How to bluff your way into Zero Trust

Peter van Eijk

WHY2025

ClubCloudComputing

No implicit trust, only ‘allow’ rules

- “All breaches happen inside an allow rule” – John Kindervag
- “We trust everything in our datacenter” -> Insider threat
- “Describe trust explicitly”
- .. But how? The devil is in the details
- .. Technology neutral?
- .. Abstract, versus actual rules? So it can be understood by non-tech?

History and status

- John Kindervag @ Forrester (Chewy Centers - 2010)
- US Executive order 14028, from 2021
  - Not (yet) cancelled, in contrast to the AI order
- Yearly DoD online symposium 2024:2000 -> 2025:4000 attendance
- Lots of public DoD material
- US Air Force detailed strategy and roadmap
- CCZT: Certificate of Competence in Zero Trust
- Products? Yes, but not one to rule them all

Major use cases & benefits

- VPN replacement
- Cloud (and micro-segmentation)
- DevOps
- Admin access
- Business process
- 3rd party / supply chain

- (YMMV)
- Better security
- Cheaper security
- Easier compliance
- Better UX
  - Less re-authentication
  - Faster provisioning
- Cheaper licensing
- …

But really …

- Trust is a vulnerability
- It is about reducing assumed trust

- Does the packet adequately represent the person (or agent)?

Fine grained allow rules on the protect surface

ZT: Attack versus protect surface

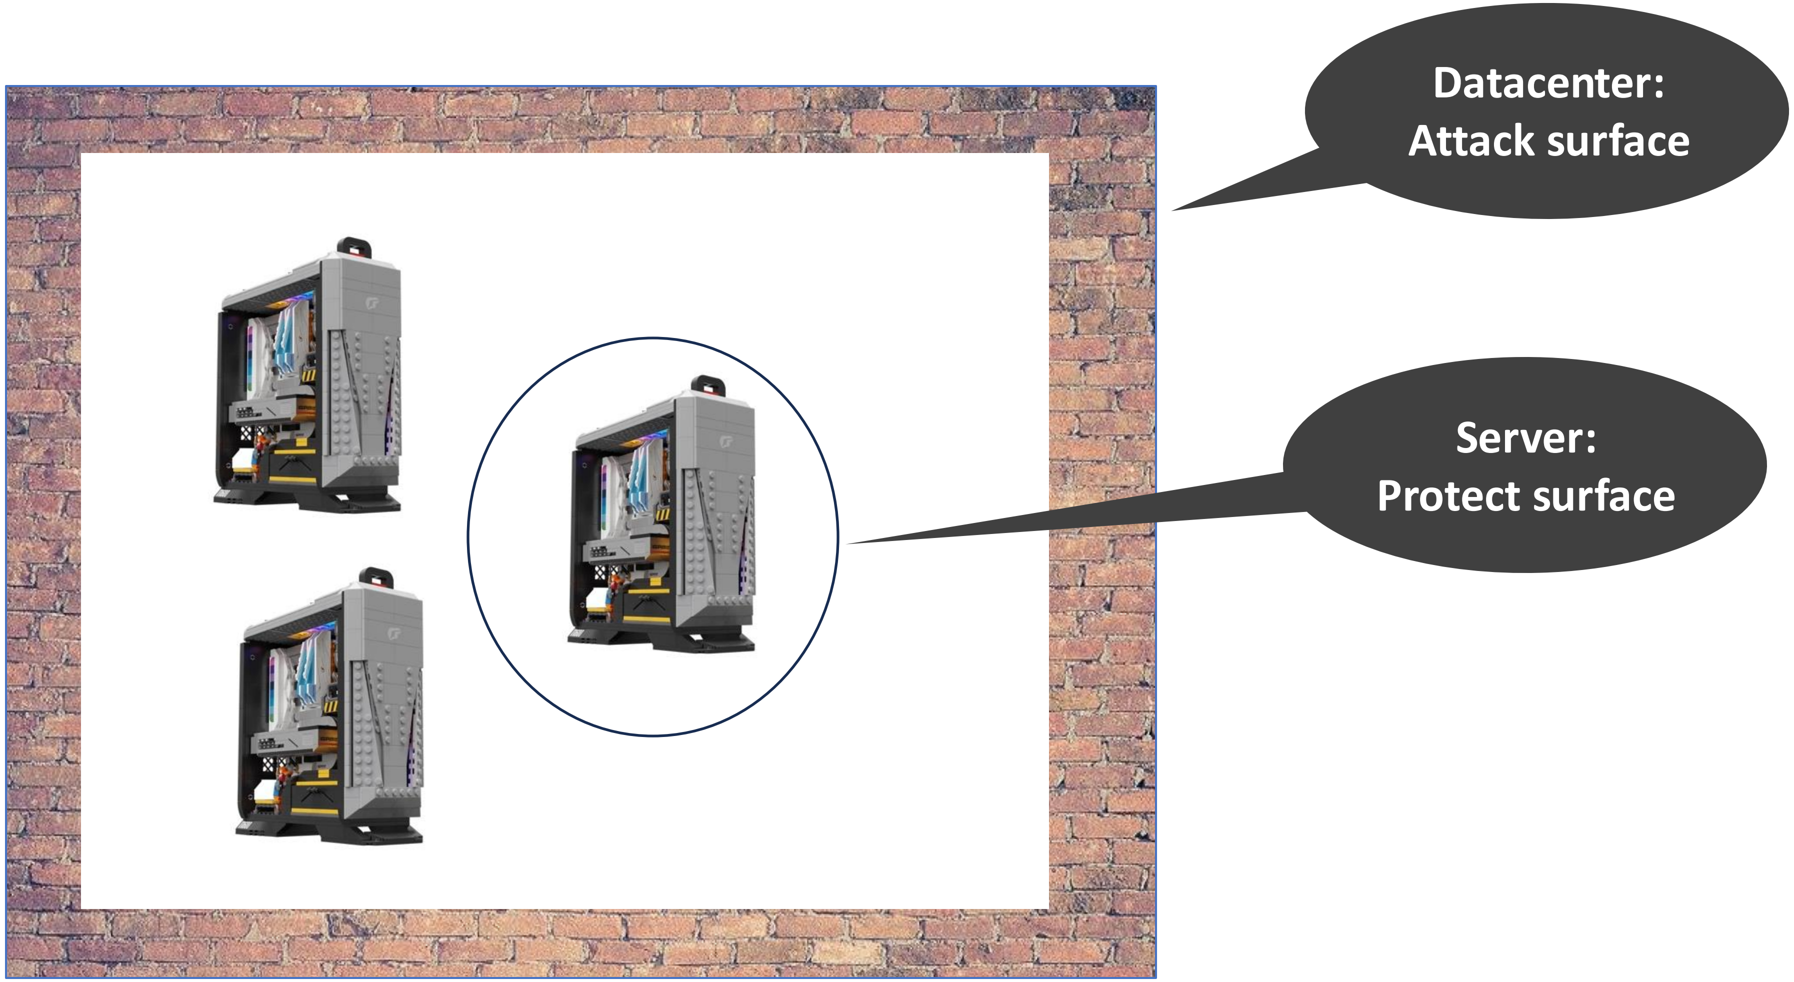

Where does ZT live?

- in the workloads
- in the hypervisors
- in the SDN controllers
- in the network fabrics
- … and possible more places

Buzzwords: DAAS elements is what you want to protect

- Data – what you want to protect
- Applications – holds data, and controls access
- Assets (e.g. devices) – runs code, possibly malicious code (i.e. unauthorized applications)
- Services – externalised applications

Key concepts PDP and PEP

- Policy Decision Point
- Policy Enforcement Point
- .. And subdivisions of those
- Various historical sources ..

- RFC2748 (2000) https://datatracker.ietf.org/doc/html/rfc2748
- https://en.wikipedia.org/wiki/Common_Open_Policy_Service
  - Introduces outsourcing and provisioning model between PDP and PEP
  - Originally for QoS policies
- Cisco & Checkpoint use this in their network product documentation, history unknown.

PDP and PEP

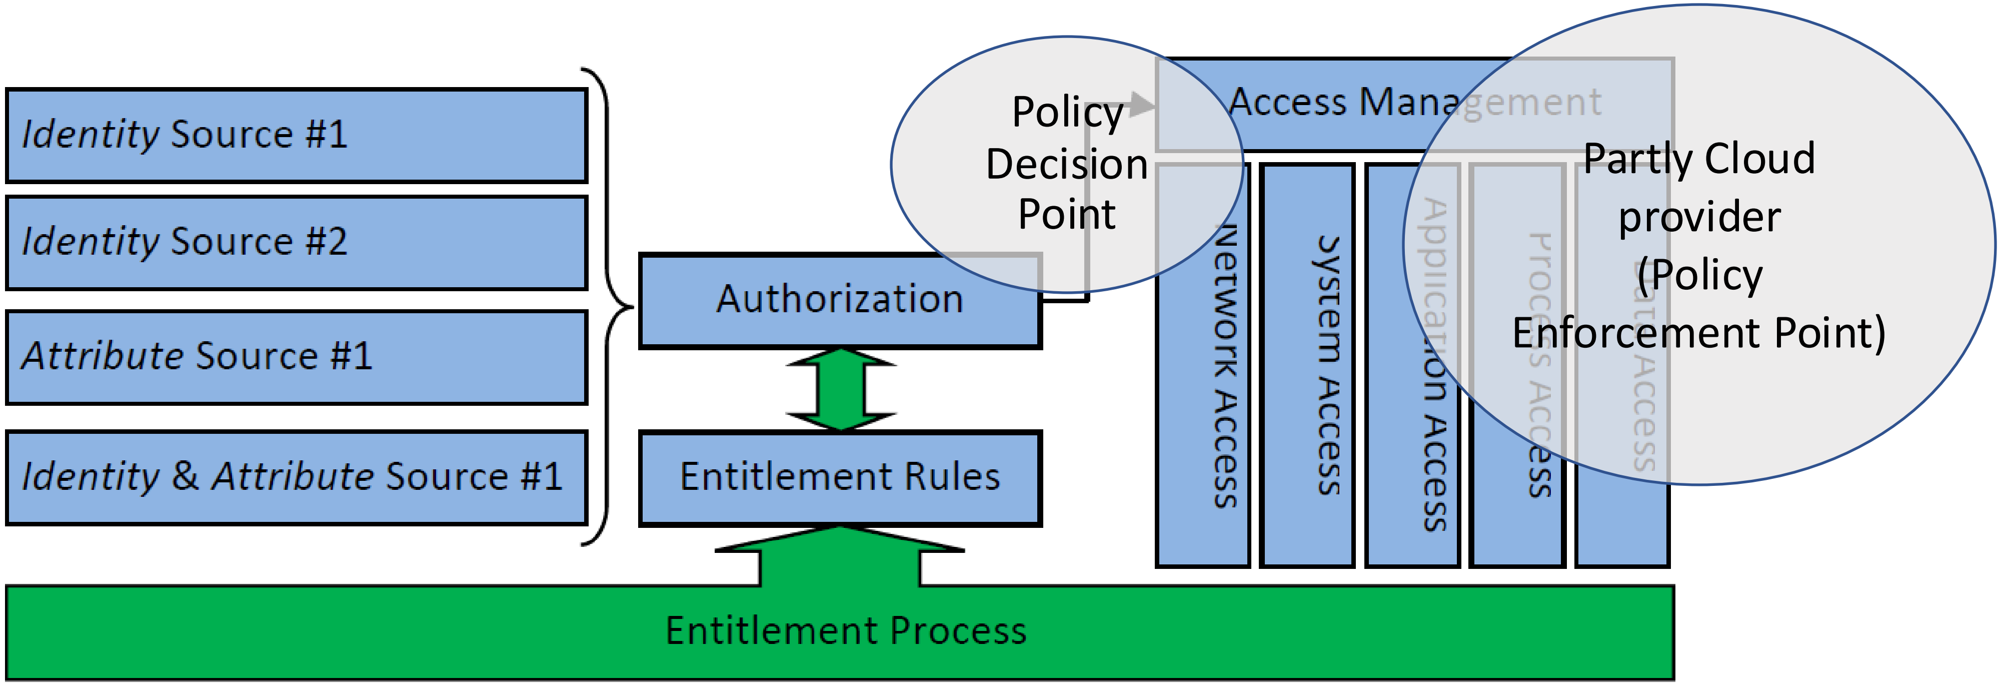

Figure 1: Generic Identity, Entitlement & Access Management System

Source: CSA guidance v3

Why separate PEP and PDP?

1. Distributed PEPs
   - Latency
   - Volume
2. Different PEP technologies
3. Cost
4. …

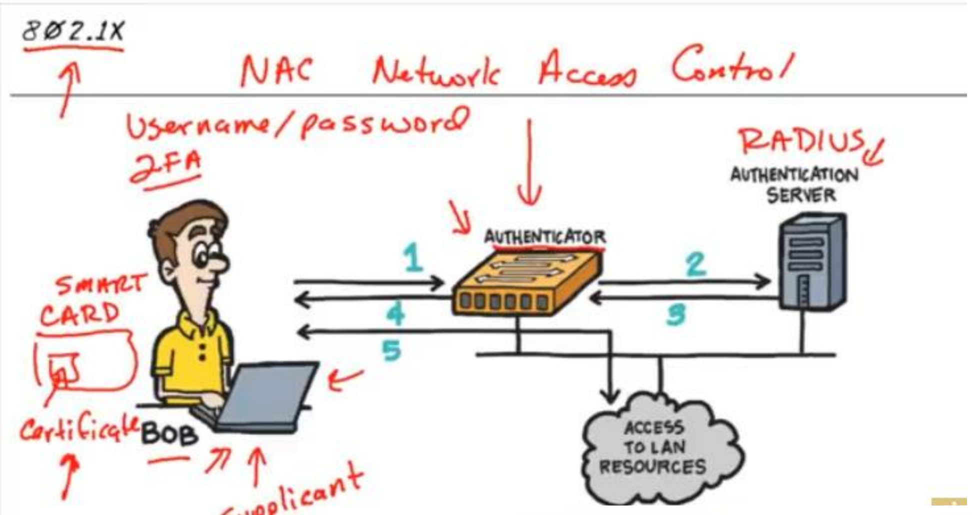

PEP = Authenticator
PDP= Authentication server (RADIUS)

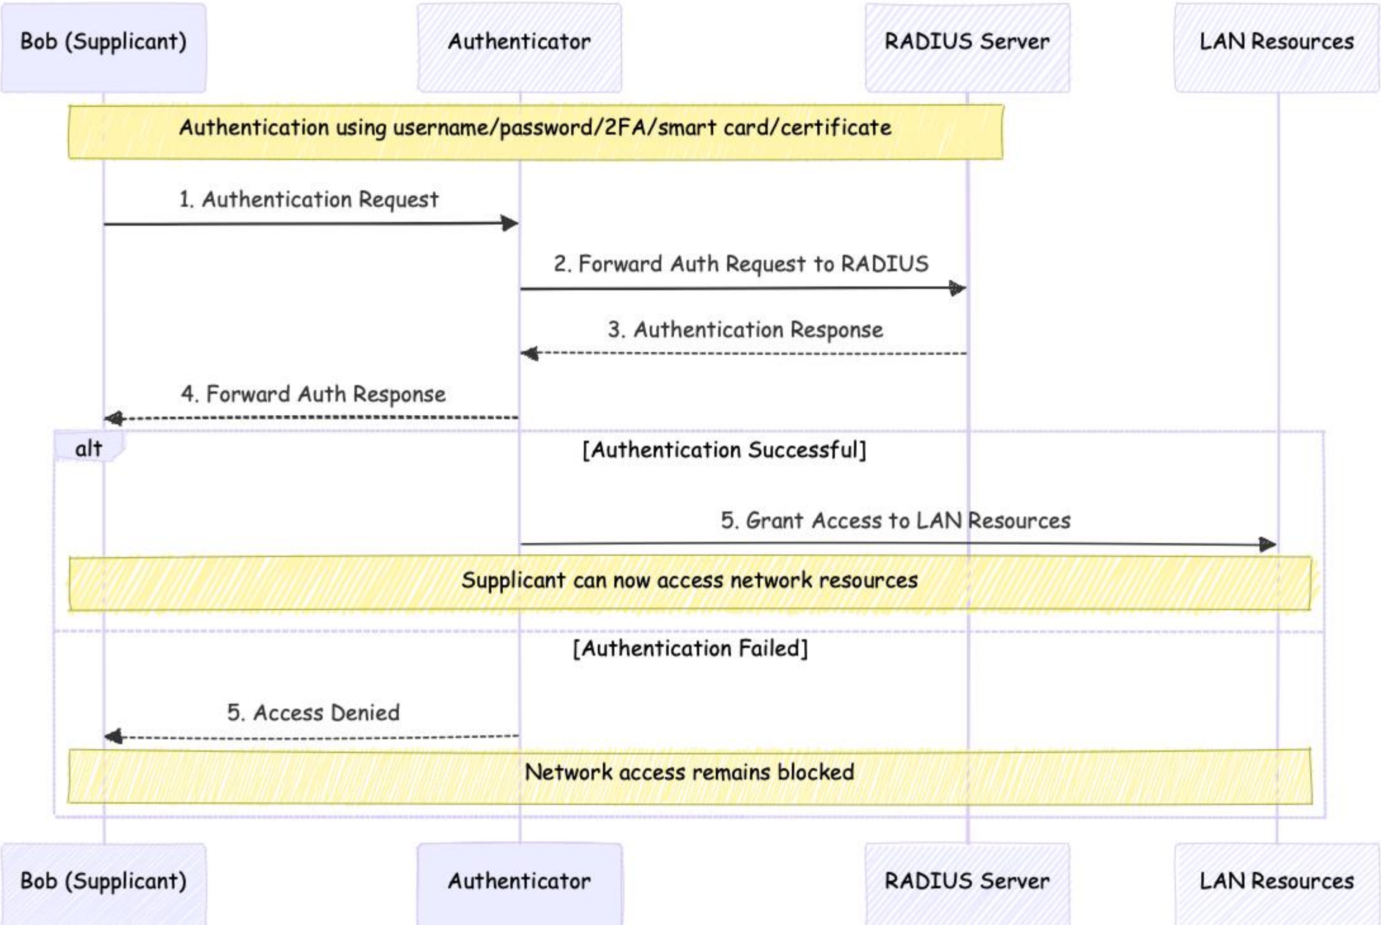

Bob (Supplicant)

Authenticator

RADIUS Server

LAN Resources

Authentication using username/password/2FA/smart card/certificate

1. Authentication Request

2. Forward Auth Request to RADIUS

3. Authentication Response

4. Forward Auth Response

alt

[Authentication Successful]

5. Grant Access to LAN Resources

Supplicant can now access network resources

[Authentication Failed]

5. Access Denied

Network access remains blocked

Bob (Supplicant)

Authenticator

RADIUS Server

LAN Resources

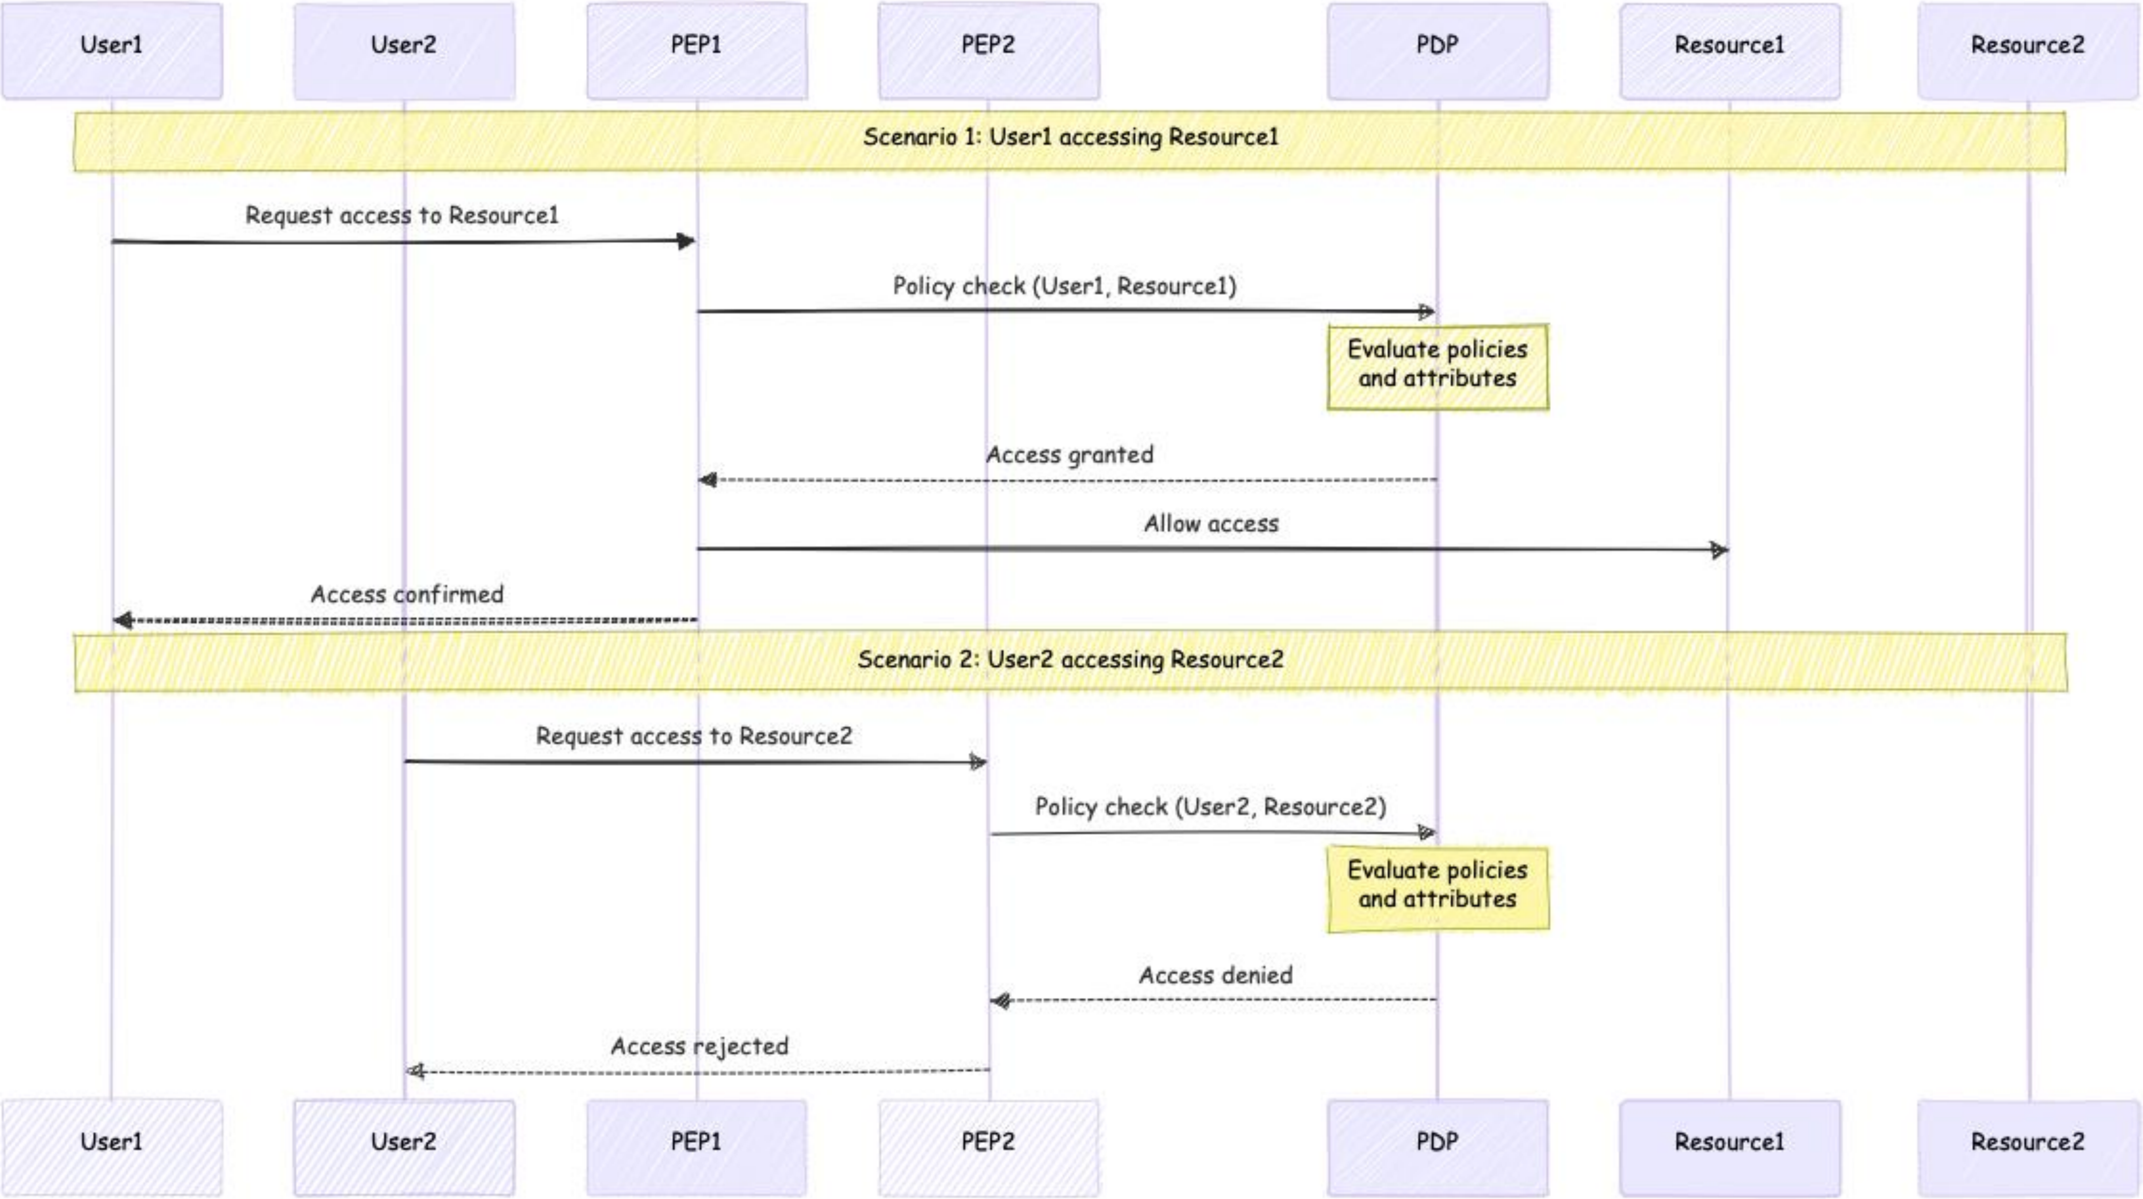
User1
User2
PEP1
PEP2
PDP
Resource1
Resource2
Scenario 1: User1 accessing Resource1
Request access to Resource1
Policy check (User1, Resource1)
Evaluate policies and attributes
Access granted
Allow access
Access confirmed
Scenario 2: User2 accessing Resource2
Request access to Resource2
Policy check (User2, Resource2)
Evaluate policies and attributes
Access denied
Access rejected
User1
User2
PEP1
PEP2
PDP
Resource1
Resource2

Software Defined Perimeter & Zero Trust Architecture

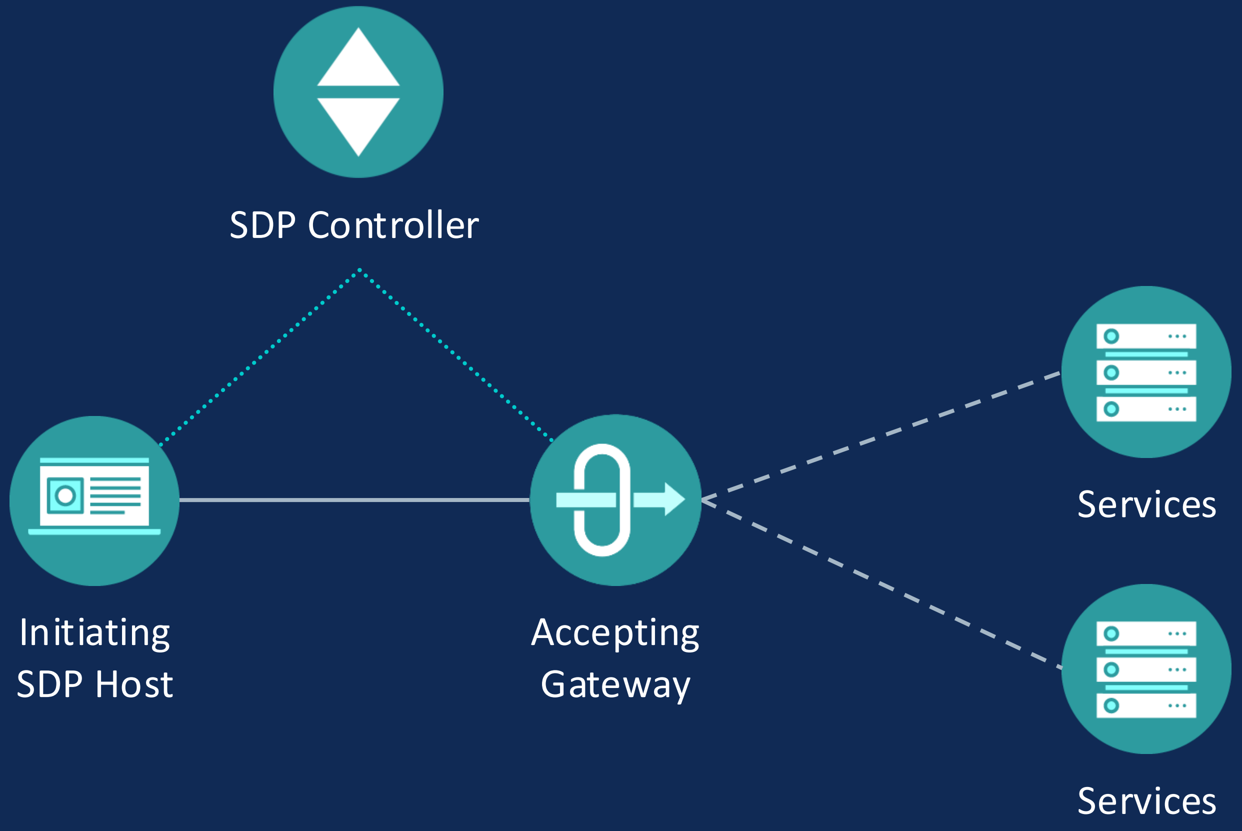

SDP is an approach to Zero Trust

- Controllers online
- Mutual TLS to controller
- Mutual TLS to controller
- List of authorized accepting hosts determined
- Accept communication from initiating host
- Receive list of IP’s of accepting hosts
- Mutual TLS tunnels

········ Control Channel

—— Secure Data Channel

- - - Insecure Data Channel

BUT, how do we write the rules?

Detailed entitlement matrices are a control, with multiple attributes as input for access decisions.

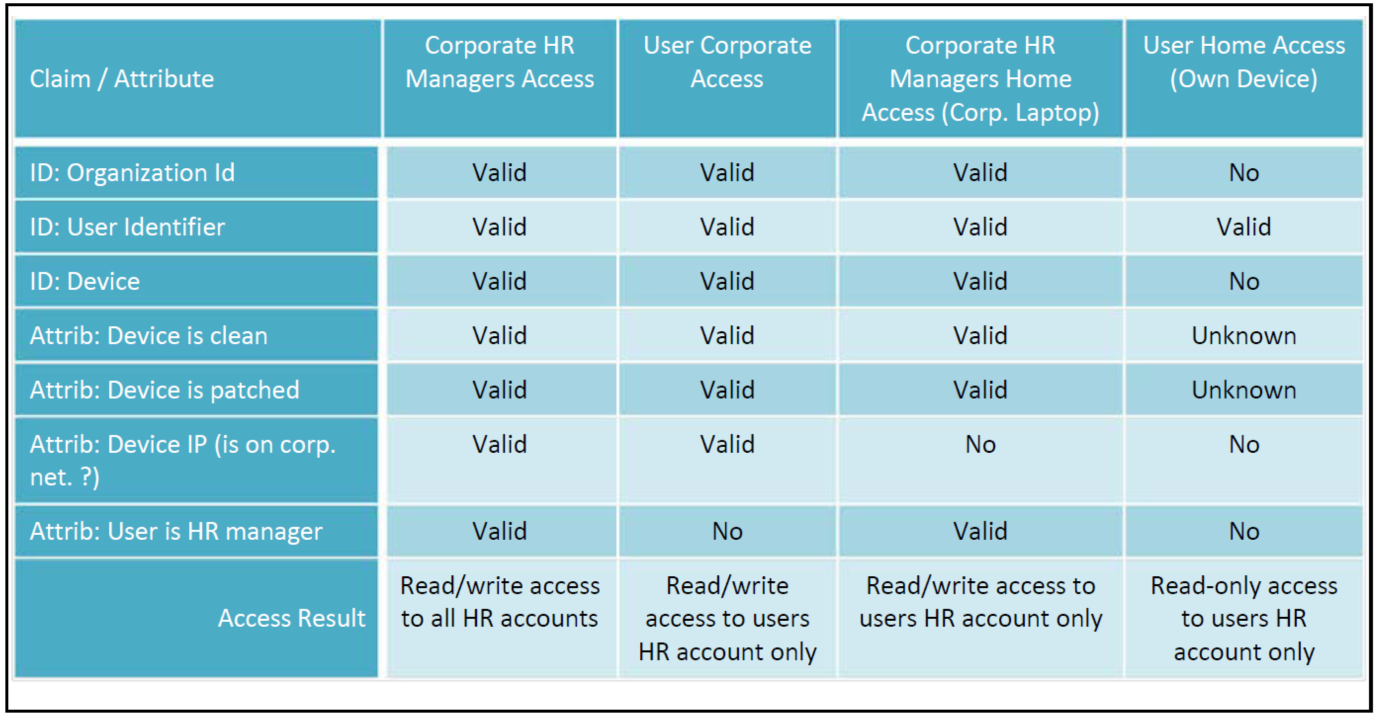

Claim / Attribute	Corporate HR Managers Access	User Corporate Access	Corporate HR Managers Home Access (Corp. Laptop)	User Home Access (Own Device)
ID: Organization Id	Valid	Valid	Valid	No
ID: User Identifier	Valid	Valid	Valid	Valid
ID: Device	Valid	Valid	Valid	No
Attrib: Device is clean	Valid	Valid	Valid	Unknown
Attrib: Device is patched	Valid	Valid	Valid	Unknown
Attrib: Device IP (is on corp. net. ?)	Valid	Valid	No	No
Attrib: User is HR manager	Valid	No	Valid	No
Access Result	Read/write access to all HR accounts	Read/write access to users HR account only	Read/write access to users HR account only	Read-only access to users HR account only

Sample HR application entitlement matrix. From CSA Guidance version 3.
Read this as follows. Check all ID and Attributes (Valid or not). Select first column that matches. The bottom row will give the authorization verdict.

Kipling Method

- The Kipling Method in the context of Zero Trust refers to using Rudyard Kipling's "six honest serving men" (Who, What, When, Where, Why, and How) as a framework for creating Zero Trust policies.

NSTAC report: Kipling Method

- A method for Zero Trust policy creation.
- A Layer 7 (application) technology determines what traffic can transit the micro-perimeter at any point in time and prevents unauthorized access to the defined protect surface.
- Describes the Who, What, When, Where, Why, and How of resource access:
  - Who should be allowed to access a resource?
  - What application is the asserted identity allowed to use to access the resource?
  - When is the asserted identity allowed to access the resource?
  - Where is the resource located?
  - Why is the user (the Who) allowed to access the resource?
  - How should traffic be processed as it accesses a resource?

That is all it says. No examples.

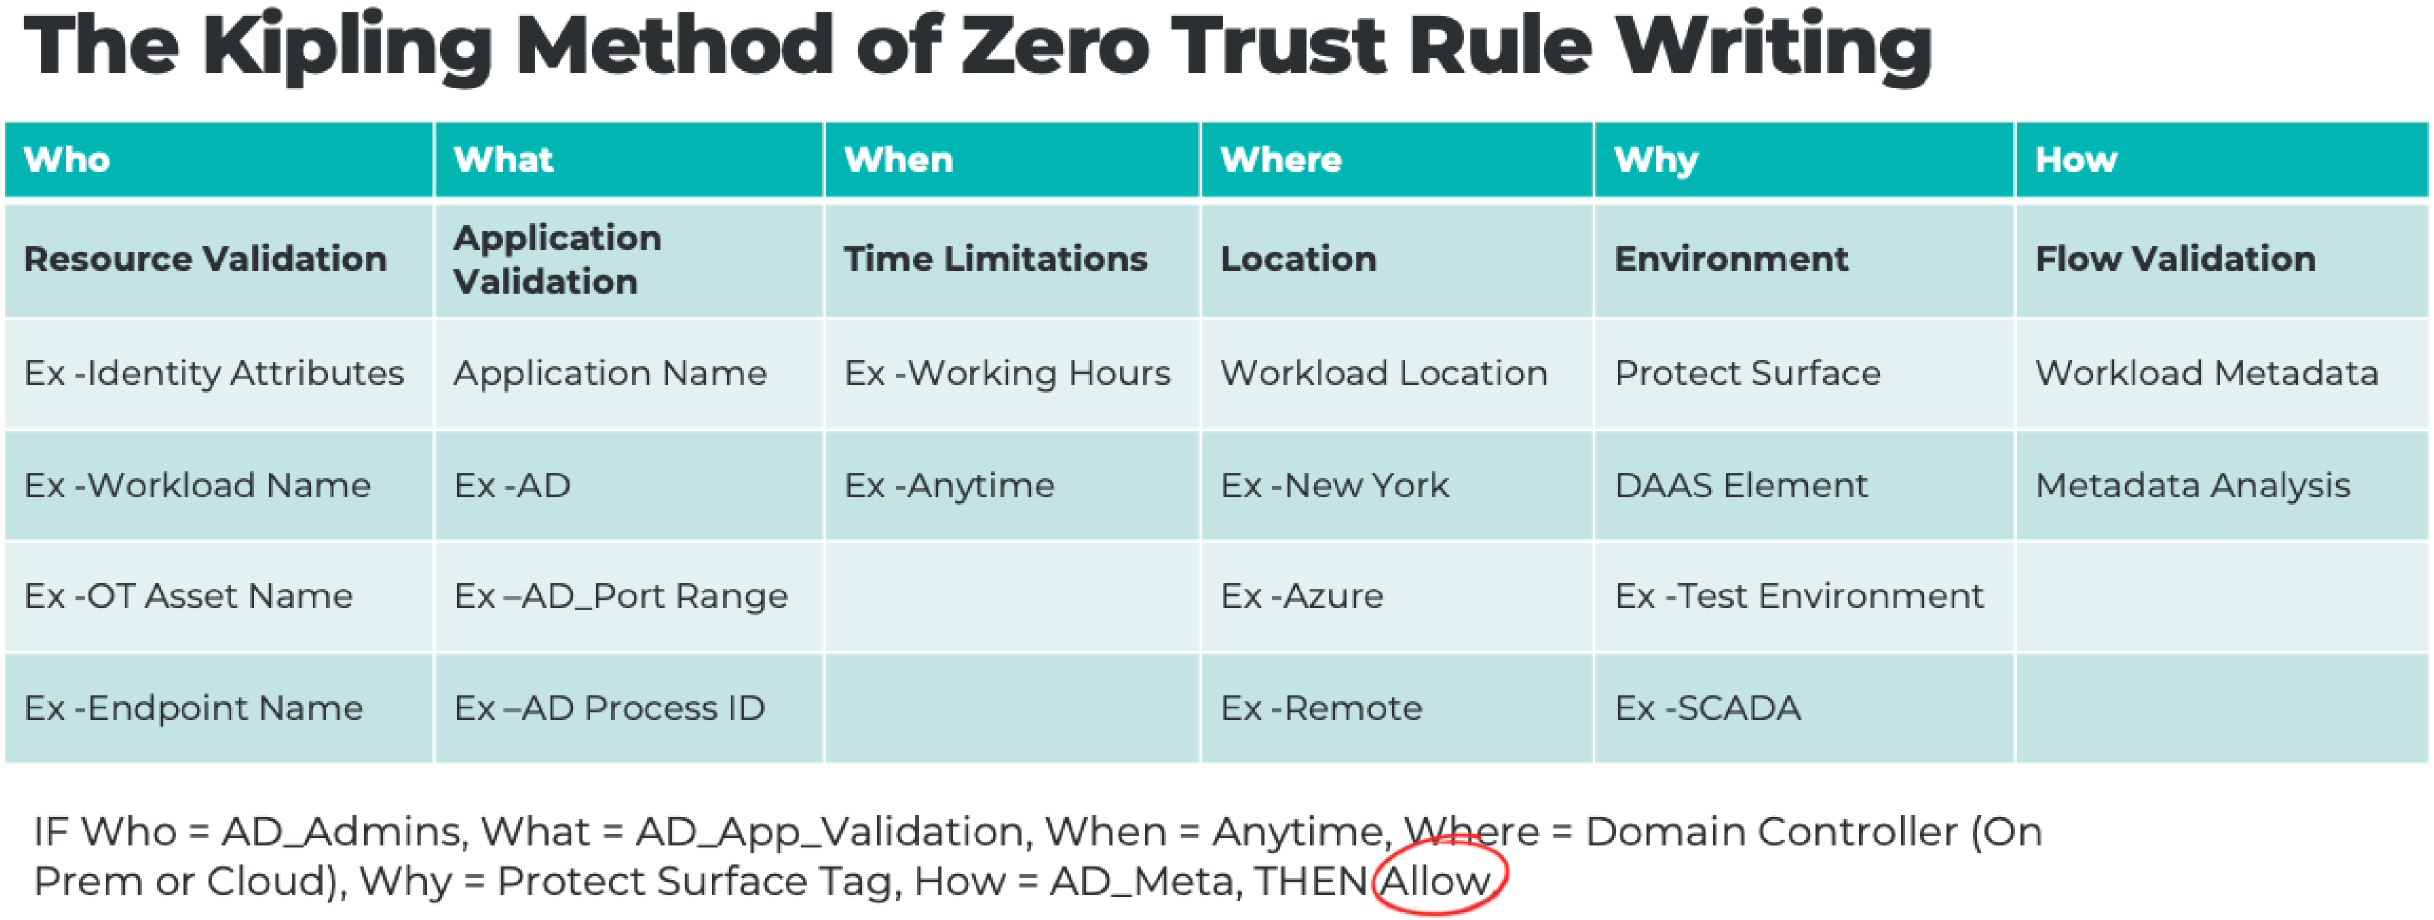

The Kipling Method of Zero Trust Rule Writing

Who	What	When	Where	Why	How
Resource Validation	**Application Validation**	**Time Limitations**	**Location**	**Environment**	**Flow Validation**
Ex -Identity Attributes	Application Name	Ex -Working Hours	Workload Location	Protect Surface	Workload Metadata
Ex -Workload Name	Ex -AD	Ex -Anytime	Ex -New York	DAAS Element	Metadata Analysis
Ex -OT Asset Name	Ex –AD_Port Range		Ex -Azure	Ex -Test Environment	
Ex -Endpoint Name	Ex –AD Process ID		Ex -Remote	Ex -SCADA	

IF Who = AD_Admins, What = AD_App_Validation, When = Anytime, Where = Domain Controller (On Prem or Cloud), Why = Protect Surface Tag, How = AD_Meta, THEN Allow.

Zero Trust Masterclass Zurich 2024.
John Kindervag, and Illumino

OK, but where do we enforce this?
'why' not consistent

Zero Trust Security by Jason Garbis

This book identifies:

- Subject criteria (who)
- Action (what)
- Target (where)
- Condition (when, who, …)

Example Garbis

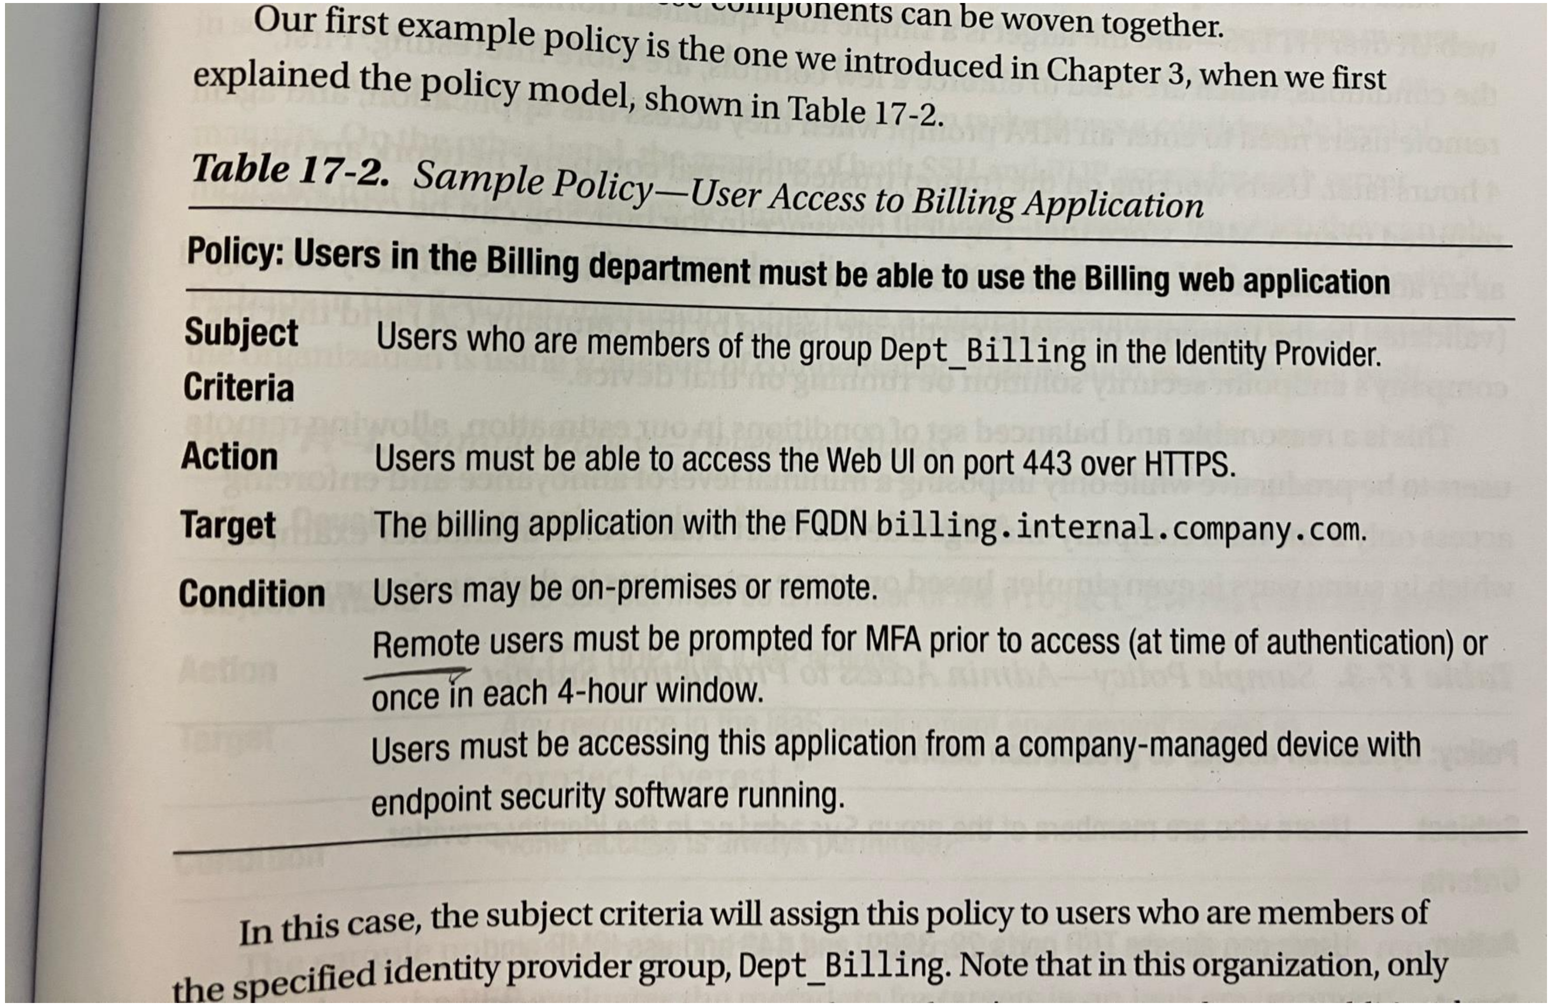

components can be woven together.

Our first example policy is the one we introduced in Chapter 3, when we first explained the policy model, shown in Table 17-2.

Table 17-2. Sample Policy—User Access to Billing Application

Policy: Users in the Billing department must be able to use the Billing web application	
Subject Criteria	Users who are members of the group `Dept_Billing` in the Identity Provider.
Action	Users must be able to access the Web UI on port 443 over HTTPS.
Target	The billing application with the FQDN `billing.internal.company.com`.
Condition	Users may be on-premises or remote. Remote users must be prompted for MFA prior to access (at time of authentication) or once in each 4-hour window. Users must be accessing this application from a company-managed device with endpoint security software running.

In this case, the subject criteria will assign this policy to users who are members of the specified identity provider group, `Dept_Billing`. Note that in this organization, only

In more detail

Element	Description
Subject	The entity performing (initating) actions. Subject must be authenticated identities
Criteria	Criteria designate the subjects to whom this policy applies
Action	The activity performed by the subject to whom this policy applies
Target	The object (resource) that the action is being performed upon
Condition	The circumstances under which the subject is permitted to perform the action upon the target

Atttributes, which are input to criteria and conditions, can be derived from identities, devices, and target resources, amongs others.

The overarching question then always is: how do we know these things, at the moment of truth (i.e. the PEP).

Flashback time

```
-rw-r--r-- 1 root root 1701 Aug 27  2023 /etc/passwd
```

Element	Description
Subject	Effective Userid, Grouid
Criteria	User id, group membership
Action	Read, write, execute
Target	/etc/passwd
Condition	The circumstances under which the subject is permitted to perform the action upon the target

Kubernetes example: allow rule

```
# networkpolicy2.yaml
kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: web-allow
spec:
 podSelector:
  matchLabels:
   app: busybox1
 ingress:
  - from:
    - podSelector:
        matchLabels:
          app: busybox2
```

Rule name

Destination

From all pods that match this label, deny the rest

PEP and PDP are both inside Kubernetes, though the Policy Administrator could reside in e.g. ArgoCD

Unifying the terminology, …

- CCZT: talks about PEP & PDP
- SDP talks about Controllers and Accepting Hosts
- What else…
    - Kipling? Attributes?

- Zero Trust Architecture as defined by NIST, originally in the SP 800-207, and then further elaborated in SP 1800-35B draft: talk about PEP PDP etc.

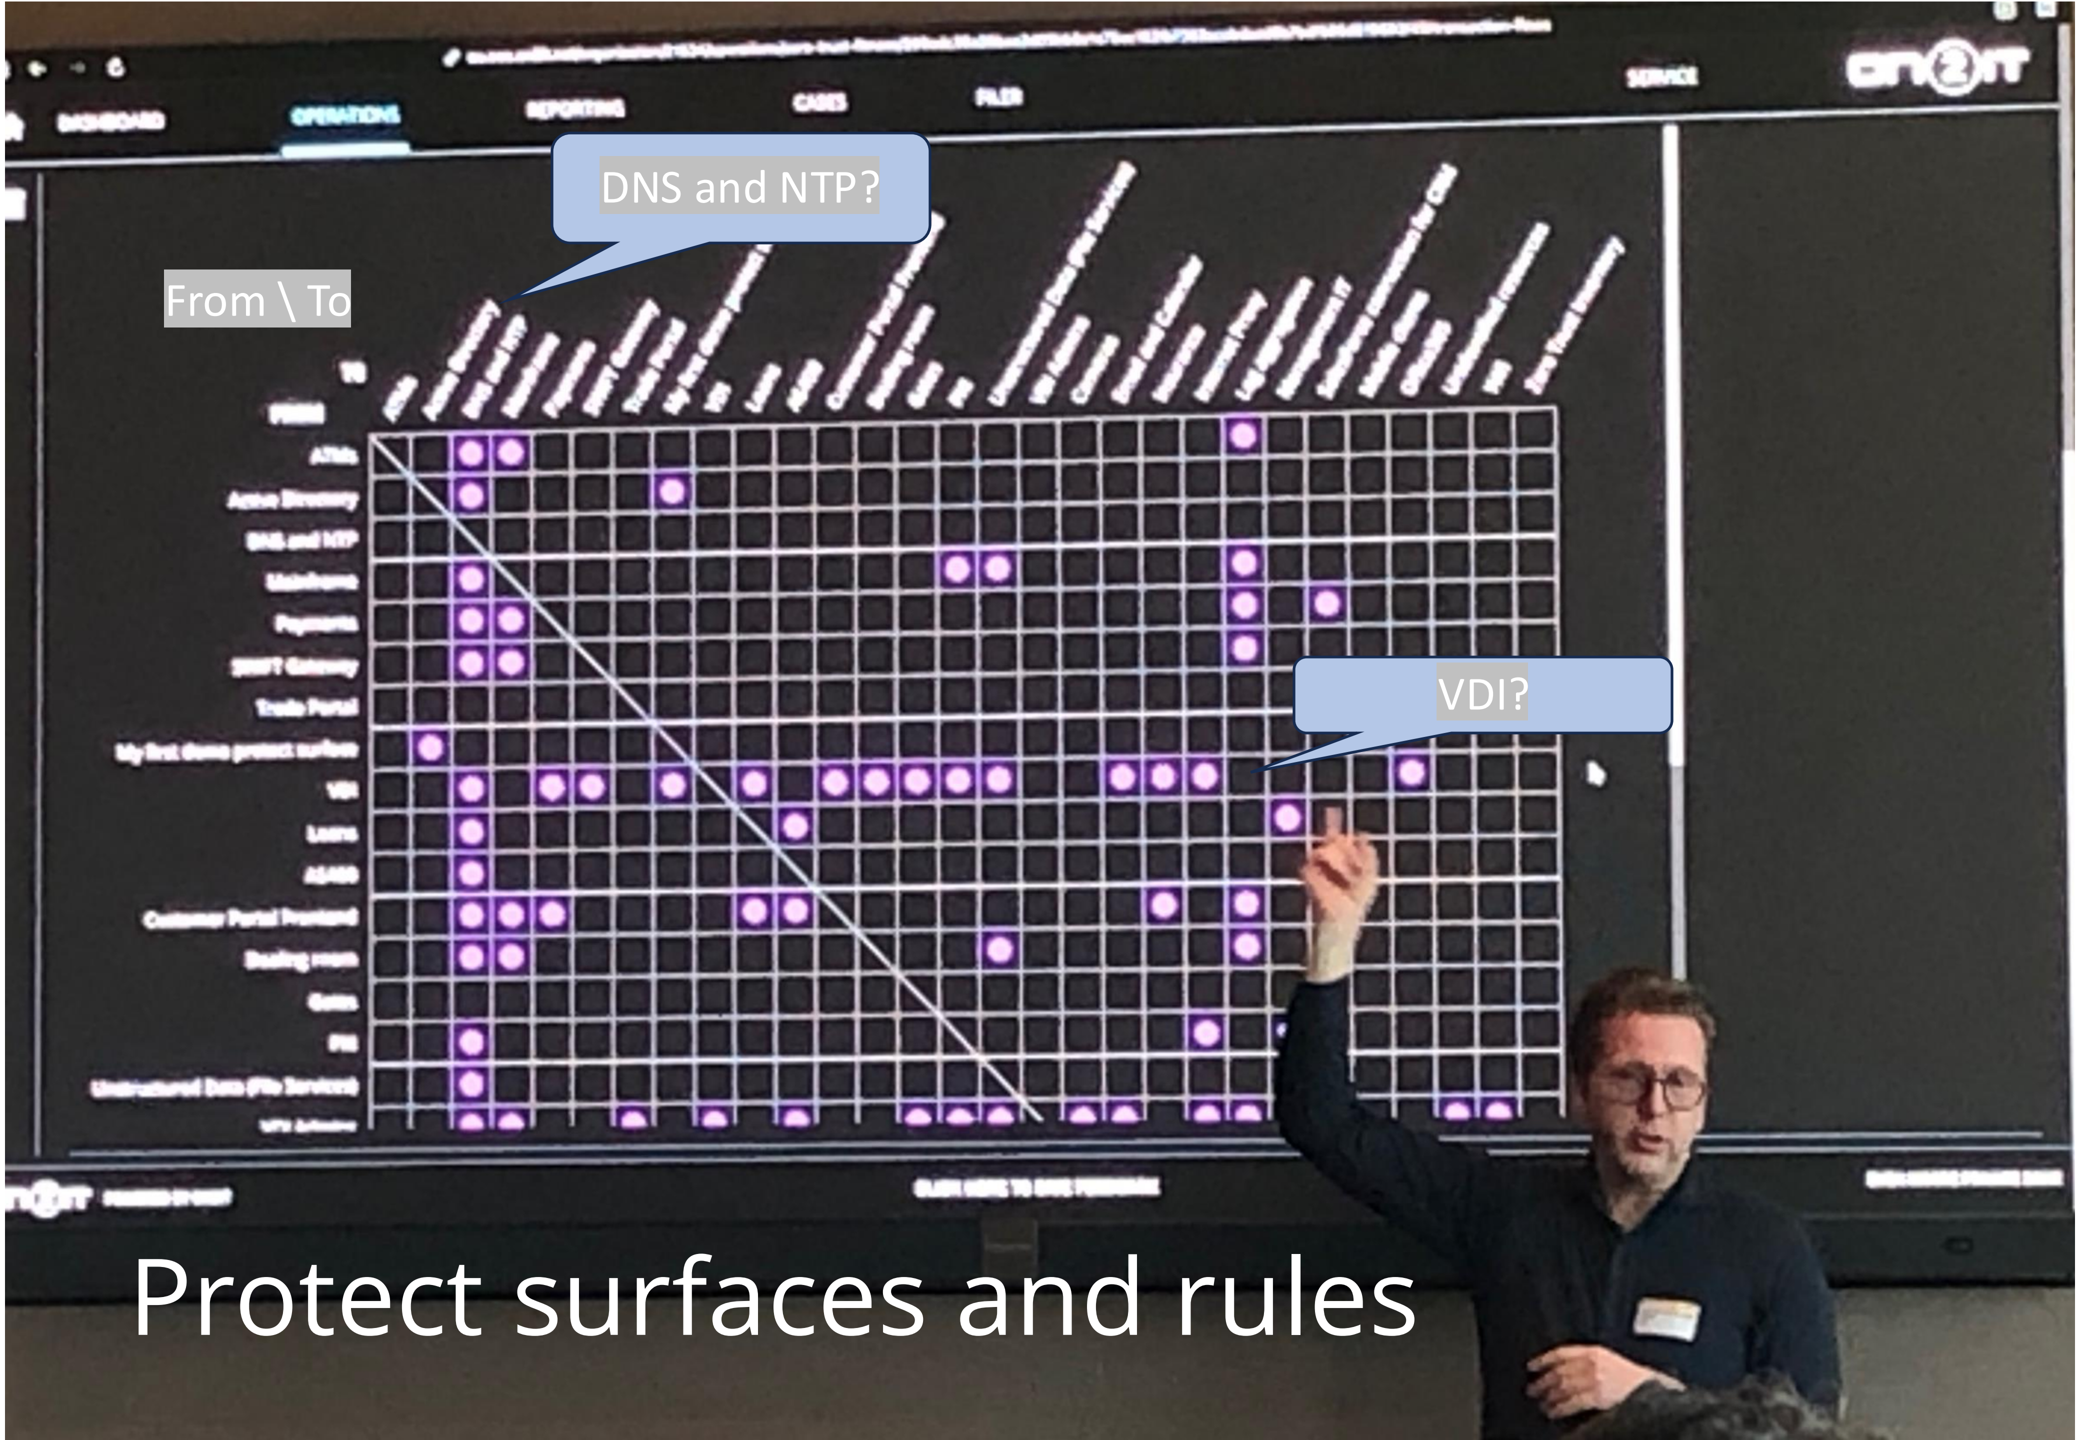

From \ To

DNS and NTP?

VDI?

Protect surfaces and rules

More rules of thumb?

- Get rid of specific deny rules.
  Only keep deny all. If you ever see the word *block* or *deny* without ‘all’, you have an opportunity to reduce implicit trust.
- Monitor the blocked stuff.
  Don’t trust a test (or rule) that you have not seen fail (principles of test-driven design).

AI and Zero Trust

- Never trust what comes out of an LLM
- Never trust an LLM with any credentials

- E.g. My claude code accesses my github only through gh, and cannot access the credentials that gh uses for that.

Retrofitting Zero Trust

Note, each control boundary can be a PEP (Policy Enforcement Point)

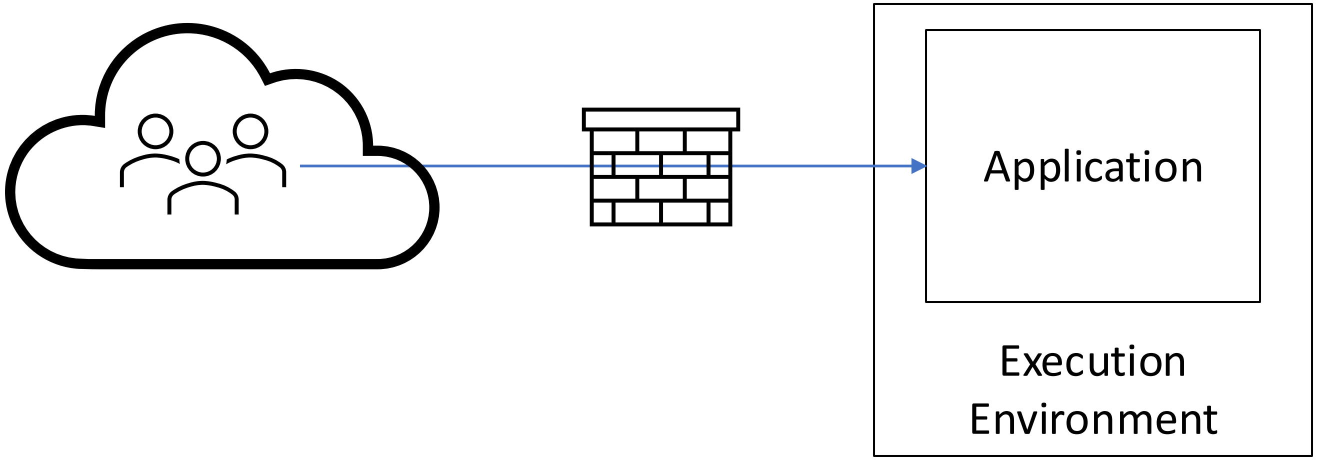

Additional controls - exfil

- Who? 4-eyes for large downloads?
- When? Time of day
- Where? Source IP filtering?
- Why? Only allow specific users

Note, each control boundary can be a PEP (Policy Enforcement Point)

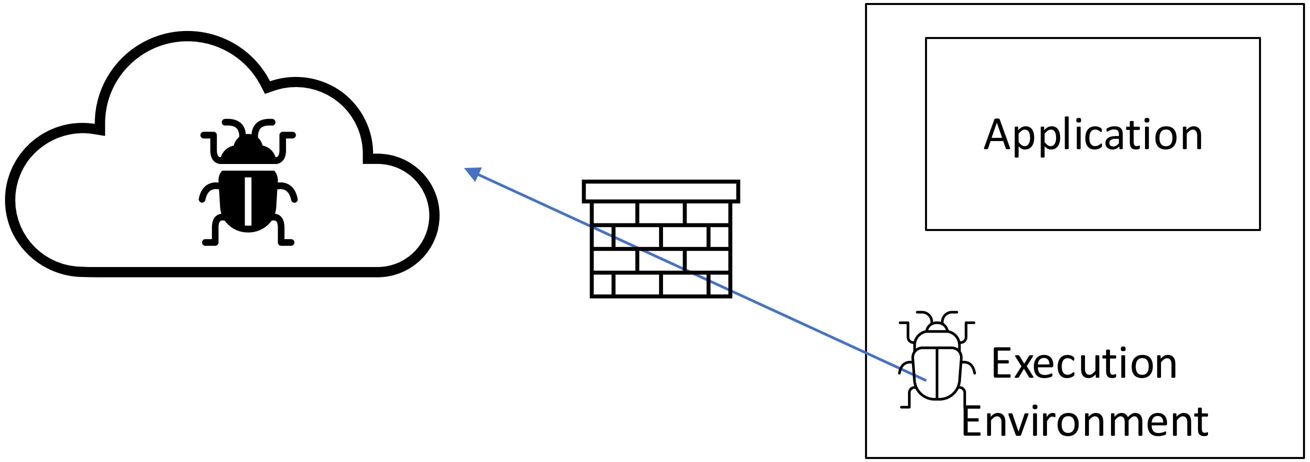

Additional controls – reverse allow

- Who? To which server?
- When? Note: logging and monitoring happens all the time
- Where? In the firewall
- Why? Exfil of sensitive data

Implementation steps

- The CCZT courseware (free) has a lot of good stuff here
- Review also Maturity Models

Identity

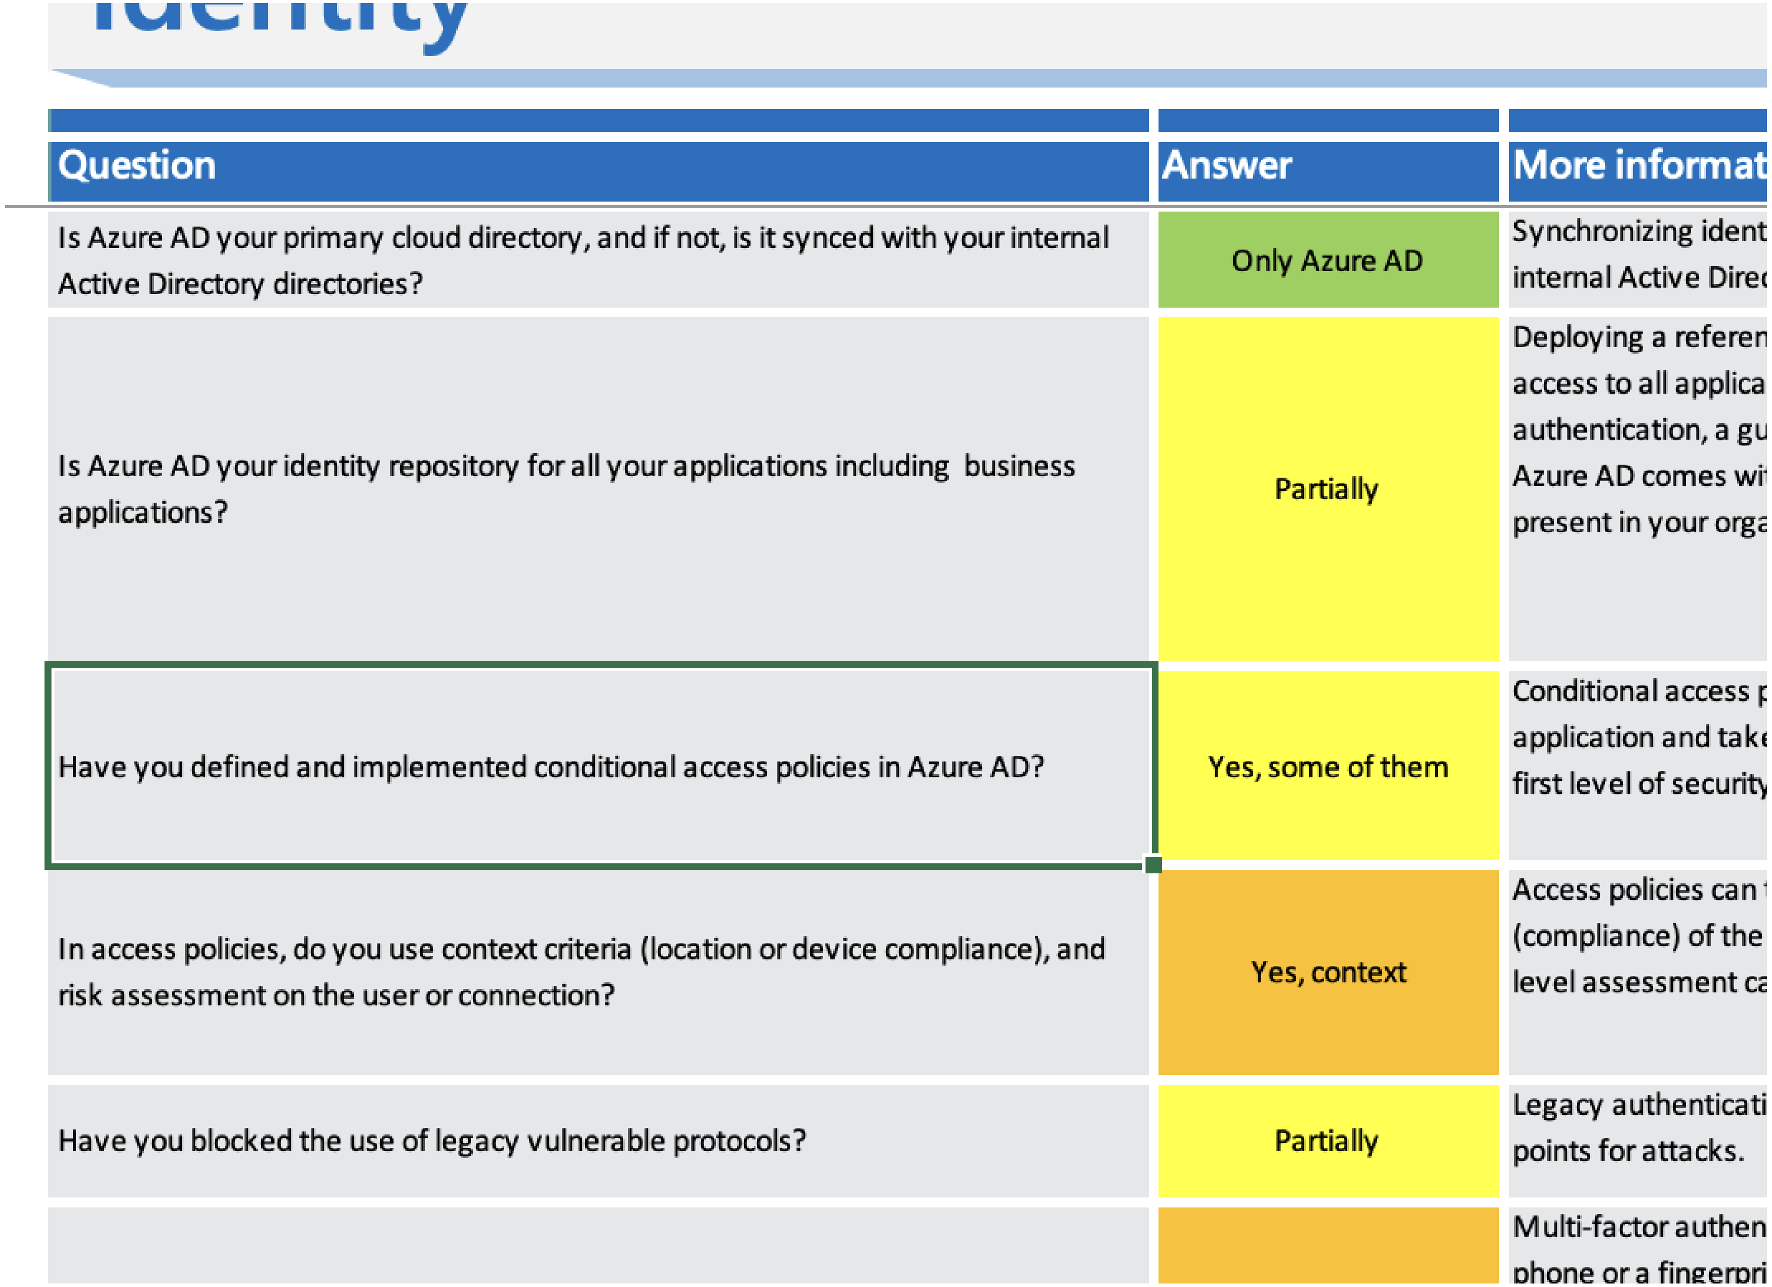

Question	Answer	More informat
Is Azure AD your primary cloud directory, and if not, is it synced with your internal Active Directory directories?	Only Azure AD	Synchronizing ident internal Active Direc
Is Azure AD your identity repository for all your applications including business applications?	Partially	Deploying a referen access to all applica authentication, a gu Azure AD comes wi present in your orga
Have you defined and implemented conditional access policies in Azure AD?	Yes, some of them	Conditional access p application and take first level of security
In access policies, do you use context criteria (location or device compliance), and risk assessment on the user or connection?	Yes, context	Access policies can (compliance) of the level assessment ca
Have you blocked the use of legacy vulnerable protocols?	Partially	Legacy authenticati points for attacks.
		Multi-factor authen phone or a fingerpri

Maturity scores

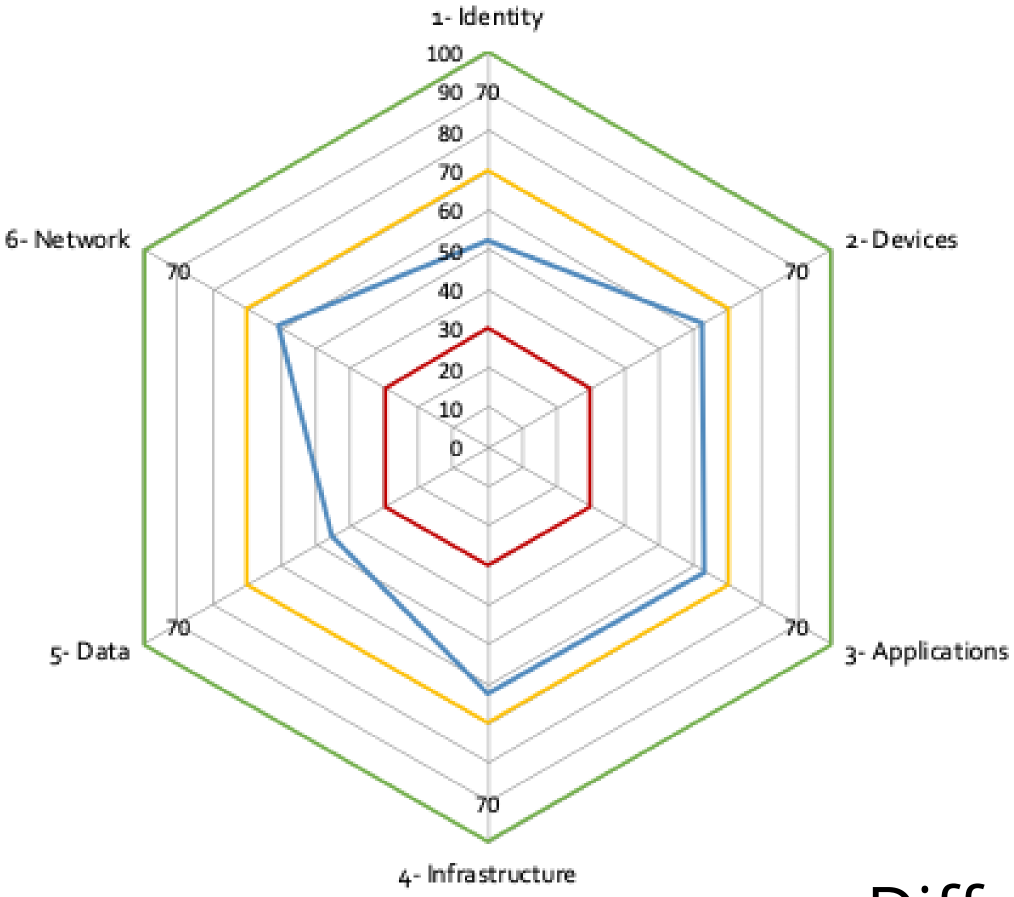

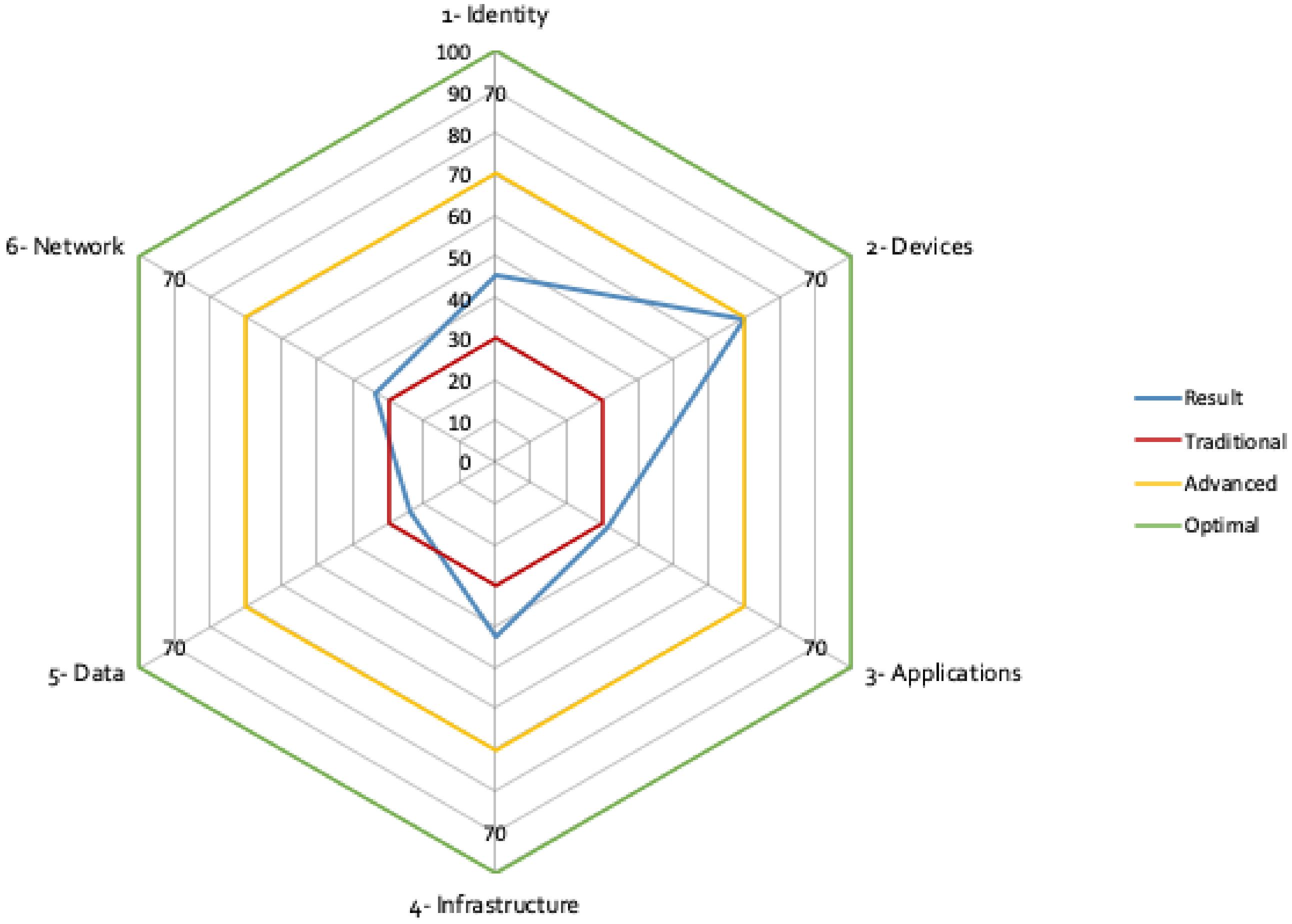

Different organisations score differently, 'Data' is consistently low

ClubCloudComputing

Tactics for Zero Trust

ZT Design Principles

- Focus on business outcomes
- Design from the inside out
- Determine who/what needs access
- Inspect and log key traffic

Foundational Step of ZT Design

- **Step 1:** Define Your Protect Surface(s)
- **Step 2:** Map the Transaction Flows
- **Step 3:** Build a Zero Trust Architecture (ZTA)
- **Step 4:** Create ZT Policy
- **Step 5:** Monitor and Maintain the Network

Zero Trust – risks and fears

- Control loss to Security Dictators?
- Organisational paralysis?
- We never had control?
- It will work as designed, dehumanizing the collaboration between workers?

Zero Trust for Cloud Infrastructure & Networks

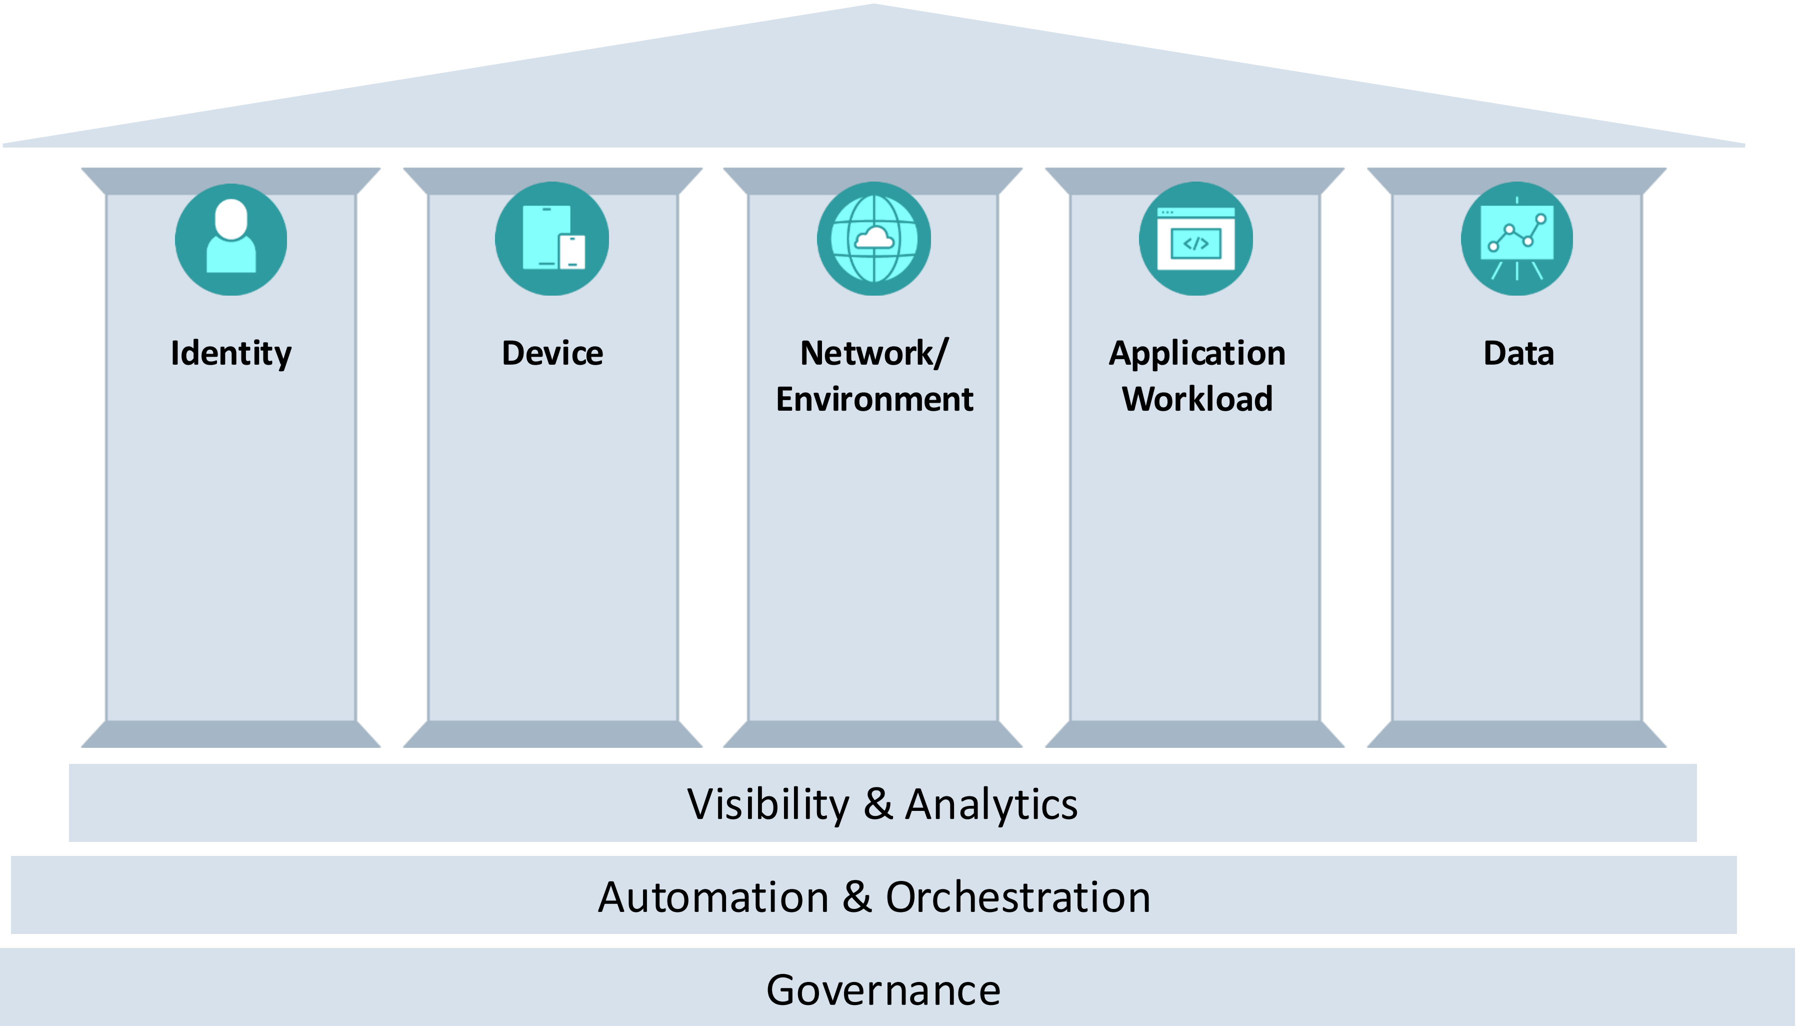

Further reading

CCZT™

Certificate of
Competence in Zero Trust

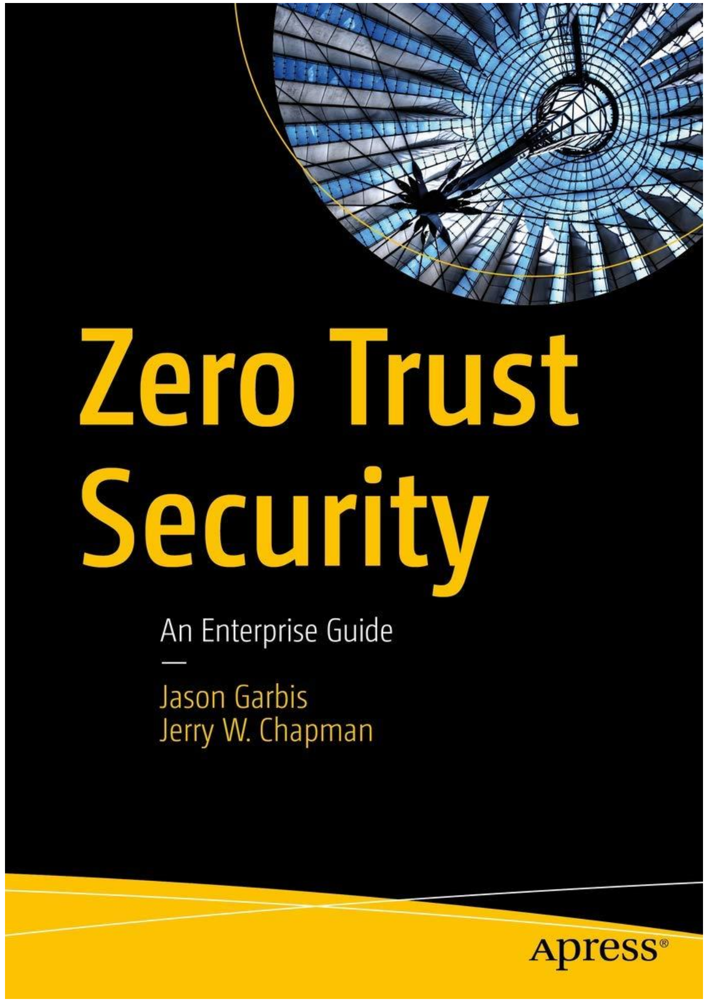

Zero Trust Security

An Enterprise Guide

—

Jason Garbis
Jerry W. Chapman

Apress®

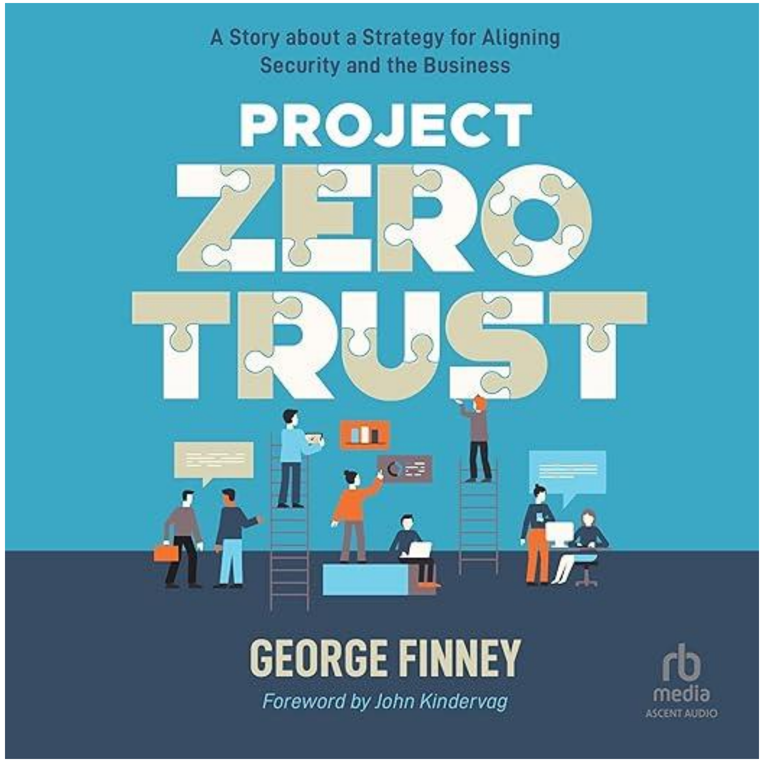

A Story about a Strategy for Aligning Security and the Business

PROJECT ZERO TRUST

GEORGE FINNEY

Foreword by John Kindervag

rb media
ASCENT AUDIO

Up next

- **Using deployment diagrams to explain architecture and security to everybody**
- **Monday 2025-08-11 20:00–20:25, Cassiopeia**

- **Thomas Fricke:**
- **Can we trust the Zero in Zero trust?**
- **Tuesday 2025-08-12 15:00–15:50, Delphinus**

Shameless plug

Peter van Eijk
p@d1g.nl

Thank you!

YouTube channel: ClubCloudComputing

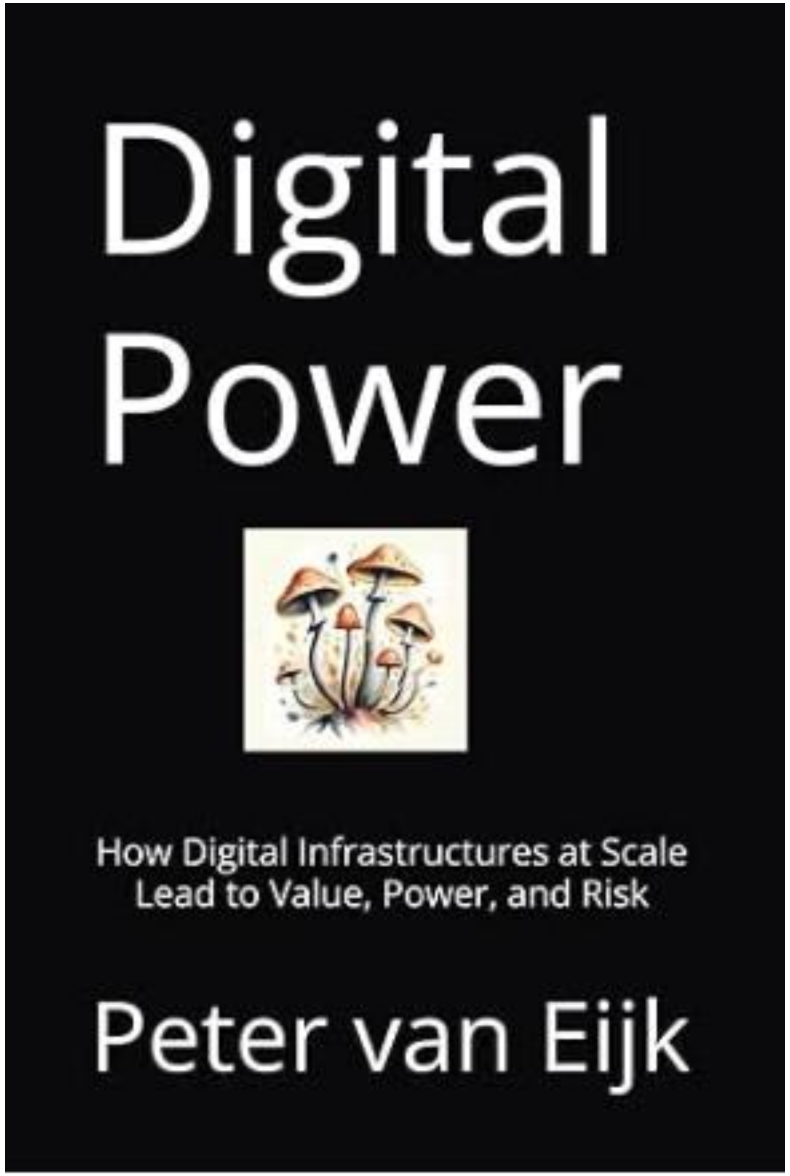

Digital Power

How Digital Infrastructures at Scale Lead to Value, Power, and Risk

Peter van Eijk

ClubCloudComputing