Carnegie Mellon University
Software Engineering Institute

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

MAY 31, 2023

Matthew Nicolai
Trista Polaski
CERT Situational Awareness

Document Markings

Agenda

- Introduction
- Area 1: Agree on a Generally Accepted Set of Basic ZT Definitions
- Area 2: Establish a Common View of ZT
- Area 3: Establish Standard ZT Maturity Levels
- Area 4: Explain How to Progress Through ZT Maturity Levels
- Area 5: Ensure ZT Supports Distributed Architectures
- Area 6: Establish ZT Thresholds to Block Threats
- Area 7: Integrate ZT and DevSecOps
- Area 8: Set Business Expectations for ZT Adoption
- Conclusion

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Introduction

Carnegie Mellon University
Software Engineering Institute

Introduction

- On March 1, 2023, the Biden Administration released the National Cybersecurity Strategy
- As part of the strategy, the Biden administration committed to improving federal cybersecurity through the implementation of a Zero Trust Architecture (ZTA) strategy and the modernization of information technology (IT) and operational technology (OT) infrastructure

What is Zero Trust?

- According to NIST SP 800-207, Zero Trust is an “evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources”
- Zero Trust Architecture (ZTA) “uses zero trust 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)”*

What is Zero Trust? Cont.

- Under ZTA, “authentication and authorization (both subject and device) are discrete functions performed before a session to an enterprise resource is established”
- ZT largely emerged in response to “enterprise network trends that include remote users, bring your own device (BYOD), and cloud-based assets that are not located within an enterprise-owned network boundary”
- Zero trust emphasizes the protection of resources, such as assets, services, workflows, network accounts, etc. rather than the protection of network segments

Overview

Why apply a zero trust strategy for cybersecurity?

Zero trust is a security model that John Kindervag and his team from Forrester Research, Inc. developed in 2009

Goals:

- Remove implicit trust (*Zero trust* is the associated buzzword)
- Move security from the network to users, applications, and workloads

Food for Thought:

- The zero trust strategy applies to personnel and physical security
- The Department of Defense (DoD) has applied zero trust to these areas for years

ZT Principles

- Ensure all resources are accessed securely, regardless of location
- Adopt a least privilege strategy and strictly enforce access control
- Inspect and log traffic necessary to support continuous auditing
- Ensure all components support application programming interfaces (APIs) for event and data exchange
- Automate actions across environments and systems driven by context and events

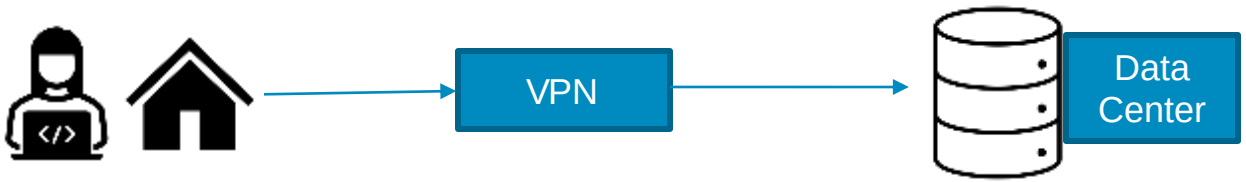

[Garbis 2021]

Working Definitions

A zero trust system employs an *integrated security solution* that uses *contextual information* from (1) identity, security, and IT infrastructure and (2) risk analytics tools to inform and enable the *dynamic enforcement of security policies uniformly across the enterprise* [Garbis 2021].

Physical security analogy

Zero trust shifts security from an ineffective perimeter-centric model to a *resource- and identity-centric model*. As a result, organizations can continuously adapt access controls to a changing environment, resulting in improved security, reduced risk, simpler and more resilient operations, and increased business agility [Garbis 2021].

NIST Zero Trust Architecture Components

The *control plane* is the part of the network that controls whether or not data is forwarded. The data plane is the actual forwarding process.

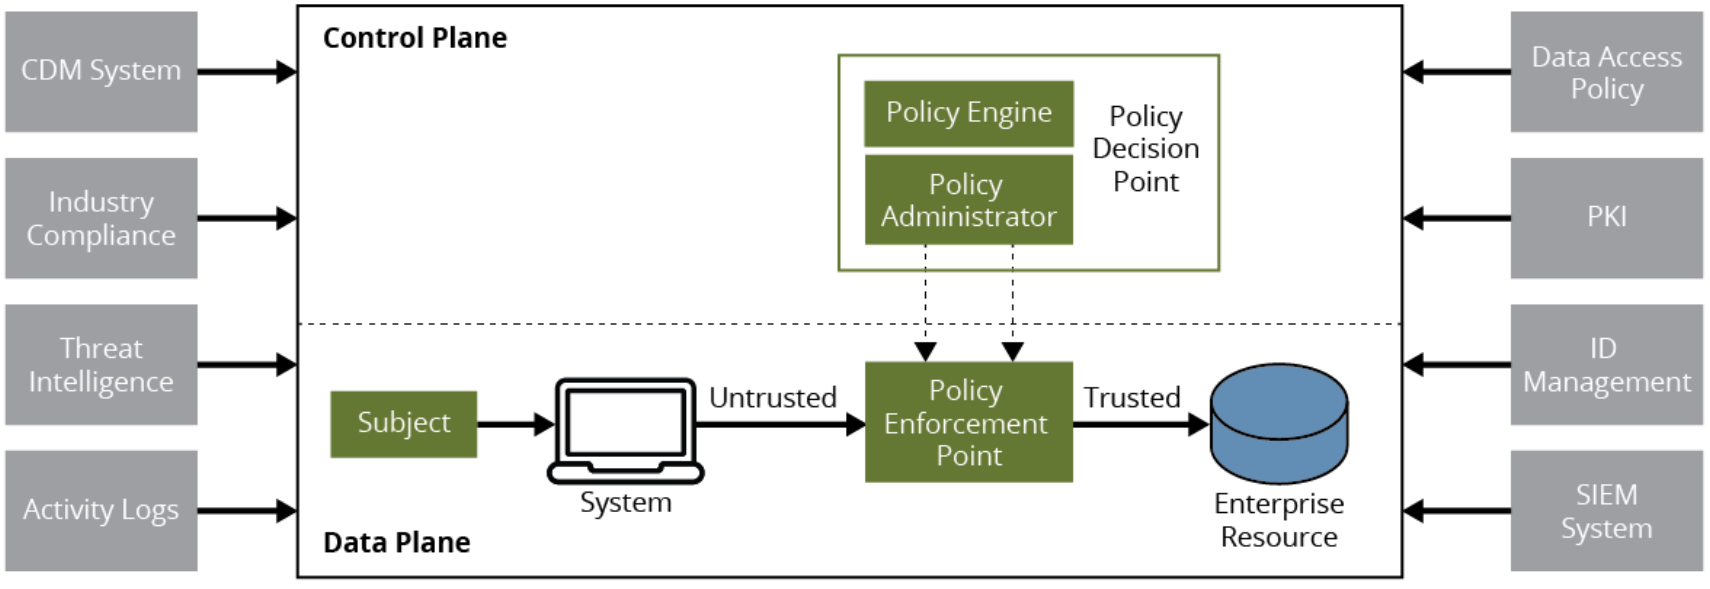

A *subject* can be a person, a non-person entity (NPE), or a machine-to-machine communication channel.

NPEs include software-based agents.

Boxes are functions, not products.

[Rose 2020]

Policy Enforcement Point Types

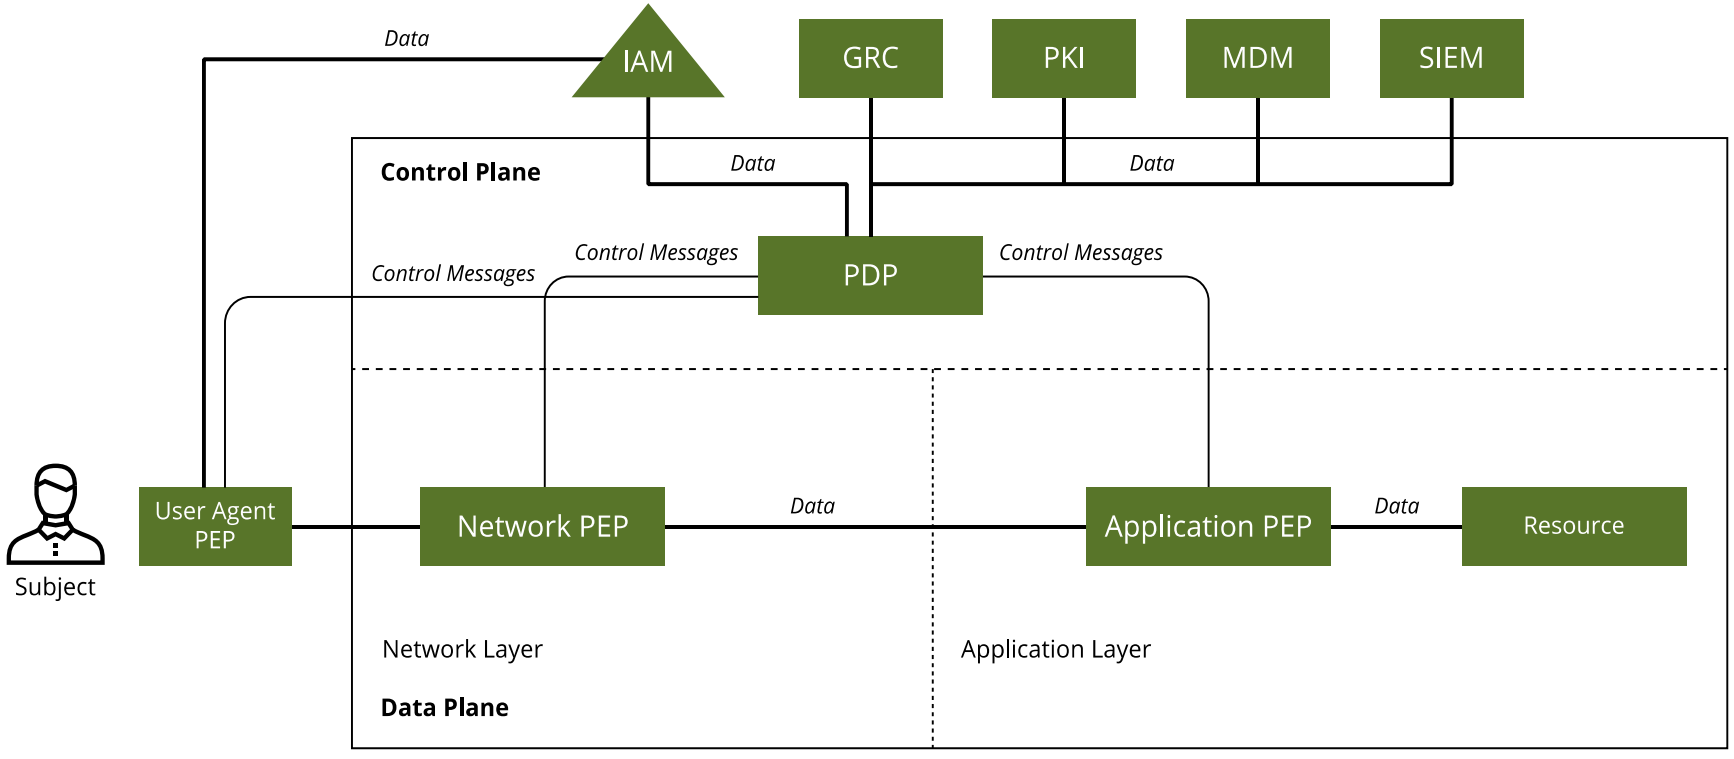

[Garbis 2021]

Resource-Based Deployment Model

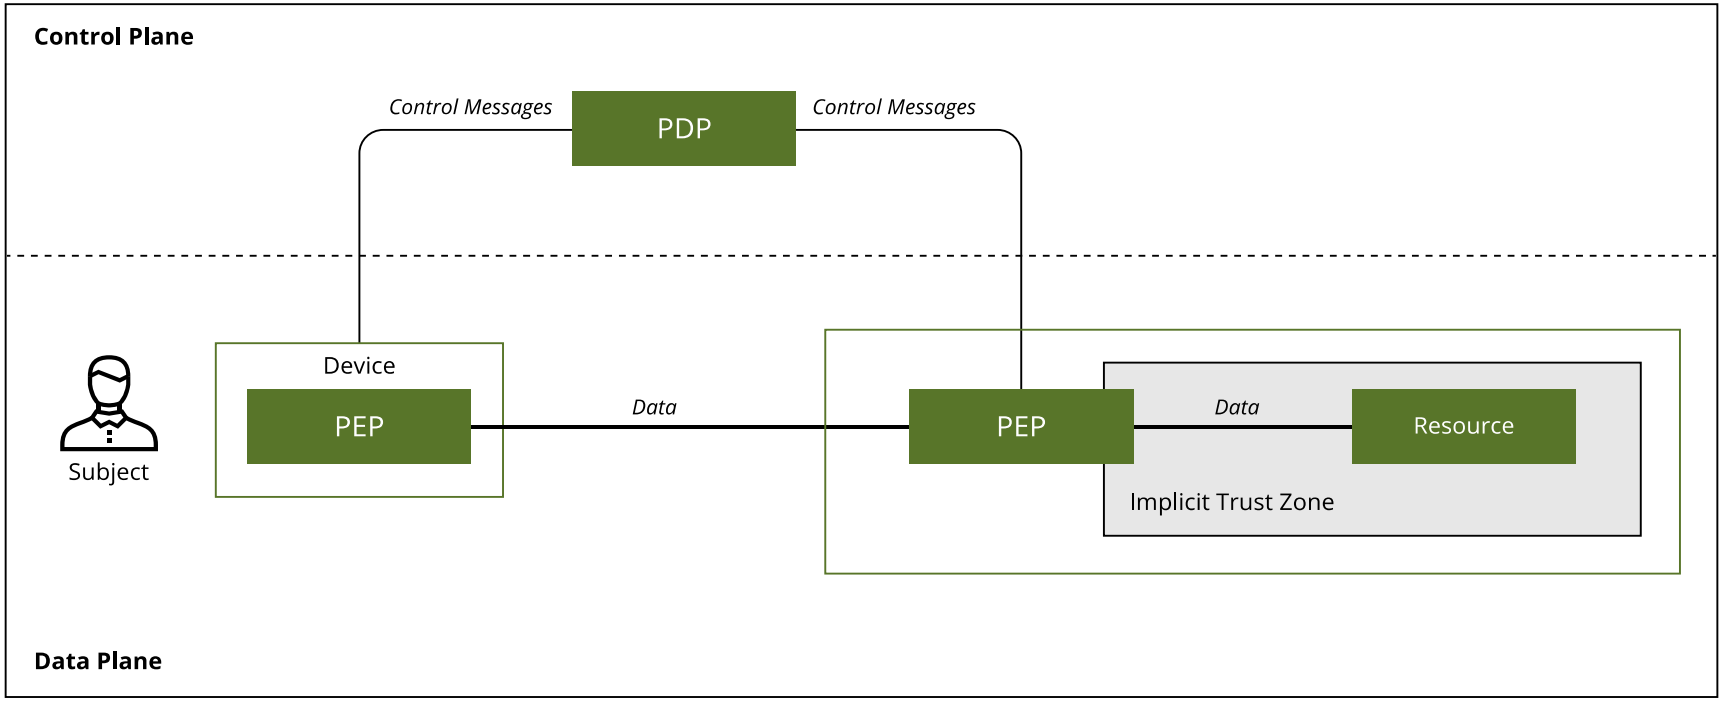

[Garbis 2021]

Enclave-Based Deployment Model

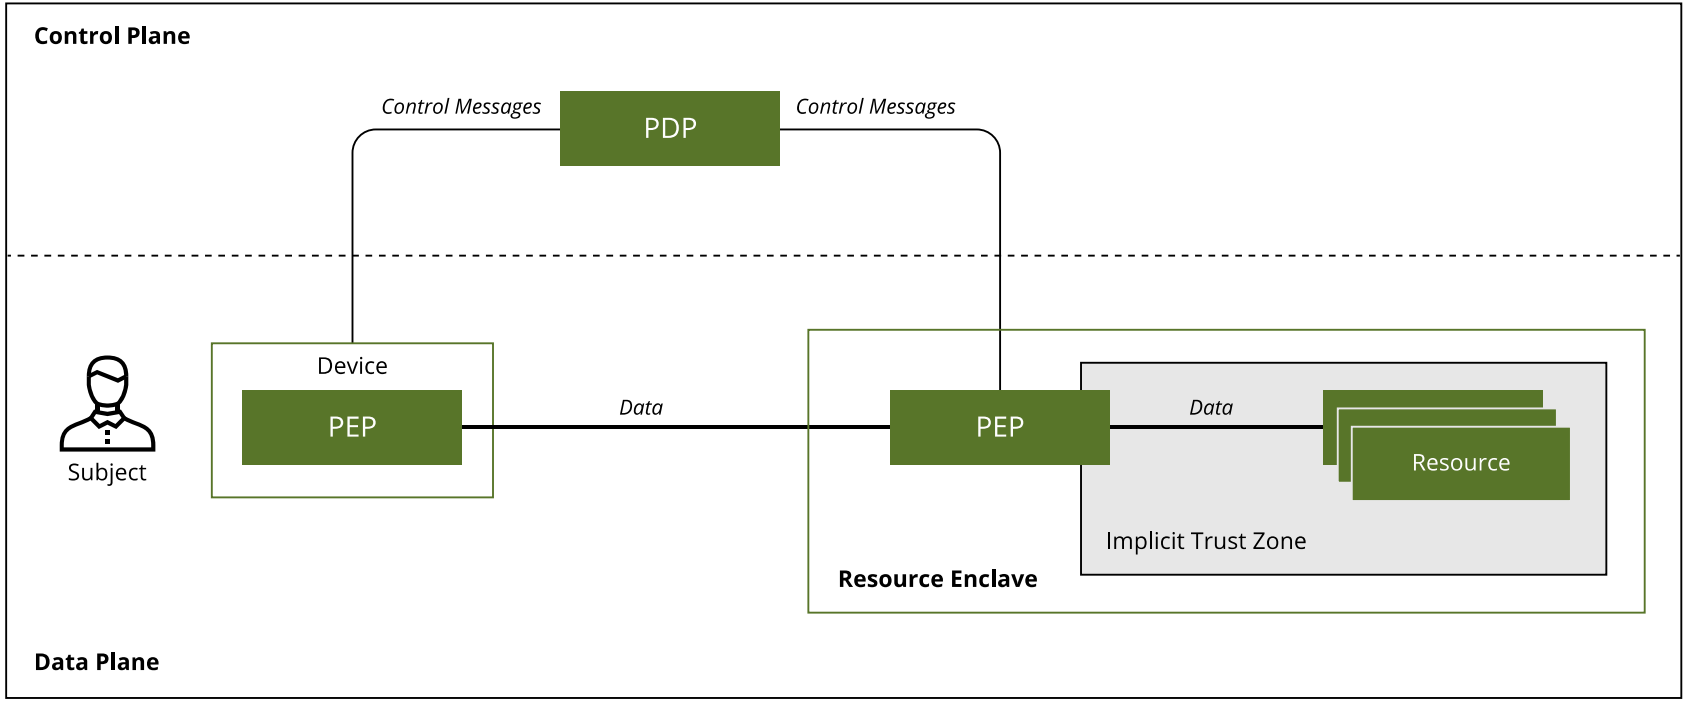

Garbis 2021]

Cloud-Routed Deployment Model

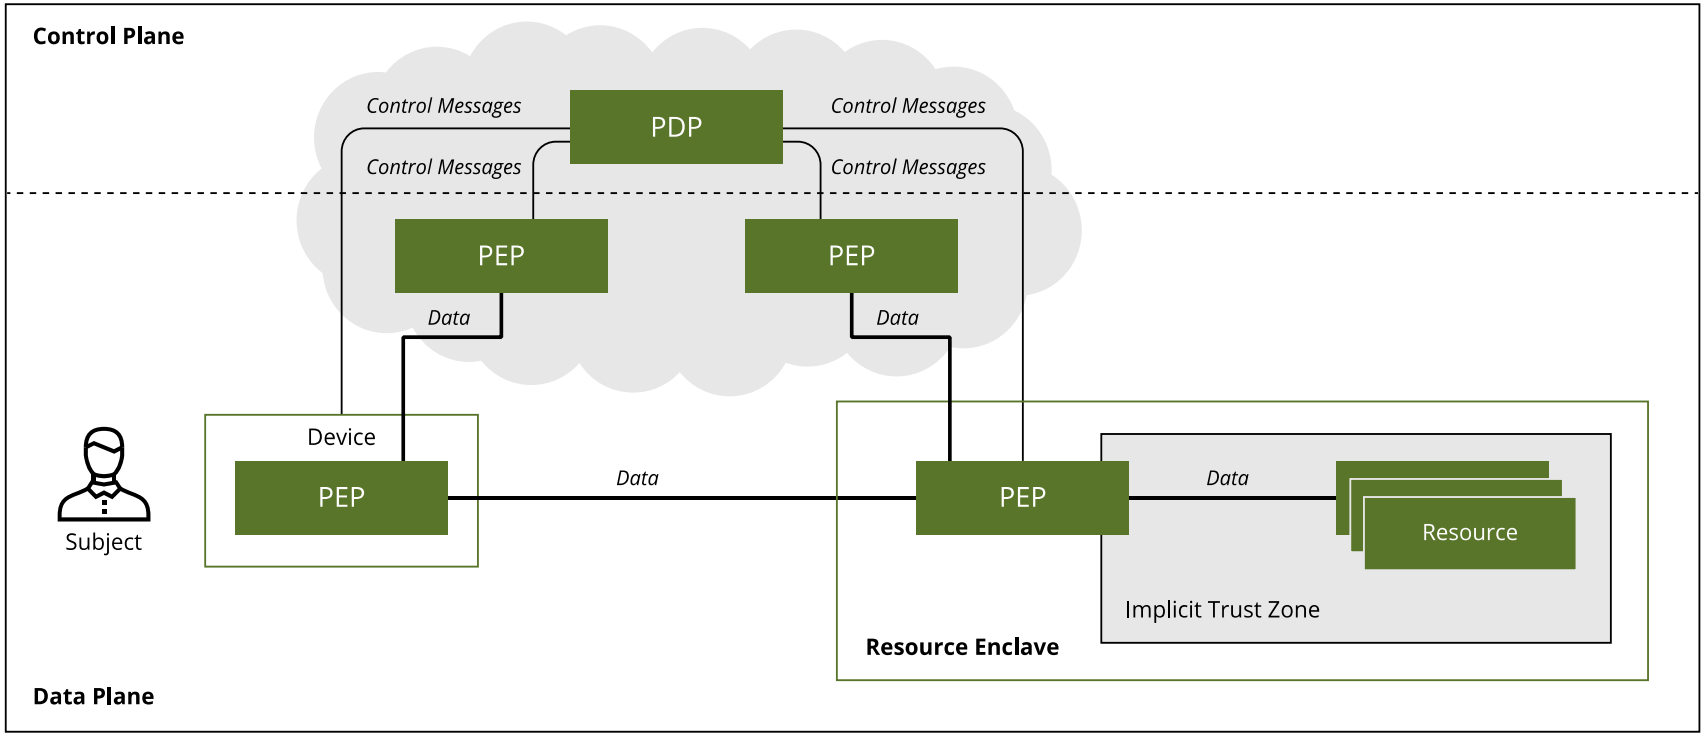

[Garbis 2021]

Microsegmentation Deployment Model

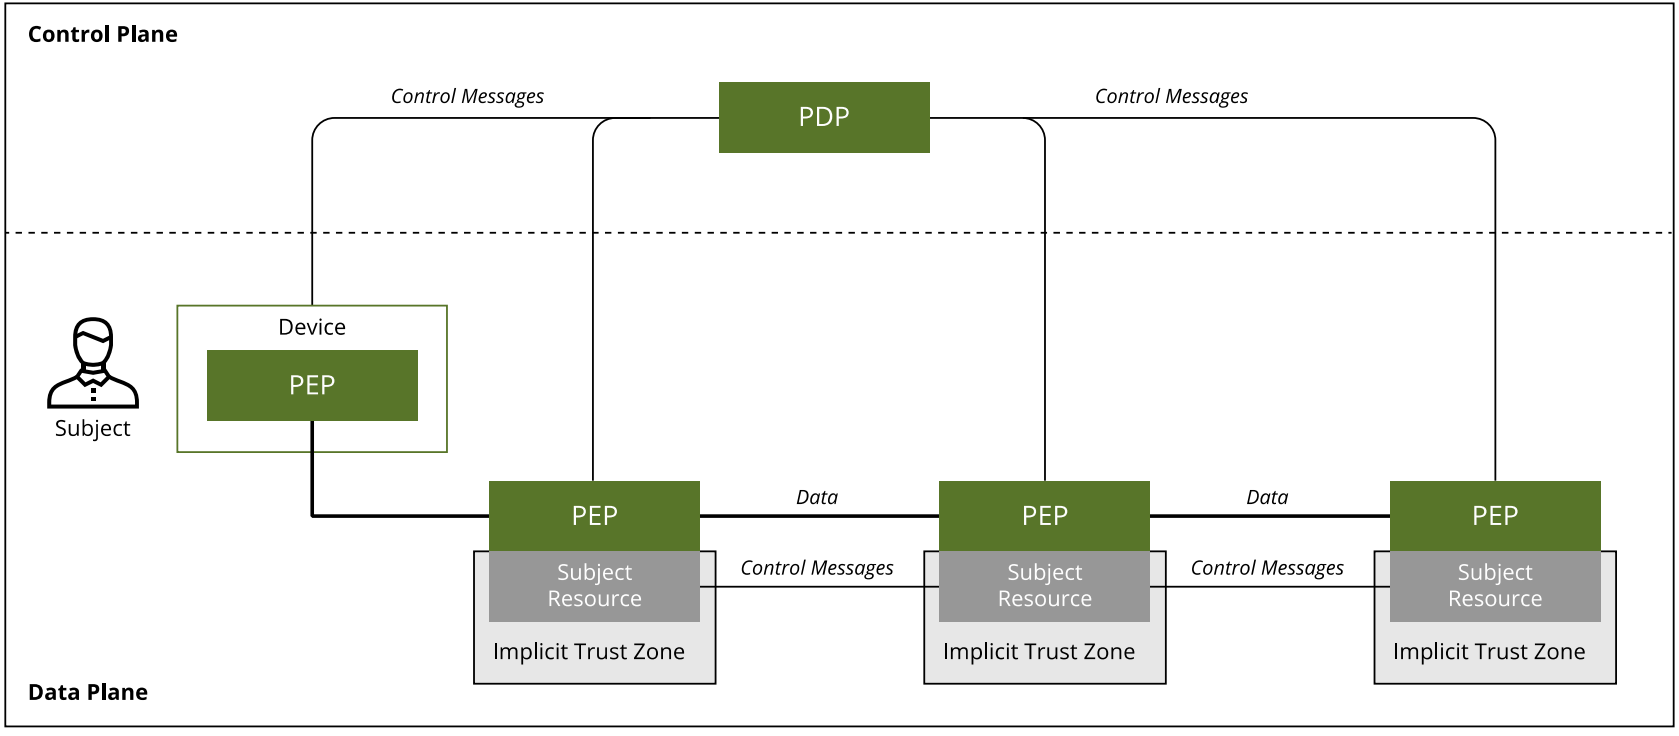

[Garbiis 2021]

CISA Zero Trust Maturity Model, v2

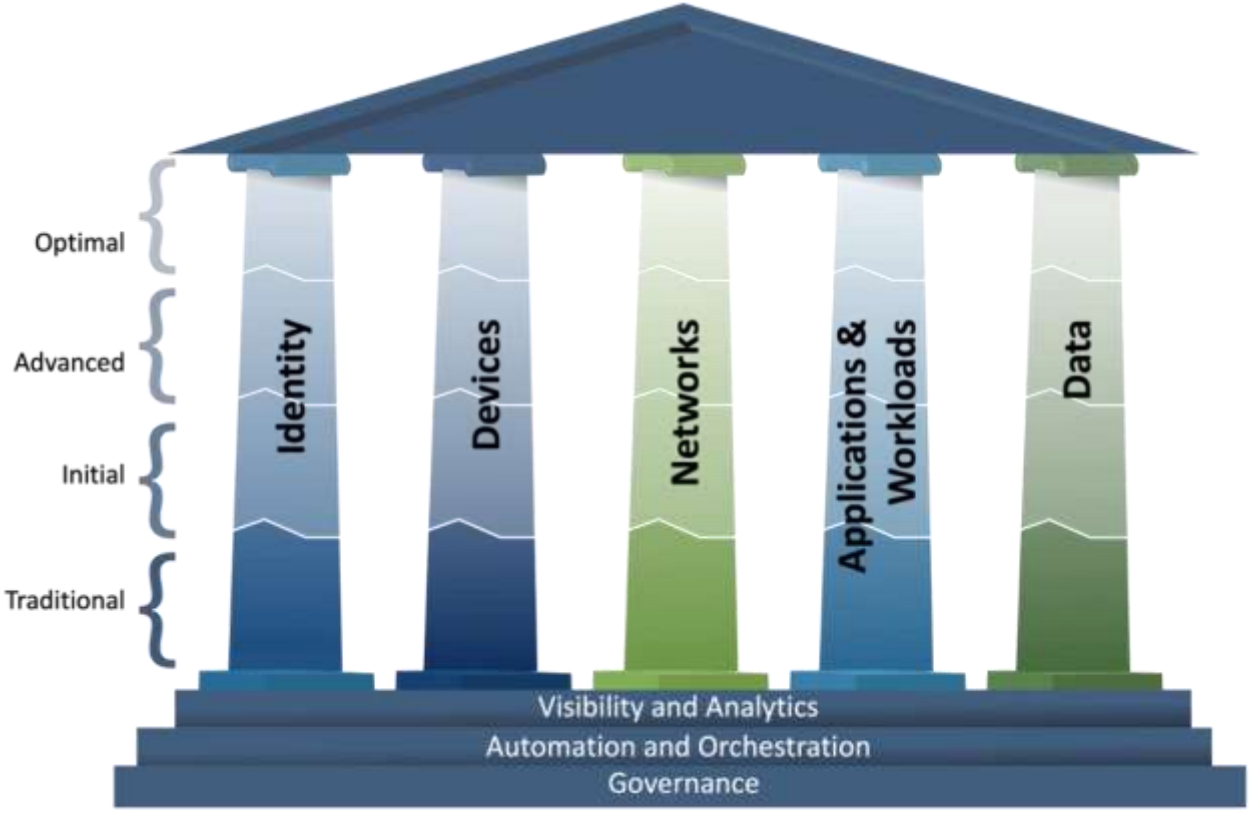

[CISA 2023]

CERT and Zero Trust

- In 2022, we hosted Zero Trust Industry Days, which featured keynote addresses, presentations from Zero Trust vendors, Q&A sessions, and expert panel discussions
- During these discussions, participants identified ZT-related issues that could benefit from additional research
- By focusing on these areas, organizations can develop solutions that streamline and accelerate ongoing ZTA transformation efforts
- In the following sections, we will discuss eight areas of potential research that will likely benefit from further exploration through public-private partnerships

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 1: Agree on a Generally Accepted Set of Basic ZT Definitions

Carnegie Mellon University
Software Engineering Institute

Agree on a Generally Accepted Set of Basic ZT Definitions

Carnegie Mellon University Software Engineering Institute

- According to NIST SP 800-207, ZT access decisions are made on a per-session basis
- However, there are several definitions of the term “session”
- Panelists at the Zero Trust Industry Day 2022 (”Industry Day”) event emphasized the importance of defining “session” and other terms, including *per session*, *per-request access*, and *per-request logging*
- Industry Day Panelists are not alone in their thinking; NIST also released CSWP 20, which explicitly states that “the unit of ‘session’ can be nebulous and differ depending on tools, architecture, etc.”

What needs to be done?

- Defining, standardizing, and reinforcing concepts such as “session” will help to solidify the industry’s overall understanding of ZT tenets and describe how they will look in practice
- Finding common ground is essential, as these definitions will need to be generally accepted
- Codifying these Zero Trust definitions will allow government, academia, and private industry to remain on the “same page”, promoting effective collaboration and innovation

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 2: Establish a Common View of ZT

Carnegie Mellon University Software Engineering Institute

Standardization Beyond Definitions

- Establishing an open-source standard for defining event communication among ZT components would allow greater integration and communication among individual components of a ZT environment
- One panelist mentioned OpenID Foundation's Shared Signals and Events (SSE) Framework to standardize and streamline the communication of user-related security events among different organizations and solutions
- The Policy Decision Point (PDPs) is a central element in Zero Trust Architecture that makes "authorization decisions for itself or for other system entities that request such decisions" (CNSSI 4009-2015 from NISTIR 7657)
- For access-related decisions, the PDP relies on policies, logs, intelligence, and machine learning (ML) or may leverage unique workflows to develop instruction sets or operating parameters
- There is little discussion, however, about how these factors might work in practice and how they should be implemented

Areas of Future Effort

- To encourage uniformity and interoperability, security organizations could develop a standardized language for PDP functionality, similar to the STIX/TAXII2 standards developed for cyber threat intelligence
- Promoting interoperability between security solutions may bolster innovation, streamline transformation, and enable more efficient and effective Zero Trust operations

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 3: Establish Standard ZT Maturity Levels

Carnegie Mellon University
Software Engineering Institute

Current State of Maturity Models

- Existing ZT maturity models do not provide granular control or discussion of the minimal baselines required for effective shifts to ZT
- It is important to consider how to develop a maturity model with enough levels to help organizations identify exactly what they must do to meet ZT standards for basic security
- One panelist emphasized the need to define the minimum baseline requirements necessary for ZTA in the real world
- It is critical to establish a standard of technical requirements for ZT maturity so that organizations can identify and audit their progress toward digital trust
- According to the World Economic Forum, digital trust represents “individuals’ expectation that digital technologies and services – and the organizations providing them – will protect all stakeholders’ interests and uphold societal expectations and values.”

Areas of Future Effort

- Since the date of our blog and paper publication, CISA has released Version 2.0 of the CISA Zero Trust Maturity Model
- The new version is improved as a whole, but there is still a notable gap between the maturity model and real-world implementation guidance
- Documents such as NIST SP 1800-35A ”Implementing a Zero Trust Architecture” are working towards bridging the gap between maturity models and real-world implementation
- However, continued R&D efforts are needed to ensure parity between maturity models and implementation resources
- It is important to note that implementation resources are relatively sizeable, complex, and subject to change due to an evolving tech stack

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 4: Explain How to Progress Through ZT Maturity Levels

Carnegie Mellon University
Software Engineering Institute

Zero Trust Maturity Progression

- Building on Area 3, organizations in the security space must identify the minimum steps required to implement ZT at some level while also demonstrating how those steps might look in practice
- Once an organization has begun implementing ZT, it can work toward higher levels of ZT maturity, with the ultimate goal of achieving digital trust.
- According to ISACA, digital trust refers to the “confidence in the integrity of the relationships, interactions and transactions among suppliers/providers and customers/consumers within an associated digital ecosystem.”

Areas of Future Research

- ZT serves as the foundation for interaction among entities from a cybersecurity perspective, and digital trust encompasses all the interactions between internal and external entities more comprehensively
- Government and related entities must actively collaborate with private-sector organizations to align models, standards, and frameworks with real-world products and services
- Public-private partnerships in this domain must focus on identifying:
  - (1) What a security offering can and cannot do
  - (2) How each offering can integrate with others to achieve a specific level of compliance
- This approach provides end users with useful information about how a particular product can leverage ZT strategies to achieve digital trust

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 5: Ensure ZT Supports Distributed Architectures

Carnegie Mellon University
Software Engineering Institute

CDNs and related systems

- With the increasing adoption of cloud solutions and distributed technologies (e.g., content delivery networks, or CDNs), it is necessary to develop security frameworks that account for applications and data moving away from a central location and closer to the user
- When developing frameworks and standards for the future of ZT, it is important to consider that offsite data storage is being moved closer to the consumer, as demonstrated by the prevalence of CDNs in modern IT infrastructures

Areas of Future Effort

- One panelist suggested exploring this topic in the context of new security frameworks, since many existing frameworks take a centralized data center/repository approach when describing security best practices
- This approach underserves CDN-oriented organizations when they are developing and assessing their security posture and architecture
- CDNs and edge systems must be considered in the context of Zero Trust Architecture, as these systems generally align with NIST's rationale of addressing security risks beyond the traditional "enterprise-owned network boundary"

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 6: Establish ZT Thresholds to Block Threats

Carnegie Mellon University
Software Engineering Institute

Establishing Minimum Thresholds

- In a ZT environment, it is important to understand what constitutes the minimum amount of information required to effectively isolate and block an activity or piece of malware
- Identifying this information is essential since a growing number of ransomware attacks are using custom malware
- To defend against this threat, organizations must improve their ability to detect and block new and adapting threats
- An important aspect of ZT is using multiple strategies to detect and isolate attacks or malware before they spread or cause damage

Areas of Future Effort

- A properly implemented zero trust architecture should not trust unknown software, updates, or applications, and it must quickly and effectively validate unknown software, updates, and applications
- ZT can use a variety of methods (e.g., sandboxes and quarantines) to test and isolate new applications
- These results must then be fed into the PDP so that future requests for those applications can be approved or denied immediately
- This domain also aligns with our previous discussion of formalized standards and definitions, as greater interoperability could bolster ZTA capabilities in response to emerging and real-time threats

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 7: Integrate ZT and DevSecOps

Carnegie Mellon University
Software Engineering Institute

DevSecOps

- In the development process, it is important to use as many security touchpoints as possible, especially those related to ZT
- It is also important to understand how to emphasize security in an organization's development pipeline for both conventional and emerging technologies
- These considerations lead us into the realm of DevSecOps, which refers to a "set of principles and practices that provide faster delivery of secure software capabilities by improving the collaboration and communication between software development teams, IT operations, and security staff within an organization, as well as with acquirers, suppliers, and other stakeholders in the life of a software system." (SEI)

ZTA and DevSecOps

- As automation becomes more prevalent, DevSecOps must account for the possibility that a requestor is automated
- ZTA uses the identity of the workloads that are attempting to communicate with one another to enforce security policies
- These identities are continuously verified; unverified workloads are blocked and therefore cannot interact with malicious remote command-and-control servers or internal hosts, users, applications, and data
- When developing software, everyone historically assumed that a human would be using it
- When security was implemented, therefore, default authentication methods were designed with humans in mind

Areas of Future Effort

- As more devices connect with one another autonomously, software must be able to use ZT to integrate digital trust into its architecture
- To enable the ZT strategy, DevSecOps must be able to answer the following questions:
  - Is the automated request coming from a trusted device?
  - Who initiated the action that caused the automated process to request the data?
  - Did an automated process kick off a secondary automated process that is now requesting the data?
  - Does the human who configured the automated processes still have access to their credentials?

Carnegie Mellon University
Software Engineering Institute

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Area 8: Set Business Expectations for ZT Adoption

Business Alignment

- Security initiatives are frequently expensive, which contributes to the organization's perception of security as a cost center
- It is important to identify inefficiencies (e.g., obsolescence) during the ZT transformation process
- It is also crucial that organizations understand how to use ZT to maximize their return on investment
- ZT is a strategy that evaluates and manages the risk to an organization's digital assets
- A ZT approach shifts the defenses from the network perimeter to in-between digital assets and requires session authentication for all access requests

Potential Action Items

- Many ZT strategies can be implemented with a reasonable amount of effort and at a low cost to the organization
- Examples include micro-segmentation of the network, encryption of data at rest, and user authentication using multi-factor authentication
- However, some solutions (e.g., cloud environments) require a lengthy transition period and incur ongoing costs
- Since organizations have unique risk tolerance levels, each organization must develop its own ZT transformation strategy and specify the initial phases
- Each of these strategies and phases will have different costs and benefits
- Developing further guidance materials in this domain may be beneficial

8 Areas of Future Research in Zero Trust: Key Takeaways from Government and Industry

Conclusion

Carnegie Mellon University
Software Engineering Institute

Parting Thoughts

- The SEI's Zero Trust Industry Day 2022 was designed to bring vendors in the ZT field together and offer a shared platform for discussion
- This approach allowed participants to objectively demonstrate how their products could help organizations with ZT transformation
- Discussions included several areas that could benefit from further exploration
- By highlighting these areas of future research, we are raising awareness, promoting collaboration among public and private-sector organizations, and accelerating ZT adoption in both government and industry
- We hope our discussion has inspired creative thinking within this domain

References

[CISA 2023]
Cybersecurity Infrastructure Security Agency, Cybersecurity Division. *Zero Trust Maturity Model, Version 2.0*. 2023.

https://www.cisa.gov/zero-trust-maturity-model

[Garbis 2021]
Garbis, J. & Chapman, J. *Zero Trust Security: An Enterprise Guide*. Berkeley, CA: Apress. 2021. https://link.springer.com/book/10.1007/978-1-4842-6702-8

[Rose 2020]
Rose, S.; Borchert, O.; Mitchell, S.; & Connelly, S. *NIST Special Publication 800-207: Zero Trust Architecture*. 2020. https://csrc.nist.gov/publications/detail/sp/800-207/final

Thank you very much!

Please feel free to reach out with any questions or comments!

Matthew Nicolai
menicolai@cert.org

Trista Polaski
tjpolaski@cert.org