The following bullets provide additional detail on the decision points, components, and capabilities that are depicted within Figure 11. The capabilities identified below are representative of an end-state ZT implementation. Controlling access to resources based on the

July 2022

risk of the user and devices is the baseline requirement for ZT and is possible without implementation of all identified capabilities.

- **Enterprise Identity Service:** which includes Federated Enterprise Identity Service (FEIS), Automatic Account Provisioning (AAP) and a Master User Record (MUR), identifies and manages the roles, access privileges, and the circumstances in which users are granted or denied privileges.
 - FEIS: The Federated Enterprise Identity Service aggregate's identity credentials and authorizations and shares among a federated group of organizations so users/NPE can access services in other domains.
 - AAP: Provides identity governance services such as user entitlement management, business role auditing and enforcement and account provisions and deprovisioning based on identity data produced during DoD people-centric activities such as on and off-boarding, continuous vetting, talent management and readiness training.
 - MUR: Enables DoD-wide knowledge, audit, and data rollup reporting of who has access to what system or applications. MUR will also provide support in identifying insider and external threats.
- **Client and Identity Assurance:**
 - Authentication Decision Point: This evaluates the credential issuance, identity of the user, NPE, and or device as access is attempted to applications and data. Devices may also be evaluated as to whether they are managed or unmanaged. Additional use cases for non-user NPE and user assisted NPE are available in the ICAM Reference Design.
 - Authorization Decision Point: A system entity that makes authorization decisions for entities that request such access decisions. It examines requests to access resources and compares them to the policy that applies to all requests for accessing that resource to determine whether specific access should be granted to the requester who issued the request under consideration. The user and device authorizations are the first stage in conditional access to resources, applications, and ultimately the data.
 - ICAM Service: The ability to create trusted digital identity representations of individuals and NPEs, bind those identities to credentials that may serve as a proxy for the individual or NPE in access transactions, and leverage the credentials to provide authorized access to an agency 's resources.
 - Capabilities provided by ICAM:
 - Continuous Authentication: An authentication concept that utilizes multiple compatible authentication strategies to verify users and NPEs identities in an ongoing, near real-time basis, as they attempt to access resources and data.
 - Conditional Authorization: The ability to grant authorization to a resource contingent upon the continued trustworthiness of the

supplicant. This trustworthiness can be affected by the device hygiene, user and entity behavior, and other factors.

- Comply-to-Connect (C2C) Service: framework of tools and technologies operating throughout the network infrastructure to discover, identify, characterize, and report all devices connecting to the network. The C2C capability will orchestrate multiple tools to prevent non-compliant and unauthorized devices and personnel from connecting to the network, thus maintaining the secure configuration to the network and protecting the information in accordance with established standards and configurations
 - Capabilities provided by C2C:
 - Device Hygiene: The ability to inspect the state of devices, checking for malware or vulnerabilities, and compliance status with security controls, of managed and unmanaged assets in order to determine risk level of allowing the device access to resources and data.
- **Data-Centric Enterprise:**
 - Resource Authorization Decision Point: This is an intermediary decision point which will evaluate the combined NPE and user to authorize the request for access. Like previous decision points, this will leverage the confidence level and defined policies to determine if access is warranted.
 - Application Authorization Decision Point: This decision point which will evaluate the combined user and NPE to authorize the request for access. Like previous decision points, this will leverage the confidence level and defined policies to determine if access is warranted.
 - Capabilities:
 - Securing Application Workload: The ability to secure and manage the application layer as well as compute containers and virtual machines. The ability to identify and control the technology stack to facilitate more granular and accurate access decisions.
 - Securing Supply Chain: The ability to prevent or act on software supply chain attacks, which occur when a cyber threat actor infiltrates a software vendor's network and employs malicious code to compromise the software before the vendor sends it to their customers.
 - Data Authorization Decision Point: Data owners use ZT measures to apply tagging of data via orchestration or DLP/DRP servers. Data tagging will be used to ensure proper access controls are met for all data.
 - Capabilities:
 - Securing Data: Processes and technical controls to identify, classify, securely handle, retain, and dispose of data.

- Data Discovery and Classification: The ability to discover, classify, label, and report all data including sensitive and at-risk data within your databases.
- Dynamic Data Masking: The ability to provide a column-level security feature that uses masking policies to selectively mask tables and columns at query time.
- **Automation and Orchestration:**
 - Policy Engine & Automation (SOAR): These terms are used to define technologies that handle threat management, incident response, policy enforcement and security policy automation. A ZT Architecture will require dynamic policy enforcement and automation. SOAR will work in concert with analytics and policy engines to develop confidence levels and automate the delivery of policy to enforcement points.
 - Capabilities:
 - Software-Defined Enterprise: The ability to create a virtualized layer over physical infrastructure, and centrally manage it in an automated manner, utilizing a policy-based access control to dynamically create, configure, provision, and decommission virtualized network functions, system functions, security functions, and workflows.
 - Cybersecurity Orchestration: The ability to coordinate and automate disparate ZT activities and interface and coordinate them with core systems.
- **Monitoring and Analysis Services:**
 - Analytics & Confidence Scoring: This system analyzes event and incident logs via systematic analysis of data via statistics or other defined functional filters or computations to obtain confidence scores. These scores indicate the probability/percent value, within a specific range of error, with which the estimation of a statistical parameter for a given set of analytic data is determined to be true. Specifically, in ZT, this represents the probability that a user or NPE is who they assert themselves to be.
 - Capabilities
 - Analytics: The ability to systematically apply statistical and /or logical techniques to describe and illustrate, condense and recap, and evaluate data.
 - Logging utilizing Security Information and Event Management: Activity data is aggregated and stored within the SIEM which provides both a security information management (SIM) and security event management (SEM) capability.
 - Capabilities:

- Auditing/Sensors and Telemetry – The ability to directly verify, such as by inspection, examination or computation, an activity or device, in order to ensure compliance to security requirements. Entities include users and NPEs, sensor reliability, compliance programs, and shared services.

4 USE CASES

ZT requires gradual implementation of capabilities, technology solutions, process changes and policy development. ZT is not intended to be a single blanket architecture, rather it is customized to suit each organization’s needs. Each environment has differing requirements, structures, and security policies in place or in the pipeline. The following use cases have been developed with this diversity in mind.

The Tenets of ZT in Section 2.2 and the high-level architecture principles discussed in Section 2.4 are the core concepts from which the use cases below have been developed. While all Pillars may interact together, each use case focuses on specific technologies and their interactions throughout the architecture.

4.1 Data Centric Security Protections (OV-1)

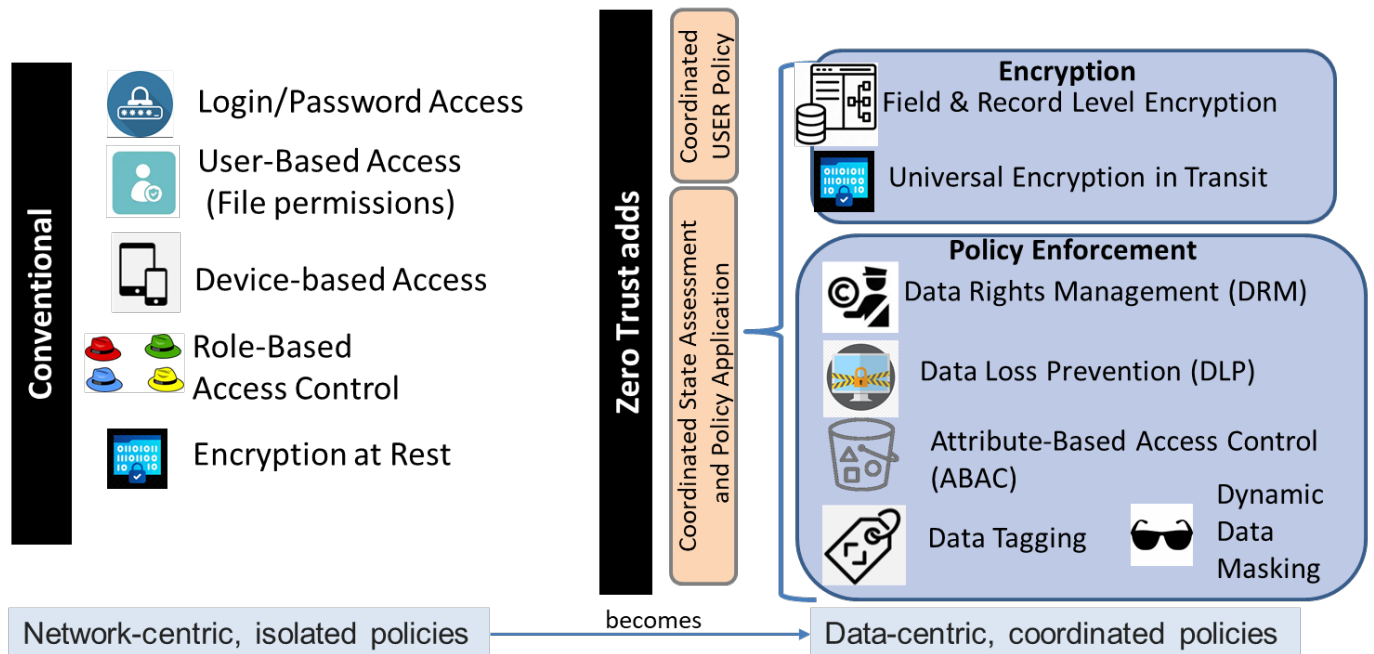


Figure 12 Data Centric Security Protections (OV-1)

July 2022

Today's approach to data security is based off legacy, isolated network-centric policies and methods. Data is vulnerable in a network-centric security model as data is protected by basic security practices such as username/password, user/device-based access and encryption only at-rest with standard Role-Based Access Control (RBAC) that is rarely updated or validated. Threat actors can circumvent these basic protections. Tomorrow's approach to data security will be under a unified ZT Framework with the focus on data-centric policies and protections that are coordinated through continuous assessment. Data-centric technologies such as encryption will help secure and protect the data at rest with additional layers of encryption within the fields and records. Data in-transit must also be encrypted. Data Tagging will feed DRM and DLP solutions which will allow creation of additional dynamic policies utilizing Attribute-Based Access Control (ABAC).

4.2 Data-Centric Security Protections (OV-2)

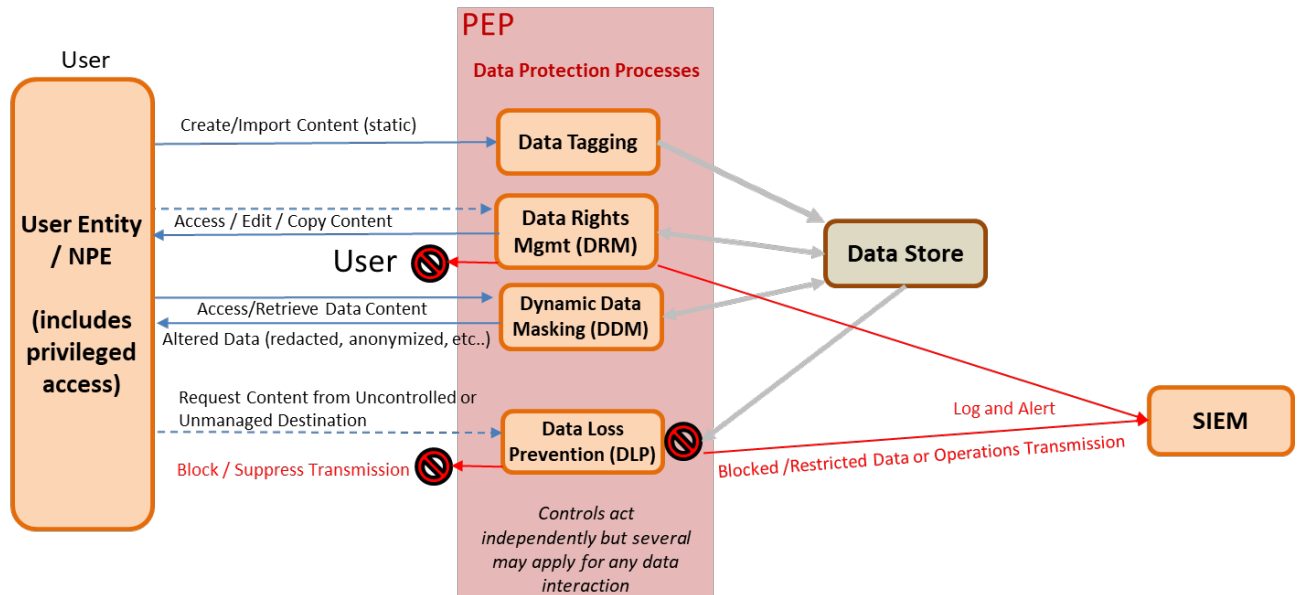


Figure 13 Data-Centric Security Protections (OV-2)

Data-Centric protections within the ZT Architecture aim to enhance protection of data regardless of where the data resides or who it is shared with. It is critical for the organization to know what data they have, the characteristics of the data, and what privacy and security requirements are needed to meet the standards for proper data protection.

Most of these protections are based around the data on the Data Store. Data Tagging on creation or import of documentation will allow the organization to categorize data with a variety of attributes. These attributes can be used in the classification of data for things like PII and sensitive data. After Data Tagging has been applied, DRM (Data Rights Management) and DLP (Data Loss Prevention) work with the SIEM and Data Store to collect and analyze access and changes to any data being accessed. DRM will allow and block access, editing or copying of data while DLP can block access and transmission of data. If a User/Endpoint is deemed trustworthy and access to the data has been granted, DDM (Dynamic Data Masking) will mask and alter the data while the data is being accessed and transmitted. These four protections,

July 2022

along with cryptographic techniques such as encryption mentioned in the section above, provide strong protections for data for a Data-Centric ZT Architecture.

4.3 Data Encryption Protections (OV-2)

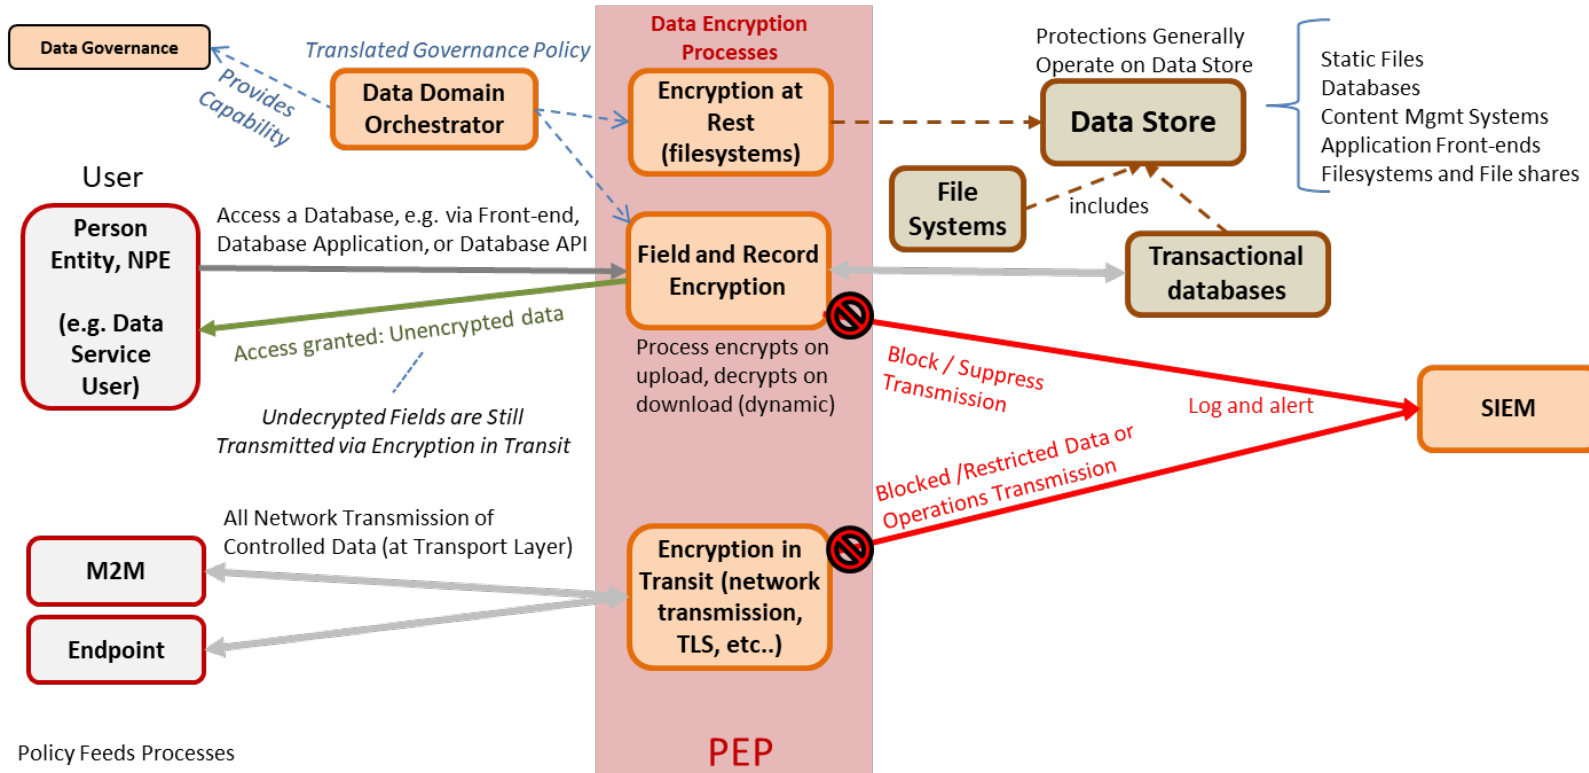


Figure 14 Data Encryption Protections (OV-2)

July 2022

Encryption is a key part of the ZT Architecture. Without encryption files are left in plain text which can expose sensitive information. With modern encryption, data is inaccessible without proper authorization. The processes and policies of Data-Centric protections are all part of Data Governance.

As a user or NPE requests access to data from an encrypted source, the request is sent through PEPs to the transactional database. If policy allows for the decryption of data to be made available for the user/device, the access to unencrypted data is granted. If policy is not met, access is blocked, and data remains encrypted.

Simultaneously to the process above, the request is recorded and analyzed by the SIEM in as near real time as possible. If the SIEM analyzes the request and deems it suspicious, it will trigger on events to be resolved by the SOAR. The SOAR, following incident response procedures can deploy mitigation policy to terminate existing sessions, reencrypt data and update policy on the PEPs to deny future requests.

4.4 Coordinating Policy for Data-Centric Security Protections (OV-2)

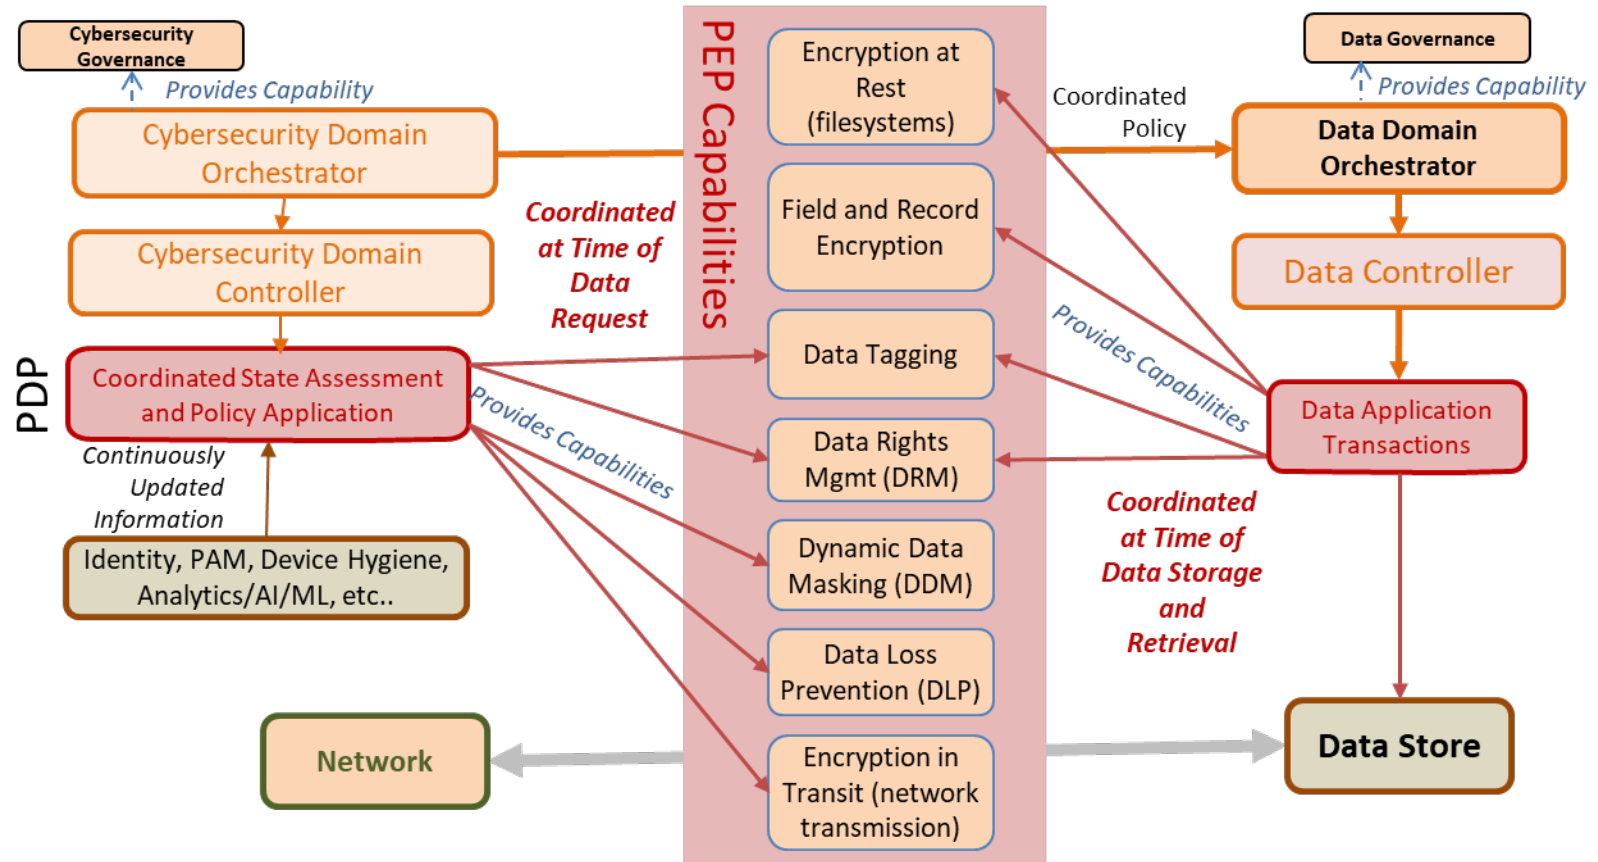


Figure 15 Coordinating Policy for Data-Centric Security Protections (OV-2)

The primary advantage of this architecture is its focus on the security of the data, not just the perimeter around the data. Data requests are routed through a policy decision point (PDP). PDP policies are kept up to date in real time through device hygiene, privileged access management (PAM), and various analytics. Users and devices cannot access the environment if policies are not met. For existing connections to DAAS, PEPs can terminate any existing connections based on PDP policy changes.

July 2022

Data access is continuously protected through numerous Policy Enforcement Points (PEPs). Data remains encrypted at rest and through transit while layers of security such as data tagging, dynamic data masking (DDM), and data loss prevention (DLP) are used.

Coordination of policy between the various components of ZT Architecture implement a defense in depth solution to maintain data integrity, availability, and confidentiality in modern architecture.

4.5 Data Analytics & AI (OV-1)

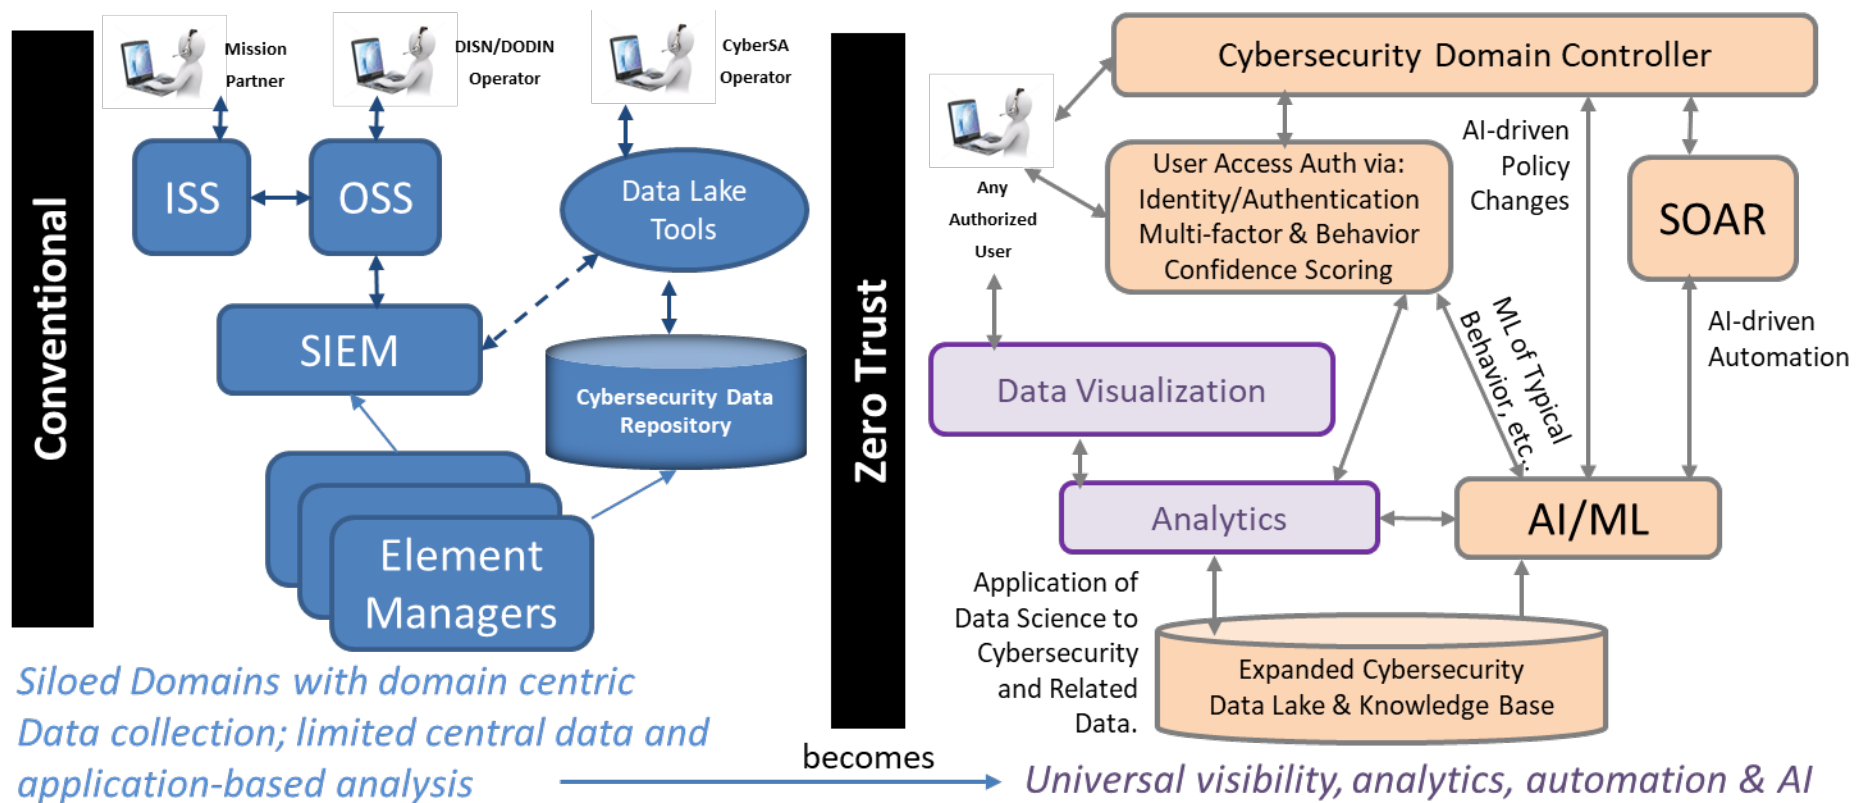


Figure 16 Big Data Analytics & AI (OV-1)

Siloed domains are normal in today's conventional architectures and cause security risks with inconsistent policies, data, logs, and analytics. The discrepancies this creates between the siloed domains make it nearly impossible to collect uniform and complete data that can be analyzed and applied into meaningful, dynamic data structures. Each siloed domain contains a subset of the data, such as the security of a device or the login location of a user at a single time. This data is fragmented across siloed domains and causes slower analysis of the data that must be optimized manually into larger relevant data.

ZT intends to make siloed domains obsolete and use data analytics and AI to create a systematic data collection architecture that can identify data types, find correlations between datasets, and observe knowledge or actionable insights using language processing. With Big Data comes the ability to accelerate the automation of data preparation tasks of gathering data, discovering, and assessing the data, cleaning, and structuring the data, transforming, and enriching the data, and then finally publishing and storing the data. What this means for ZT is the ability to have consistent policies, data, logs, and analytics to allow uniform and cohesive collection of data which in turn greatly enhances threat detection and mitigation across the architecture.

4.6 Data Analytics & AI (SV-1)

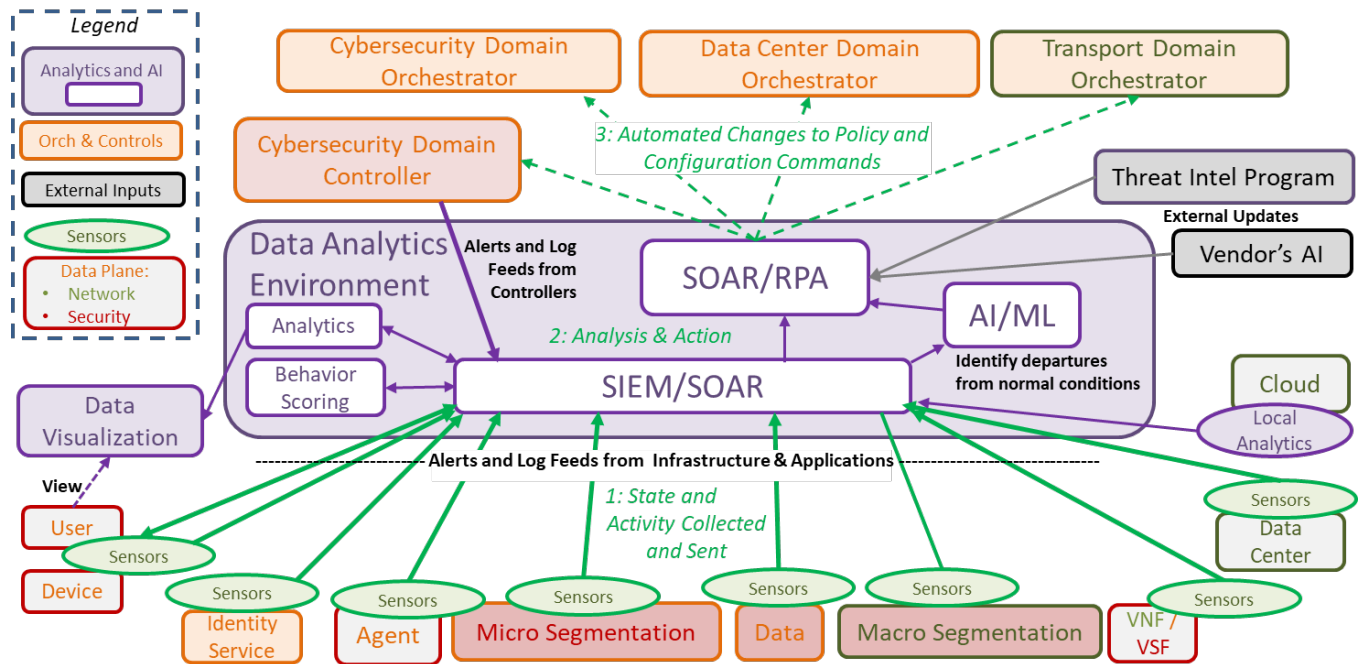


Figure 17 Data Analytics & AI (SV-1)

Big Data Analytics and AI within ZT dramatically increases visibility, insight, and automation into the environment. Data is centrally collected from all aspects of the environment and analyzed. The amount of data being collected in a ZT model is far larger than traditional architecture due to data required to power automation, and thus requires more advanced tools.

Sensors collect data from all components of the environment and send it to a SIEM where it is initially processed and analyzed for threats and abnormalities. Threat and anomaly data processed from the SIEM will be forwarded to the SOAR and further analyzed with the assistance of AI. Confirmed threats will be mitigated by the ZT controller through automation. This information is recorded and stored for future ML and AI purposes that include User/NPE confidence scoring, advanced threat detection, creating and modifying baselines, and working with external intel programs and other AI to assist with automation and orchestration. As a subset of ZT architecture, big data analytics and AI greatly enhance the security of modern environments from robust data collection, advanced analysis, and automated threat mitigation.

4.7 Centralized Orchestration & Policy Management (OV-1)

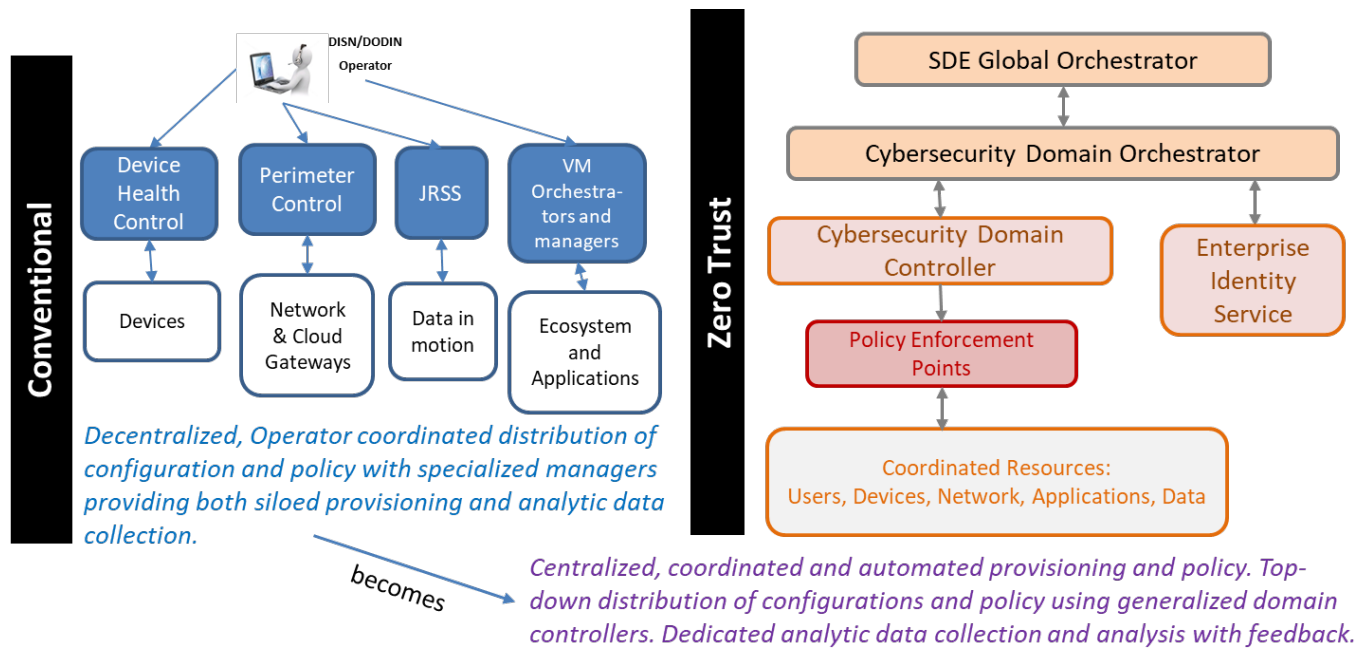


Figure 18 Centralized Orchestration & Policy Management (OV-1)

Traditional architecture has been previously based around administrators applying configuration and policy changes within their specific domains of influence with little regard to other control areas. This has brought around the issue of non-cohesive policies and configurations in the enterprise. ZT is in the process of shifting the paradigm around to a centralized orchestration of Comply-to-Connect methodologies with not only policy creation but also the policy deployment and continued validation of those policies looks to change that. In the new ZT access to and through resources will be coordinated from a centralized control structure. Policy will be able to change and adapt quickly to new threats in the environment as well as to allow automation to deploy those changes more efficiently and quickly to enforcement points in the field.

As the controlling agent for security in the enterprise, the Cybersecurity Domain Orchestrator will interrogate a SDE Global Orchestrator for the desired state of the environment to determine the delta in security policies. Depending on configuration, one or more orchestrators may be used. In its evaluation, the Cybersecurity Orchestrator will also need to be aware of enterprise logging data about users and NPEs throughout the enterprise. As big data becomes available for analysis and use, the Cybersecurity Orchestrator will be able to utilize that data in the creation and deployment of Policy. Updated information will be posted to the Enterprise Identity Service that Cybersecurity will also use in its policy creation. Security Policy is then pushed down the stack to The ZT Policy Controller and down to the actual Policy Enforcement Points. This will allow a coordinated plan and action around the resources available for the enterprise's security posture and authorization.

4.8 Centralized Orchestration & Policy Management (OV-2)

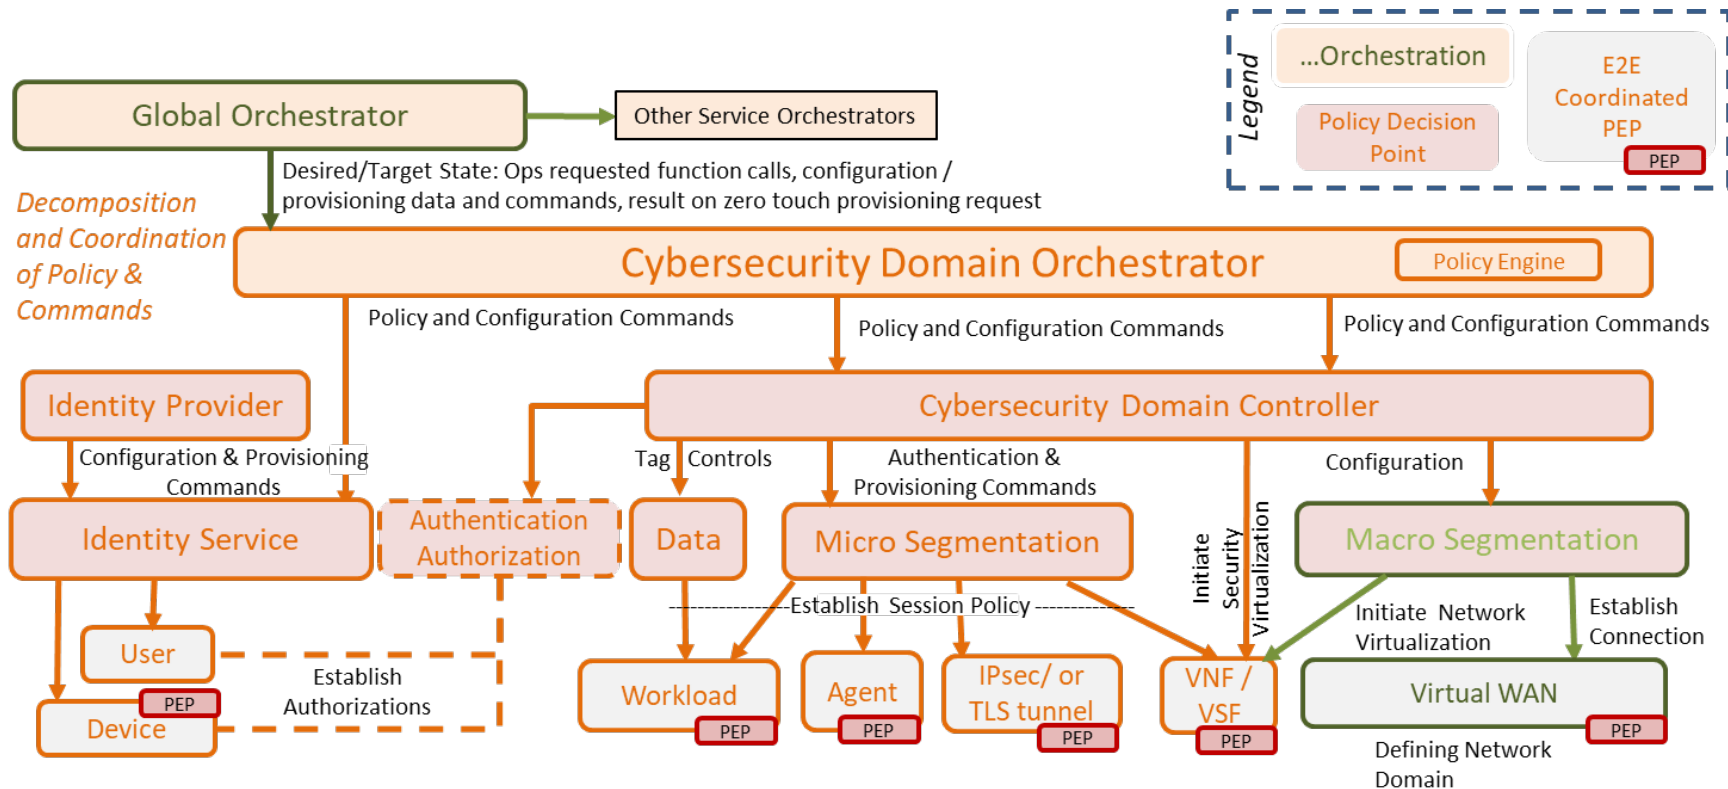


Figure 19 Centralized Orchestration & Policy Management (OV-2)

July 2022

As Policy changes are created, ZT will utilize one or more Central Orchestrators to distribute and verify changes down through the enterprise. The Cybersecurity Domain Orchestrator (CDO) is dependent on reviewing what is to be the desired/target state of the environment from originating organizational provision tools such as a Service Global Orchestrator. The CDO will be responsible for matching up desired/target state with security policy changes, resolve conflicts, and communicate changes down to a controller interfacing layer. The policy that is brought into this policy engine layer will be disseminated out to policy enforcement points specific to the area of influence that policy would be effective at. Ultimately, the controlling policy would allow for specific and a unified control structure to be applied which would allow for an exact control of entities, traffic, applications and data permissions. The CDO will also have to interface with an Identity Services platform to update records and permissions available for Users and Non-Person Entities. As a user or NPE tries to access data, the authorization of that user will now have to be vetted by a unified and cohesive security posture.

4.9 Dynamic, Adaptive Policy Feedback Loop (OV-1)

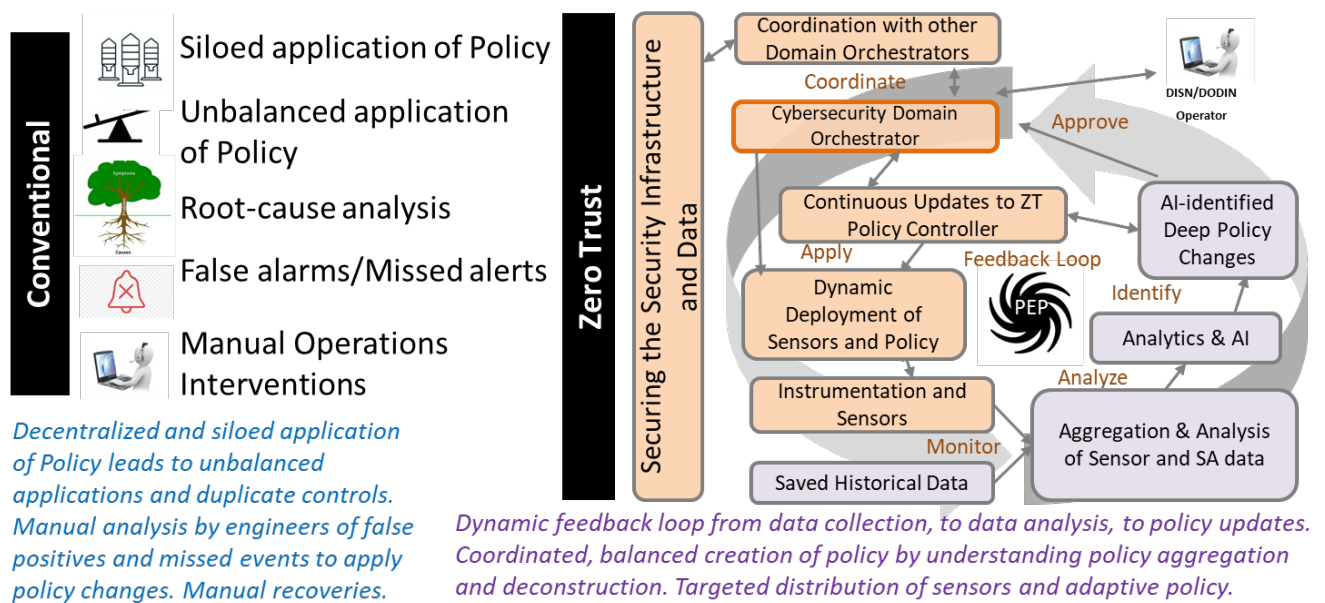


Figure 20 Dynamic, Adaptive Policy Feedback Loop (OV-1)

Environments currently house inefficiencies in their procedures and important security data is often missed, misdiagnosed, not understood or not in the sphere of influence. This requires constant manual interventions by the operations and security teams leading to slow and sometimes incorrect changes to the environment. The ZT approach is to have a unified Adaptive Policy feedback loop. The desired state will be for the policy enforcement points to be continually refined and monitored to more accurately protect users, devices, infrastructure, applications and ultimately data. A created policy will be deployed, monitored, analyzed to identify required

July 2022

changes. Evaluation of analytics, and as technology progress the incorporation of first out-of-band Artificial Intelligence (AI) and later in-band AI, will generate policy for review or for immediate stopgap implementation. These changes will be approved in a coordinated fashion to then be reapplied back to the PEPs to begin the process all over again. Information gathered from multiple parts will allow for more data points of the environment and enable a wholistic understanding of effectiveness of the applied policy and the changes. More data sources will improve AI via Machine Learning (ML). Having a single point of coordination will allow for a unified view of what is applied and how changes might affect other areas that a siloed system would not be so quickly aware.

4.10 VPN-Less Implementation (OV-1)

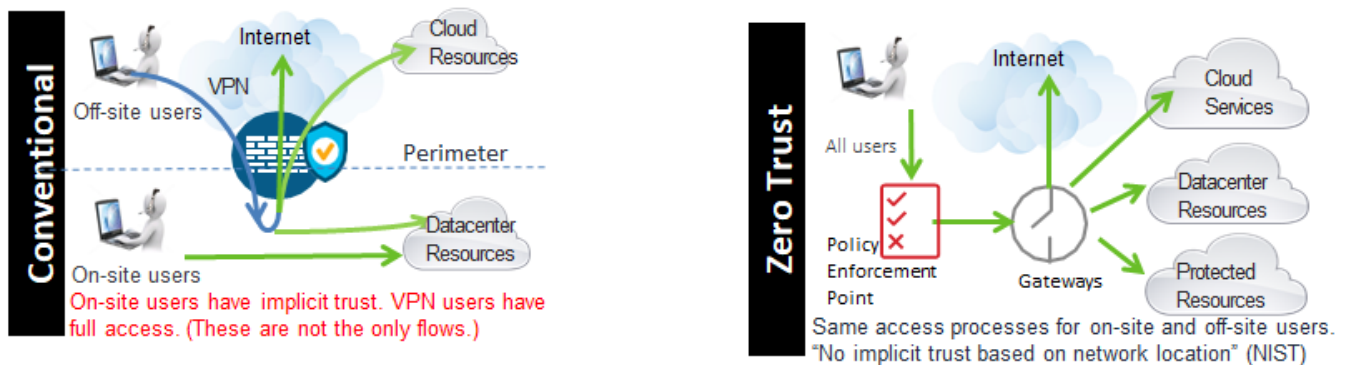


Figure 21 VPN-Less Implementation (OV-1)

A ZT environment dispenses with the distinction between “internal” and “external” users. An internal user should have no implicit trust associated with it than an external user. All users are untrusted. One outcome that can follow is the removal of VPN. In a ZT environment, all users are effectively “external” or untrusted and therefore must undergo the same rigorous authentication and authorization processes.

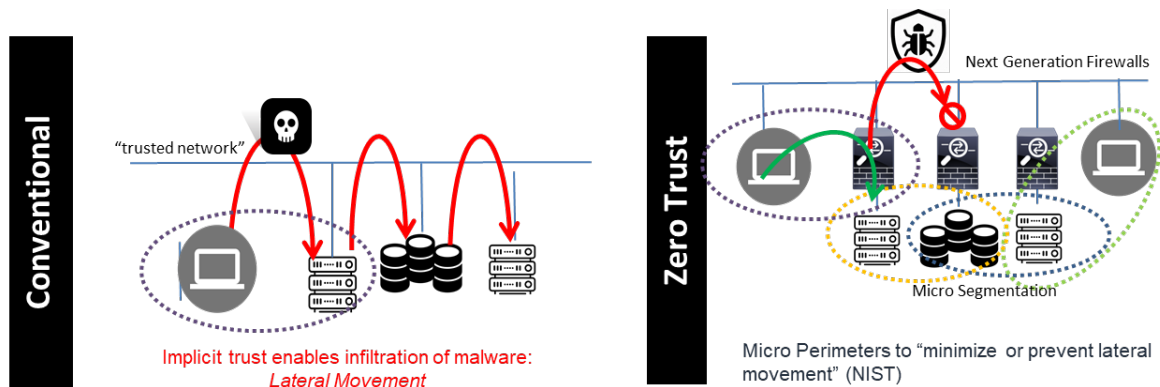
In the conventional approach, off-site users connect to the internal network via a VPN, which effectively places them on the “internal” network with on-site users. If the external user accesses external or Internet resources, traffic first passes through the enterprise perimeter before heading back out. This increased traffic flow requires continuous bandwidth and can create significant latency issues. Additionally, VPNs pose a threat to enterprise security. They create a path in the network perimeter and provide access to network resources after authentication. The conventional approach cannot provide a method to intelligently confirm the identities of users and entities attempting to access the network or provide adaptive policy enforcement based on authentication.

In ZT, all users and NPEs pass through the same Policy Enforcement Point and gateways before they can access resources with Comply-to-Connect, many of which will reside in

July 2022

datacenter resources and cloud services accessible via the Internet. All requests for access will be highly scrutinized using continuous multi-factor authentication and the concept of least-privilege. In this model, formerly external users do not incur additional latency by hair-pinning through a VPN.

4.11 East-West Segmentation (OV-1)



- **Prevent and Curtail Lateral Movement**
- **Explicitly Allow Authorizations (Deny by Default)**

Figure 22 East-West Segmentation (OV-1)

Security states of previous deployments of application and server stacks have had issues involving implicit trust in communication between systems. This trust has allowed malicious users and devices the ability to traverse through the environment with relative ease. Once through the perimeter controls malicious users and software can move laterally across to infect or attack systems and data within the area of influence. ZT aims to enhance the security posture of static DMZ network configuration by only allowing the specific communication that is required for the applications to work and implement ever evolving controls. Micro-segmentation will require communication between devices to be limited with just enough access to complete the intended task of communication between servers, devices and applications. Communication will be controlled not only at the network level between hosts, but also from process to process and in the application stack through API Micro-segmentation. Additional Authentication and Authorization will be part of each step of the process towards the data layer.

4.12 Global Uniform Device Hygiene (OV-1)

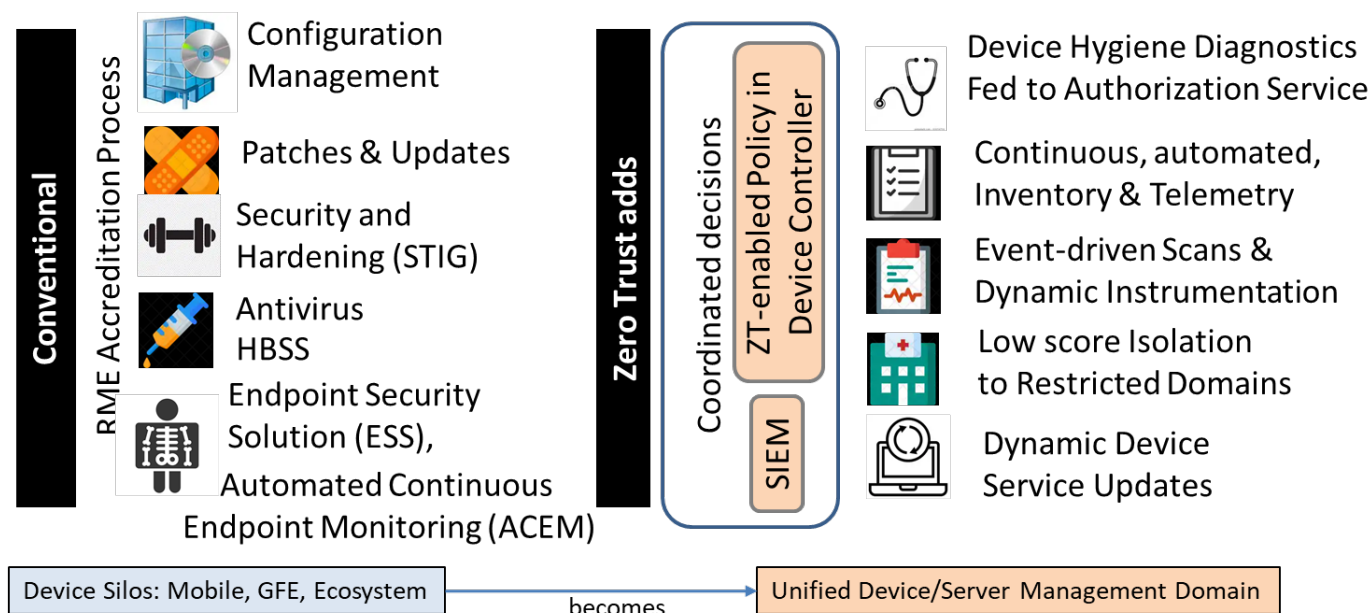


Figure 23 Global Uniform Device Hygiene (OV-1)

While not the only vector for attack, individual unpatched non-compliant systems allow for an attacker to gain a beachhead into any environment. Previous iterations of security have focused on areas or silos of control and do not give a unified view of the security of a device. A system may be completely patched, but a single change to another part of the system may cause that system to become more susceptible to compromise. ZT aims to change this by giving a unified centralized evaluation of systems and a coordinated response to not only system status but event analysis and action over the environment and on individual devices.

Conventional device hygiene has been focused around hitting checklists, being at certain version numbers and general event monitoring of a system. If a system hits a certain number of Security Technical Implementation Guide (STIG checks,) is up-to-date and is not currently flagged as being infected, it has been considered safe to be on the network.

ZT aims to take what was done in the past and add to it by incorporating a more unified strategy and inclusion of event data to inform decisions. A system will be under the same patching that was required before, but hygiene will now be a part of the authorization to specific information. It will also be continuously checked not only by the patching systems, but also from other hygiene tools in the environment. Metrics will be collected about systems so that baselining can be configured for devices. ZT will aim to not only create baseline what a normal device would look like in the environment, but also what the patterns are of individual machines. Discrepancies between current actions and previous patterns can have different policies applied to the device. This pattern recognition will allow for event-driven triggers and scanning to take place.

July 2022

Triggering on the detection of a system will initiate a unified and coordinated policy to be provisioned and later applied to policy enforcement points through the environment. Dependent upon the severity, policy can roll out a gradual change or an instant termination affect as required to protect data. Additionally, Policy Controllers can pull confidence level scoring, that will be available from big data and analysis, to assist in the creation and enforcement of policy in the environment. ZT also looks to add confidence scoring not only to users but also devices. Erratic systems could have their score affected by network behavior, process behavior, or other defining characteristics.

4.13 Global Uniform Device Hygiene (OV-2)

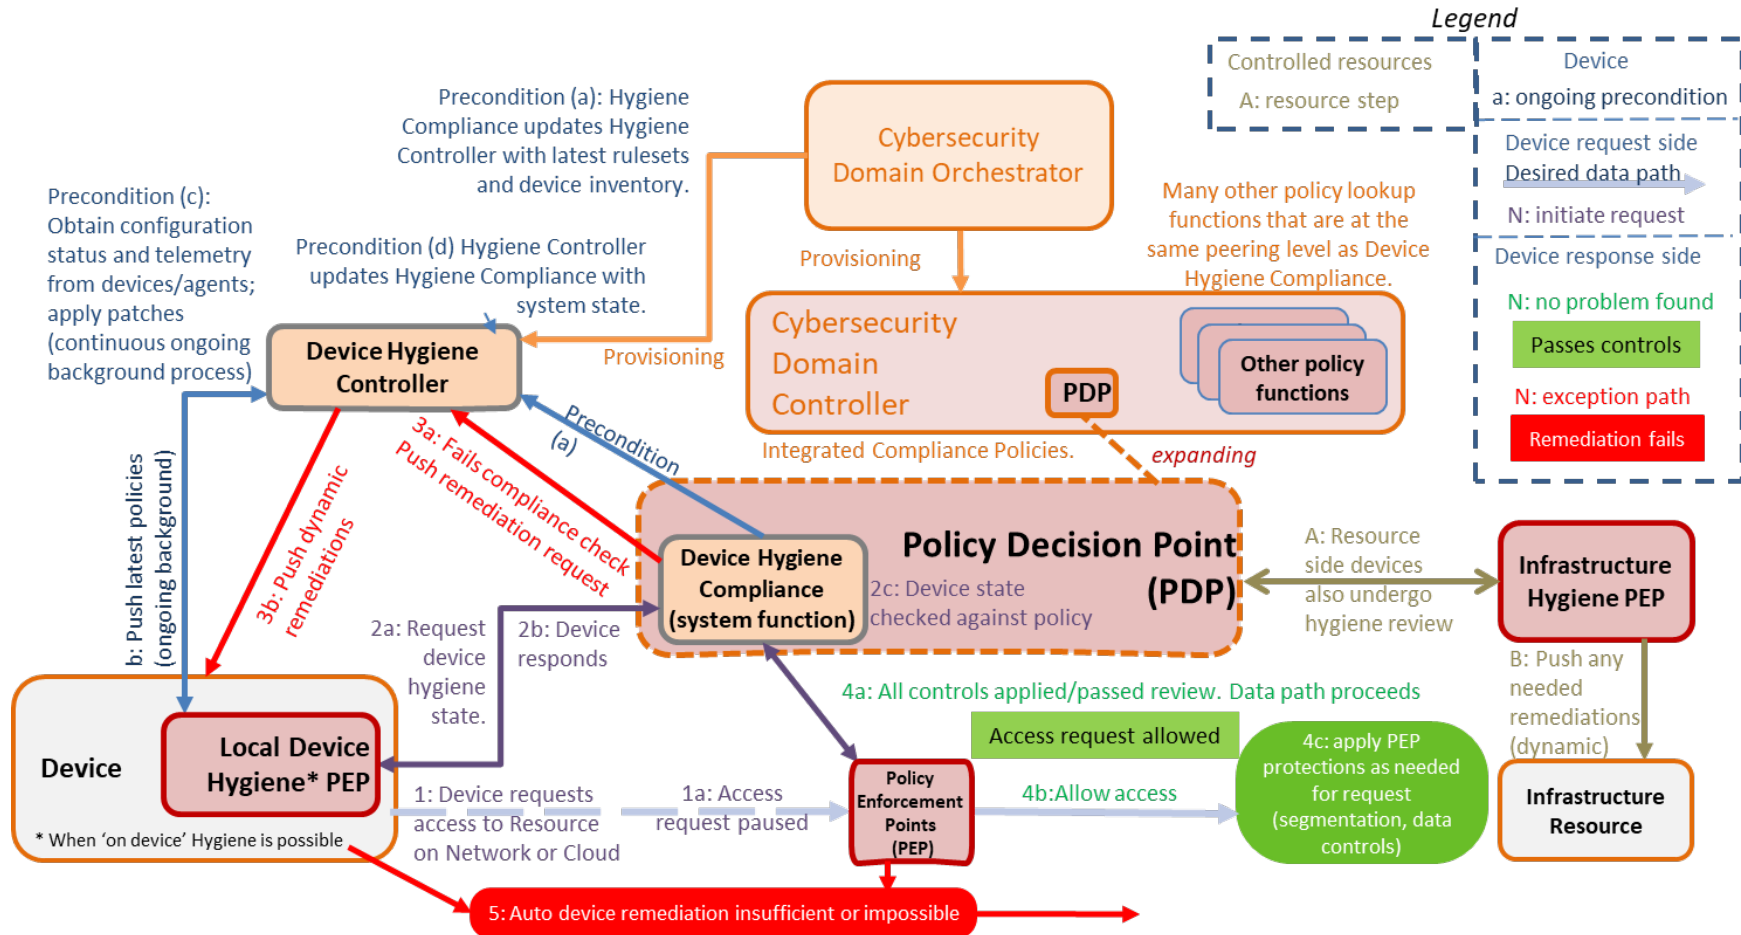


Figure 24 Global Uniform Device Hygiene (OV-2)

July 2022

Global Uniform Device Hygiene is different than today's methodology as it doesn't only include software patches or STIG checklists. Instead, it relies on ZT policies that guide decisions to achieve logical outcomes. The policy's intent is to be a statement that is implemented as a protocol or procedure that can assist with automated decision making. These guides to decision making are based off the Event Condition Action structure:

- The Event specifies the signal or criteria that invokes the rule for condition.
- The Condition is the logical test that causes an action to be executed based on if the condition is met or true.
- The Action consists of updating policy in accordance with the condition on network and data access level.

ZTA validates in real time to ensure any vulnerabilities or rulesets that apply to the device to be validated and corrected to ensure it's in compliance with the applied policy at the time it tries to access any resource. It is constantly checked against any possible exploits and if any exist, it attempts to remediate and if that's not possible it will remove it from the environment to mitigate any exploitation.

4.14 Dynamic, Continuous Authentication (OV-1)

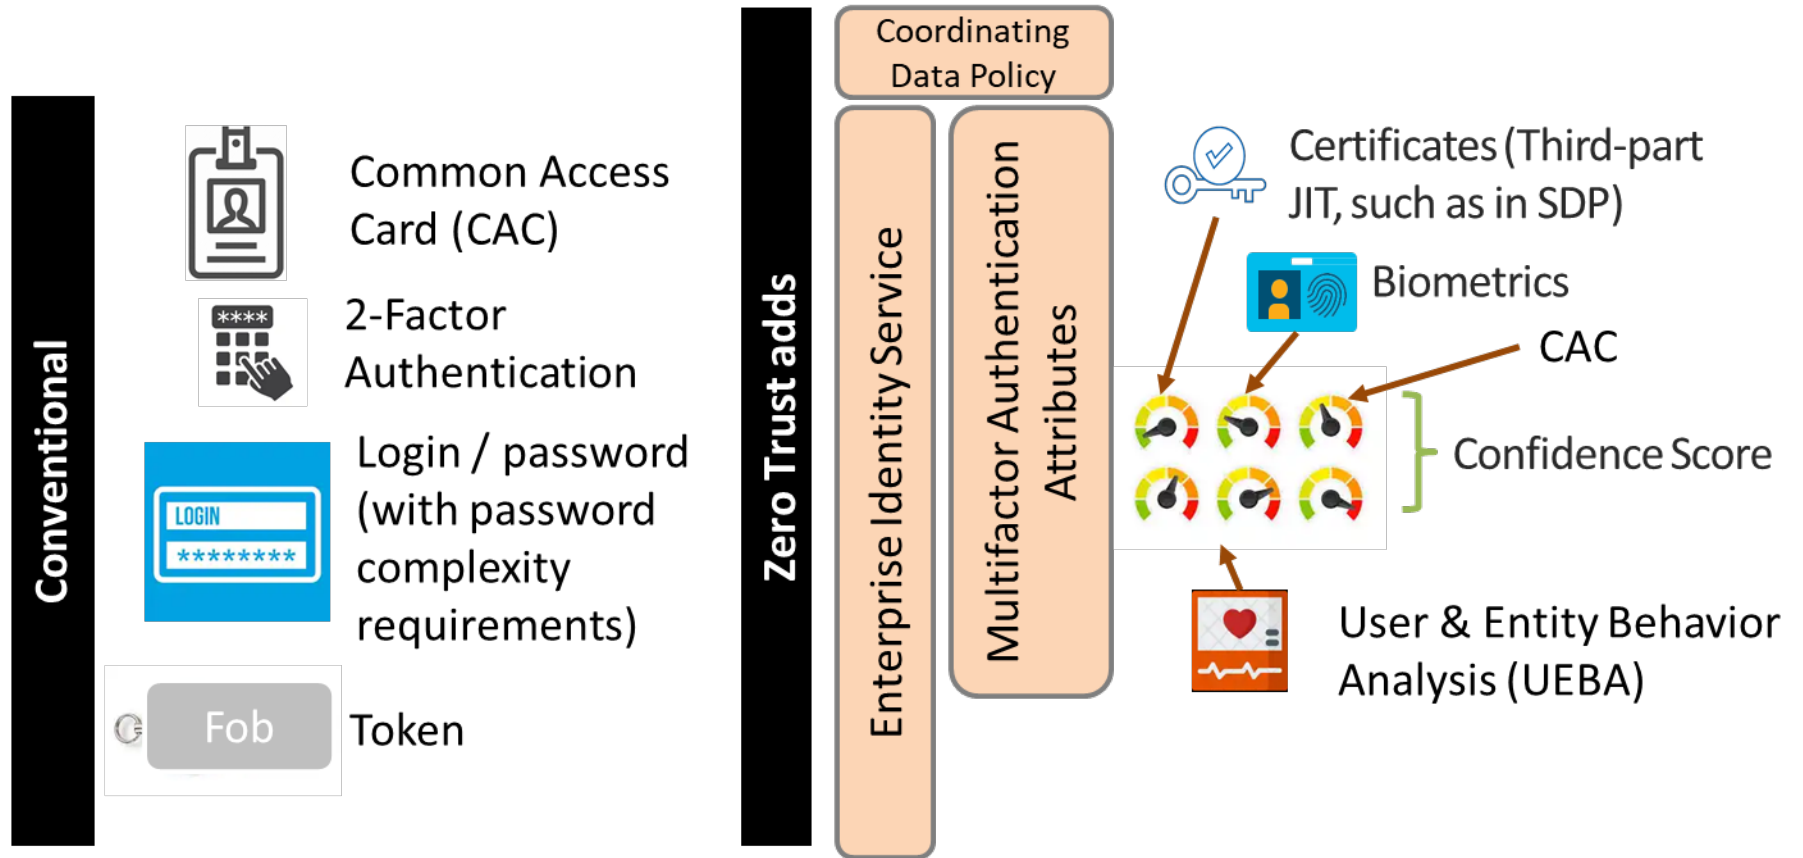


Figure 25 Dynamic, Continuous Authentication (OV-1)

July 2022

ZT is not eliminating the all conventional means of authentication. ZT security model shifts to a more focused use of multi-attribute-based confidence levels to enable authentication and authorization policies based on the concept of least privileged access. The conventional use of persona-based identities, credentials, and attributes are not dynamic or context aware. Current methods tie to a user's physical location. After authentication, every person entity and non-person entity is treated the same. Two factor authentications, authentication tokens, and username and password login have not kept pace with the industry's multi-factor authentication advances. Conventional authentication methods do not address non-persona entities such as bots, hardware devices, or software applications.

ZT draws on technologies such as multifactor authentication (MFA), enterprise identity service, and user/entity behavior analysis (UEBA) to enable continuous and dynamic authentication. These tools evaluate the identity of the user or NPE (non-person entity) in real time as access to applications and data are requested. User and entity transactions are continuously monitored for anomalous behavior, which is then flagged, and the user/entity is then restricted access. To fully enable ZT, enterprise identity service and multifactor authentication are both critical. Use of a single unified platform with integrated identity and access management to provide MFA is ideal to avoid any gaps in security or any hurdles to implementing ZT, enabling thorough and continuous monitoring.

4.15 Dynamic, Continuous Authentication (OV-2)

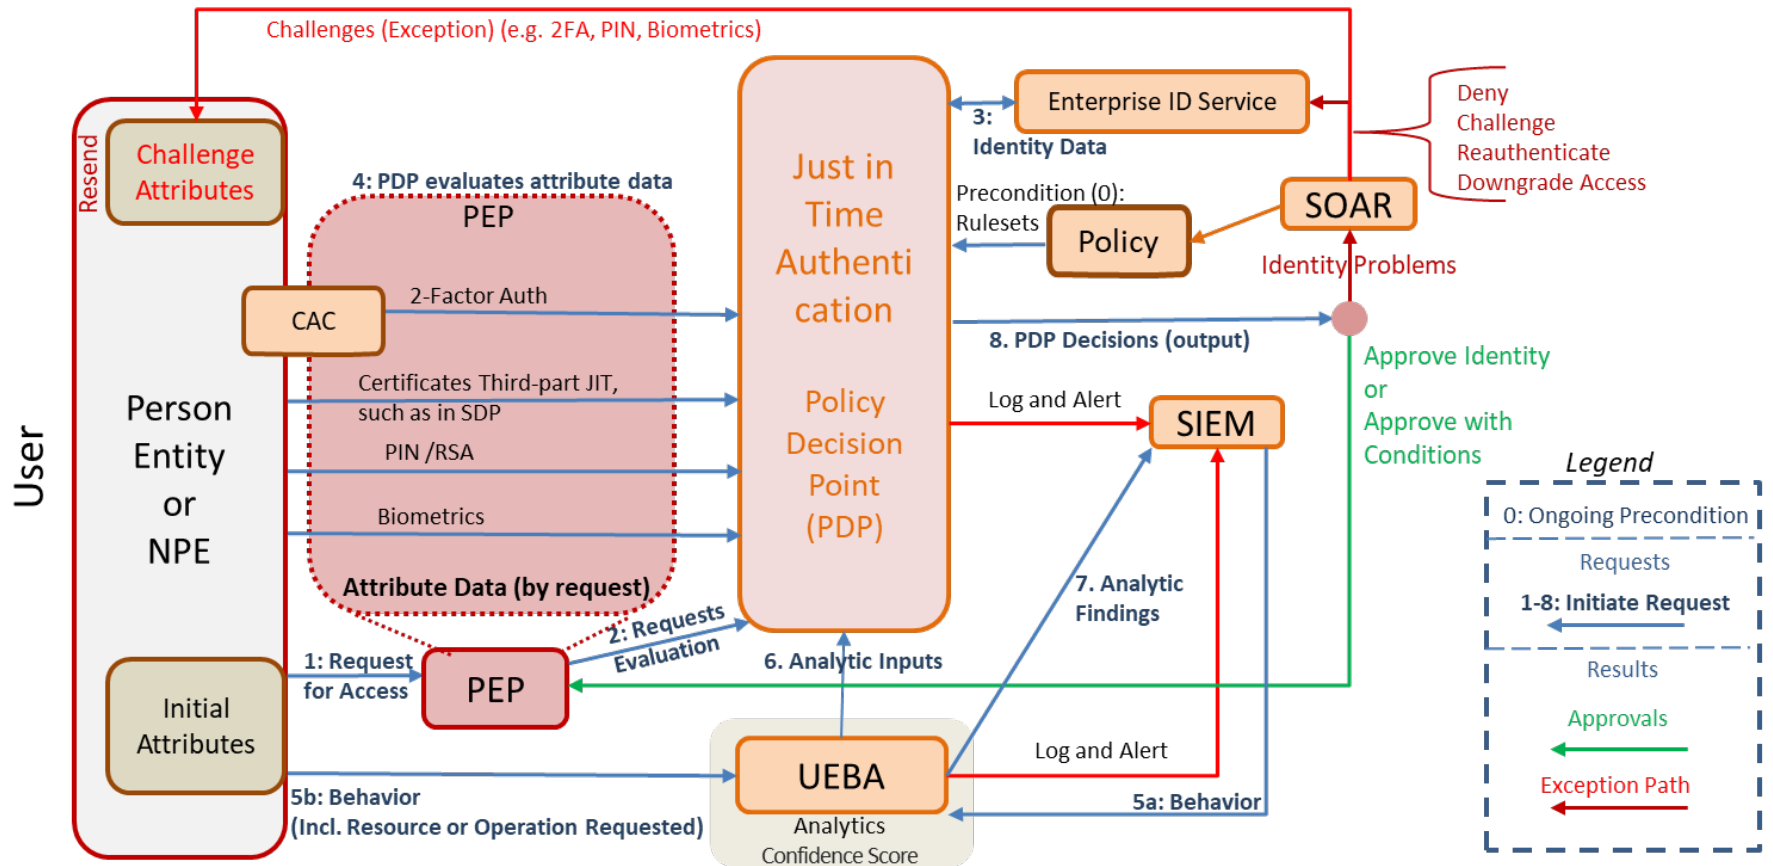


Figure 26 Dynamic, Continuous Authentication (OV-2)

July 2022

The process of dynamic, continuous authentication begins with a user's/NPE's request for access. Attribute data such as a CAC and certificate or biometric will be provided to the identity agent for validation. Throughout the authentication process, behavior data such as time of day, resource or operation requested is collected at policy decision points and are logged to the SIEM which then feeds analytics to the UEBA engine for analysis. The UEBA engine uses analytics to develop confidence score. Confidence levels are developed for each access request and then are distributed to policy points for enforcement. The attribute is then checked for validity at the policy decision point (PDP) and the decision is to either approve the identity or sent to the SOAR to deny, challenge, re-authenticate or downgrade access.

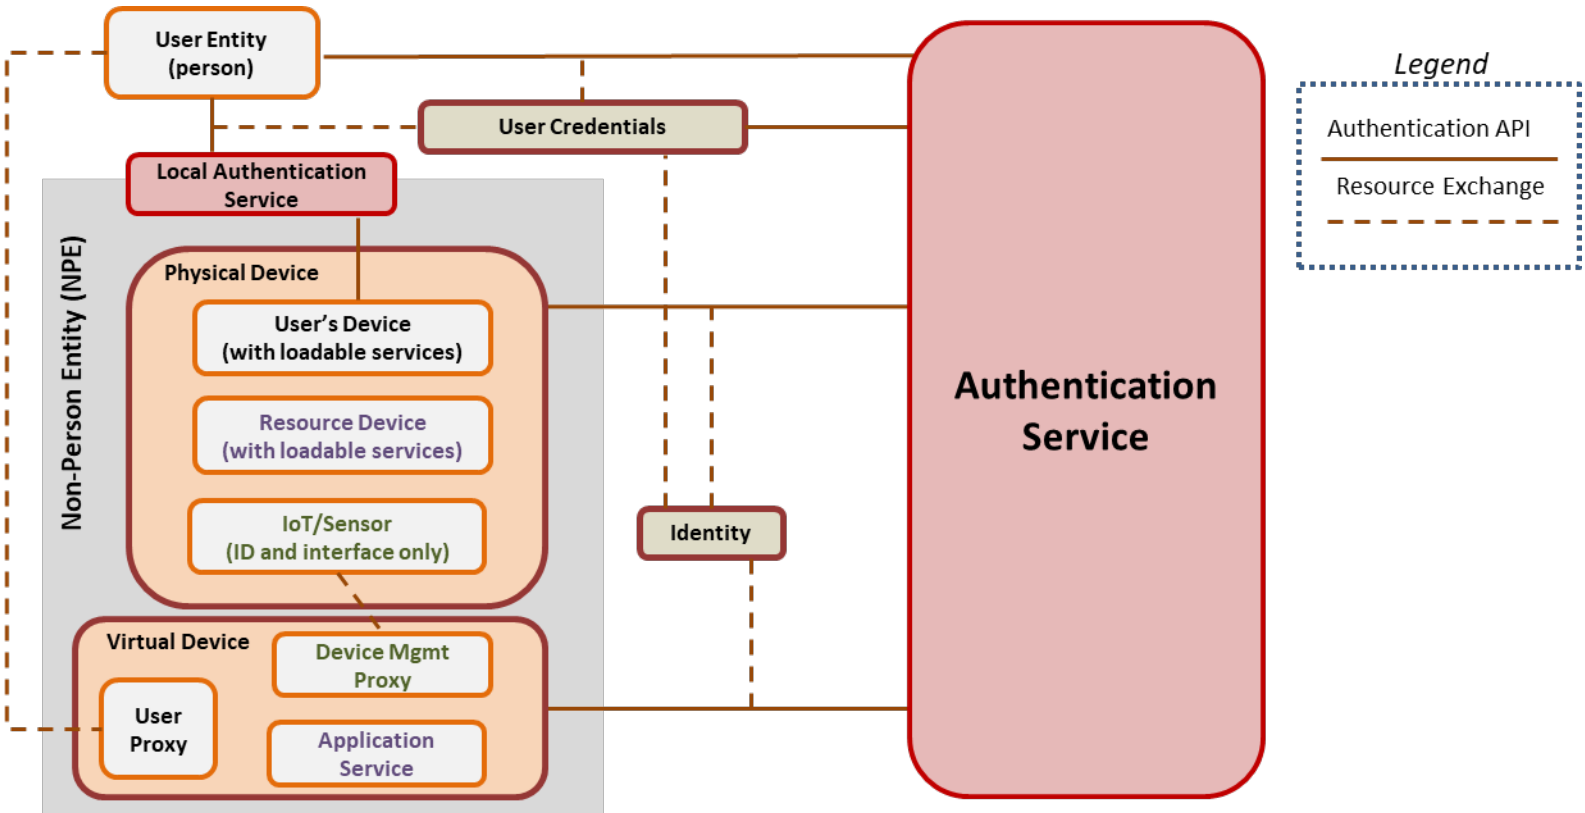


Figure 27 Performers Requiring Authentication

All User and Endpoint entities must be authenticated before an established connection to any resources can be authorized. Users will provide credentials that proves the identity of that user and, if validated, is authorized to access the resource. A Non-Person Entity is everything else which also maintains an identity that seeks to establish a connection with resources and needs to be authorized and authenticated. Authentication can also be assisted by utilizing a proxy for both the user and NPEs.

July 2022

Examples of NPE Devices that require authentication:

- User Device (with loadable services): Device used by a user to initiate and use a session for access to resources
- Resource Device (with loadable services): Infrastructure that hosts either an application resource(s); or provides networking.
- IoT/sensor (ID and interface only): A sensor or IoT device that has a unique ID and that can only run embedded services that respond to a manager.

Examples of additional services that can associate unique IDs to authentication services.

- User Proxy: Application that can stand in for the user to an Authentication Service.
- Device Management Proxy: The representation of a unique device ID to the authentication service by a device manager.
- Application Service is software running on an OS that maintains unique information profile and can stand in for a specific unique instance ID.

4.16 Conditional Authorization (OV-1)

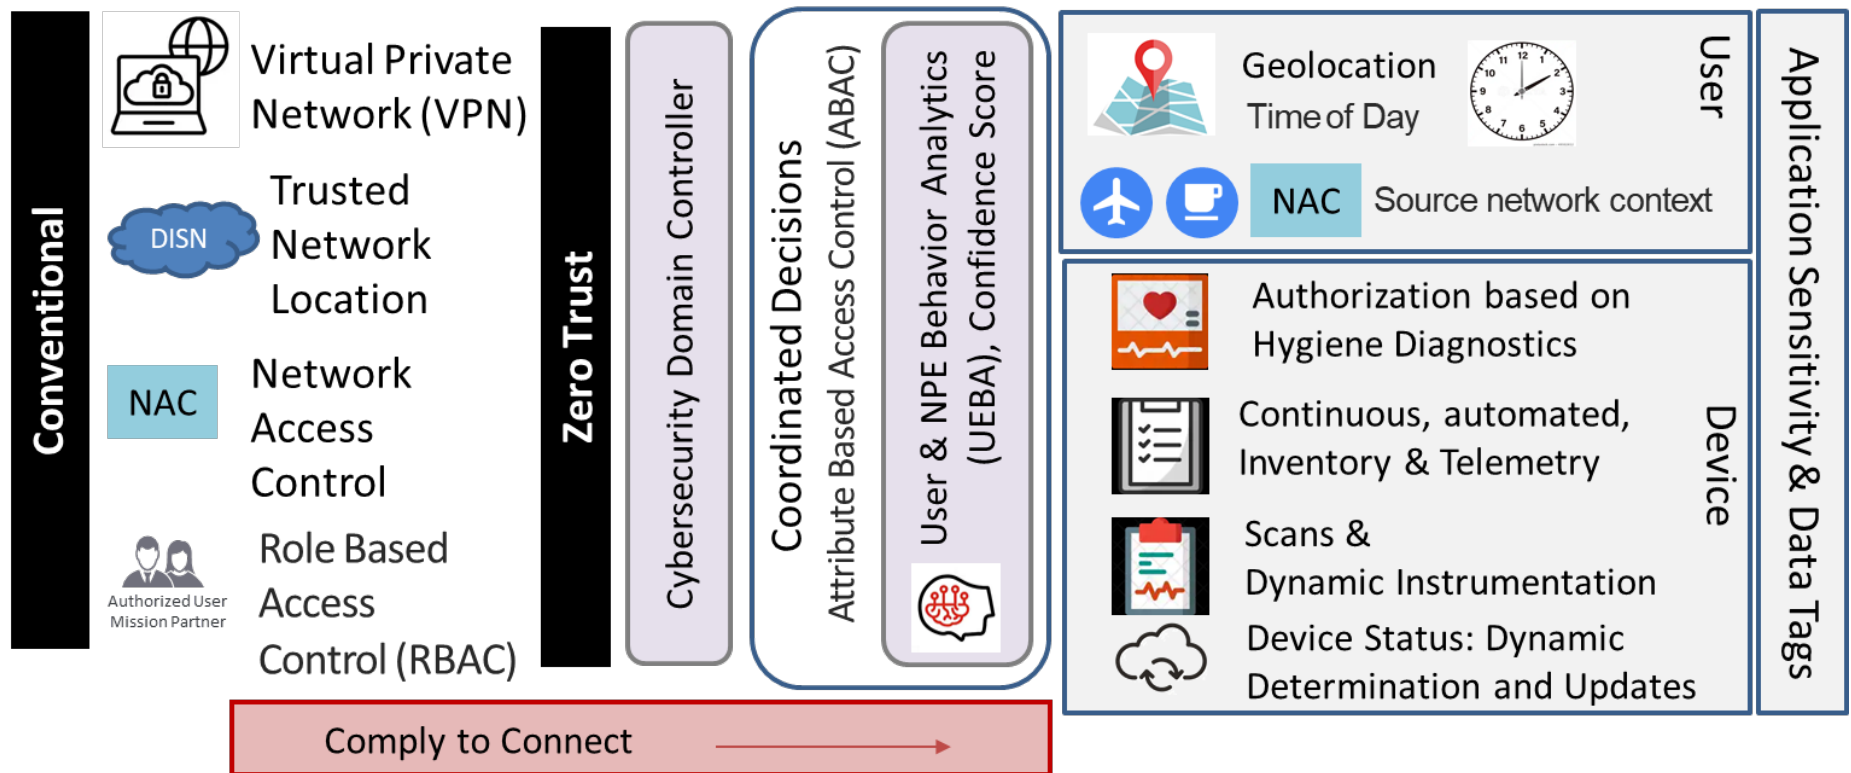


Figure 28 Conditional Authorization (OV-1)

In a conventional approach, authorization is accessed on network location, user or entity role, and authentication methods such as login/password, PKI/CAC, and even two-factor authentication. ZT architecture applies a more comprehensive authentication process that also considers dynamic policy, context, and multifactor attributes such as device health, location, time, and behavior. Activities are logged at the SIEM and User and NPE Behavior Analytics are used to develop a confidence score. Confidence scores are

July 2022

accessed for each individual and NPE and aggregated for policy enforcement. The policy engine develops policy based upon the confidence score of a user or NPE. Authentication and authorization activities occur at focused policy enforcement points throughout the enterprise. Policy enforcement points (PEP) are responsible for enabling, monitoring, and terminating connections within the enterprise. All activity throughout the enterprise is continuously monitored for anomalies in accounts, devices, network activity and data access.

4.17 Conditional Authorization (OV-2)

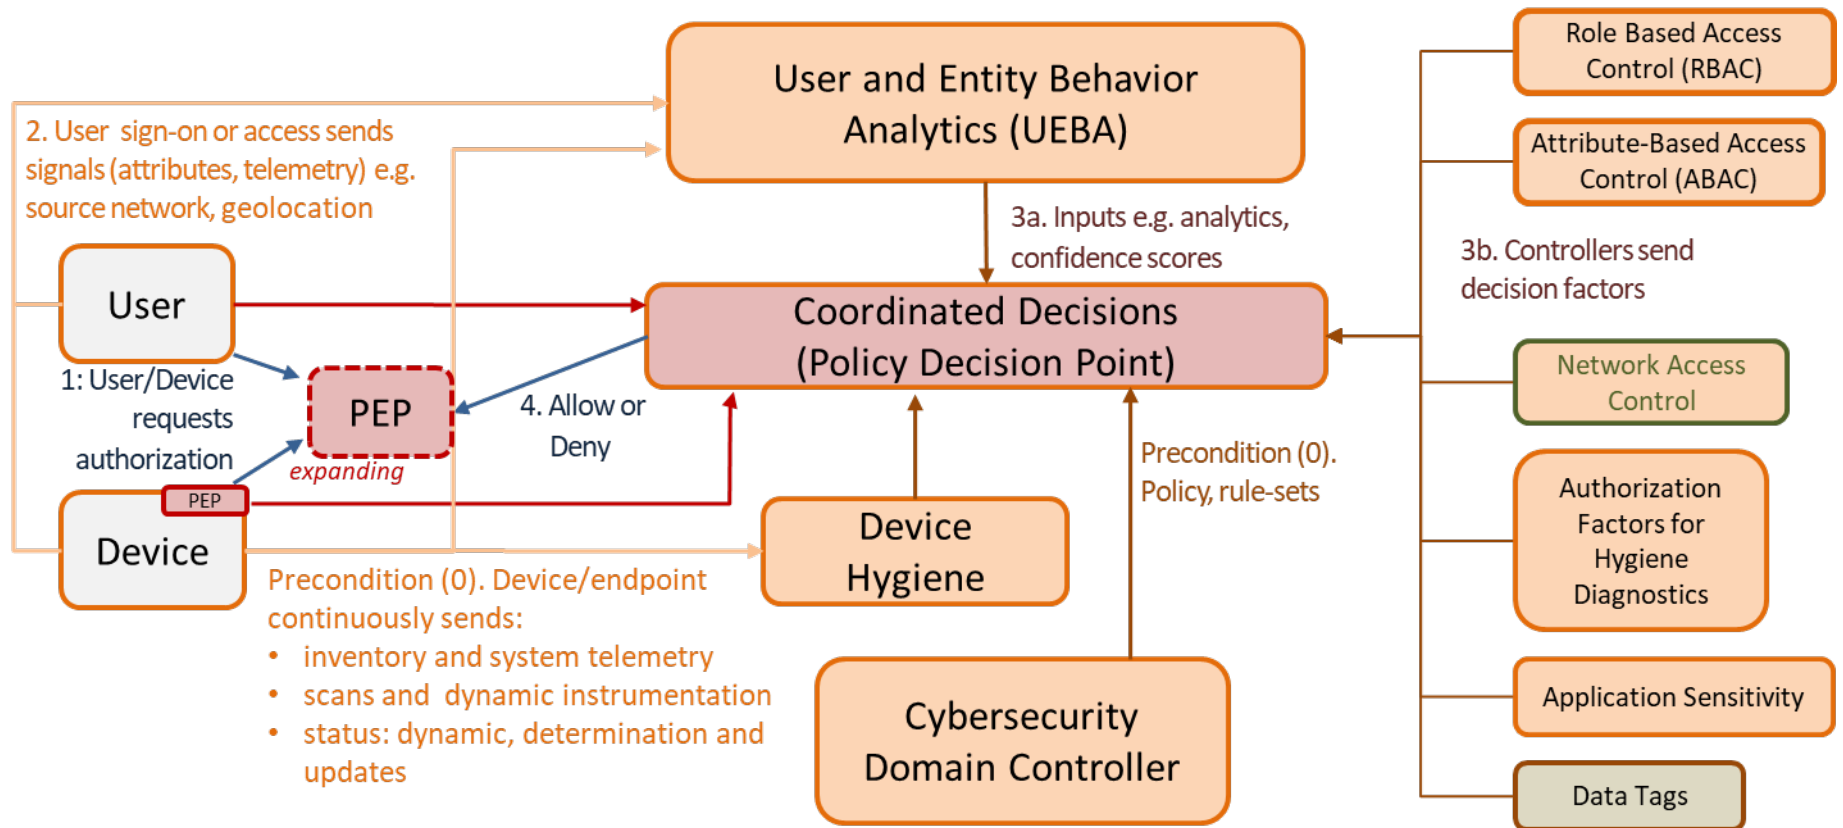


Figure 29 Conditional Authorization (OV-2)

July 2022

During Conditional Authorization, at step 0 the device itself will continuously send inventory and system information, scans and dynamic instrumentation, and the status of the device (needed updates) to the Policy Decision Point. Continually at same time (also step 0) the ZT Policy controller is constantly sending policy to the PDP so specific “rule-sets” are to initialize the request for authorization. The user will then send off a request for authorization from their device and if the device has passed the step 0 check the user/device will hit the PEP. The user will then send their sign-on access to try and gain authorization. In this step we are looking at the User and using NPE Behavior Analytics to prove that the user/device is who they say they are. During this step we are checking a multitude of information such as Role Based Access Control (RBAC), Attribute-Based Access Control (ABAC), Comply to Connect (C2C), Network Access Control, Authorization based on Hygiene Diagnostics, Application Sensitivity, and Data Tags. Each of these controllers send back a “score” of sorts based on these checks to the PDP to come up with a final score for the user. Once the score comes back at the level deemed fit by the organization then the device/user will be granted authorization.

5 TECHNICAL POSITIONS

5.1 Emerging Technologies

ZT requires incremental technology changes to achieve an ZT end-state. Emerging technologies can be used as technical opportunities. Examples are provided in the OV-2s and SV-1s above. Emerging technologies are solution capabilities that, as they evolve, likely will meet one of more conceptual capabilities associated with the ZT Pillars. It will be up to the Reference Design (RD) and the Reference Implementation (RI) architecture to specify how these technical capabilities are used. ZT is not reached at some arbitrary future point but instead, it is a evolution of technology and operational approaches which over time evolve with the threat environment it is seeking to ameliorate.

5.2 Standards, Associated Architectures and Guides



Figure 30 Standards Profile for DoD Zero Trust Architectures

There is no single set of cybersecurity standards that will define a ZT architecture. Instead, all cybersecurity technical standards developed by the Standard Development Organization (SDO) at both national and international levels play a role in developing and implementing a ZT architecture. The set of ZT tenets drive the development of the ZT architecture principles leading to right cybersecurity architecture and policy development and deployment complying with the DoD ZT requirements. DoD calls for adoption of industry open cybersecurity standards driving cybersecurity technology maturity and providing foundation for cybersecurity systems interoperability.

Some of the industry leading SDOs developing the cybersecurity standards include ISO-IEC JTC1/SC27 WGs, National Institute of Standards and Technology (NIST), Cloud Security Alliance (CSA), Internet Engineering Task Force (IETF), and other open standards development organizations. Other than the technical cybersecurity standards that provide technical capability and ensure end-to-end interoperability, the laws, regulations and policies (LRP) issued by US federal government and Department of Defense set the cybersecurity policies for DoD to develop and deploy ZT architecture across the department. Both existing and mature cybersecurity standards and emerging standards are essential in ZT architecture development and guaranteeing interoperability across the DoD enterprise.

DoD is aware of Federal Associated standards and designs including the GSA and HLS ZT Frameworks.

July 2022

The Federal ICAM Architecture framework provides an external guide for the development of DISA's ICAM.

5.3 Linkages to Other Architectures

5.3.1 DoD Cybersecurity Reference Architecture (CS RA) Integration

5.3.1.1 Architecture Description

The CS RA describes the capabilities, services, activities, principles, functions, and technical infrastructure necessary to successfully operate and defend the Department of Defense Information Network (DoDIN). This RA is not static but provides a baseline (or standard) list of cybersecurity capabilities. Technology and architecture will be configured to support any interim, transitional, or objective cyberspace command and control (C2) model selected for implementation by the DoD. The CS RA will serve as a primary source of guidance for RAs, solution architectures, and programs necessary to achieve the vision of the Joint Information Environment (JIE) and will be used to assess compliance of security architecture to established standards.

5.3.1.2 Architecture Usage

The CS RA will be used by DoD Components as the basis for development of Component-specific solution architectures, engineering documentation, and implementation plans. This document will serve as a source of input for funding justification, acquisition planning documents, testing and evaluation plans, and information technology portfolio management decisions. The CS RA should also be considered for relevance to existing and new programs.

5.3.1.3 Linkage

The CS RA provides an architectural frame of reference for implementations but does not currently incorporate ZT (as of version 4.1). As a result, DoD ZT Reference Architecture will be authoritatively referenced in the ZT addendum of the DoD CS RA which will include other non-infrastructure considerations. Updated versions of the CS RA will infuse ZT Principles and Pillars migrating from a Perimeter Centric Architecture to a ZT Architecture. This RA will account for critical security considerations around identity, automation and data security while the CS RA will account for higher level security and engineering concepts. This coverage of ZT within the CR SA will allow ZT RA and CS RA to align. This document will account for critical security considerations around identity, automation and data security while the CS RA will account for higher level security and engineering concepts.

July 2022

5.3.1.4 Artifact Availability

The CS RA is accessible on SIPRNet via the Warfighting Mission Area Architecture Federation and Integration Portal.⁷

5.3.2 DoD ICAM Reference Design (RD)

5.3.2.1 Reference Design Description

The purpose of this Identity, Credential, and Access Management (ICAM) Reference Design (RD) is to provide a high-level description of ICAM from a capability perspective, including transformational goals for ICAM in accordance with the Department of Defense (DoD) Digital Modernization Strategy. As described in Goal 3, Objective 2 of the DoD Digital Modernization Strategy, ICAM “creates a secure and trusted environment where any user can access all authorized resources (including services, information systems, and data) to have a successful mission, while also letting the Department of Defense (DoD) know who is in the environment at any given time.”⁸ This objective focuses on managing access to DoD resources while balancing the responsibility to share with the need to protect. ICAM is not a single process or technology but is a complex set of systems and services that operate under varying policies and organizations.

5.3.2.2 Reference Design Usage

This document is not intended to mandate specific technologies, processes, or procedures. Instead, it is intended to:

- Aid mission owners in understanding ICAM requirements and describing current and planned DoD enterprise ICAM services to enable them to make decisions ICAM implementation so that it meets the needs of the mission, including enabling authorized access by mission partners.
- Support the owners and operators of DoD enterprise ICAM services so that these services can effectively interface with each other to support ICAM capabilities.
- Support DoD Components in understanding how to consume DoD enterprise ICAM services and how to operate DoD Component, COI, or local level ICAM services when DoD enterprise services do not meet mission needs.

Each mission owner is responsible for ensuring ICAM is implemented in a secure manner consistent with mission requirements. Conducting operational, threat representative cybersecurity testing as part of ICAM implementation efforts is a mechanism that needs to be used to check secure implementation.

⁷ Cybersecurity Reference Architecture, Version 4.0, July 2016

⁸ DoD ICAM Reference Design, June 2020

July 2022

5.3.2.3 Linkage

The DoD ZT RA leverages concepts and lexicon from the ICAM RD to provide a unified and consistent approach to implementing ZT Architecture. This document will not include exhaustive references to ICAM use cases but will acknowledge critical concepts as enablers to ZT. References to the ICAM RD are included throughout the DoD ZT RA, however more in depth ICAM specific use cases are only available in the ICAM RD.

5.3.2.4 Artifact Availability

The ICAM Reference Design is accessible on via the DoD CIO Library at:

https://dodcio.defense.gov/Portals/0/Documents/Cyber/DoD_Enterprise_ICAM_Reference_Design.pdf

5.3.3 NIST Special Publication 800-207 Zero Trust Architecture

5.3.3.1 Architecture Description

Zero Trust (ZT) is the term for an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources. A Zero Trust architecture (ZTA) uses ZT principles to plan industrial and enterprise infrastructure and workflows. Zero Trust assumes there is no implicit trust granted to assets or user accounts based solely on their physical or network location (i.e., local area networks versus the internet) or based on asset ownership (enterprise or personally owned). Authentication and authorization (both subject and device) are discrete functions performed before a session to an enterprise resource is established. ZT is a response to enterprise environment trends that include remote users, bring your own device (BYOD), and cloud-based assets that are not located within an enterprise-owned network boundary. ZT focus on protecting resources (assets, services, workflows, network accounts, etc.), not network segments, as the network location is no longer seen as the prime component to the security posture of the resource. This document contains an abstract definition of ZTA and gives general deployment models and use cases where ZT could improve an enterprise's overall information technology security posture.⁹

5.3.3.2 Linkage

The DoD ZT RA leverages concepts and lexicon from the NIST guidance to provide a unified and consistent approach to implementing ZT Architecture. References to the NIST 800-207 are included throughout the DoD ZT RA.

5.3.3.3 Artifact Availability

The NIST Special Publication 800-207 Zero Trust Architecture is available from the NIST Computer Security Resource Center:

⁹ NIST SP 800-207 Zero Trust Architecture, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

6 SECURITY ASSESSMENT

6.1 Governance

Similar to isolation among security capabilities, today's governance policies are also individually aligned to data policies and cybersecurity policies. The adoption of ZT principles encourages the convergence of these areas for a unified data and security governance approach.

The analytics focuses on behavior and data residency to determine the impact of the ZT security policies. This drives changes to the policies to improve the security posture along with defensive cyber operations and incident response. As data continues to be more dispersed throughout the enterprise environment, the overarching governance policies need to modernize to align.

6.2 Data Governance (OV-2)

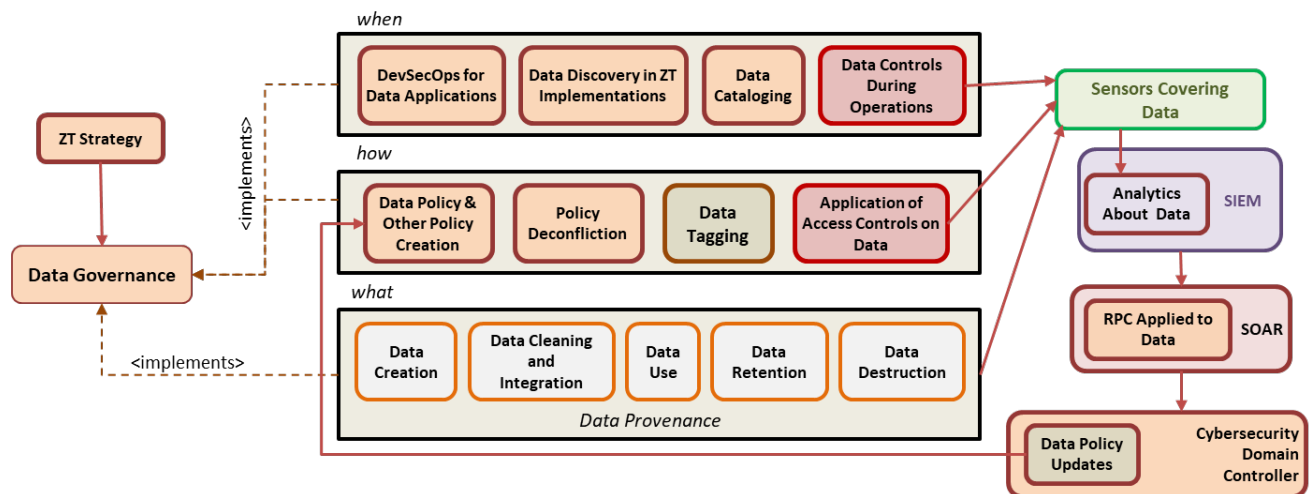


Figure 32 Data Governance: Applying Data Policies (OV-2)

ZT strategy should explain the goals and objectives of data governance in ZT implementations. These will apply to the creation of ZT compliant applications, the application of ZT to existing IT networks and systems, and to operations in a ZT environment. These will be organization and implementation specific. However, some of the architectural concerns follow.

Within ZT, Data Governance is concerned with maintaining proper security and authorized access to data while other data governance policies are applied. In the federal government system, specific access and rights of openness apply to much of the data. This potentially can conflict with the technologies and approaches to data access and security embedded in the core of ZT. These must be reconciled in the ZT architecture so that the data governance technical

July 2022

policies, data access roles, and proper retention policy are implemented within the methods of the ZT access and security.

The first step in applying governance under ZT is understanding the critical data within the environment along with the impact if the data is compromised. ZT implies the understanding that the more sensitive the data, the more potential damage caused by data compromise, the stronger the access controls and other data protections that must be applied. The earlier that risk assessment and data classification can be applied in the software supply chain, the more mature the ZT application.

In operational use, data discovery, risk assessment, and data classification provide a baseline for the application of data and security governance policies. Yet these must occur within an access framework of ZT. Since access controls, implemented at a session level, show what data even can be seen; the functions of data discovery, and the principle of data openness, must be reconciled with the protections of ZT. This likely will be done via an organization's Data Catalog. The data catalog, itself accessed via ZT controls, will contain descriptions and meta data about the data without itself holding that data. Therefore, data cataloging is a companion function to ZT. However, the address/location of data in the organization must be tightly controlled if captured by the Data Catalog.

Once the initial data assessment policy is complete, ZT security rules will be implemented to protect the data and enact segregation of duties. Much of this will be implemented via data tagging. Specific ZT policies will govern how data is tagged and ensure the data is tagged. These tags will be used by multifactor authorization, under ZT policy. Data tags will likely have a much wider organizational function, applying to other data governance needs such as openness and retention. Data tagging is used by ZT, and is a function of mature ZT, but will have business use beyond ZT.

ZT data protection should apply to the full scope of data including sensor data, status information, inventory, and static data properties. It applies to data at rest and data in transit. It also applies to ZT policy statements when policy is data external to the program code. ZT should be applied throughout the full lifecycle of data, from data generation, through normal use, and into retention periods, ending only when/if data is destroyed. ZT data protections will apply to all logs of data use.

The monitoring of data and security governance is incorporated into ZT analytics. Analytics should evaluate the compliance and effectiveness of the governance. As with other ZT policies, the SIEM, SOAR, Controller feedback loop is used to apply automation and refine data policy via actual use experiences. With ZT, this is most important in adapting policy to evolving data attacks.

6.3 Securing Supply Chain (OV-2)

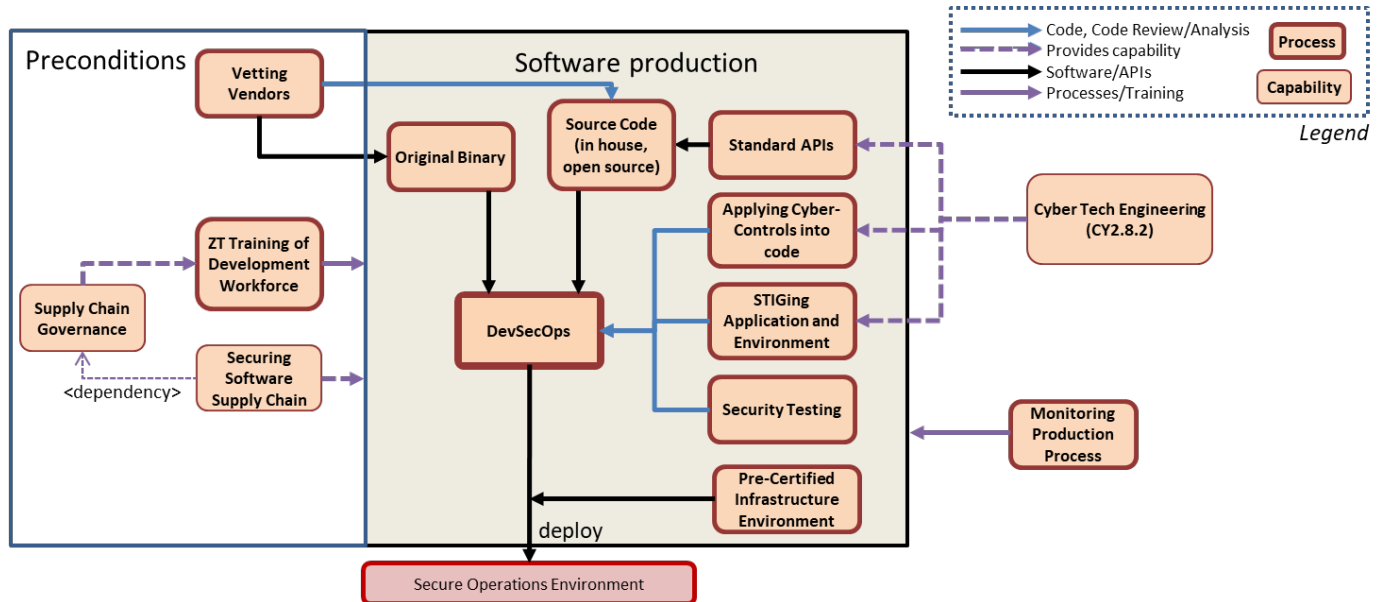


Figure 31 Securing the Supply Chain (OV-2)

The ZT tenet of never trust, always verify, applies beyond the conditional access of data to include the development of applications. The combination of improperly secured applications introducing lateral movement risks and critical data being generated and stored within application necessitates a focus on supply chain.

As applications are built, the source code and binaries need to be vetted throughout the development process. A DevSecOps continuous integration, continuous deployment process includes numerous steps to ensure proper application security. Binaries are evaluated for CVEs and whether they are being incorporated from a trusted DoD source repository. Static code analysis is used in the source code evaluation process to perform dynamic vetting as the application is built. These two security processes ensure the ingredients used in the application development are secure which limits the ability to misuse for lateral movement or other nefarious activities.

The development of binaries and source code should incorporate security hardening such as STIG guidance for the application. Additionally, continuous monitoring of both source elements and the application once promoted to production will provide behavior indicators. A baseline analytic will be developed for application behavior and the security hardening activities can be adjusted in the build process to maintain a robust security posture.

When in the procurement and validation step, the hardware and software need to be given a confidence level. This score is vital and gives a better idea on if the device or data is safe to

July 2022

utilize. If a received device or piece of data has been tampered with, then this lowers the confidence level putting the overall operation at risk. With putting the applications on tampered hardware there is significant potential that the information will be compromised giving the adversaries access to secure information. Depending on where the compromised hardware sits, it potentially could infect other hardware within the building leading to more devices and data being compromised.

7 ARCHITECTURE PATTERNS

7.1 Architecture Patterns (CV-4)

Architecture Pattern	Artifact Type	Architecture Pattern uses or supplies these capabilities
Domain Policy Enforcement for Resource Access	SV-1 SoS	Continuous Authentication, Conditional Authorization, Macro/Micro Segmentation
Software Defined Perimeter	OV-2	Conditional Authorization, Micro segmentation, Encryption
ZT Broker Integration	SV-1 SoS	Continuous Authentication, Conditional Authorization, Device Hygiene
Micro-segmentation via NGFW and SDDC	SV-1 SoS (multiple)	ZT Enabling Infrastructure, Software Defined Networking
Macro-segmentation via SD-WAN	SV-1 SoS	ZT Enabling Infrastructure, ZT Orchestration, SDN, SDDC

Table 2 Design Pattern Table (CV-4)

7.1.1 Domain Policy Enforcement for Resource Access

(SV-1)

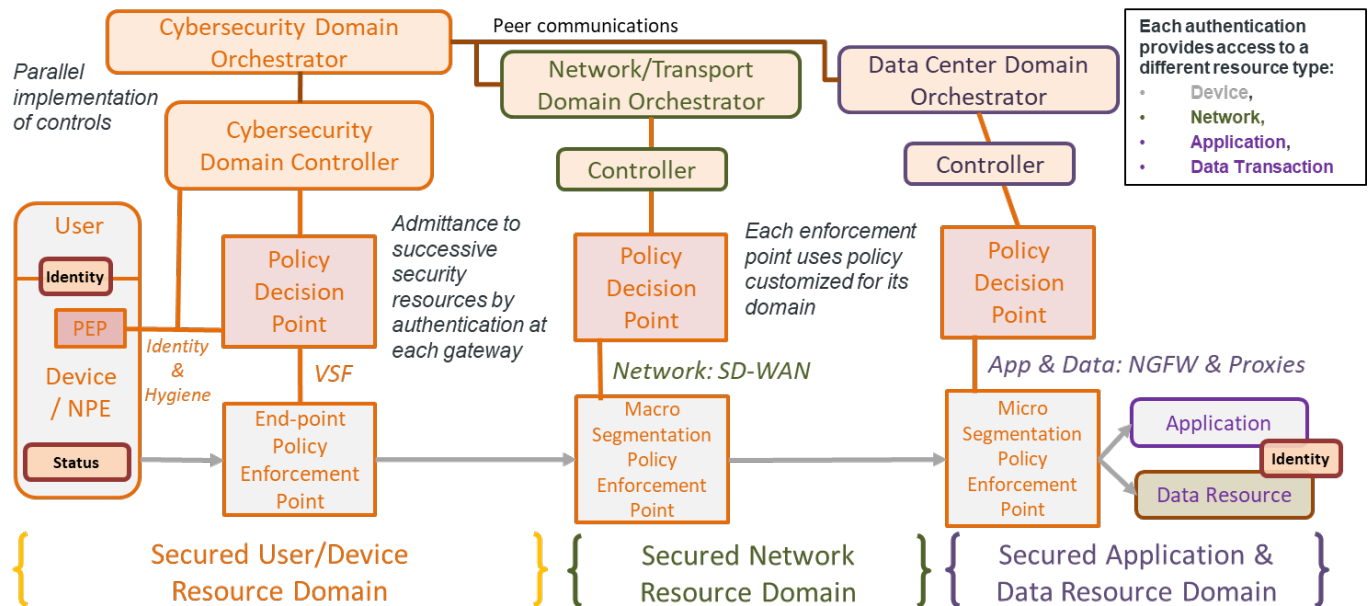


Figure 32 Domain Policy Enforcement for Resource Access (SV-1)

The ZT architecture policy is enforced through parallel domain orchestration. Each security domain provides custom policy orchestration and automated response via its dedicated controller. Network/Transport Domain Orchestrator controls network traffic and grants or denies access to network resource through predefined policies. Similarly, through predefined policies the Data Center Domain Orchestrator is responsible for granting or denying access to application and databases. The collected network and application data is also fed to the Cybersecurity Domain Orchestrator where it is further evaluated for anomalies. Any detected threats will be mitigated by the Cybersecurity Domain Orchestrator via automated policy changes.

7.1.2 Software Defined Perimeter (OV-2)

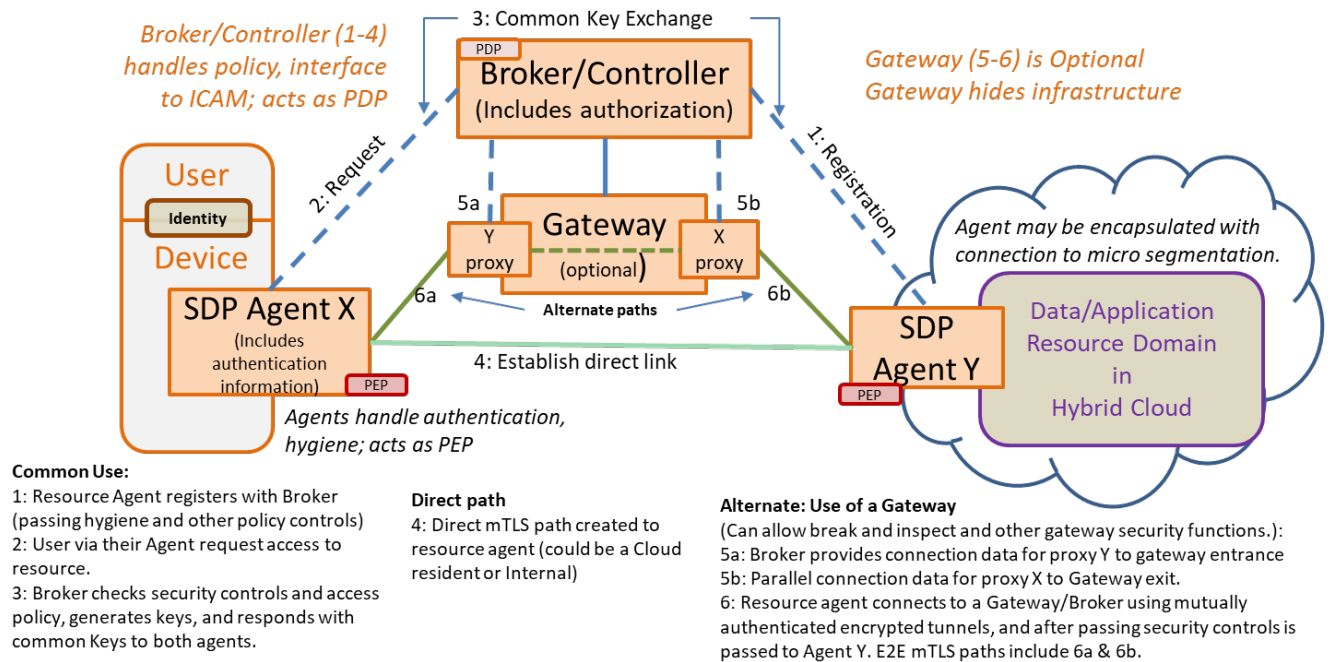


Figure 33 Design Pattern: Software Defined Perimeter (OV-2)

Software Defined Perimeter will move away from the strong network perimeter concept and move towards conditional authorization with micro-segmentation and encryption. While creating an end-to-end encrypted communication path, all data and applications will have direct visibility removed from the public internet. Devices wanting to access resources would be required to pass a ZT enabled SDP. During requests, all communication will be assumed untrusted and require conditional access based on device identity, device hygiene, and user identity with confidence level scoring. These abilities are enabled by agents installed on both the request and receiving end and a ZTA broker with policy enforcement points. An optional but highly recommended piece to include would be a gateway to broker these communications. This gateway would enable the ability to break and inspect traffic to view traffic for malicious actions and data loss. The ability to monitor user behavior, session duration and bandwidth consumption would be vital in providing accurate user and device confidence scores.

7.1.3 ZT Broker Integration (SV-1)

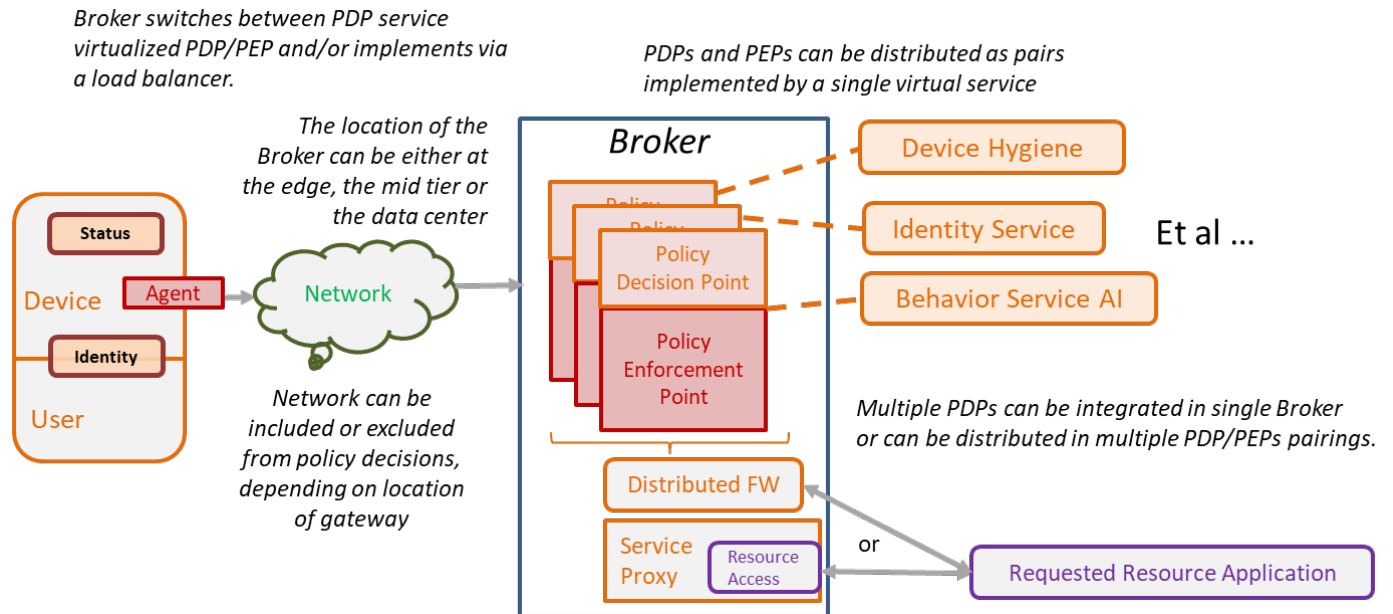


Figure 34 SoS Design Pattern: Zero Trust Broker Integration (SV-1)

In a ZT environment, all applications are hidden from the end user network and require the user or device contact and connect to a trust broker to facilitate connections. This broker can provide continuous authentication and conditional authorization by ingesting device hygiene, identity services, and other factors provide by the ZT big data environment. As traffic flows through or terminates at the broker, access to resources can be denied without the need for sessions to expire. The positioning of the broker can be positioned at the edge or closer to the application as required by the program.

7.1.4 Micro Segmentation (SV-1)

Micro segmentation increases security by breaking down networks into smaller components and enhancing the control of network and process traffic through unified policy enforcement driven by the ZT policy controller. It is important to note that micro segmentation can continue to break down to smaller and smaller components, defining process to process micro segmentation and evolving to API micro segmentation. These micro segmentation patterns below will only touch on network based microsegmentation.

In a scenario where a user makes a request to a three-tier web application, traffic would flow to the PEP of the web server, and if the traffic meets the policy enforced by the ZT policy controller, it would be passed to the application tier. Again, the PEP of the application tier would evaluate the traffic, and if access is granted it would be passed on to the database tier and be evaluated once more before sending the request back to the user.

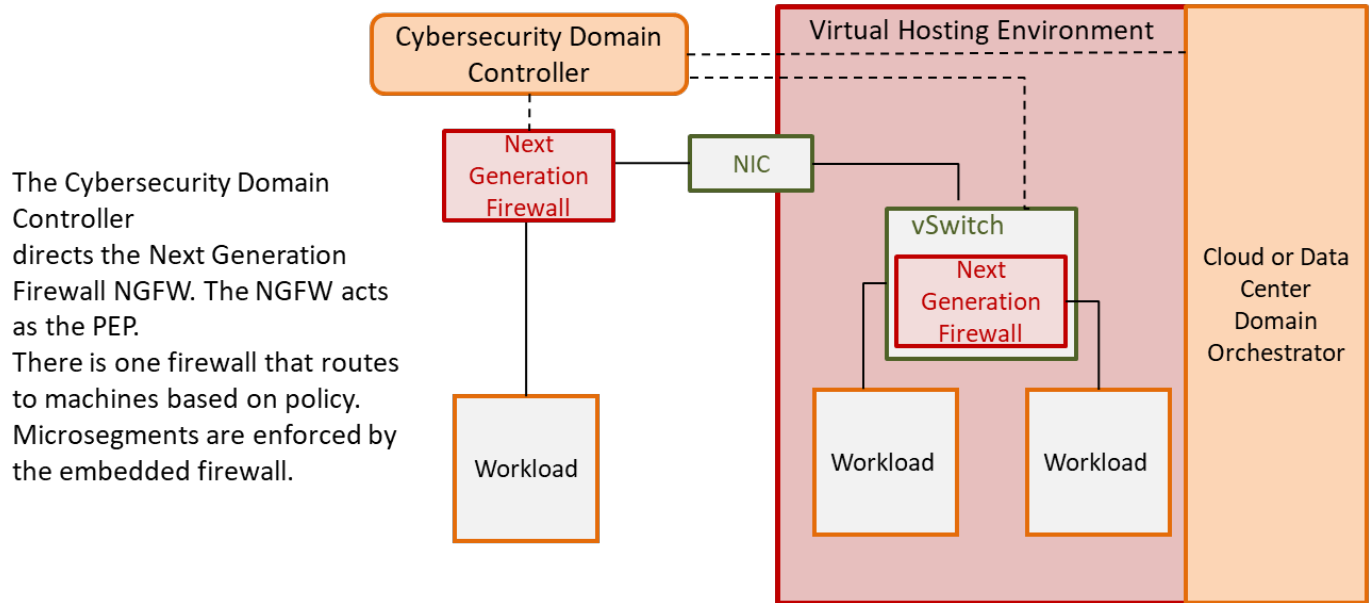


Figure 35 SoS Micro Segmentation (SV-1)

In this design pattern micro segmentation is achieved at the network level through advanced boundary protection capabilities such as a NGFW acting as a PEP. This NGFW provides additional capabilities over a traditional firewall to include intrusion prevention, application firewall functionality, in-line deep packet inspection, malware detection, and policy enforcement. In this architecture all traffic must flow through the NGFW before reaching its destination microsegment. Policy is centrally applied to the NGFW from the policy controller.

The Cybersecurity Domain Controller directs the agents/Micro Firewalls on the virtual machines. The Micro Firewalls are the PEP
The Micro segment is realized by the policy statements to the micro firewalls by the ZT policy controller.

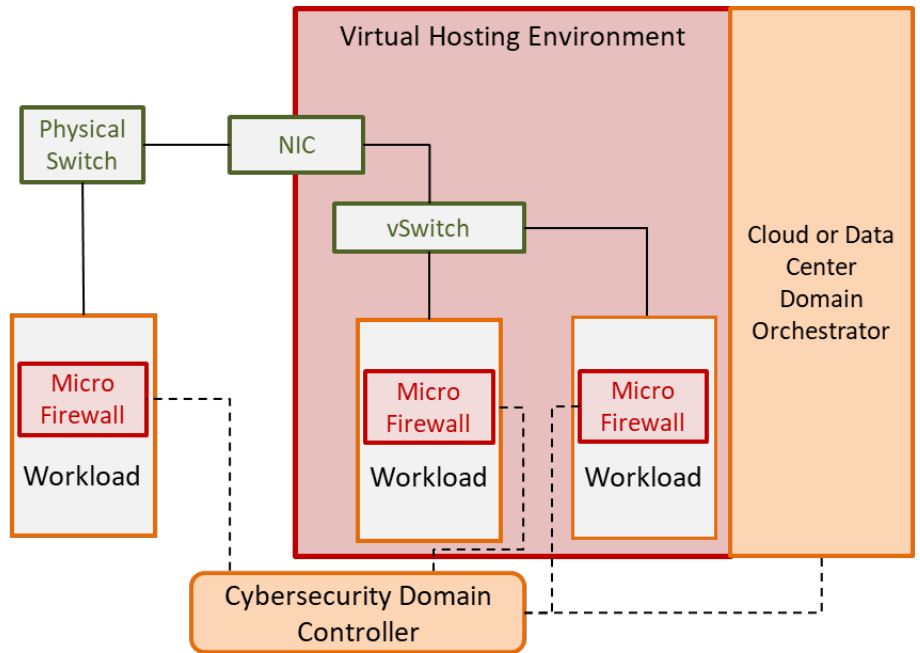


Figure 36 SoS Micro Segmentation (SV-1)

In this architecture micro segmentation is achieved at the hypervisor level through micro firewalls. Each virtual machine has a dedicated micro firewall which filters all traffic before reaching its destination workload. This model provides enhanced threat mitigation of east-west traffic flow and unifies policy enforcement via the ZT policy controller across all micro firewalls.

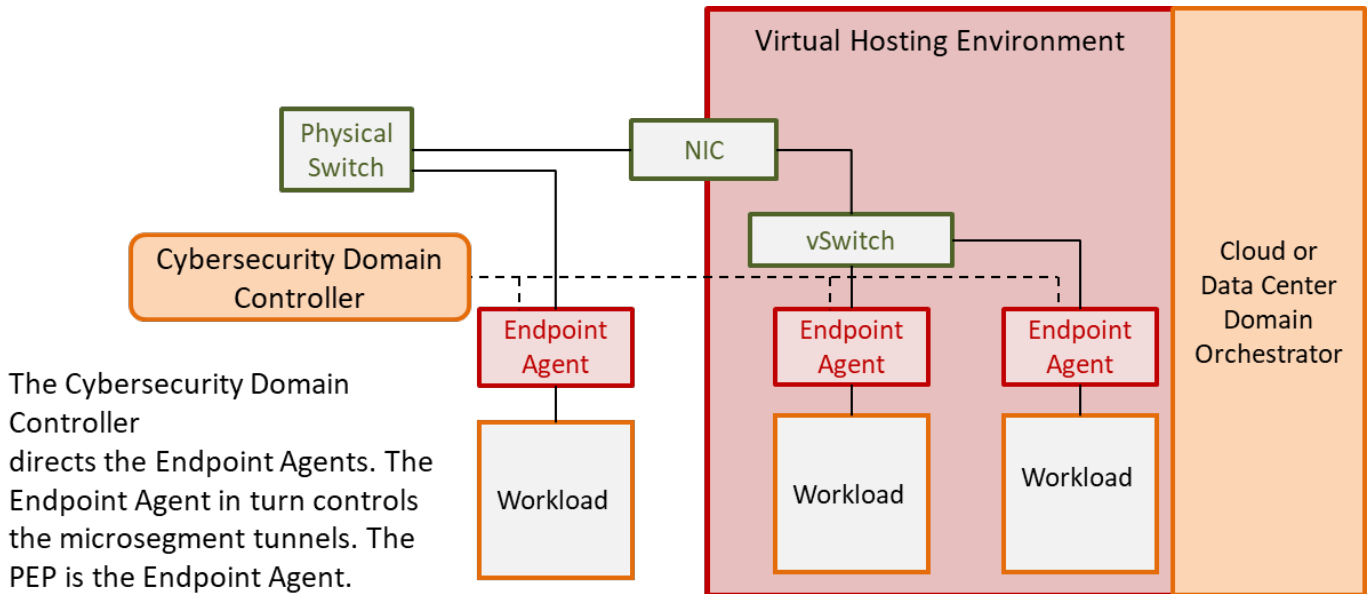


Figure 37 SoS Micro Segmentation (SV-1)

Micro segmentation is established in this architecture through host-based agent endpoints. Each host has its own agent that extends monitoring to the application layer, providing enhanced threat mitigation down to the individual process level. All traffic will be inspected by the endpoint agent before reaching the workload on the destination host. This architecture includes robust protection against lateral movement, granular access control, and unified policy enforcement from the ZT policy controller. In addition, this model is independent of the underlying infrastructure making it flexible to deploy in the cloud or on-premises.

7.1.5 Macro Segmentation (SV-1)

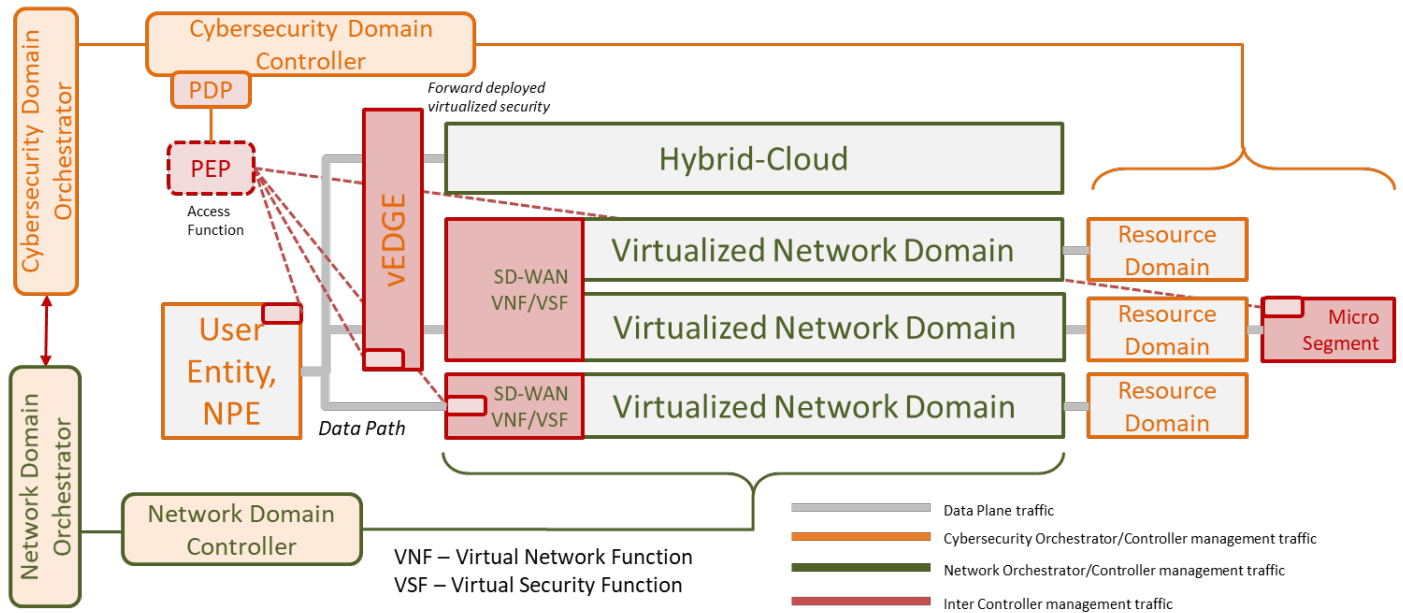


Figure 38 Design Patterns: SoS Macro Segmentation (SV-1)

Macro Segmentation as of today exists with usual VLAN or some sort of broad scope of segmentation done by managed switches in a manual method and usually only provides a perimeter protection model. ZTA expands on that with providing security against devices located within the environment by validating the device, user or NPE on each attempt of accessing a remote resource before it can connect thus providing protection within the perimeter.

For a user, device or NPE to access any resource domain, the entity would need to authenticate and be authorized by a type of access control that houses Cybersecurity policies, such as EIS. If allowed, it will transverse through other ZT capabilities such as SD-WAN/VNF/VSF which also have cybersecurity and separate network policies to ensure the traffic is allowed at that current time and if so, will reach the resource domain.

7.2 External Services

7.2.1 SvcV-1: External Services(SvcV-1)

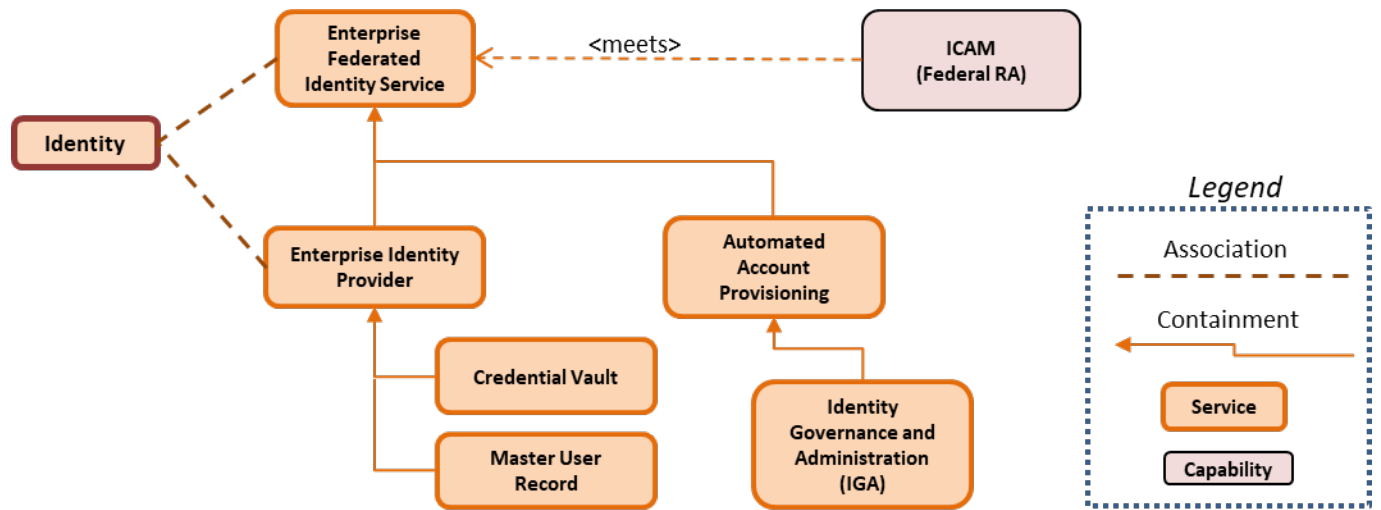


Figure 39 External Services (SvcV-1)

Enterprise Federated Identity Service (EFIS) is a fundamental component of a ZT environment. EFIS is fundamental for the transformation to a data-centric identity-based access management architecture that is required in a future-state ZT Architecture. The Automated Account Provisioning (AAP) will provide identity governance services such as user entitlement management, business role auditing and enforcement, and account provisioning and de-provisioning based on identity data produced during DoD person-centric activities such as on and off-boarding, continuous vetting, talent management, and readiness training.

The identity provider is the system that creates, maintains and manages identity information and provides authentication services based on an individual's identity information. The Identity provider will use the Master User Record (MUR) to enable knowledge, audit, and data rollup to report who has access to what system or applications. The MUR will collect and correlate attribute and entitlement information for person entities that have access to enterprise resources. Enterprise ICAM is the DoDs version of an Enterprise Federated Identity Service.

July 2022

7.2.2 SvcV-2: Enterprise Federated Identity Service (SvcV-2)

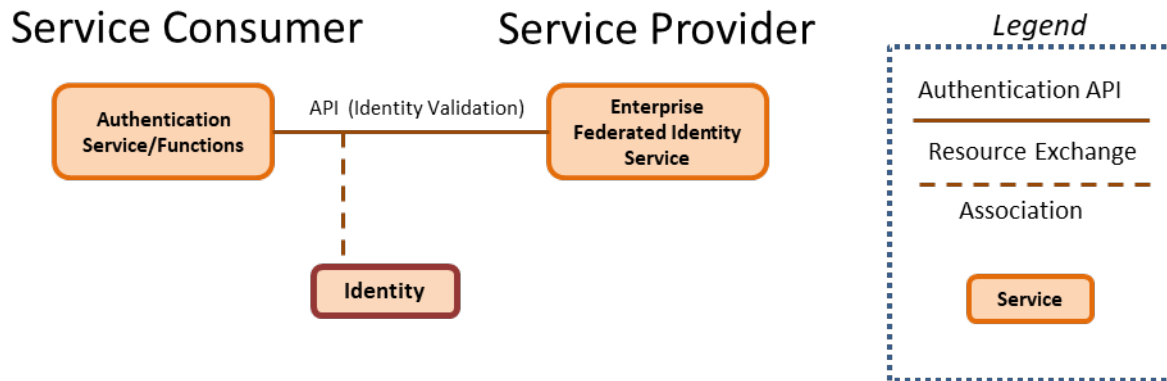


Figure 40 Enterprise Federated Identity Service (SvcV-2)

The Enterprise Federated Identity Service has been developed to distribute persona and personnel attributes for access control using a standards-based access control such as Security Assertion Markup Language (SAML). Using this connection approval process provides an individual's identity and attributes for the purpose of enabling Attribute Based Access Control (ABAC) to individuals who have justification and the need-to-know.

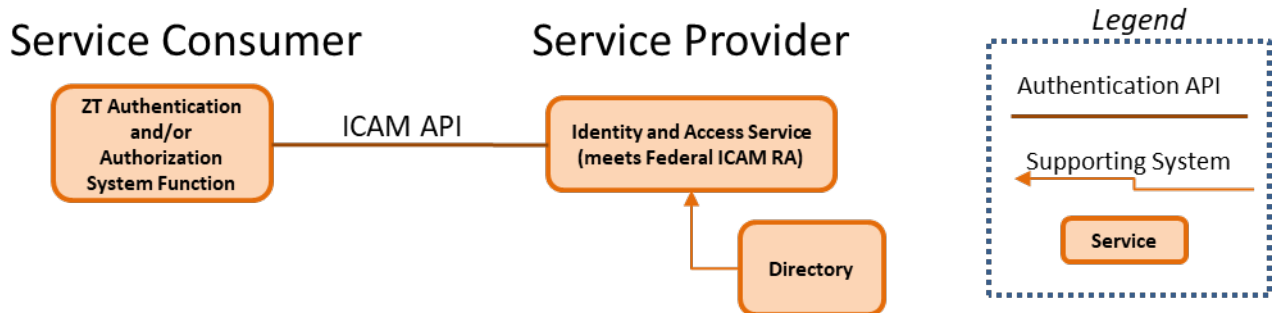


Figure 41 ICAM Service (SvcV-2)

Both Federated ID Service and ICAM have their own external architecture, so noted here as service interfaces used by ZT.

Required external services used by most ZT authentication and authorization components.

8 TRANSITION ARCHITECTURE PLANNING (FFP)

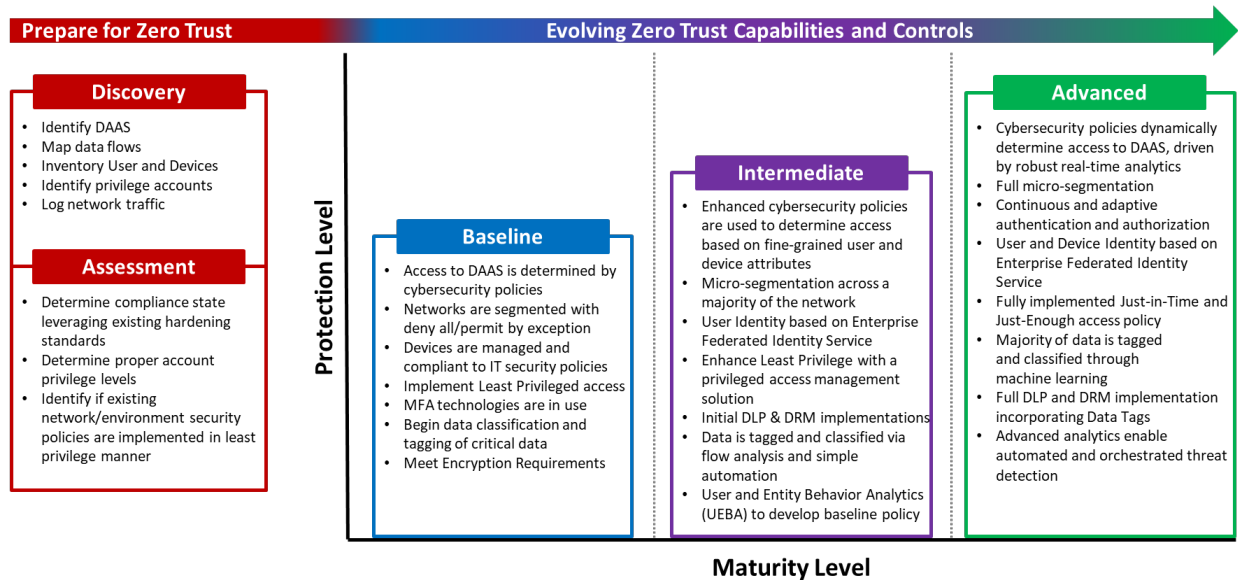


Figure 42 Maturity Model (FFP)

8.1 Maturity Model (FFP)

Zero Trust Maturity is the logical progression of an as-is security model to an advanced Zero Trust architecture. Although one’s Zero Trust journey will continue to mature beyond this FFP, this Maturity Model serves as a reference to help information system owners conceptualize their migration from their as-is to their to-be architecture. The approach to full Zero Trust implementation begins with preparatory discovery and assessment tasks. The initial discovery process will identify critical DAAS as well as access and authorization activity existing within the architecture. The tasks within the “Prepare for Zero Trust” section are critical, as the focus of Zero Trust is to protect DAAS. To do this, the relationships between workloads, networks, devices, and users must be discovered.

An advanced Zero Trust Architecture requires the implementation of security policies tied back to specific authorization attributes and the confidence level of the user and entity. Prerequisite assessment of the environment will determine the compliance state, privilege account levels and validate implementation of existing security controls. Advanced, as depicted in the FFP does not mean an end to maturing Zero Trust. Zero Trust will continue to be refined as AI and ML continue to refine security controls within the architecture as depicted in Figure 20.

8.2 Baseline (OV-1)

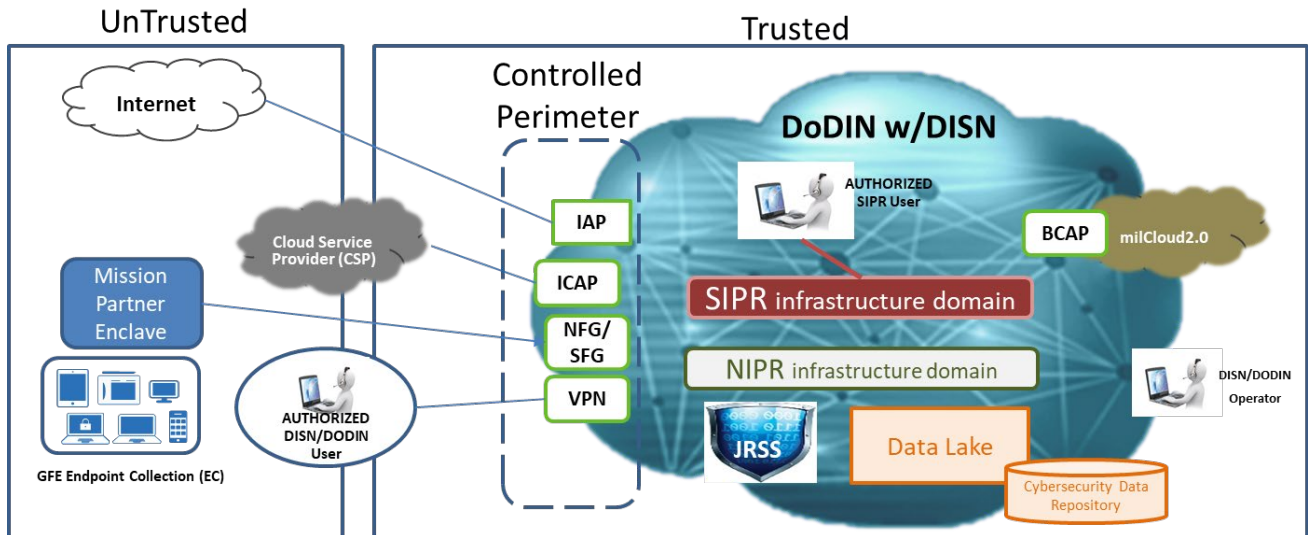


Figure 43 Transition Architecture Baseline (OV-1)

The baseline architecture implemented today within DoD commonly consists of untrusted and trusted zones with security elements siloed in endpoint capabilities, perimeter capabilities and mid-tier capabilities. The design of the architecture assumes threats only exist in the untrusted segment while communication flows are more widely allowed in the trusted segment.

The endpoint elements support numerous device types, however conditional access to data, applications, assets and services is not consistently implemented. Devices are protected with host-based security systems and connected to the DoDIN via VPNs as trusted and authorized entities.

The perimeter consists of internet access points (IAPs), cloud access points (CAPs) and federated gateways (FGs). These perimeter capabilities are implemented to keep adversaries out of the environment while also managing access to public or external resources.

Mid-tier security capabilities are based on the Joint Regional Security Stack (JRSS) which consists of firewall, intrusion detection, intrusion prevention and virtual routing. The complicated mid-tier design, while featuring significant security capabilities, is subject to performance challenges and is not well integrated with endpoint and perimeter capabilities.

Data has shifted from being solely located in the datacenter to cloud infrastructure and endpoints, which are not always managed, to complicate the protect surface. The iterative implementation of ZT architecture solves numerous challenges within today's cybersecurity architecture.

8.3 Transition (OV-1)

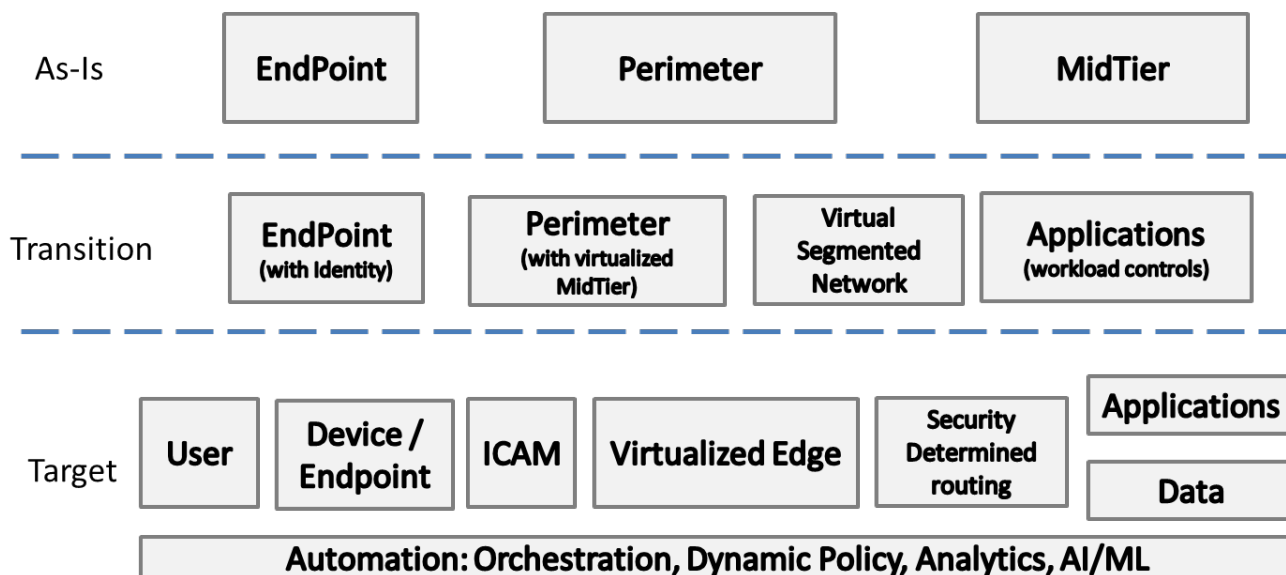


Figure 44 Transition Architecture Transition (OV-1)

The transition from the as-is baseline architecture to the target ZT end-state is best defined as a journey filled with capability advancement, integration activities, virtualization, automation, least privilege policy development and analytics development. Legacy implementations within the DoD will vary greatly based on organizations, purpose and age. They will each require a transition architecture specific to the changes required to reach that objective since no one size fits all approach can exist.

Transition focuses on modernizing endpoint security with integration of identity and attribute-based authentication and authorization. Users must be validated prior to network access along with compliance of the host which is being used to connect. The combination of identity, user attributes and device attributes enable a risk based conditional access decision which was not previously possible. The second element of the transition is integrating and virtualizing perimeter and mid-tier capabilities. Consolidation of these capabilities leads to streamlined capabilities and least privileged access policies. Less complication eases management burden, improves performance, enables integration with endpoint and application capabilities. The third element in transition is the adoption of application security capabilities to improve cybersecurity posture and limit lateral movement. Micro-segmentation implemented on the application host limits the machine-to-machine communication to only necessary ports, protocols, and processes.

The target architecture continues to evolve these operational capabilities with the adoption of new technology such as SDP, integration with enterprise ICAM services, data protections, and robust risk-based analytics. The integration of enterprise ICAM services ensures the same identity and attributes are leveraged throughout authorization decisions to provide better federation of ZT solutions. Adoption of SDP provides an additional abstraction layer to the DoDIN as inspection and segmentation are performed prior to user/device connecting the

July 2022

resource. SDP combined with data tagging and protections such as data loss prevention and data rights management provide strong protection against data exfiltration. Each of these modern security capabilities provides the visibility necessary to perform analytics required to develop the risk scores used in authorization decisions.

The transformation from the as-is architecture to the transition architecture and eventually to the final state architecture is an iterative journey to establish new capabilities, security policies and analytics. The maturity model further provides a view of capabilities and policies to implement throughout the transition process.

9 APPENDIX (AV-2)

Appendix provides context dependent ontology of semantic classification and meaning of the acronyms, terms and definitions of architecture elements used within the subject area. It enables a common understanding of terms and consistency of definitions used across the subject area. It includes acronyms, a taxonomy of terms, and definitions that are used in the Reference Architecture and relevant to solution architectures.

9.1 Systems

Name	Short Name	Description
Analytics & Confidence Scoring	N/A	This system analyzes event and incident logs via systematic analysis of data via statistics or other defined functional filters or computations to obtain confidence scores. These scores indicates the probability/percent value, within a specific range of error, with which the estimation of a statistical parameter for a given set of analytic data is determined to be true. Specifically in ZT, this represents the probability that a user/NPE is who they assert themselves to be.
Comply to Connect	C2C	Comply-to-Connect (C2C) is the identification, protection, and detection of DoDIN connected devices to ensure a continuous secure configuration. C2C enables the conduct of Defensive Cyber Operations in response to detected and prevailing threats by providing critical enabling information for the development of a Common Operating Picture. C2C standards are based on a framework of managing access to the network and its information resources by restricting or limiting access to those devices that do not comply with the standards.
Cybersecurity Domain Controller (w/ZT)	CSDC	The Cybersecurity Domain Controller administers (directs & controls) policy for all cybersecurity functions. It ensures coordination of policy implementation and consistency of policy application. For Zero Trust, the Cybersecurity Domain Controller acts either as a Policy Decision Point for Zero Trust functions or delegates that PDP function to a specific subdomain controller.
Cybersecurity Domain Orchestrator (w/ZT)	CSDO	The CDO administers provisioning and workflow for all subordinate cybersecurity controllers. The CDO ensures coordination and integration of all provisioning and policy. For Zero Trust, the CDO has functional provisioning for all Zero Trust functions. The CDO will accept intent policy requests, identify impacted component controllers and partition the customized policy to each controller. Intent policy requests include customer requests, automatic scale-up/scale-down of resources, responses to security incidents, response to new TI or vulnerability information, changes to user/group/application access parameters, new policies and changes to existing policies, etc.

Data Lake	N/A	<p>A data lake is a centralized repository that allows you to store all your structured and unstructured data at any scale. You can store your data as-is, without having to first structure the data, and run different types of analytics—from dashboards and visualizations to big data processing, real-time analytics, and machine learning to guide better decisions.</p> <p>A data lake is a centralized repository designed to store, process, and secure large amounts of structured, semistructured, and unstructured data. It can store data in its native format and process any variety of it, ignoring size limits.</p>
Domain Controller	DC	Domain Controller directs/programs the behavior of the domain resources using well defined interfaces. This system accepts commands from the Domain Orchestrator and coordinates policy and provisioning. The domain controller subsumes and extends the traditional functions of element managers.
Domain Orchestrator	DO	The automated coordination and management of IT systems within a specific functional or resource domain. Orchestrators arrange, sequence and automate implementation of tasks, rules and policies to coordinate logical and physical resources in order to meet a customer or on-demand request to create, modify or remove infrastructure resources.
Endpoint	N/A	Endpoint is a role given to any devices capable of initiating or terminating a session on a network. They are often described as end-user devices, such as mobile devices, laptops, and desktop PCs; although hardware such as servers in a data centers are also considered endpoints. Devices such as zero clients, virtualized systems, and infrastructure equipment (i.e. routers and switches) are considered endpoints.
Endpoint Agent	N/A	Client software installed on a network endpoint that communicates or is controlled by a centralized system.
Federated, Identity, Credential, and Access Management	FICAM	<p>FICAM is the Federal Government’s enterprise approach to design, plan, and execute common ICAM processes.</p> <p>The FICAM Architecture is a framework for an agency to use in ICAM program and solution roadmap planning. The FICAM Architecture focuses on enterprise identity processes, practices, policies, and information security disciplines.</p>
Global Orchestrator	GO	The Global Orchestrator manages the SDN control information for all the many different systems comprising the DoDIN. The GO is the software-based coordinator that plays the central role of creating the requested services by coordinating and orchestrating services among Domain Orchestrators (DOs) and Domain Controllers. The GO controls and manages the domain orchestrators and controllers through sets of Representational State Transfer (REST) Application Programming Interfaces (APIs).

Information Sharing System	ISS	The Information Sharing System component resides outside the protective boundary of the core DISN OSS components and provides human and machine information sharing services to the DISN Service / Product managers that do not sit within the DISN OSS enclave, the DISN customers, the DoD NetOps community, DISA business systems, and commercial service providers.
Micro Firewall		Micro Firewall are virtualized NGFWs within a microsegment used to isolate workloads.
Next Generation Firewall	NGFW	NGFW goes beyond ports, protocols, and IP addresses, providing standard policy-based protection, while including more advanced tools such as intrusion prevention systems, application filtering, uniform resource locator (URL) filtering, and geo-location blocking.
Non-Person Entity	NPE	An entity with a digital identity that acts in cyberspace but is not a human actor. This can include an autonomous service or application, hardware devices (e.g. IOTs), proxies, and software applications (e.g. Bots).
Operations Support System	OSS	The DISN OSS is All the systems that provide operations, administration, maintenance, and provisioning (OAM&P) management functions: Service Fulfillment Functions: Inventory and Configuration Management, Order Management, Network Activation and Service Provisioning, Service Assurance Functions: Alarm and Fault Management, Performance Management, Incident Management, and Release Management.
Person Entity	N/A	The role a human actor (i.e. User) performs when accessing IT assets with a specific identify.
Policy Administrator	N/A	This component is responsible for establishing and/or shutting down the communication path between a subject and a resource (via commands to relevant PEPs). It would generate any session-specific authentication and authentication token or credential used by a client to access an enterprise resource. It is closely tied to the PE and relies on its decision to ultimately allow or deny a session. If the session is authorized and the request authenticated, the PA configures the PEP to allow the session to start. If the session is denied (or a previous approval is countermanded), the PA signals to the PEP to shut down the connection. Some implementations may treat the PE and PA as a single service; here, it is divided into itstwo logical components. The PA communicates with the PEP when creating the communication path. This communication is done via the control plane.

Policy Decision Point	PDP	Mechanism that examines requests to access resources, and compares them to the policy that applies to all requests for accessing that resource to determine whether specific access should be granted to the particular requester who issued the request under consideration.
Policy Enforcemnt Point	PEP	This system is responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource. The PEP communicates with the PA to forward requests and/or receive policy updates from the PA. This is a single logical component in ZTA but may be broken into two different components: the client (e.g., agent on a laptop) and resource side (e.g., gateway component in front of resource that controls access) or a single portal component that acts as a gatekeeper for communication paths.
Policy Engine	N/A	This component is responsible for the ultimate decision to grant access to a resource for a given subject. The PE uses enterprise policy as well as input from external sources (e.g., CDM systems, threat intelligence services described below) as input to a trust algorithm (see Section 3.3 for more details) to grant, deny, or revoke access to the resource. The PE is paired with the policy administrator component. The policy engine makes and logs the decision (as approved, or denied), and the policy administrator executes the decision.
SDP Agent	N/A	A software agent that interacts with the SDP Controller and other agents to activate within the SDP, connect and authenticate to the SDP Controller, initiate connections to other agents, or accept connections from other agents.
SDP Broker/Controller	N/A	An appliance or process that secures access to isolated services by ensuring that users are authenticated and authorized, devices are validated, secure communications are established, and user and management traffic on a network remain separate.
SDP Gateway	N/A	An SDP gateway provides authorized users and devices with access to protected processes and services. The gateway can also enact monitoring, logging, and reporting on these connections.

Security Domain Orchestrator	N/A	The Security DO is the software-based coordinator that creates the requested security services by coordinating and orchestrating services among Security Domain Controller and legacy management. The Security Domain Orchestrator controls and manages the security domain controller and legacy management through sets of Representational State Transfer (REST) Application Programming Interfaces (APIs) and legacy management interfaces.
Security Incident and Event Manager	SIEM	The SIEM aggregates security and event data from across the environment.
Sensors	N/A	An intrusion detection and prevention system component that monitors and analyzes network activity and may also perform prevention actions. [https://csrc.nist.gov/glossary/term/sensor] A sensor collects information on devices throughout the network in order to determine 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.
Virtual Hosting Environment	N/A	A virtual hosting environment lets you run multiple guest operating systems on a single host computer at the same time. Host software virtualizes the following resources: CPU Memory Disk Network Local devices
Virtual Private Network	VPN	A restricted-use, logical (i.e., artificial or simulated) computer network that is constructed from the system resources of a relatively public, physical (i.e., real) network (such as the Internet), often by using encryption (located at hosts or gateways), and often by tunneling links of the virtual network across the real network.
VPN Gateway	N/A	VPN gateways provide secure connectivity between multiple sites, such as on-premises data centers, Virtual Private Cloud (VPC) networks, and VMware Engine private clouds. Traffic is encrypted because the VPN connections traverse the internet. Each VPN gateway can support multiple connections. When you create multiple connections to the same VPN gateway, all VPN tunnels share the available gateway bandwidth.

ZT Domain Controller	N/A	The Domain Controller covering the Zero Trust. This system accepts commands from the ZT Orchestrator and coordinates policy and provisioning. It acts as the element manager for Zero Trust devices both physical and virtual.
ZT Domain Orchestrator	N/A	The Orchestrator dynamically adjust policy based on administrator input and events within the network as identified by the Analytics Engine.

9.2 Services

Name	Short Name	Description
Automated Account Provisioning	AAP	User Account Provisioning service provides an automated identity management process that grants and manages access to applications, systems and data within an organization. Provisioning automation in ZT occurs through using policy to allocate privileges and permissions to users, protecting the security in the enterprise via least best fit user/NPE access.
Credential Vault	N/A	A credential vault is a repository that holds the assigned credentials (certificates, user IDs, tokens, and passwords) for accounts and resources (Users & NPE). As a service, the Credential Vault allows state determination and storage of a credential via API. Management access to the credential vault requires a privileged administrator.
Enterprise Identity Provider	N/A	A service which provides state/status determination and access to Identity and Credential information. It may also provide baseline user/NPE access roles.
Enterprise Identity Service	EIS	The Enterprise Identity Service provides data controlling for IdAM and ICAM to a client system requesting confirmation of identity credential and permissions to access functional services.

Federated Enterprise Identity Service	FEIS	The Federated Enterprise Identity Service aggregates identity credentials and authorizations and shares among a federated group of organizations so users/NPE can access services in other domains.
Federated Identity, Credential, and Access Management	FICAM	Federated ICAM provides for sharing of information on Identity, credentialization, and Access among agency/regional/organizational based systems. As a service FCIAM provides APIs for exchange of identity determination and access to services. FICAM has components that realize policies, standards, and APIs.
Identity Governance and Administration	IGA	Identity governance and administration system supports automated service provisioning of access certifications, access requests, password & token management following pre-established governance polies.
Identity, Credential, and Access Management	ICAM	The set of security disciplines that allows an organization to enable the right entity to access the right resource at the right time for the right reason. It is the tools, policies, and systems that allow an organization to manage, monitor, and secure access to protected resources. These resources may be electronic files, computer systems, or physical resources such as server rooms and buildings
Master User Record	MUR	The user master record contains the assignment of roles to the user where each role may be associated with corresponding authorizations for services and activities. Associated with the MUR are Master Device Record (MDR) which provides the same function for a non person entity (NPE) endpoint and Master System Record (MSR) covering a set of hardware and software.

9.3 General Terms

Name	Description
Access Management	Access Management is how an agency authenticates enterprise identities and authorizes appropriate access to protected services.
Capability	The ability to achieve a desired effect under specified standards and conditions through combinations of means and ways to perform a set of tasks.
Continuous	Occuring periodically without interruption during the ordinary performance of services.
Credential Management	Credential Management is how an agency issues, manages, and revokes credentials bound to enterprise identities.
Dynamic	Occuring in near-real-time under conditions then present.
Enterprise assets	Enterprise assets include end-user devices, network devices; non-computing/Internet of Things (IoT) devices; and servers) connected to the infrastructure physically, virtually, remotely, and those within cloud environments.
Identity Federation	Federation is the technology, policies, standards, and processes that allow an agency to accept digital identities, attributes, and credentials managed by other agencies.
Identity Management	Identity Management is how an agency collects, verifies, and manages attributes to establish and maintain enterprise identities for employees and contractors.
Just in Time	Using the current values of all indicators and analytics as input to a policy decision or enforcement.
Permission	Authorization to perform some action on a system. [https://csrc.nist.gov/glossary/term/permission]
Resource	Data, Information, Performers, Materiel, or Personnel Types that are produced or consumed.
Security Technical Implementation Guide (STIG)	Based on Department of Defense (DoD) policy and security controls. Implementation guide geared to a specific product and version. Contains all requirements that have been flagged as applicable for the product which have been selected on a DoD baseline. [https://csrc.nist.gov/glossary/term/security_technical_implementation_guide]
Signature (virus)	A recognizable, distinguishing pattern associated with an attack, such as a binary string in a virus or a particular set of keystrokes used to gain unauthorized access to a system. [https://csrc.nist.gov/glossary/term/signature]
Telemetry	Telemetry is the automated collection of measurements or other data at remote points and their automatic transmission to receiving equipment for monitoring. [https://www.afcea.org/content/next-generation-cybersecurity-telemetry-offers-promise]

Unified Profile for DoDAF/MODAF (UPDM)	The Unified Profile for DoDAF/MODAF (UPDM) is a visual modeling standard that supports the development of architectures that comply with the USA Department of Defense Architecture Framework (DoDAF) and the UK Ministry of Defence Architecture Framework (MODAF).
Workload	The virtualized environment that includes network, compute, storage, data, application services, and security services that provide a ability to perform a specific set of tasks.

9.4 DIV-1

Resource Exchanged	Description/Definition
Alerts	Data that indicates some trigger or threshold passing event has occurred and which is transmitted from the managed device/service to the managing service. A notification that a specific attack has been detected or directed at an organization's information systems. https://niccs.cisa.gov/about-niccs/cybersecurity-glossary
Analytics	Information resulting from the systematic analysis of data or statistics. This analysis includes discovering, interpreting, and communicating significant patterns in data.
Application Sensitivity	A relative measure of how sensitive a service with access to specific data is to the well being and security of the organization.
Behavior	Aggregate data from logs and reports that provides packet, flow, file and other types of information, as well as certain kinds of threat data to figure out whether certain kinds of activity and behavior are likely to constitute a cyberattack. https://www.techopedia.com/definition/32366/user-and-entity-behavior-analytics-ueba
Biometrics	A biometric is a measurable physical characteristic or personal behavioral trait used to recognize the identity, or verify the claimed identity, of an applicant. Facial images, fingerprints, and iris scan samples are all examples of biometrics. (FIPS 201)
Certificates	A certificate provides authentication of the identity claimed. Within the National Security System (NSS) public key infrastructure (PKI), identity certificates are used for both authentication and digital signatures. A set of data that uniquely identifies an entity, contains the entity's public key and possibly other information, and is digitally signed by a trusted party, thereby binding the public key to the entity identified in the certificate. Additional information in the certificate could specify how the key is used and the validity period of the certificate. https://csrc.nist.gov/glossary/term/certificate
Challenge	Additional or secondary question response from a user to confirm identity or further authenticate.

Confidence Scores	A confidence level indicates the probability/percent value, within a specific range of error, with which the estimation of a statistical parameter for a given set of analytic data is determined to be true. Specifically in ZT, this represents the probability that a user/NPE is who they assert themselves to be.
Configuration	The conditions, parameters, policy, and specifications with which an information system or system component is described in order to provide the services and behavior desired by a management application. https://csrc.nist.gov/glossary/term/configuration
Configuration Commands	Data by which a managing service provides a managed device/service a set of policy and threshold data changes.
Credential	An object or data structure that authoritatively binds an identity - via an identifier or identifiers - and (optionally) additional attributes, to at least one authenticator possessed and controlled by a subscriber. https://csrc.nist.gov/glossary/term/credential
Data (at rest)	Digital Information contained in a system medium as encoded bites.
Data (in transit)	Information being transmitted from one service point to another service point as a series of encoded spectrum communications.
Data Tag Controls	Policy on how to associate data tags, a form of association by label, with specific data information. Tags include security domain characteristics and classifications by use and by type. Tags will be utilized by workloads in determining access to information.
Data Tagging	The ability to associate a data object with characterizing metadata for a defined purpose.
Device Hygiene	Information on the state of compliance to policies, configurations and state of use of a device.
Dynamic Instrumentation	Dynamic Instrumentation or Dynamic Binary Instrumentation (DBI) is debug code or other injected program statements that enable techniques to execute code within a process to examine its internals. https://ieeexplore.ieee.org/document/9449226 https://blogs.blackberry.com/en/2021/04/malware-analysis-with-dynamic-binary-instrumentation-frameworks
Dynamic Remediations	Data which results in changes to policy and/or configuration of a device or service that brings the end state of the device/service closer to governance policy.
Geolocation	The data that identifies the geographical location of a person or NPE by means of digital information dynamically captured via network access, GPS, reckoning, or other analytical means.
Identity	The set of attribute values (i.e., characteristics) by which an entity is recognizable and that, within the scope of an identity manager's responsibility, is sufficient to distinguish that entity from any other entity. https://csrc.nist.gov/glossary/term/identity
Logs	Digital information that provided a history of events and states of a specific system or device.

Inventory	Returned or stored data that captures enterprise assets within a selected domain and comprising the physical, virtual, remote, and cloud infrastructure that needs to be monitored and protected within the enterprise. https://www.cisecurity.org/controls/inventory-and-control-of-enterprise-assets
Key	A parameter used in conjunction with a cryptographic algorithm that determines the specific operation of that algorithm. https://csrc.nist.gov/glossary/term/key
Privileged Access Management (PAM)	A class of solutions that help secure, control, manage and monitor privileged access to critical assets.
Patches and Updates	A patch, upgrade, or other modification to code that corrects security and/or functionality problems in software. https://csrc.nist.gov/glossary/term/update
PIN/RSA key	Additional information used as 'what you know' in two factor authentication.
PKI	The architecture, organization, techniques, practices, and procedures that collectively support the implementation and operation of a certificate-based public key cryptographic system. Framework established to issue, maintain, and revoke public key certificates. https://csrc.nist.gov/glossary/term/pki
Policy	Statements, rules or assertions that specify the correct or expected behavior of an entity. For example, an authorization policy might specify the correct access control rules for a software component. https://csrc.nist.gov/glossary/term/policy
Provisioning data	Information that describes what changes are needed in network and services to support a specific task, stakeholder, or access update.
Rule set	The capture of policy in a collection of Event/Condition/Action, or other forms of assertive statements, that can be interpreted by an algorithm.
Sensor Data	Information that describes the state and/or history of activity in a specific part of the infrastrucute or in a service. Usually includes logs, alerts, transactions, MIB contents and other captures of ongoing infrastructure and service activity.
Status information	The capture of data, usually as metrics, that characteristics the state of infrastructure or service at a specific instance or range of time.
Target State	Information that describes the intent of a change that is captured in characteristics of service and infrastructure provisioning.
Threat Intelligence	Threat information that has been aggregated, transformed, analyzed, interpreted, or enriched to provide the necessary context for decision-making processes. https://csrc.nist.gov/glossary/term/threat_intelligence
Token	Something that the Claimant possesses and controls (typically a key or password) that is used to authenticate the Claimant's identity. A portable, user-controlled, physical device (e.g., smart card or memory stick) used to

	store cryptographic information and possibly also perform cryptographic functions. https://csrc.nist.gov/glossary/term/token
Two-Factor	The second factor transmitted for meeting authentication that requires two or more factors to achieve authentication. Factors include: (i) something you know (e.g., password/personal identification number [PIN]); (ii) something you have (e.g., cryptographic identification device, token); or (iii) something you are (e.g., biometric). https://csrc.nist.gov/glossary/term/2fa

9.5 StdV-1-2 References

Document Name	Description	Reference URL
Zero Trust Architecture, Special Publication (NIST SP) - 800-207	Contains an abstract definition of Zero Trust architecture (ZTA) and gives general deployment models and use cases.	https://www.nist.gov
Guide to Attribute Based Access Control (ABAC) Definition and Considerations	Provides Federal agencies with a definition of attribute based access control (ABAC).	https://csrc.nist.gov/
Embracing a Zero Trust Security Model	Explains the Zero Trust security model and its benefits, as well as challenges for implementation.	https://media.defense.gov
Cloud Native Access Point (CNAP) Reference Design	Describes and defines the set of capabilities, fundamental components, and data flows within a CNAP. It presents logical design patterns and derived reference implementations for deploying, connecting to, and operating a CNAP.	https://software.af.mil
DoD Enterprise Identity, Credential, and Access Management (ICAM) Reference Design	Provides a high-level description of ICAM from a capability perspective, including transformational goals for ICAM in accordance with the Department of Defense (DoD) Digital Modernization Strategy	https://dodcio.defense.gov
DoD Enterprise DevSecOps Reference Design	Describes the DevSecOps lifecycle, supporting Pillars, and DevSecOps ecosystem; lists the tools and activities for DevSecOps software factory and ecosystem.	https://dodcio.defense.gov
SDP Specification v1.0	Specifies the base architecture for Software Defined Perimeter (SDP)-compliant systems.	https://cloudsecurityallian
Cloud Security Technical Reference Architecture	Illustrates recommended approaches to cloud migration and data protection, as outlined in Section 3(c)(ii) of Executive Order 14028.	https://www.cisa.gov
Federal ICAM Architecture	Describes the basics of ICAM, the FICAM Architecture, and how to use this information to facilitate enterprise ICAM practices.	https://playbooks.idmanagement.gov/arch/
Zero Trust Maturity Model	This document is designed to be a stopgap solution to support Federal Civilian Executive Branch (FCEB) agencies in designing their Zero Trust architecture (ZTA) implementation plans in accordance with Section 3,b,ii of Executive Order 14028, “Improving the Nation’s Cybersecurity” .	https://www.cisa.gov

Executive Order on Improving the Nation's Cybersecurity		https://www.whitehouse.gov/
Zero Trust Cybersecurity Current Trends	Assess the maturity of ZT technologies, their readiness and suitability for use in government, and the issues agencies would face if they chose to pursue ZT.	https://www.actiac.org

9.6 Capability Table

ID	Name	Short Name	Operational Definition/Discription
1	Continuous Authentication	N/A	The ability validate network users are the ones who they claim to be throughout an entire session at every step.
1.1	Continuous Multifactor Authentication	N/A	The ability to conduct authentication using two or more different factors to achieve authentication. Factors include: something you know (e.g., password/PIN); something you have (e.g., cryptographic identification device, token); or something you are (e.g., biometric), something you do. Continuous means just-in-time authentication (just -in time usually refers to authorization).
1.2	Behavioral Biometrics	N/A	Observing activities of users, information systems, and processes and measuring the activities against organizational policies and rule, baselines of normal activity, thresholds, and trends.
2	Conditional Authorization (Users, NPEs, M2M)	N/A	The ability to grant authorization to a resource contingent upon the continued trustworthiness of the supplicant. This trustworthiness can affect by the device hygiene, user and entity behavior, and other factors.
2.1	Attribute-Based Access Control (ABAC)	ABAC	An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environment conditions, and a set of policies that are specified in terms of those attributes and conditions.
2.2	Device Hygiene	N/A	The ability to determine the compliance status of managed and unmanaged assets.
2.2.1	Continuous, automated, Inventory & Telemetry	N/A	The ability to locate and identify devices connected to an environment, detect their removal/addition, to accurately know the totality of assets that need to be monitored and protected within the enterprise, and to obtain information about them. [CIS] Also support identifying unauthorized and unmanaged assets to remove or remediate.

2.2.2	Status Scans & Dynamic Instrumentation	N/A	The ability to poll devices for status, state, and configuration via remote management function or installation of agents/code on the device by management.
2.2.3	Dynamic Device Service Updates	N/A	The ability to remotely install new configurations and services on a device in order to bring the device into conformity or compliance with existing policy.
2.3	Just in Time Authorization	N/A	Just in Time Authorization allows a timed expiration of group membership. In practice, this allows administrative rights to be given at the time of need for as long as an action or duty needs them. As a result, access to administrative privileges becomes limited and abuse must be timed for when those privileges are given.
2.4	Privileged Access Management	PAM	Privileged Access Management (PAM) refers to a class of solutions that help secure, control, manage and monitor privileged access to critical assets.
3	ZT enabling Infrastructure	N/A	Infrastructure capabilities that enable ZT
3.1	Macro-segmentation	N/A	Similar in concept to physical network segmentation, macro-segmentation can be achieved through the application of additional hardware or VLANs.
3.2	Micro-segmentation	N/A	Micro-segmentation is the practice of dividing (isolating) the network into small logical segments by enabling granular access control, whereby users, applications, workloads and devices are segmented based on logical, not physical, attributes. This also provides an advantage over traditional perimeter security, as the smaller segments present a reduced attack surface (for malicious actors). In a ZT Architecture, security settings can be applied to different types of traffic, creating policies that limit network and application flows between workloads to those that are explicitly permitted.
3.2.1	Workload Definition	N/A	The ability to define the objectives, compute requirements, and communication pathways required for a specific application workload.
3.2.2	Workload Isolation	N/A	The ability to segment out an application workload so as to only allow the required connections be made between processes, network traffic, and api calls. As a subset of micro-segmentation the capability is limiting east west traffic preventing lateral movement.

3.3	Software-defined perimeter (SDP)	SDP	The ability to control access to resources based on identity and a need-to-know model in which device state and identity are verified before access to application infrastructure is granted.
4	Securing Application & Workload	N/A	The ability to secure and manage the application layer as well as compute containers and virtual machines. The ability to identify and control the technology stack to facilitate more granular and accurate access decisions.
4.1	API and Process Micro Segmentation	N/A	The ability to allow or block communication of API calls and process to process communication on both remote and local systems.
4.2	Securing Software Supply Chain	N/A	The ability to prevent or arrest software supply chain attacks, which occur "when a cyber threat actor infiltrates a software vendor's network and employs malicious code to compromise the software before the vendor sends it to their customers."
4.2.1	DevSecOps	N/A	A process capability that improves the lead time and frequency of delivery outcomes through enhanced engineering practices; promoting a more cohesive collaboration between Development, Security, and Operations teams as they work towards continuous integration and delivery.
4.2.2	API Standardization	N/A	The ability to reach agreement and publish, locally, the application programming interface for a commonly used service. Enforcement of compliance in the use of commonly agreed API's.
4.3	Application Proxies	N/A	An application proxy or application proxy server receives requests intended for another server and acts as the proxy of the client to obtain the requested service.
4.4	Risk-adaptive Application Access	N/A	In Risk-adaptive Application Access, access privileges are granted based on a combination of a user's identity, mission need, and the level of security risk that exists between the system being accessed and a user. RAdAC will use security metrics, such as the strength of the authentication method, the level of assurance of the session connection between the system and a user, and the physical location of a user, to make its risk determination.
5	Securing Data	N/A	Processes and technical controls to identify, classify, securely handle, retain, and dispose of data.
5.1	Encryption	N/A	A procedure used in cryptography to convert plaintext into ciphertext to prevent anyone but the intended recipient from reading that data.

5.1.1	Encryption In Transit	N/A	The ability to protect data if communications are intercepted while data moves between sites or services. This protection is achieved by encrypting the data before transmission; authenticating the endpoints; and decrypting and verifying the data on arrival.
5.1.2	Encryption At Rest	N/A	The ability to protect data from a system compromise or data exfiltration by encrypting data while stored.
5.2	Dynamic Policy enforcement	N/A	The ability to adapt policy and configurations, and enforce that change, in near real time based on environmental circumstances and indications of user and network behavior.
5.3	Data Rights Management (DRM)	DRM	DRM is a set of access control technologies and policies that proactively detect and protect access to data and proprietary hardware and prevent unauthorized modification or redistribution of protected data.
5.4	Data Loss Prevention (DLP)	DLP	The ability to detect and prevent the unauthorized use and transmission of information.
5.5	Dynamic Data Masking	DDM	The ability to provide a column-level security feature that uses masking policies to selectively mask tables and columns at query time.
5.6	Data Discovery & Classification	N/A	The ability to discover, classify, label, and report upon the sensitive data in your databases.
5.6.1	Data Tagging	N/A	The ability to associate a data object with characterizing metadata for a defined purpose.
6	Analytics	N/A	The ability to systematically apply statistical and /or logical techniques to describe and illustrate, condense, and recap, and evaluate data.

6.1	Data Visualization	N/A	The ability to represent information graphically, highlighting patterns and trends in data and helping the reader to achieve quick insights.
6.2	Security Information and Event Management (SIEM)	SIEM	The ability to centrally collect event and incident alerts across disparate sources, analyze them, and provide reports, situational awareness, and notifications. It is frequently used in support of incidence response, compliance, and reporting.
6.3	Big Data	N/A	The ability to enable enhanced insight, decision making, and process automation by consuming high-volume, high-velocity and/or high-variety information assets.
6.4	Sensors & Telemetry	N/A	The ability to collect status, state, and configuration of a service or device via the use of active or passive probes or other analytic activities on the device.
6.5	Continuous Monitoring	N/A	The ability to determine if the complete set of planned, required, and deployed security controls within an information system or inherited by the system continue to be effective over time in light of the inevitable changes that occur.

6.6	Machine Learning	ML	The ability to apply machine learning algorithms composed of many technologies (such as deep learning, neural networks and natural language processing), in unsupervised and supervised learning, that operate guided by lessons from existing information.
6.7	Entity and Activity Auditing	N/A	The ability to conduct "Monitoring for anomalous or suspicious behavior ...with signatures, statistical analysis, analytics or machine learning on user activity events. The analysis seeks to find patterns amongst data generated by user activity.
7	ZT Governance	N/A	A set of processes that ensures that ZT assets are formally managed throughout the enterprise. A ZT governance model establishes authority and management and decision making parameters related to ZT policies produced or managed by the enterprise.
8	ZT Orchestration	N/A	The ability to coordinate and automate disparate Zero Trust services, systems, and activities as part of of Cybersecurity Domain Orchestrator.
8.1	Automation	N/A	The ability to create and apply application technology to monitor and control the production and delivery of otherwise manual services.
8.1.1	Artificial Intelligence	AI	The capability of computer processes to perform functions that are normally associated with human intelligence such as reasoning, learning, and self-improvement.
8.1.2	Robotic Process Automation	RPA	The ability to use software tools that partially or fully automate human activities that are manual, rule-based, and repetitive.
8.1.3	Policy Administrator	N/A	A component with the ability to establish and/or shut down the communication path between a subject and a resource (via commands to relevant Policy Enforcement Points). The ability to direct Policy Enforcement Points to grant or deny access to resources based on policies created by the policy engine.
8.2	ZT Policy Engine	N/A	The ability for a component responsible for the ultimate decision to grant access to a resource for a given subject.

8.3	ZT Policy Administration	N/A	The ability to coordinate and enforce policy created by the ZT policy engine by translating it to settings and configurations at designated policy enforcement points (PEPs).
8.4	Software-Defined Enterprise	SDE	The ability to create a virtualized layer over physical infrastructure, and centrally manage it in an automated manner, utilizing a policy-based access control to dynamically create, configure, provision, and decommission virtualized network functions, system functions, security functions, and workflows.
8.4.1	Domain Orchestration	DO	The ability to coordinate services and operations, for a specific domain, across multiple types of devices and systems.
8.4.2	Domain Control	DC	The ability direct or command elements and associated systems to perform specific actions within a specified domain.
8.4.3	Software Defined Networking	SDN	The ability to separate the control and data planes and centrally manage and control the elements in the data plane.
8.4.5	Software-defined Wide-area Network	SD-WAN	The ability to virtualizethe enterprise connection of local networks into a wide-area network through the use of central routing, management, control & configuration of virtualized, distributed network and security services.
8.4.6	Network Function Virtualization / Virtual Security Function	NFV/VSF	The ability to decouple network functions (VNF) and security functions (VSF) from hardware appliances and deliver those functions as software in virtual machines.
8.5	Data Governance	N/A	A set of processes that ensures that data assets are formally managed throughout the enterprise. A data governance model establishes authority and management and decision making parameters related to the data produced or managed by the enterprise.
8.6	Risk Management Framework	RMF	Provides a comprehensive, flexible, repeatable, and measurable process that any organization can use to manage information security and privacy risk for organizations and systems and links to a suite of NIST standards and guidelines to support implementation of risk management programs to meet the requirements of the Federal Information Security Modernization Act (FISMA).

July 2022

10 REFERENCES

1. Cybersecurity Reference Architecture (CS RA) Version 4.0, July 2016 *This version of the CS RA is located on SIPRNet; link will be provided once classification of this document is completed.*
2. Identity, Credential, and Access Management (ICAM) Reference Design Version 1.0, June 2020, <https://dodcio.defense.gov/library>
3. DoD Digital Modernization Strategy, *DoD Information Resource Management Strategic Plan FY19-23*, July 2019, <https://dodcio.defense.gov/library>
4. 2018 DoD Cyber Strategy
https://dodcio.defense.gov/Portals/0/Documents/Factsheet_for_Strategy_and_CPR_FIN_AL.pdf
5. DoD IT Standards Registry (DISR) 20-2, October 2020,
gtg.csd.disa.mil/disc/dashboard.html
6. NIST SP 800-207, Zero Trust Architecture, August 2020,
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>
7. NIST SP 800-63-3, Digital Identity Guidelines, June 2017,
<https://doi.org/10.6028/NIST.SP.800-63-3>
8. DoD Architecture Framework, DoD Deputy Chief Information Officer, version 2.02, August 2010, <https://dodcio.defense.gov/library/DoD-Architecture-Framework/>
9. DoD Zero Trust Strategy, DoD CIO, July 2022, <https://dodcio.defense.gov/library>
10. Department of Defense Artificial Intelligence Strategy, DoD CIO, 2018,
<https://dodcio.defense.gov/>
11. Forrester, *developing a Framework to Improve Critical Infrastructure Cybersecurity*, April 2013, NIST RFI# 130208119-3119-01, <https://nist.gov/>
12. American Council for Technology – Industry Advisory Council (ACT-IAC), *Zero Trust Cybersecurity Current Trends*, April 2019, <https://actiac.org/zero-trust-cybersecurity-current-trends>
13. Improving the Nation's Cybersecurity: NIST's Responsibilities under the Executive Order 14028, <https://www.nist.gov/itl/executive-order-improving-nations-cybersecurity>