Zero Trust Architecture Overview & Innovations prepared for USAF Cyberspace Dominance

Jason L. Lind, USAF (Sep.)
lind@multiplex.studio
+1 608.301.1170
12 May 2020

Trust not: verify always! – Zero Trust was first coined in late 2009 by John Kindervag, a principal analyst at Forrester Research and at its core is about shrinking the verification perimeter to as close to the actual data as possible and continuously monitoring the "protect surface" for threats. Since then the concept has gained significant momentum including recently a NIST Special Publication (SP 800-207) and a host of vendors claiming to be the best of breed Zero Trust. The below graphic is Centrify's interpretation of ZT.

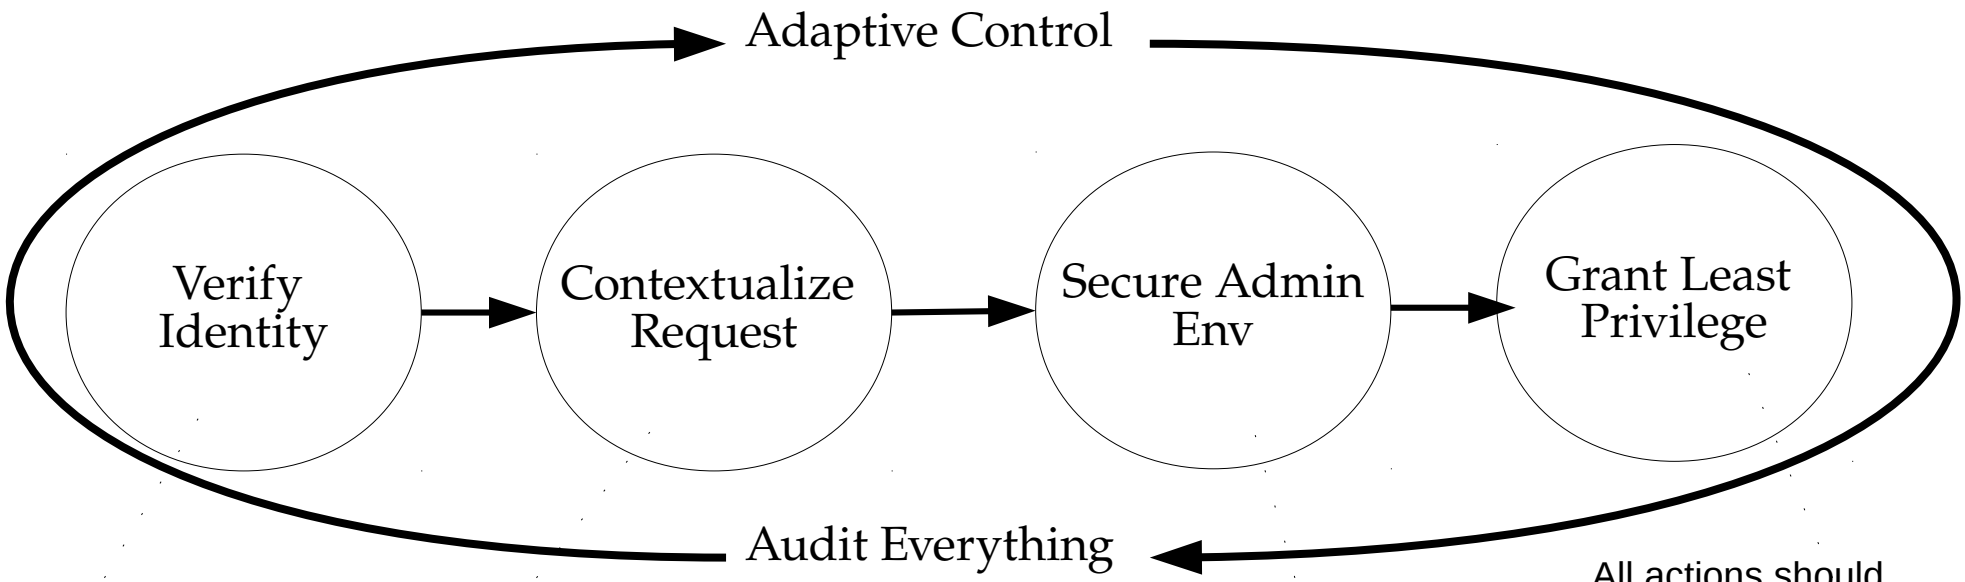

Zero Trust Tenets According to NIST SP 800-207

1. **All data sources and computing services are considered resources**
2. **All communication is secure regardless of location**
3. **Access to individual enterprise resources is granted on a per-session basis**
4. **Access to Resources is determined by dynamic policy**
5. **The enterprise ensures all owned and associated devices are in the most secure state possible and monitors assets to ensure they remain in the most secure state possible**
6. **All resource authentication and authorization are dynamic and strictly enforced**
7. **The enterprise collects as much information as possible about the current state of network infrastructure and communications and uses it to improve its security posture**

Zero Trust Tenet Model

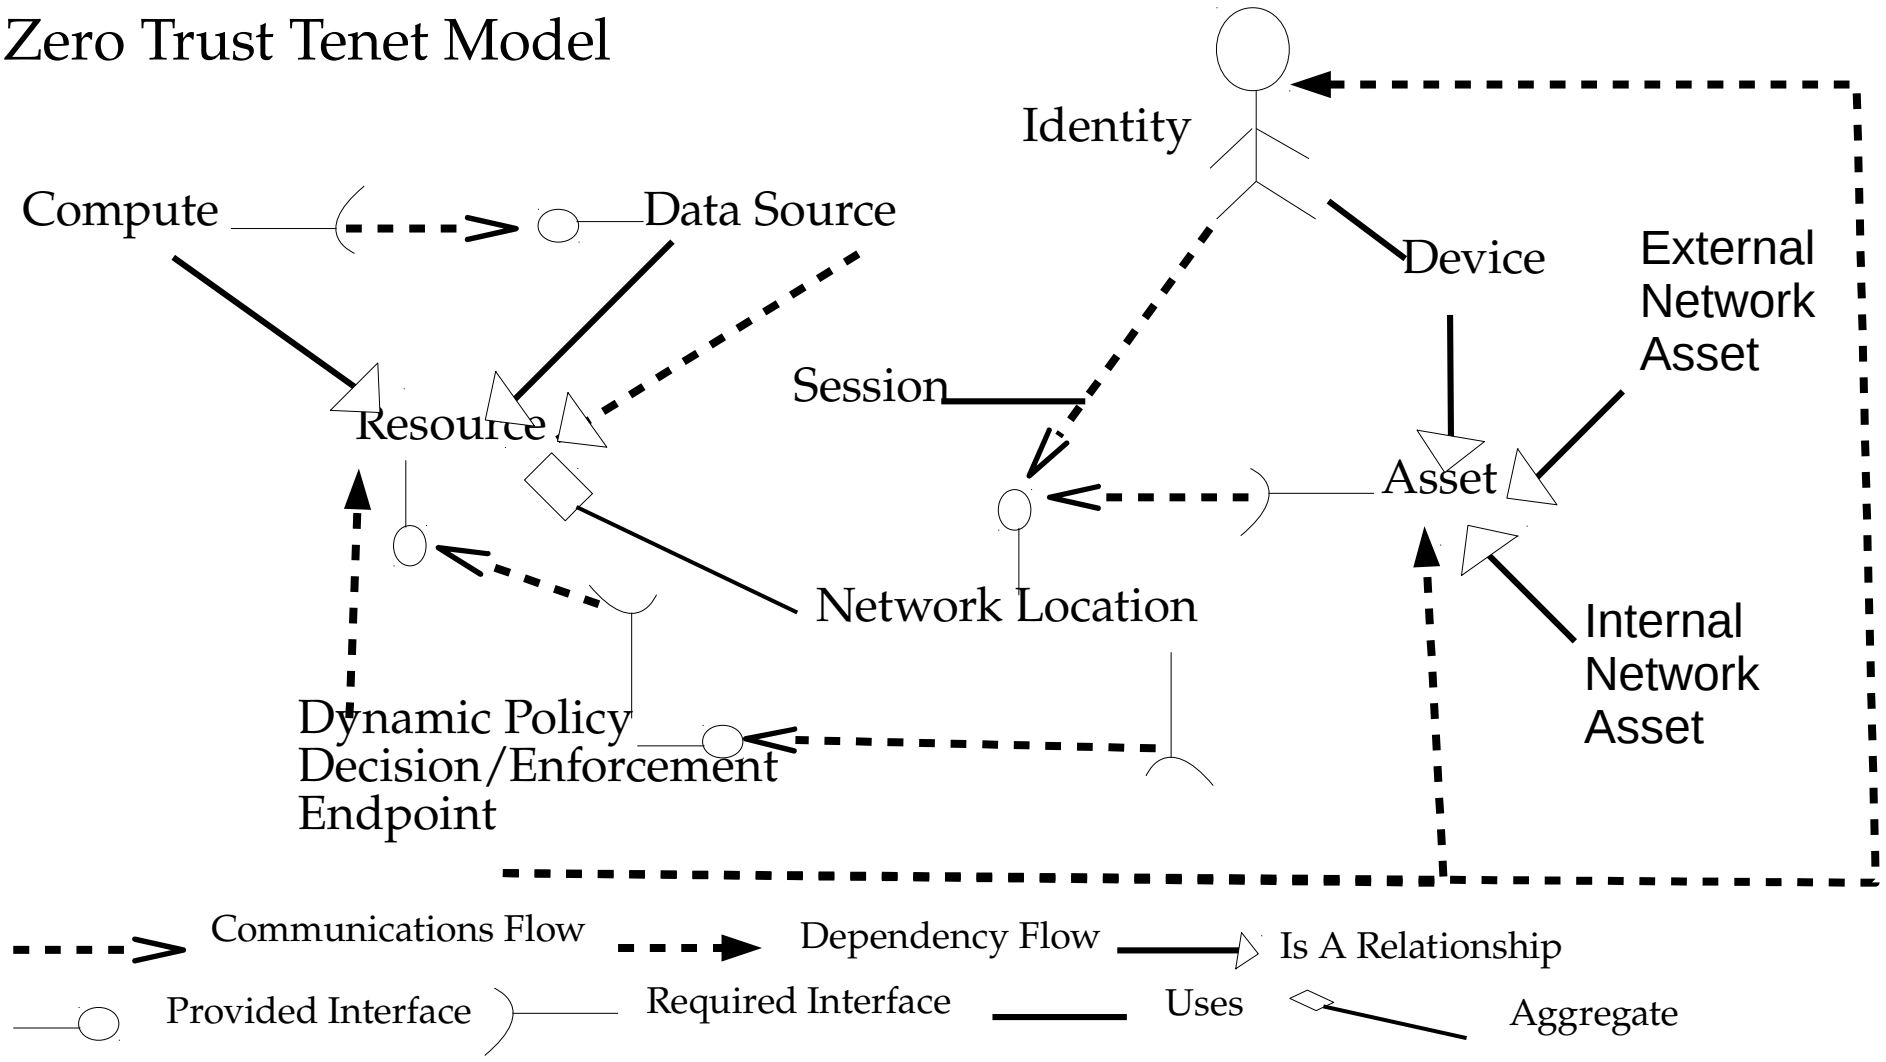

A core concept in Zero Trust is that the verification surface in continuously evaluated

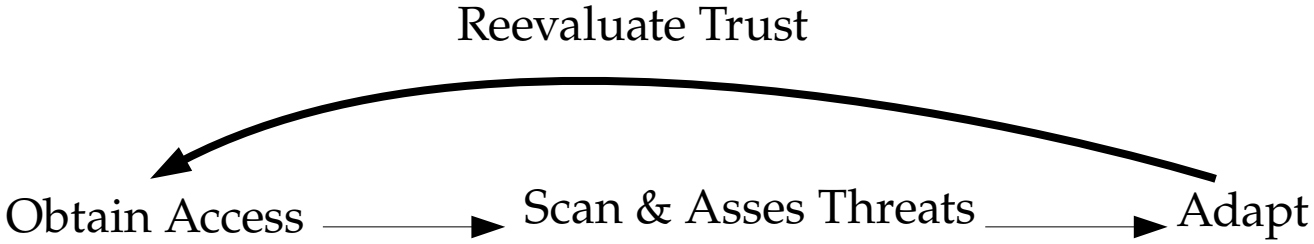

Pure Zero Trust Logical Architecture

Zero Trust logically has two core logical system: The Policy Decision Point (PDP) and Policy Enforcement Point (PEP). The PDP is composed of the Policy Engine, “the brain” having ultimate decision power for granting resource access and the Policy Administrator (PA) responsible for establishing communication paths.

The PEP is responsible for enabling, monitoring and eventually terminating connections between a subject and an enterprise resource

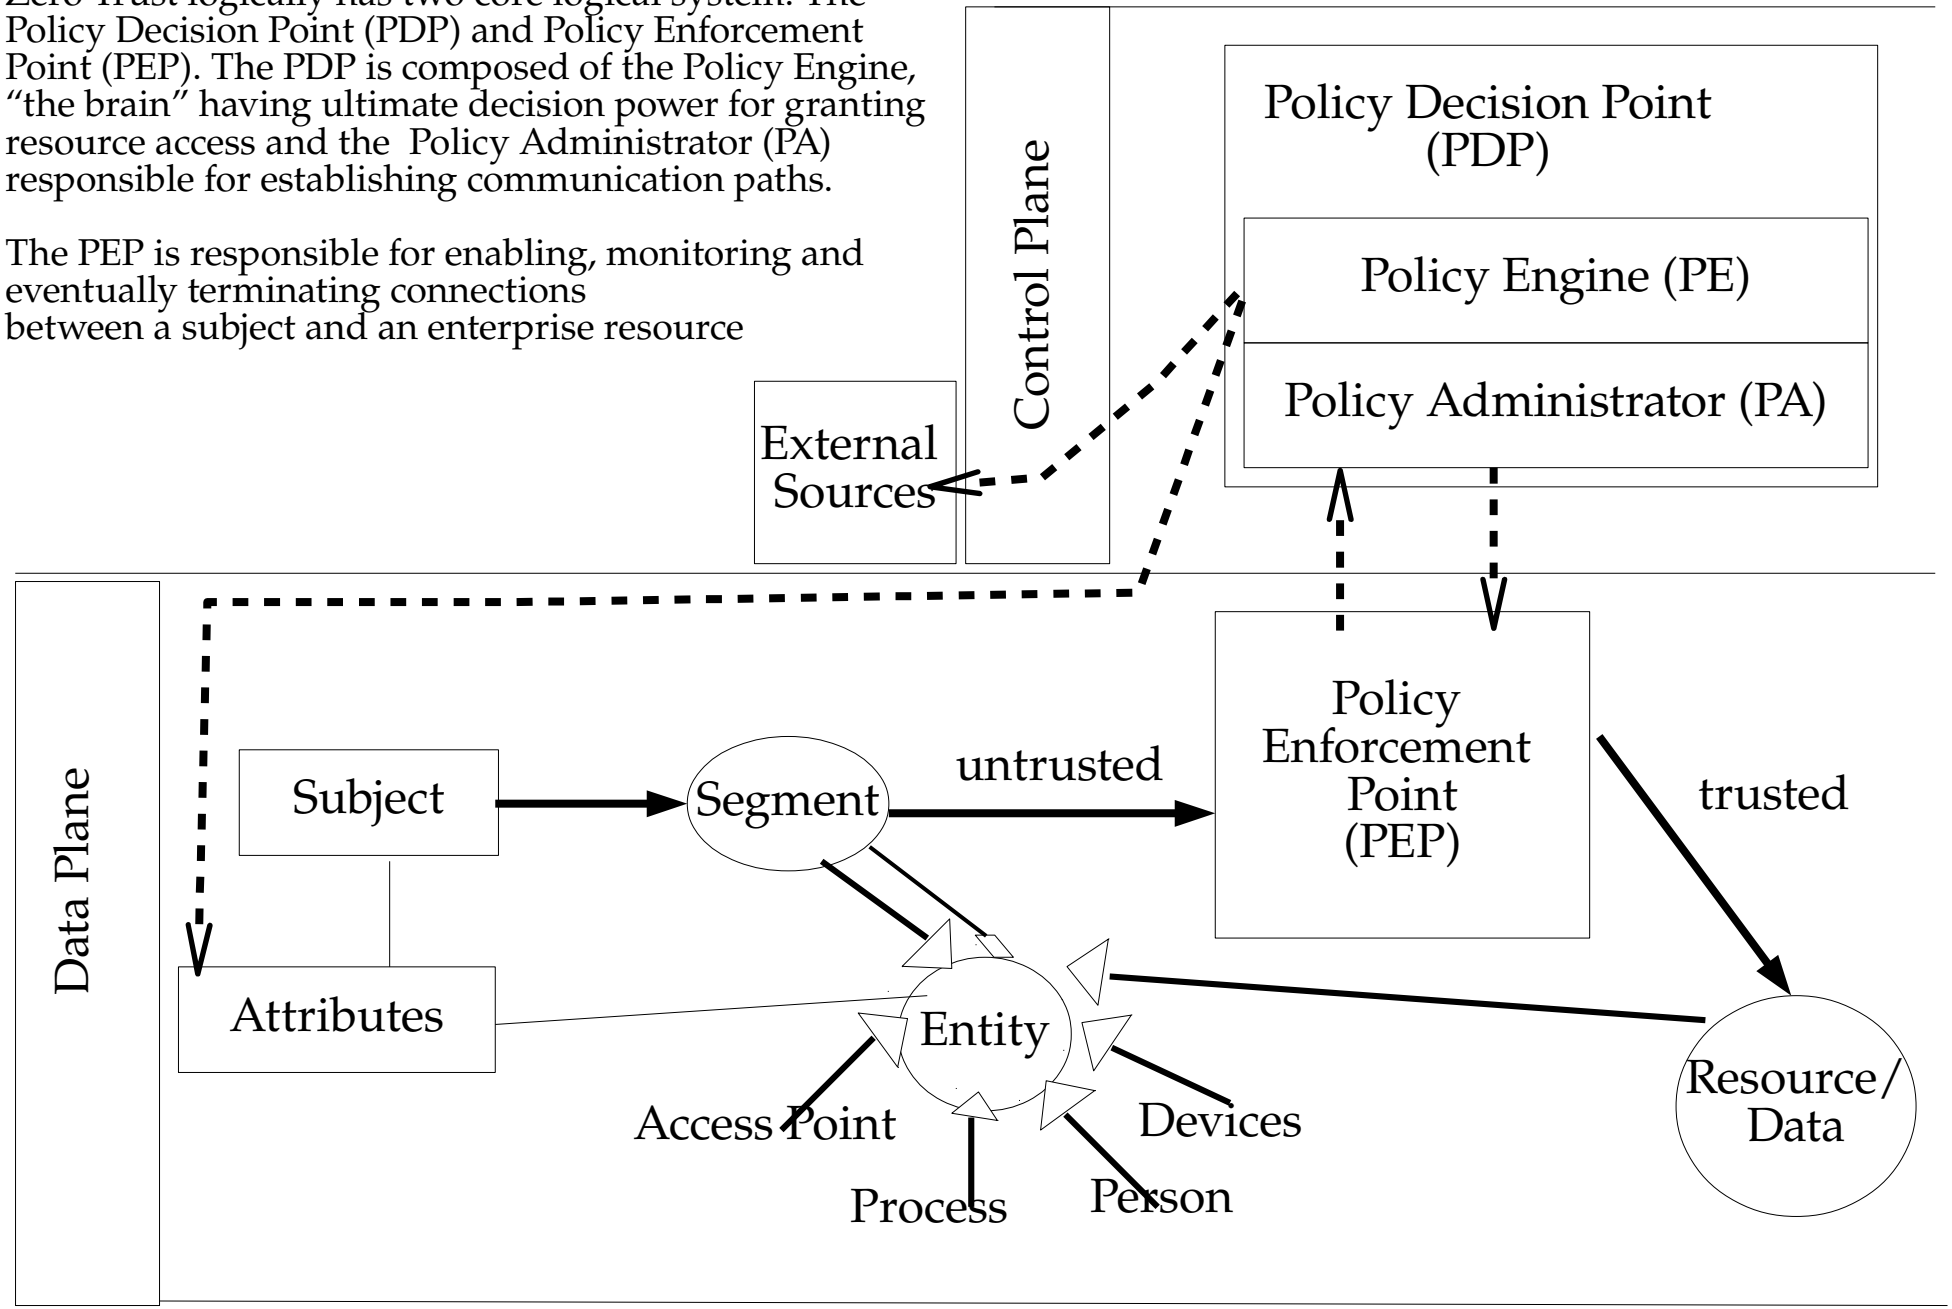

Networking Requirements

Enterprise Assets have basic network connectivity

Distinguish between assets owned or managed by the enterprise and their current security posture

Enterprise Resources unreachable without accessing a PEP

The Data Plane and Control Plane are logically separate

Enterprise Assets can reach PEP component

The infrastructure used to support the access decision process should be made scalable to account for changes in process load

Assets may not be able to reach certain PEPs due to observable factors

Policy Engine's Trust Algorithm

If the Policy Engine can be thought of as the brain of the PDP then its Trust Algorithm can be thought of as its primary thought process. Trust Algorithms have historically been static and simple, however today's evolving threats require a dynamic and complex method for authorizing resources access.

Effective trust algorithms will evaluate on at least the following criteria:

- **Access Request:** The actual request from the subject, including device information
- **User identification, attributes and privileges:** The "who" that is requesting the resource
- **Asset Database and Observable Status:** A database that contains the known status of each enterprise owned asset
- **Resource Access Requirements:** Defines minimal requirements for access to a resource
- **Threat Intelligence: Informational** feeds about general threats and active malware operating on the Internet

Policy Engine External Inputs

Continuous Diagnostics & Mitigation (CDM)
Provides PE with metrics regarding asset's configuration.

Industry Compliance System
Enures that the enterprise remains compliant with regulatory constraints (e.g. blocking a request that would violate HIPPA for a patient)

Threat Intelligence Feed(s)
Information about newly discovered attacks or vulnerabilities (e.g. blacklists, malware, reported attacks to other assets, etc)

Data Access Policies
Attributes, rules and policies about access to enterprise resources, preferably dynamically generated by the policy engine but could be encoded.

Enterprise Public Key Infrastructure (PKI)
Generation and logging of cryptography certificates

ID Management System
Creation, storage and managing of enterprise use accounts and identity records (e.g. LDAP)

Network and System Activity Logs
Aggregates asset legs, network traffic and resource access actions in order to provide real-time feedback on the security posture of enterprise information systems

Security Information & Event Management (SIEM)
Collection of security-centric information for later analysis

Legacy Perimeter Defense

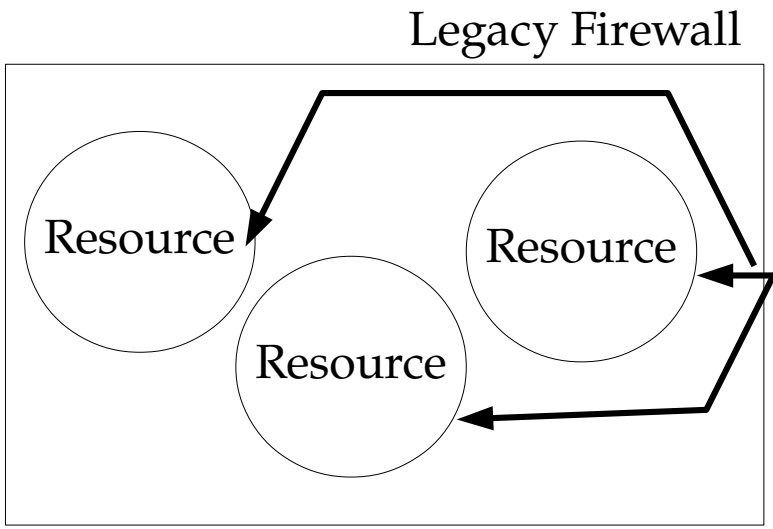

Micro-segmentation is a logical strategy to leverage virtual subnets to protect enterprise resources via isolation. A resource, or collection of resources, are protected on a virtual, or physical, network via a Next-Generation Firewall or gateway device acting as a PEP. This is sharply opposed to the legacy perimeter based model where all resources are behind the same firewall and accessible once authenticated leaving a large surface to protect in its wake. By shrinking the perimeter as close to the logical data flows as possible an enterprise can better mitigate successful attack vectors against their organization

µSegmentation Approach

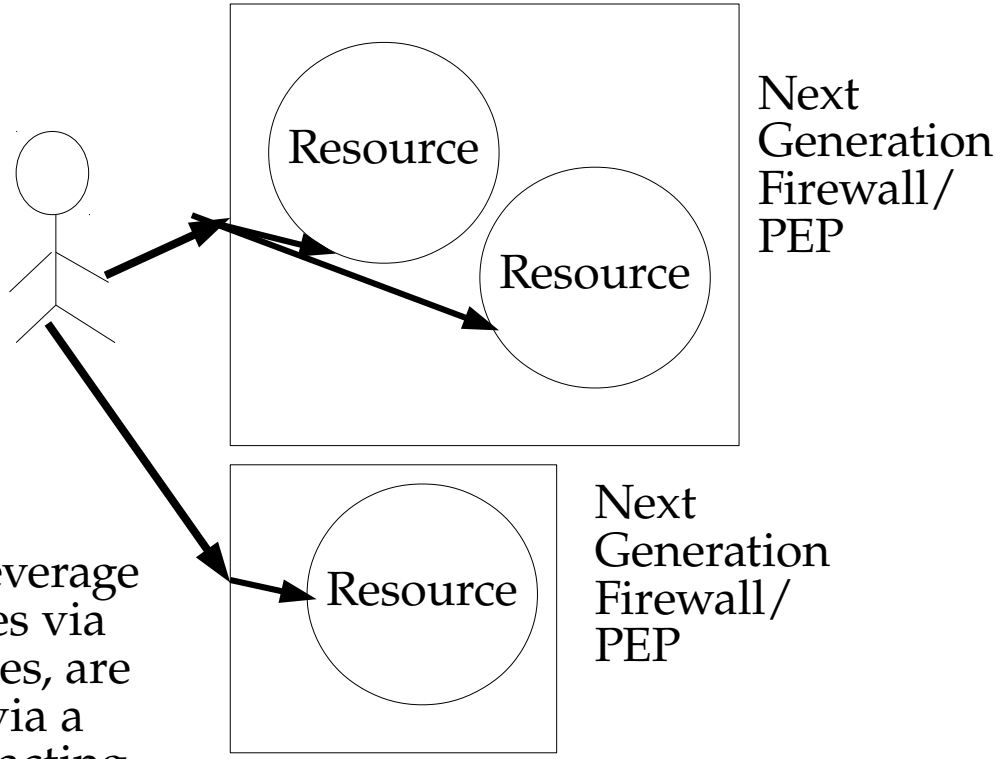

Bureaucratic Silos & Segmentation

Conway's Law notes that our applications and infrastructure are representative of how we communicate in our organization. Segmentation therefore must be carefully architected to emphasize the communication strategy of the enterprise.

Software Defined Perimeter (SDP) Strategy for ZTA

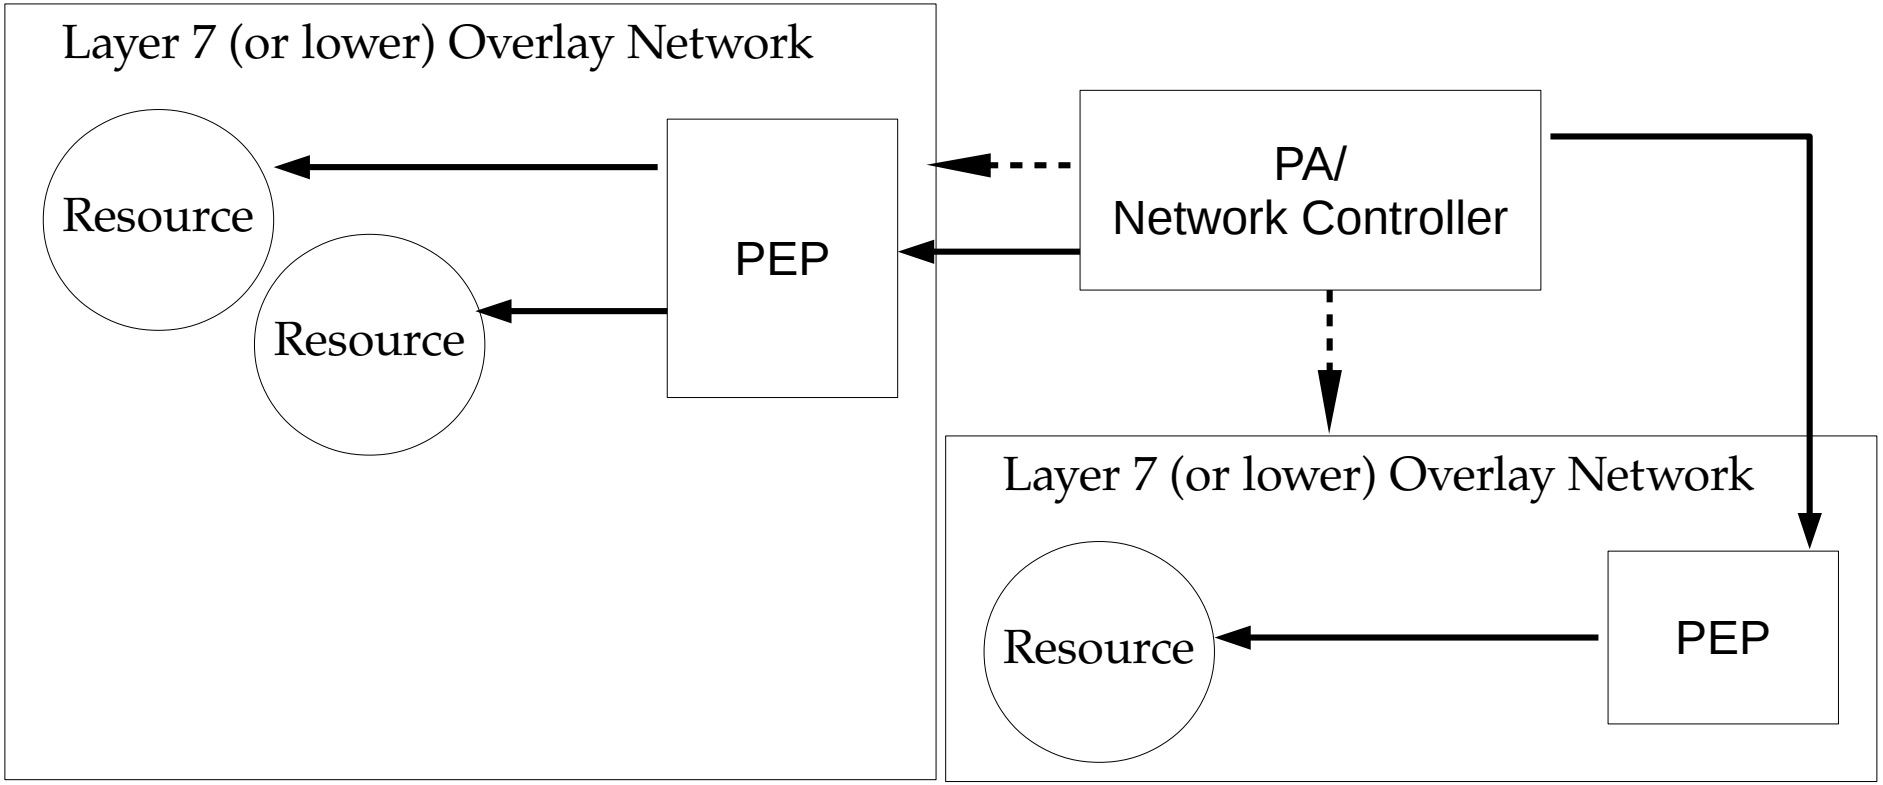

Software Defined Perimeters enable enterprises to isolate resources at the application, or lower, layer (OSI Model). Like micro-segementation the goal of this is to distribute and shrink the attack surface, in effect taking that paradigm further enabling greater granularity of control by utilizing virtual versus physical boundaries.

Dr. Chase Cunningham's Seven Pillars of Zero Trust eXtended (ZTX)

Automation & Orchestration

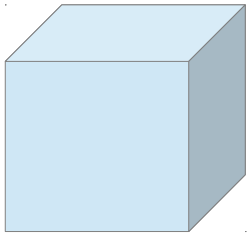

Artificial processes are a proper subset of the threat vectors as people and therefore their agents must be treated as any other actor

Data

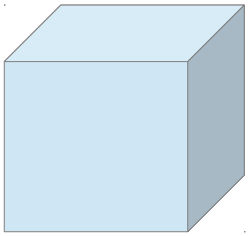

Data is centric to Zero Trust as at the end of the day data, not just systems, is what is being protected and secured.

Workloads

The conceptual and computational throughput of a process

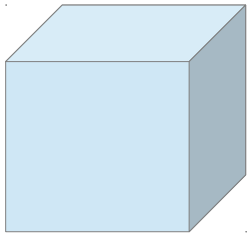

Visibility & Analytics

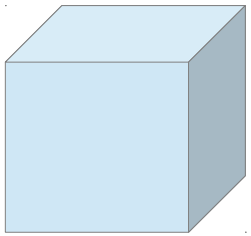

Build dashboards and analysis tools to monitor and analyze the entire `protect surface`

Devices

The devices people use define how they access the network and the kind of identity factors and presence that can be presented

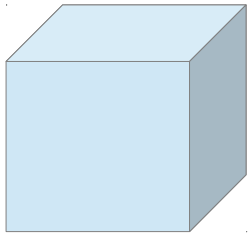

Networks

Networks are at the outermost of the `protect surface` resides, software defined perimeters and micro-segmentation are to be considered here.

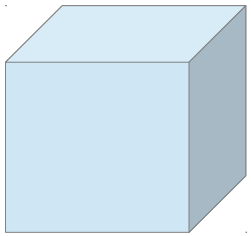

People

At the end of the day it is people who are using the system and your own users are the biggest threat to the system: Train! Train! Train!

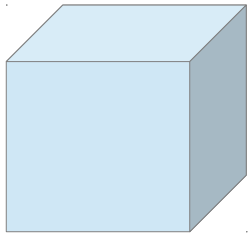

My Interpretation of ZTX Relationships