ZERO TRUST

WHAT IT IS
WHY WE NEED IT
HOW WE GET THERE

PAULA QUEIROZ

PART 1

INTRODUCTION

AGENDA

- Understand **what Zero Trust is** and how it **differs from traditional infrastructure security**.
- **Contextualise** the need to think about having and implementing a Zero Trust strategy based on the security needs in **today's world**.
- Explore **industry-recognised standards, models and reference architectures** for Zero Trust.
- Look at **security controls and techniques** which will support a Zero Trust architecture.
- Explore expected **corporate difficulties** which will be encountered when planning a shift towards Zero Trust.
- Have a **practical look** at how Zero Trust can be implemented.

INTRODUCTION

- Zero Trust is a security framework that enforces the requirement for **authentication**, **authorization**, and **continuous security configuration** as well as **posture validation** prior to access to all enterprise applications and data.

- Zero Trust aims to **secure infrastructure and data in today's world**, which includes the **decentralization of the hosting of assets and people**, with the extension or migration of the traditional data center to the cloud, employee mobility, etc.

- There are actually a number of vendor-neutral standards provided by several recognized organizations that can guide us towards a Zero Trust strategy. Here we will focus heavily on **NIST 800-207 and the American Department of Defense’s (DoD) reference architecture for Zero Trust**.

- Very important to keep in mind from the get-go is that fact that, more than a shift in technology, **Zero Trust is a shift in mindset** which, more than representing the challenge to change technology, represents the challenge of changing the way we think about technology, and how we secure it.

SUBJECT AND OBJECT (OR RESOURCE)

PRE-REQ

As we will see later (and we will refer to them as such), the understanding of the concepts of **subject** and **object** are paramount to understanding Zero Trust, as we will be continuously talking about the relationship between the two. Zero Trust will always focus on the access and interaction of a subject with an object.

Using the CISSP as a reference:

- **A subject will always be something or someone (active) which manipulates an object (passive).** Subjects will most often be people, but they can often be software or devices.
- **An object (or resource) is something (passive) which is manipulated by a subject (active).** Objects will normally ultimately refer to data (in either soft or hard format), but they can also be pieces of software, equipment, etc.

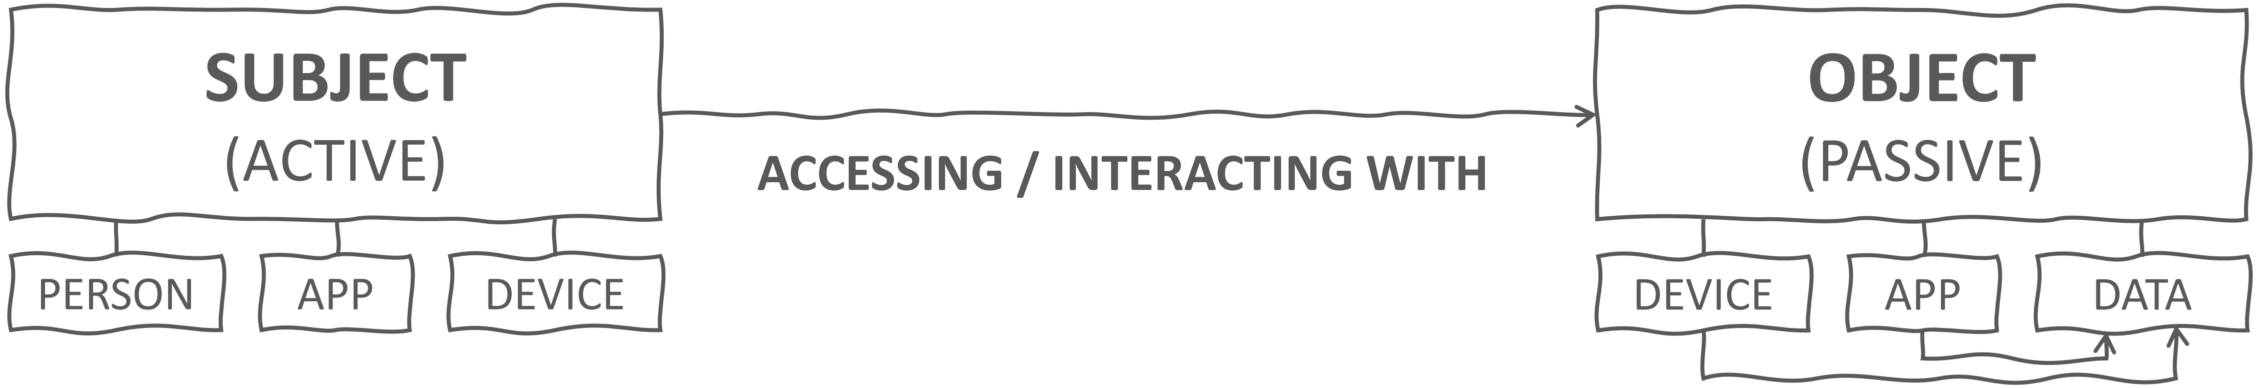

PART 2

PAST VS PRESENT

PART 2.1

PAST VS PRESENT

THE PAST

PAST VS PRESENT: TRADITIONAL NETWORK SECURITY

- Until recent years, we have widely accepted the idea what creating **a hardened security perimeter to host all our core assets** was the best way to secure those assets.
- We employed often complex defense in depth strategies that revolved around **securing that core perimeter with multiple layers of security** - but always assuming that **whatever was inside that static perimeter was safe** - and **whoever had access to that perimeter was to be trusted**.
- Note: zero trust falls in the category and obeys the security principle of defense in depth, but it involves a foundational paradigm and practical implementation rather different from what we would traditionally find.

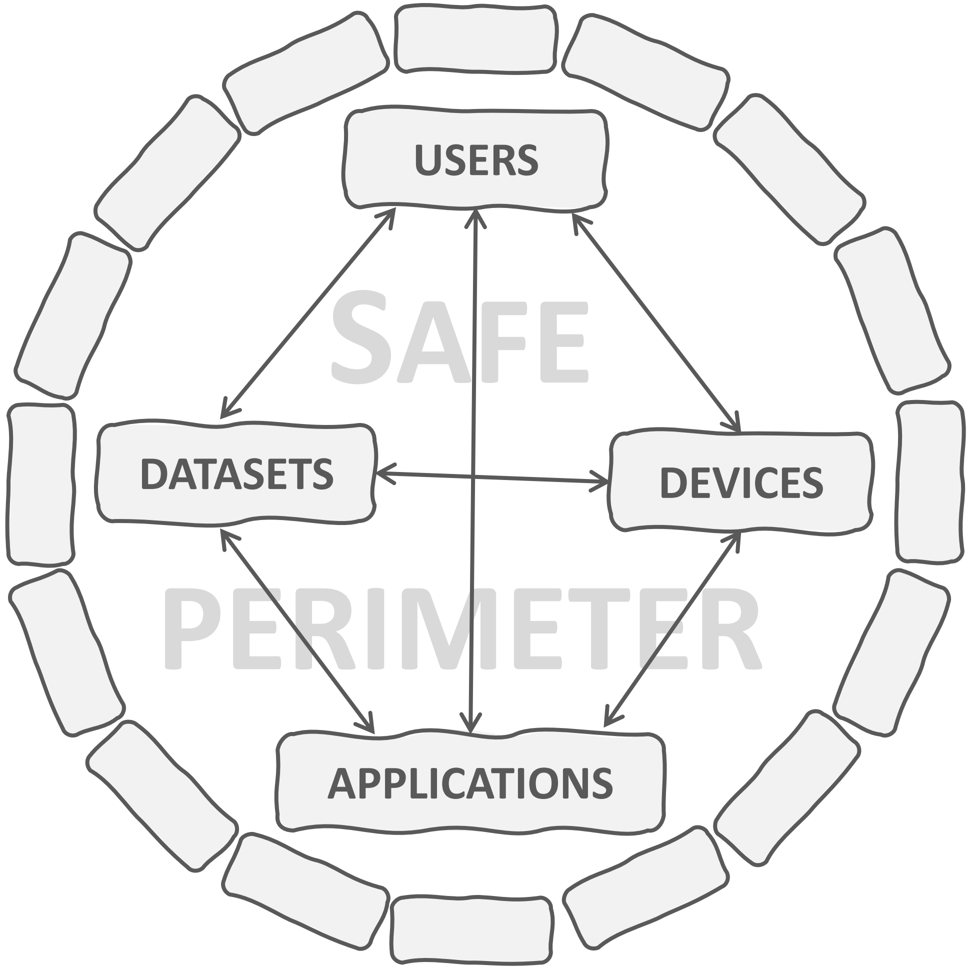

PAST VS PRESENT: TRADITIONAL NETWORK SECURITY

There are, some problems with the traditional approach, that render it unfit for today’s world:

- ✕ The traditional model assumed that **enterprise assets were hosted on-premises**, normally in the corporate datacenter. This is more often than not no longer the case. Our assets today tend to be **fully decentralised** and **hosted in hybrid environments**, often multi-cloud, and we can expect it to be increasingly so from here onwards.
- ✕ It also assumed (quite wrongly, as the statistics show us) that security **threats came from the outside**. However, we know it for a fact that this is not always the case. Whether an insider is the root of a security breach intentionally or accidentally, the fact is that **a significant number of breeches occur from the inside**.
- ✕ Add **workforce mobility**, **BYOD** and the increasing need for staff to **access their work systems and data remotely and from multiple locations** (which was drastically accelerated by the COVID-19 pandemic), and then add the security exploits that target VPN security and the **difficulties scaling up VPN solutions** on short notice, and it becomes clear that this historical way of securing our corporate assets is no longer fit for purpose.

PAST VS PRESENT: TRADITIONAL NETWORK SECURITY

So, this leaves us with the need to find a way to secure our infrastructure and data in a way that meets certain criteria:

- ✓ Our **assets must be protected regardless of where they are hosted** - these can be traditional on-premises datacenters or cloud providers, PaaS, SaaS or IaaS.
- ✓ We must be ready to **face threats that comes from the outside** as well as those that come from the inside, and insider threats can take many shapes: an insider may wish to do harm, their credentials may be stolen or, quite simply, they may make an honest but catastrophic mistake.
- ✓ Our staff must be able to **work securely outside the office premises, without difficulty doing so**.
- ✓ An **alternative to VPNs** (increasingly less secure and hard to scale massively at short notice) must be available to allow for secure remote work.
- ✓ **Data isn’t all the same** and, although corporate data always tends to have value to the enterprise (when it doesn’t, it should be destroyed), different data has different value and, therefore, **should be protected differently**.

This is where Zero Trust kicks in...

PART 2.2

PAST VS PRESENT

THE PRESENT

PAST VS PRESENT: ZERO TRUST

- The focus has changed from securing a perimeter (which, as we’ve discussed, has been dissolved), to **securing users, assets and resources**, by **continuously monitoring and authorising interactions** between them, as we start from the premise that **the network is always hostile**.
- Basically, we are now looking at **individual interactions between subjects and objects**, rather than subjects and perimeters.
- Trivia: the expressions “zero trust” and “zero trust architectures” were invented in 2010 by the analyst John Kindervag.

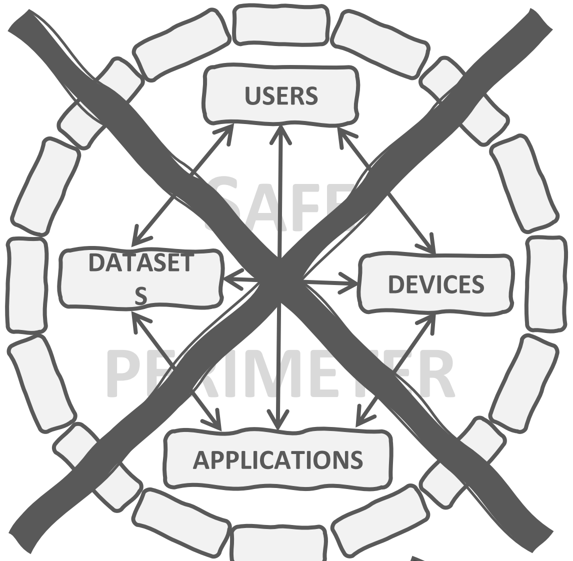

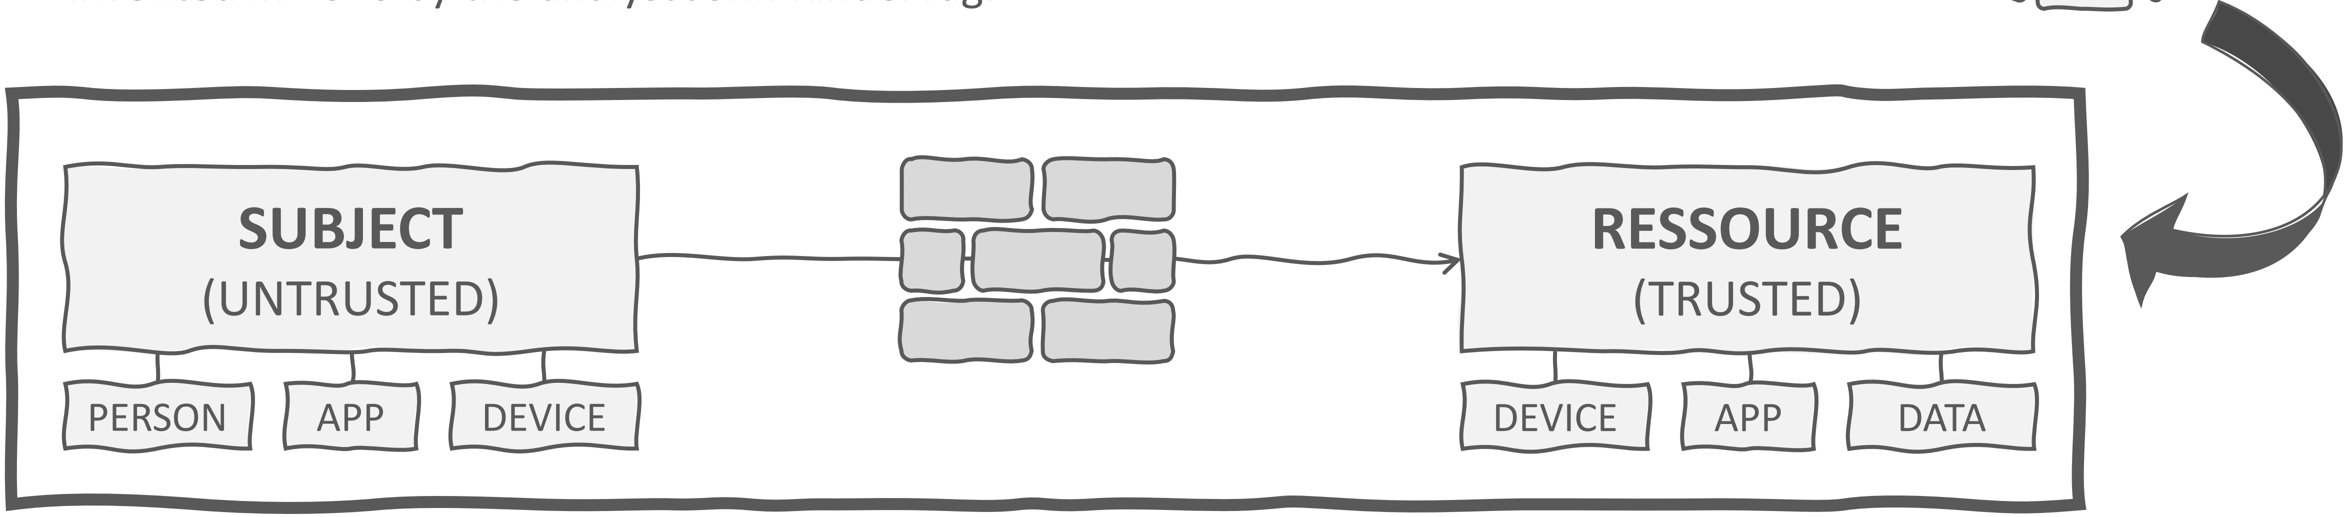

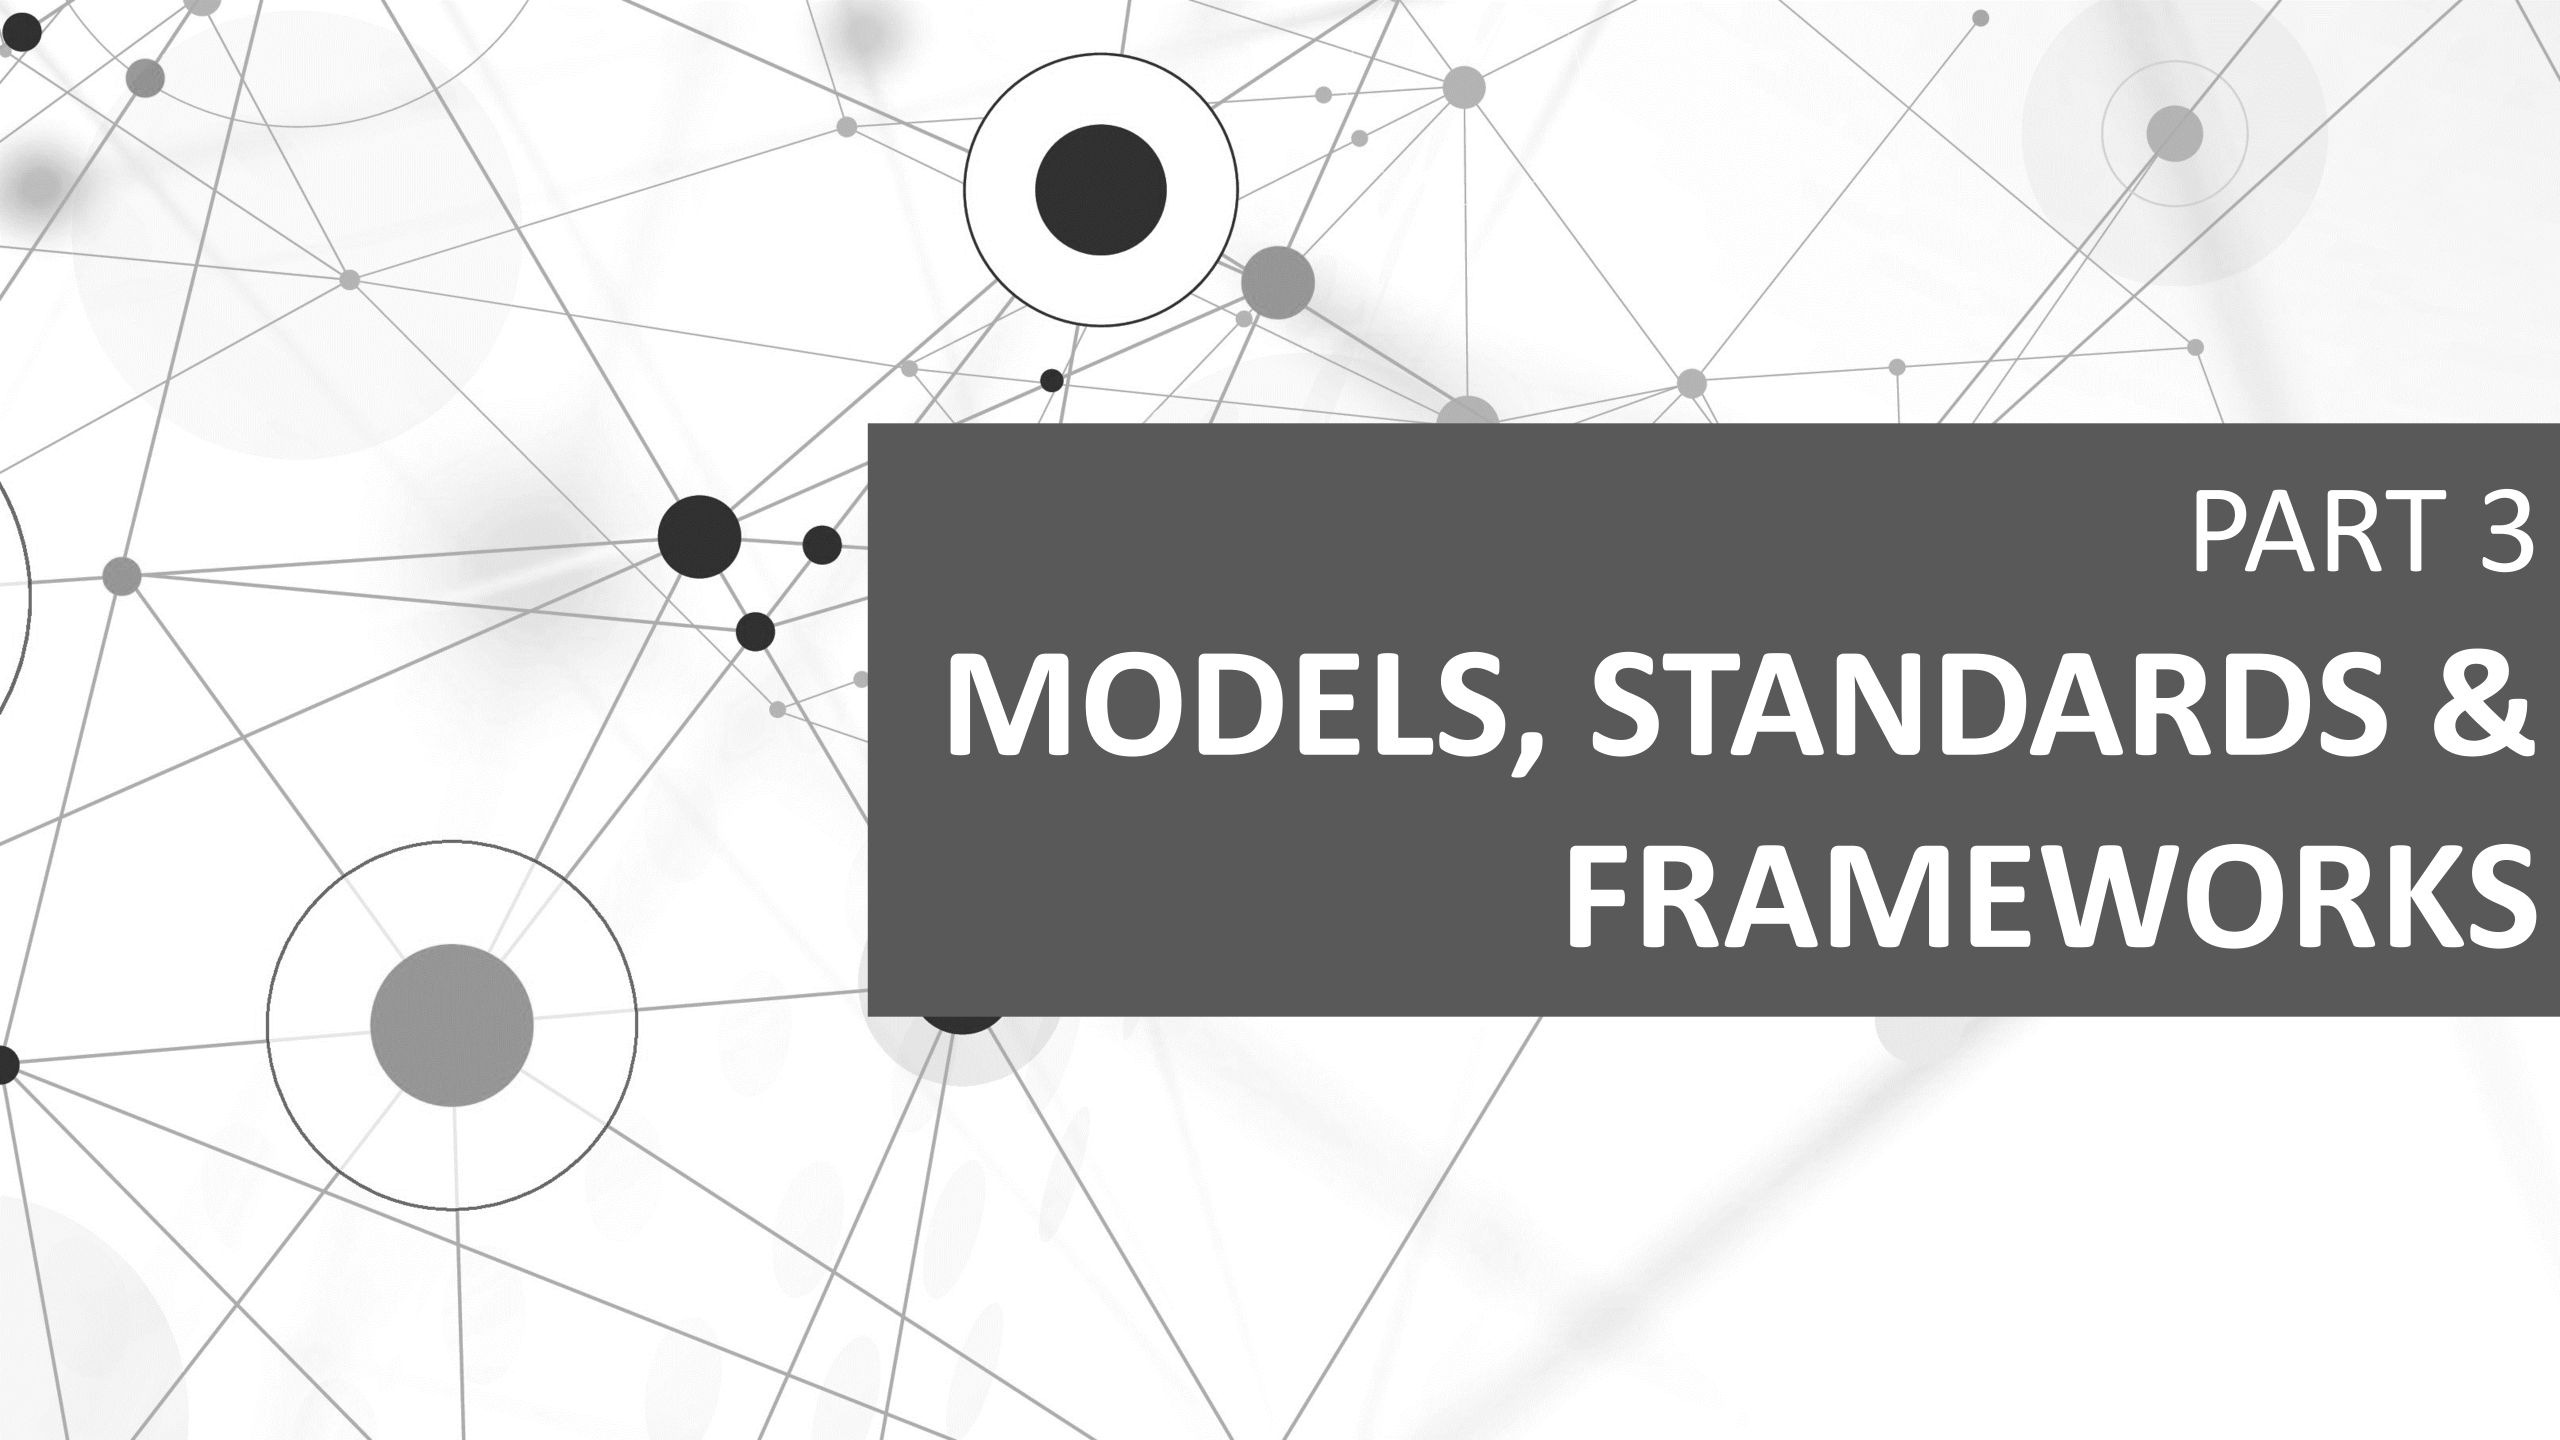

PART 3

MODELS, STANDARDS & FRAMEWORKS

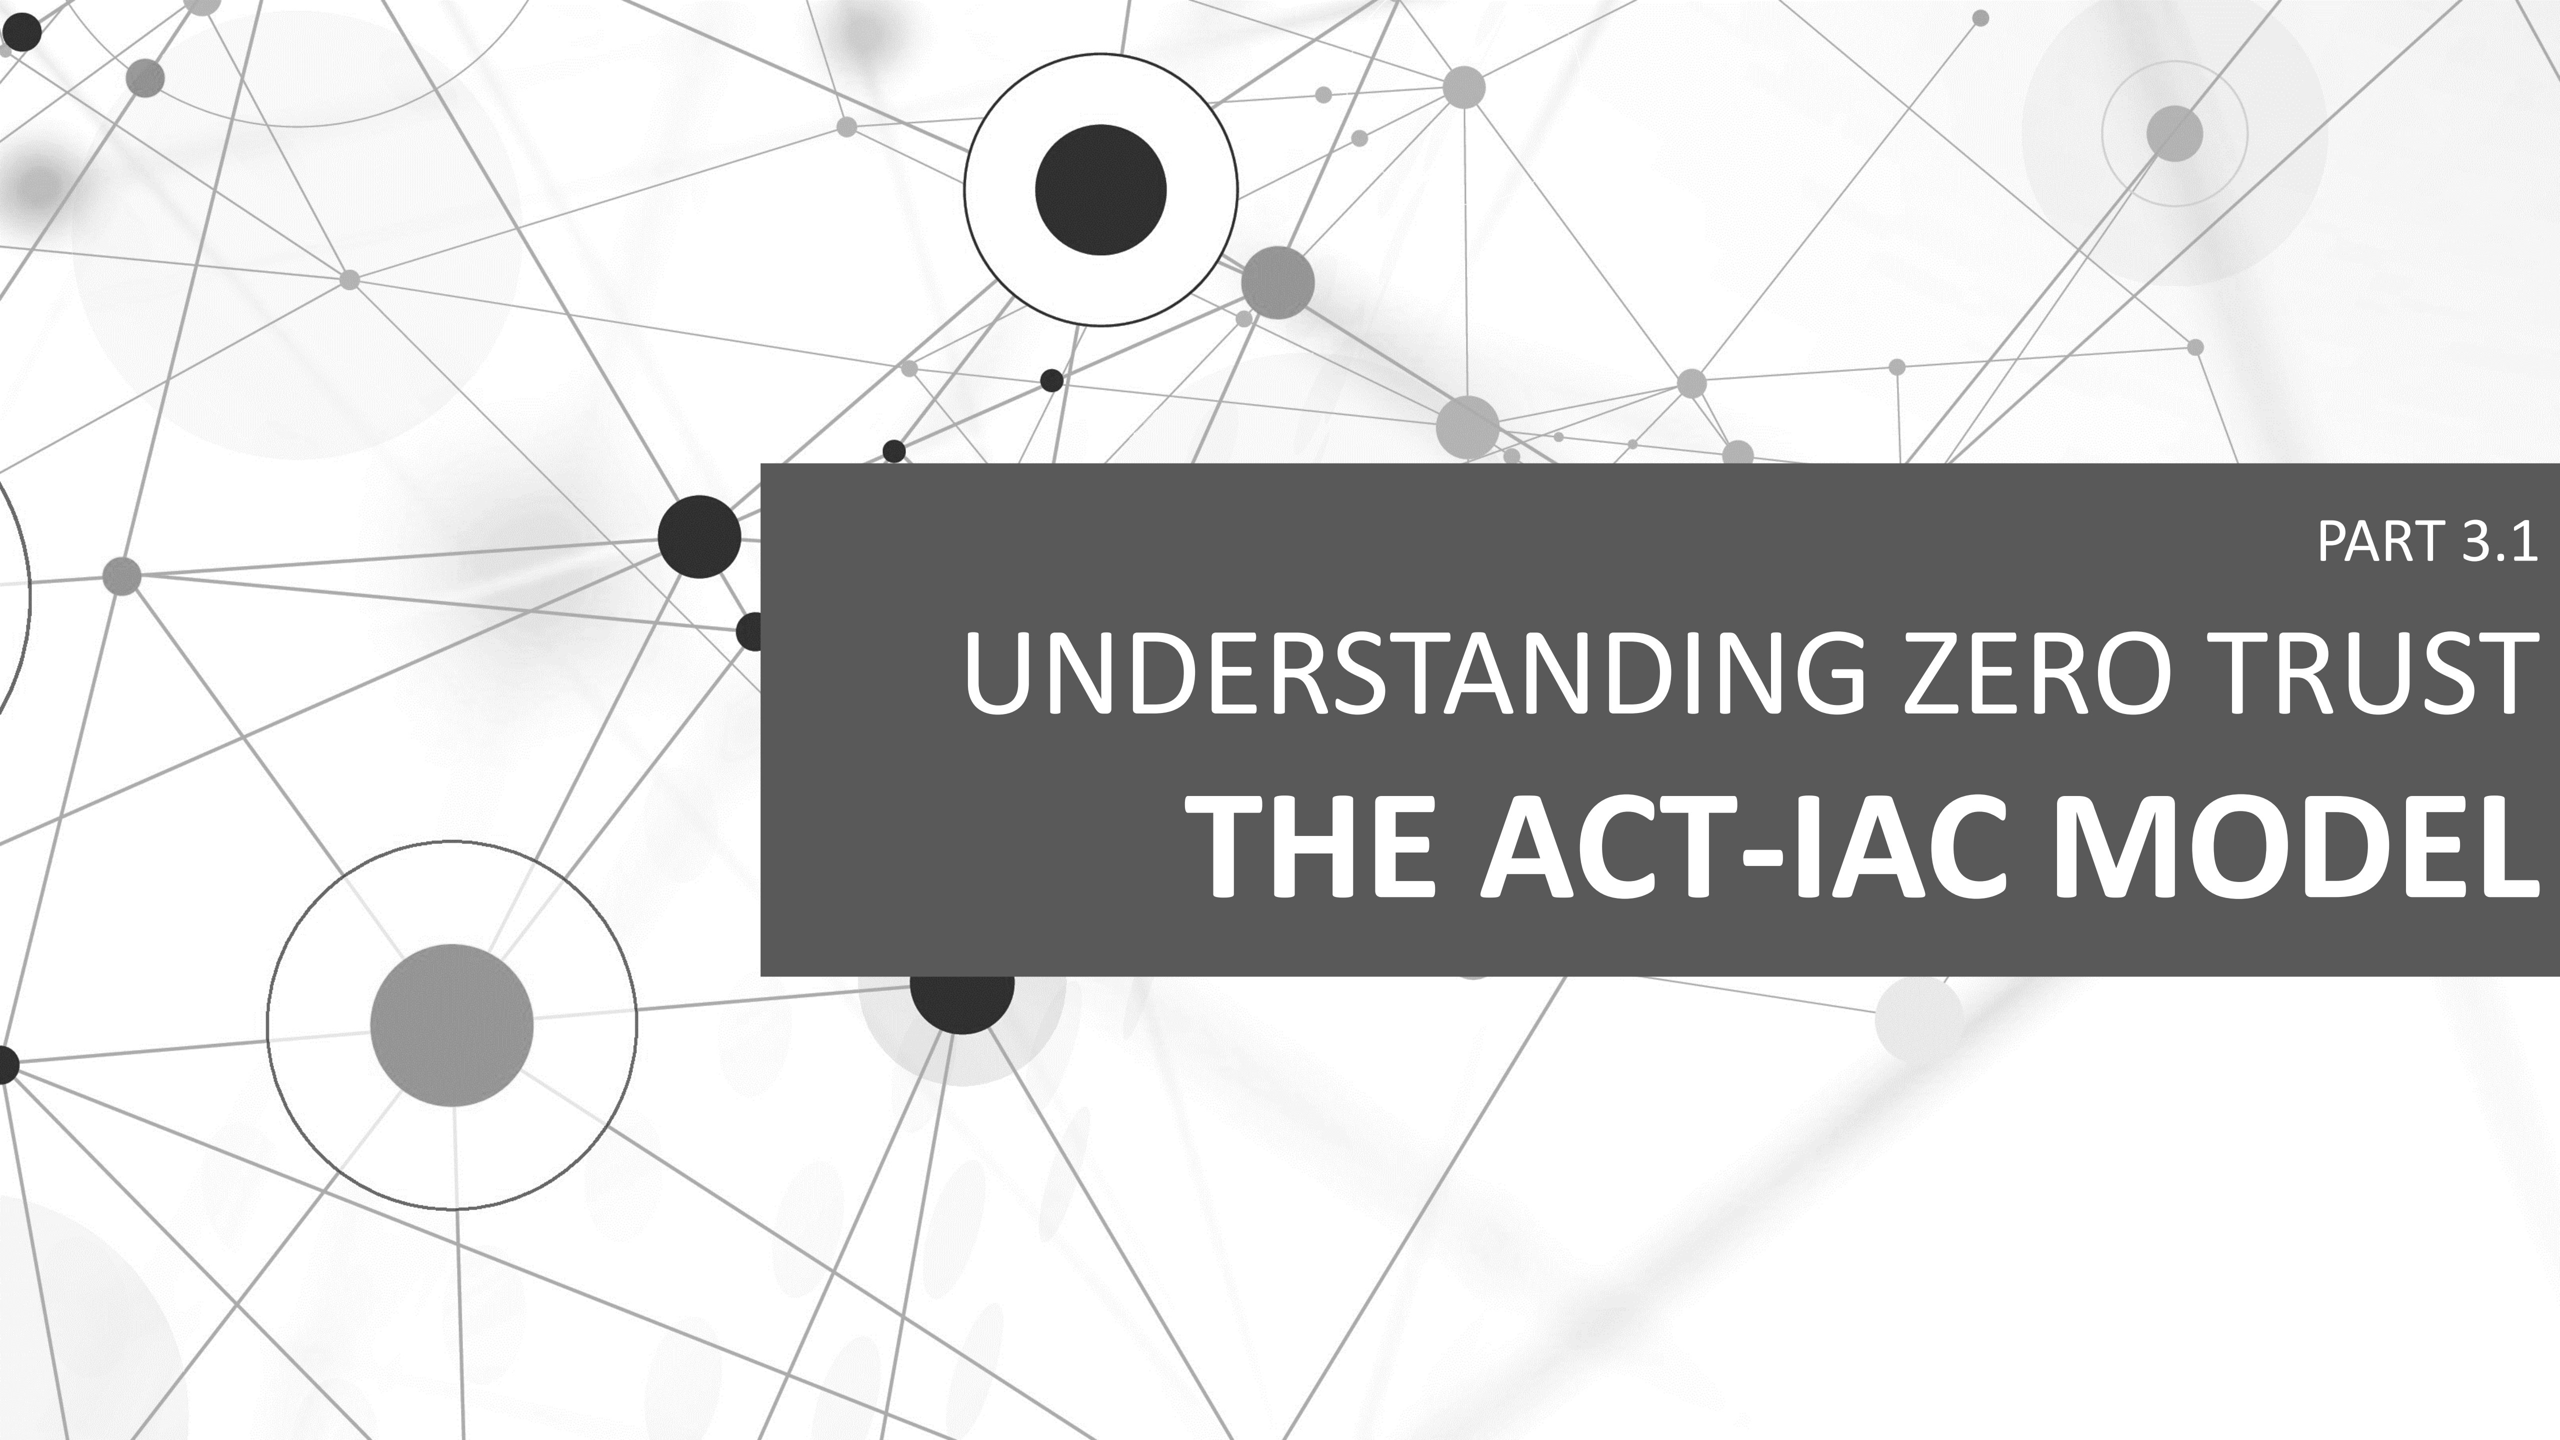

PART 3.1

UNDERSTANDING ZERO TRUST

THE ACT-IAC MODEL

THE ACT-IAC MODEL

- If this so far sounds confusing, maybe looking at the model created by the ACT-IAC prior to start discussing zero trust frameworks and architectures will help clarify what we’ll be looking at when thinking “zero trust”.
- The ACT-IAC has created a model to support Zero Trust architectures at its most fundamental level, **built upon a data foundation**, which is easy enough to understand. Six pillars are considered which Zero Trust targets in different ways to ensure secure data access.

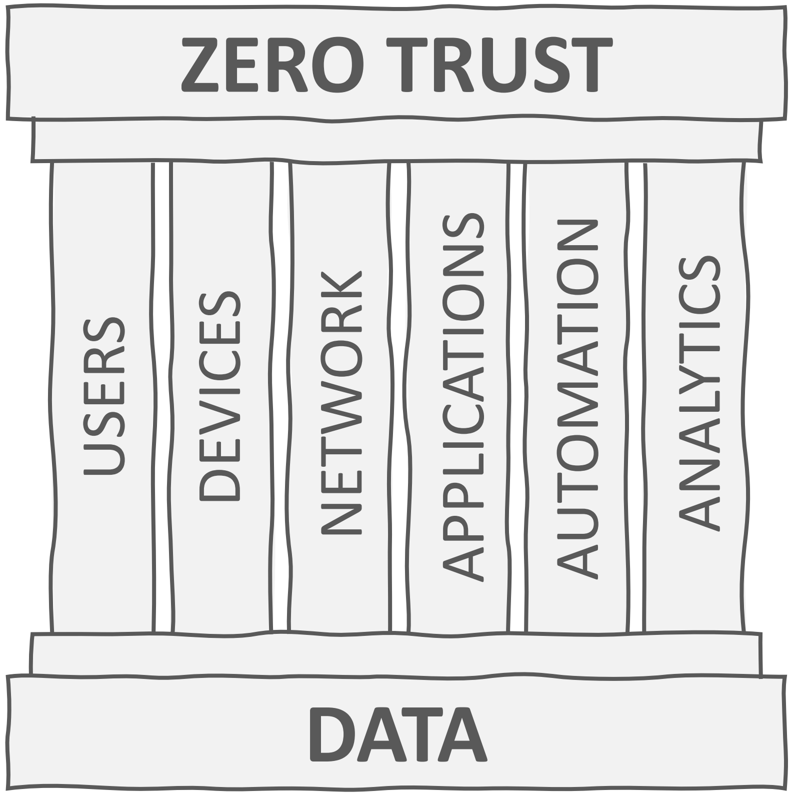

THE ACT-IAC MODEL

- **USERS**: users are **continuously validated, authenticated and monitored** - this allows for continuous validation of account privileges, rather than their assumption and static assignment upon network login.

- **DEVICES**: devices and their security posture are **continuously analysed and validated** in real time.

- **NETWORK**: network is **segmented** in ways that match **use-cases** rather than “office floors” (for instance, software-defined networks).

- **APPLICATIONS**: security must be **definable and applicable at the application layer**: an account is matched against an application and application access and privileges are **never implicit upon network logon** (note that, in this model, VMs and containers are “seen as” applications).

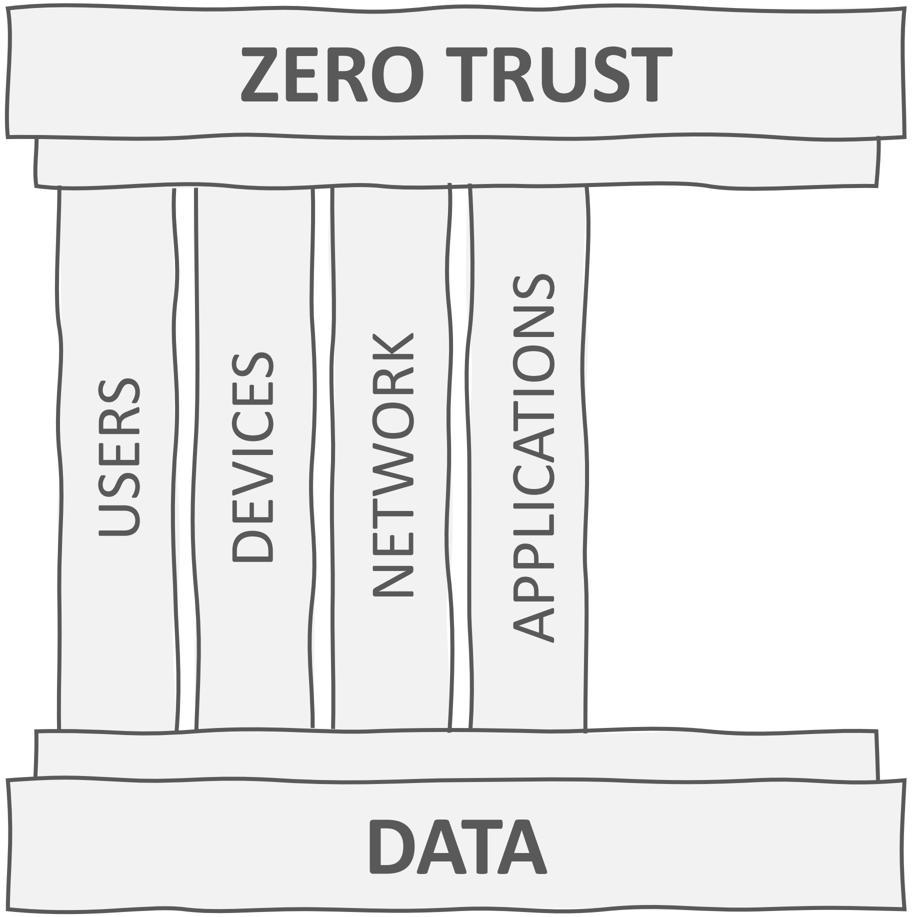

THE ACT-IAC MODEL

- **AUTOMATION**: security automation, orchestration and response (**SOAR**) **automates and orchestrates tasks**, including those related to monitoring, through **predefined workflows**.

- **ANALYTICS**: security information and event management (**SIEM**), user and entity behaviour analytics (**UEBA**) and others provide clear, real-time (or as close as possible to real-time) **visibility over what's happening** at any point in time, as well as easily visualisable **behavioural trends** that allow for more efficient steering of the zero trust implementation.

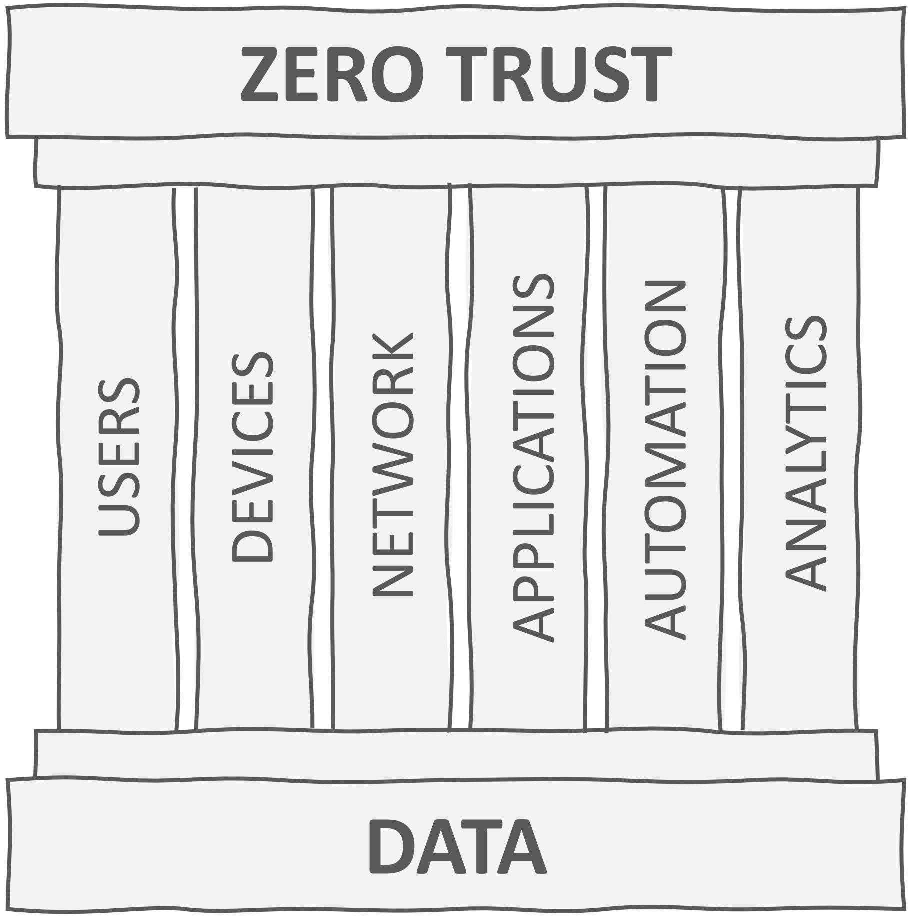

PART 3.2

UNDERSTANDING ZERO TRUST

NIST 800-207

NIST 800-207 | INTRODUCTION

- For context, given the recent number of "high profile" security breaches affecting American federal agencies (amongst others), that have taken place in recent years, the Biden administration (U.S. government) ordered in 2021 that all U.S. federal agencies align with NIST 800-207.

- As a result, the standard has undergone extensive validation and input from a range of commercial customers, suppliers and government agencies, making it what is now widely regarded as the most recognized widely recognised and accepted framework for Zero Trust.

- NIST 800-207 will consider all data sources and computing services as resources which must be taken into account when conceiving a zero trust architecture, and access to any resource will be determined on a per-session basis, based on a dynamic policy, defined using a set of attributes, which can be static and predictable (account identity, resource being accessed, etc) or behavioural and more complex (the more attributes are included in a policy, the stronger the level of security).

NIST 800-207 | DATA IS ALWAYS KEY

- Wherever we look, we will find that **data is always key** to designing an efficient zero trust strategy. Consequently, we will often come across the term “**data-centric enterprise**”.
- The more information collected (and processed - obviously, just collecting it is not enough), the bigger will be the **understanding the enterprise will have of its own resources**, and the best positioned it will be to design and implement a truly meaningful, bespoke security posture.
- Implicitly, we are reiterating the fact that the **understanding of our data is foundational to a successful zero trust strategy**. Again, **not all data has the same value** but, if the data lifecycle is adhered to (which it should be), **all data will have some value**.
- As a reminder, data is not limited to that which users produce and/or gather in the shape of human-intelligible text. **Any logging or tracing**, for instance, **will produce data, and this data** will be of immense value when it comes to a corporate zero trust strategy, namely when it comes to **security automation and orchestration**.

NIST 800-207 | THREE PRINCIPLES FOR ZERO TRUST

Zero Trust considers three principles included in NIST 800-207:

1. **Continuous verification**: Always verify access, all the time, for all resources.
2. **Limit the blast radius**: Minimize the impact if an external or internal security incident occurs.
3. **Automate context collection and response**: Integrate behavioral data and get context from the entire IT stack (identity, endpoint, workload, etc.) for the most accurate response.

More details to follow in the next slides…

NIST 800-207 | THREE PRINCIPLES FOR ZERO TRUST

1. Continuous verification

Continuous verification means that there are **no trusted zones, credentials, or devices** at any time, hence the common phrase "**never trust, always verify**".

A check that needs to be continuously applied to such a large set of assets means that several key elements need to be in place for it to work effectively:

- **Risk-Based Conditional Access**: this ensures that the workflow is **only interrupted when risk levels change**, enabling continuous verification, without sacrificing user experience.
- **Rapid and scalable deployment of a dynamic policy model**: the policy should not only consider the risks, but also **include the company's IT compliance requirements**.

NIST 800-207 | THREE PRINCIPLES FOR ZERO TRUST

2. Limiting the blast radius

If a breach occurs, it is essential to **minimize its impact**. Zero Trust **limits the scope of credentials or paths for an attacker**, giving humas and automated systems **time to react** and mitigate the attack. This means:

- **Identity-based segmentation**: traditional network-based segmentation can be operationally difficult to maintain in today's ever-changing, decentralized world.
- **Principle of Least Privilege**: whenever credentials are used, it is essential that those credentials have access to the **minimum required** to perform the task.

As tasks change, the scope of privileges must also change. Many attacks rely on privileged service accounts because they are typically unmonitored and often over-authorized.

NIST 800-207 | THREE PRINCIPLES FOR ZERO TRUST

3. Automating context collection and response

To make the most efficient and accurate decisions, **having as much data as possible is useful** as long as it can be processed and acted upon in real time. The NIST provides guidance on using information from the following sources:

- **Credentials** (human and service, including SSO credentials)
- **Workloads** (including VMs, containers and others)
- **Endpoints**
- **Network**
- **Data**
- Other sources include SIEM, SSO, identity providers, threat intelligence…

PART 3.3

UNDERSTANDING ZERO TRUST

THE AMERICAN DOD’S MODEL

DOD ZERO TRUST REFERENCE ARCHITECTURE

The American Department of Defense (DoD), with heavy reliance on the contents of NIST 800-207, has developed its own model and reference architecture guide, which simplifies the contents of NIST 800-207 quite a bit.

Its unclassified (obviously) document has been prepared by the Defense Information Systems Agency (DISA) and National Security Agency (NSA) Zero Trust Engineering Team. The link to this document can be found in the References section of this document.

The main premise we get from this document is that access to resources based on the risk of the user and devices is the baseline requirement for Zero Trust.

According to the DoD, a Zero Trust architecture will include:

1. **Defense of Enterprise Identity, Credential, and Access Management (ICAM)**
2. **Client and Identity Assurance**
3. **Data-Centric Enterprise**
4. **Dynamic Access Control Plane**
5. **Analysis & Confidence Scoring**
6. **Logging utilizing Security Information and Event Management**

More details to follow in the next slides…

DOD ZERO TRUST REFERENCE ARCHITECTURE

1. Defense of Enterprise Identity, Credential, and Access Management (ICAM):

This will include:

- **Identity Provider (IDP),**
- **Automatic Account Provisioning (AAP) and**
- **Master User Record (MUR)**

To identify and manage the roles, access privileges, and the circumstances in which users are granted or denied privileges.

DOD ZERO TRUST REFERENCE ARCHITECTURE

2. Client and Identity Assurance

Which will include:

- **Authentication Decision Point:** to evaluate the **identity** of the user, NPE, and or device **as access is attempted** to applications and data.
- **Authorization Decision Point:** a system entity that makes **authorization decisions** for entities that request such access decisions.

The capabilities of **Comply-to-Connect** (C2C) and **PAM** are expected.

DOD ZERO TRUST REFERENCE ARCHITECTURE

3. Data-Centric Enterprise

Which will include:

- **Resource Authorization Decision Point**: this is an intermediary decision point which will evaluate the combined NPE and user to authorize the request for access, including the capabilities of **Macro Segmentation** and **Application Delivery Control** (which can be a Proxy in a DMZ).
- **Application Authorization Decision Point**: this is another intermediary decision point which will evaluate the combined NPE and user to authorize the request for access.
- **Data Rights Management**: a set of **access control** technologies that prevent the unauthorized access, modification and redistribution of data.
- **Data**: the final step in the process is **access to the data and applications**. **Data tagging** will be used to ensure **proper classification levels** for all data are used to help prevent spillage.

Some capabilities that must be considered are:

- **Micro segmentation**
- **DevSecOps Application Development and**
- **Data Authorization Decision Point**

DOD ZERO TRUST REFERENCE ARCHITECTURE

4. Dynamic Access Control Plane:

To be included is:

- **Policy Engine & Automation (SOAR):** these are technologies that handle **threat management, incident response, policy enforcement and security policy automation**.

The capabilities to be expected are:

- **Automated Policy Deployment,**
- **Endpoint Detection and Response** (**EDR**, with real-time monitoring and detection of malicious events on endpoints), and
- **User Activity Monitoring (UAM).**

DOD ZERO TRUST REFERENCE ARCHITECTURE

5. Analysis & Confidence Scoring

These technologies perform **continuous assessments** of entities, attributes and configurations to adapt and risk-optimize security policy for deployments.

Confidence scores are leveraged in authorization activities.

We will consider the capabilities of:

- **Entity Behavior Analysis (with SIEM log analysis)** and
- **Data Loss Prevention (DLP)**.

6. Logging utilizing Security Information and Event Management

Used for **data aggregation and storage**, ensuring both **security information management** (SIM) and **security event management** (SEM) capabilities, allowing for **Entity Activity Auditing**.

▶ **We will see how these 6 baseline requirements fit, in practical terms, into a Zero Trust strategy later on…**

PART 4

DESIGNING AND IMPLEMENTING ZERO TRUST

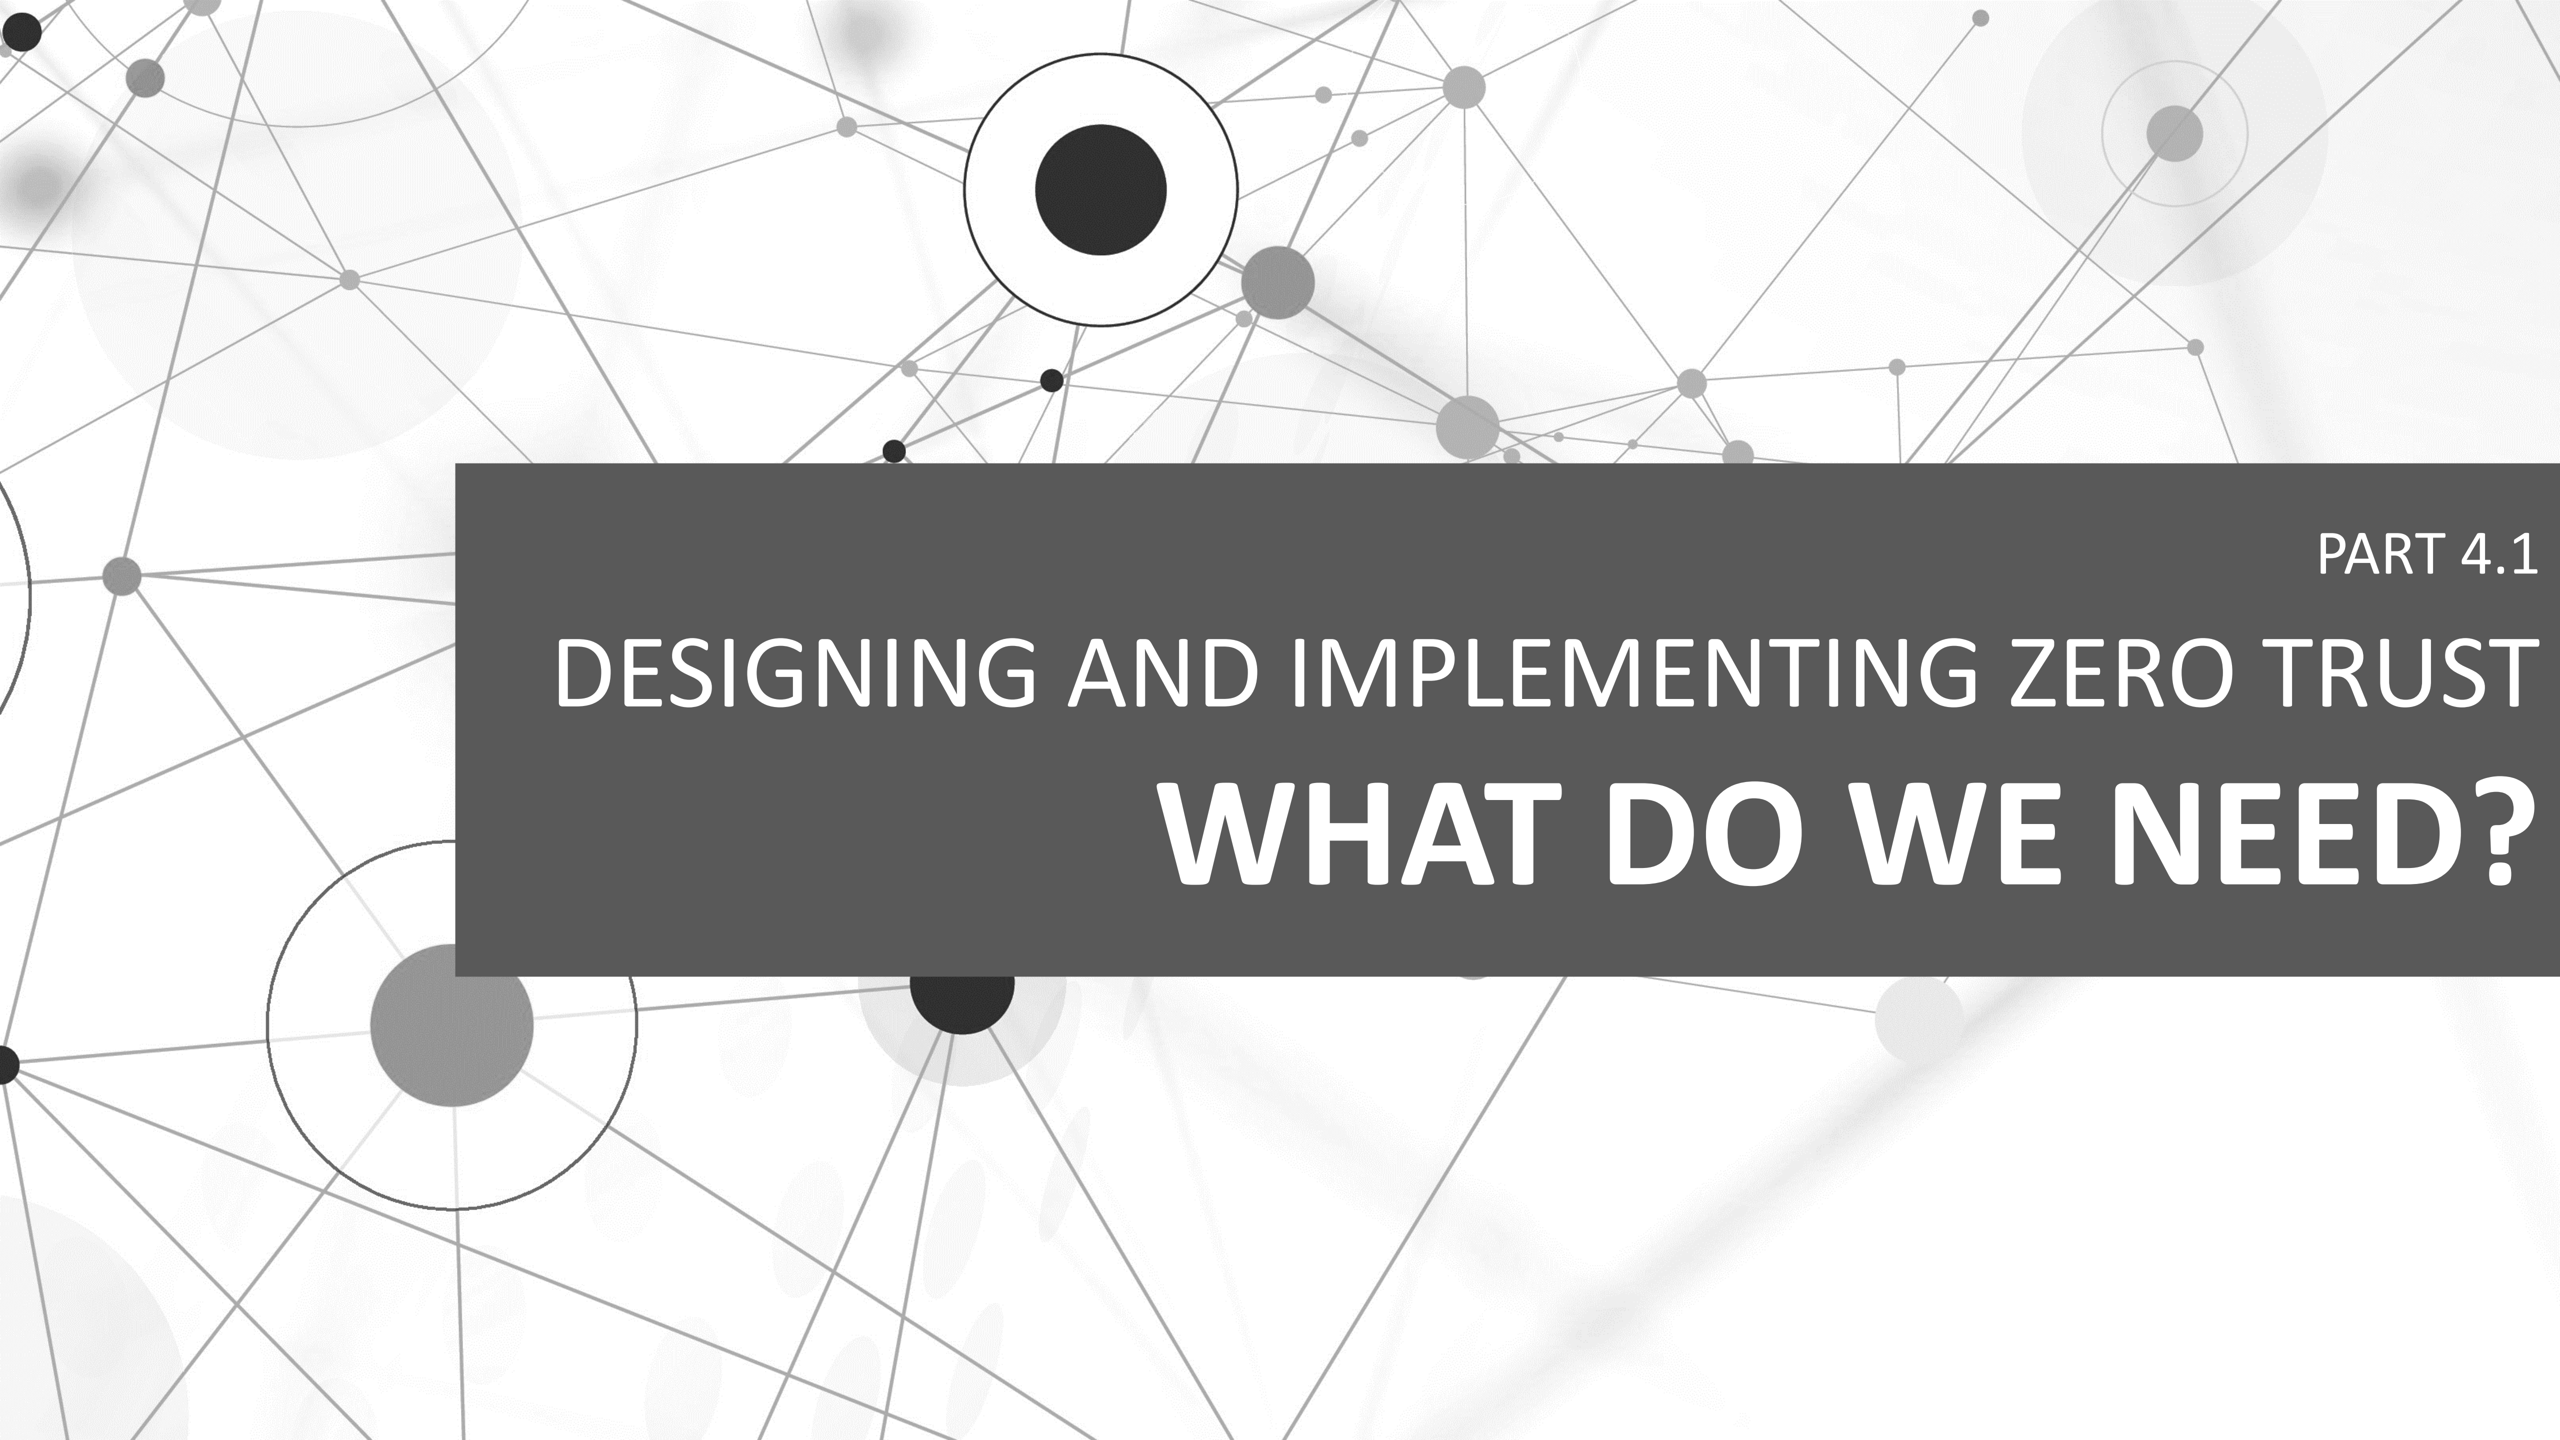

PART 4.1

DESIGNING AND IMPLEMENTING ZERO TRUST

WHAT DO WE NEED?

RECAP OF WHAT WAS DISCUSSED SO FAR

- As discussed, the concept of Zero Trust introduces a different paradigm to traditional “trust but verify” network security, where users and endpoints within the corporate network perimeter were automatically trusted.
- This traditional security model has introduced risks to the organization, ranging from malicious activity by internal users to legitimate credentials being exploited by external parties who have access to them.
- Unlike this traditional model, the Zero Trust architecture requires organizations to continuously monitor and validate that accounts and devices have the appropriate privileges and attributes, in addition to enforcing security policies and ensuring compliance with requirements, with predefined security checks before allowing any action.
- **In very simplistic words, whilst we used to validate and authenticate a subject against a network, and automatically grant that subject access to the objects (aka resources) that network contained, Zero Trust validates and authenticates the subject against the resource they are trying to access, regardless of the network.**

WHAT DO WE NEED FOR ZERO TRUST?

- There is no sole technology for the full implementation of zero trust, although several vendors are now in the position of being able to offer rich product suites that get close to covering all formal aspects of it.
- Needless to say, the human brain remains probably any enterprise’s most meaningful appliance, with a core set of tools irreplaceable - at least at this point in time - by any product suite available on the market. So it’s not just about rolling out products, it’s learning, thinking planning, rethinking and replanning before rolling them out.
- Although the paradigm is different from what we have traditionally seen in terms of securing our assets, zero trust relies mostly on existing security principles, controls and techniques, such as the principle of least privilege, the principle of compartmentalisation, segregation, IAM, PAM, CASB, IDS/IPS, CSPM/SSPM, content inspection, DNS security, etc.

WHAT DO WE NEED FOR ZERO TRUST?

- Understandably, we can now ask the question: “if we have all the above, why are we talking about zero trust? Don’t we de facto have zero trust already in place”.
- And the fair answer is probably “maybe” - which takes us back to the start: understanding what zero trust is and how it differs from traditional security architectures.
- If the above has been put in place as part of a security architecture designed to cater for centralised assets and users, where securing the network was the goal and we implicitly trusted whatever was on the network, then the answer is no.
- As a very silly but practical example: if we apply several of the above to authenticate us to our Active Directory but, once logged onto the Windows network, we are automatically granted uncontrolled access to whatever applications are available, then we clearly don’t have zero trust in place.

PART 4.1

DESIGNING AND IMPLEMENTING ZERO TRUST

CORE LOGICAL COMPONENTS

SIMPLIFIED VERSION OF CORE LOGICAL COMPONENTS OF ZTA

The following is an extremely simplified version of the “official” NIST illustration of the core logical elements of a Zero Trust architecture (which has also been adopted in the context of the CISSP anyone training for the certification (or recertification) should know it by heart):

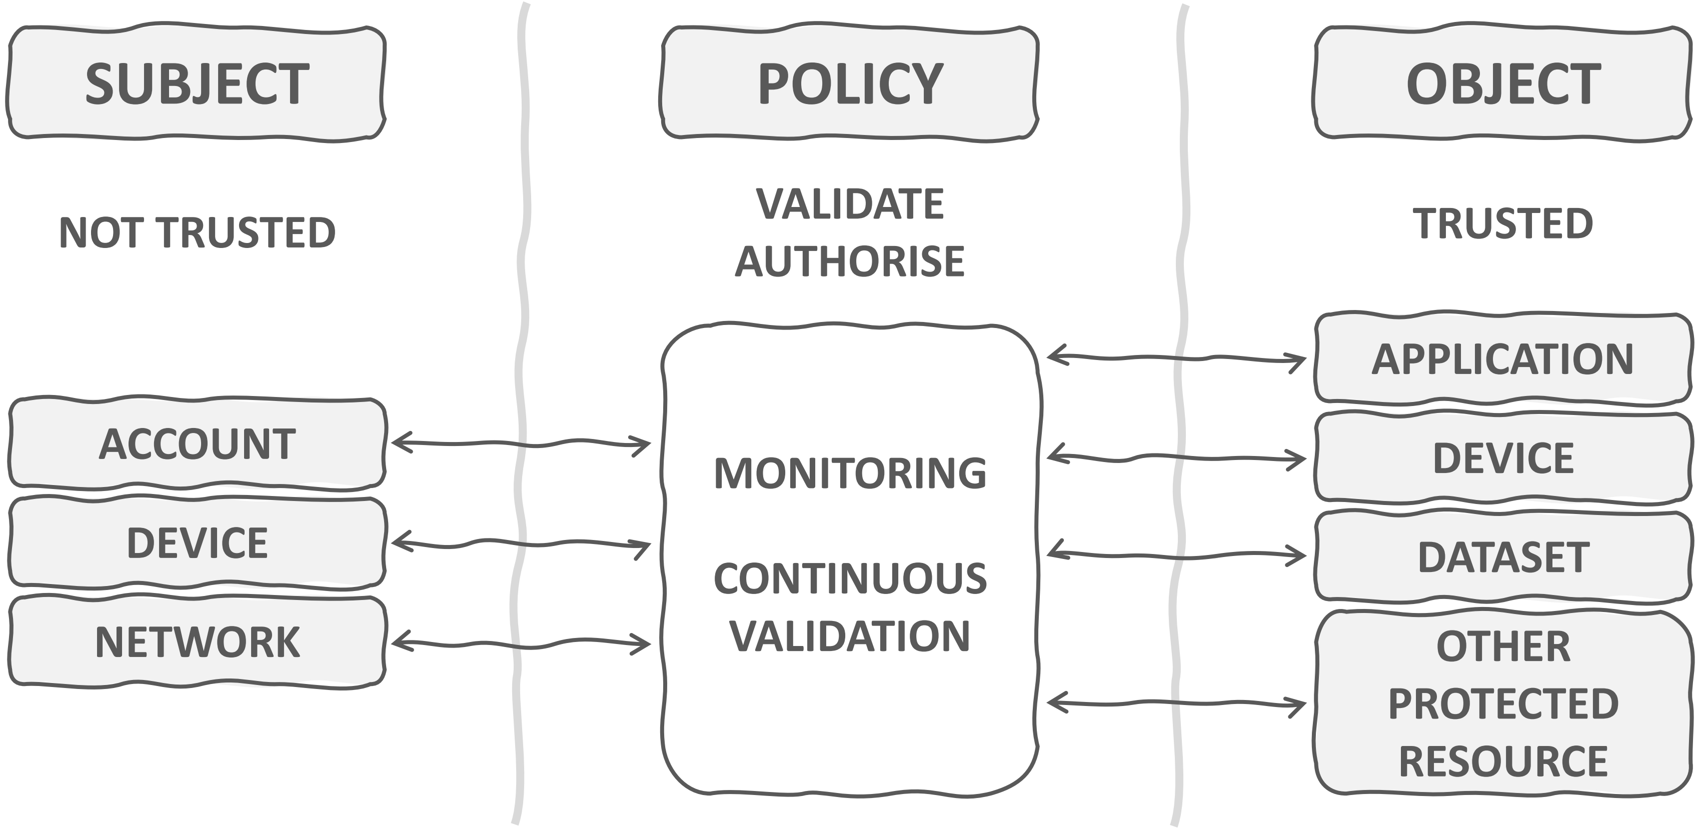

SUBJECT

NOT TRUSTED

ACCOUNT

DEVICE

NETWORK

POLICY

VALIDATE
AUTHORISE

MONITORING

CONTINUOUS
VALIDATION

OBJECT

TRUSTED

APPLICATION

DEVICE

DATASET

OTHER
PROTECTED
RESOURCE

DIVING DEEPER INTO THE CORE LOGICAL COMPONENTS OF ZTA

As we have understood the very simplified version of the diagram, we can jump to the real thing as it will help us understand and contextualise what lies in the centre and what takes a peripheral location in a Zero Trust architecture. And here is what the NIST presents us with, and the industry agrees is a fair representation of the core elements of ZTA:

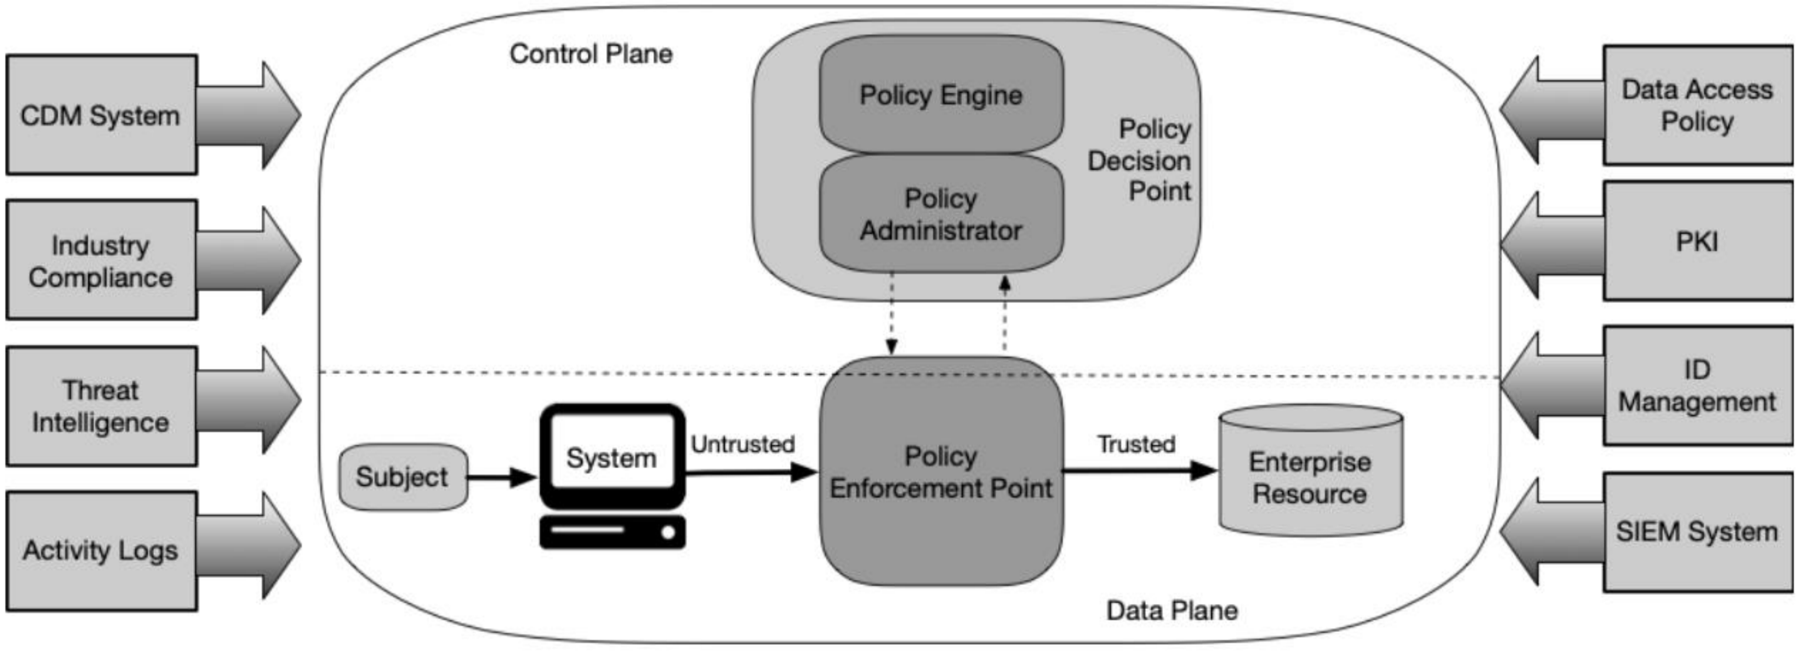

As we see, two parallel planes sit squarely in the middle of the diagram: the control plane and the data plane. The Two interact through what's named here a "policy enforcement point" (which is rather self-explanatory).

DIVING DEEPER INTO THE CORE LOGICAL COMPONENTS OF ZTA

Straight from the NIST 800-207 (meaning, copy and paste, with a few changes to make it easier to read):

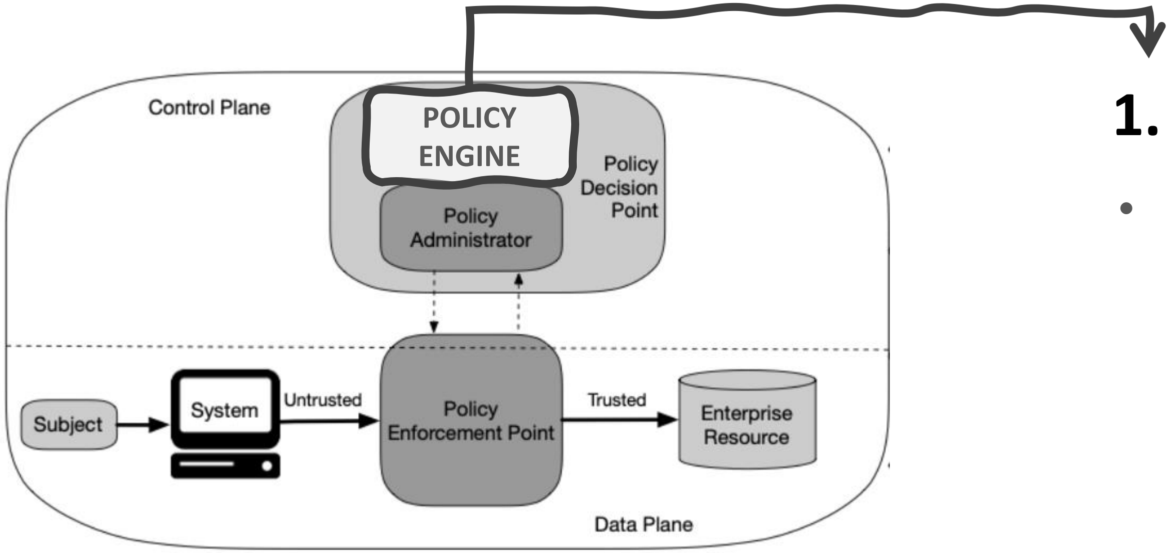

1. Policy engine (PE)

- 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 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.

DIVING DEEPER INTO THE CORE LOGICAL COMPONENTS OF ZTA

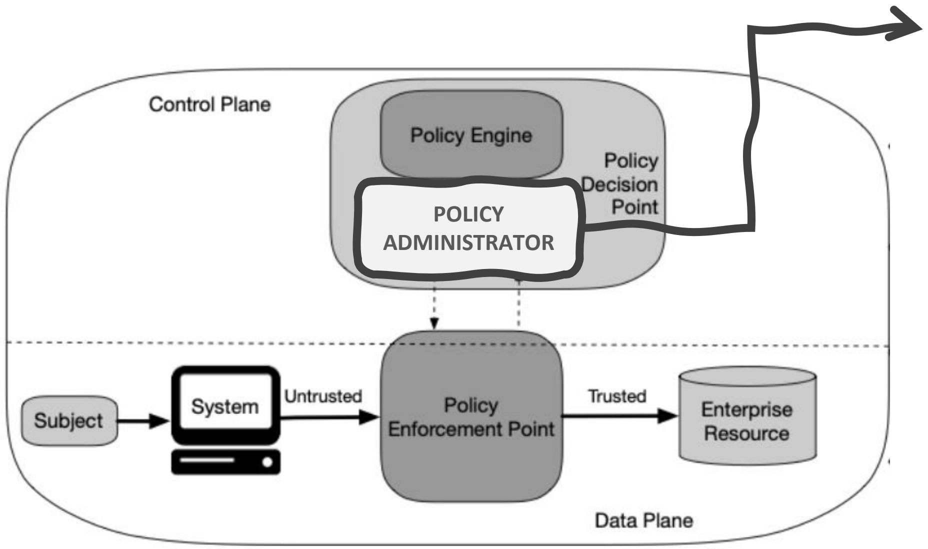

2. Policy Administrator

- 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 its two logical components.
- The PA communicates with the PEP when creating the communication path. This communication is done via the control plane.

DIVING DEEPER INTO THE CORE LOGICAL COMPONENTS OF ZTA

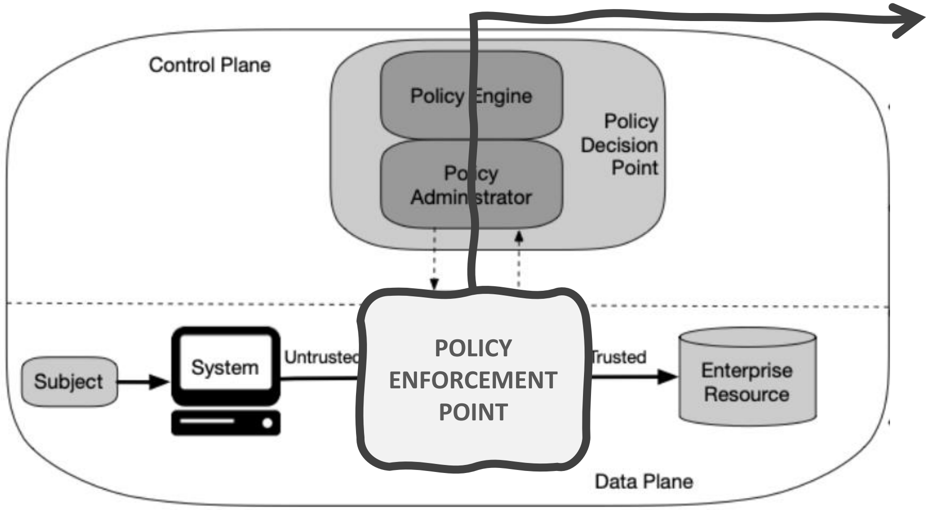

2. Policy enforcement point (PEP)

- 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.
- Beyond the PEP is the trust zone hosting the enterprise resource.

PART 4.2

DESIGNING AND IMPLEMENTING ZERO TRUST

IN CORPORATE TERMS

ZTA IN CORPORATE TERMS

So, what does this all mean in corporate terms?

The first part of the answer to that question, you probably know it already (it’s included in the introduction section): it means that, more than just a change in technology, Zero Trust is a change in mindset. And changing mindsets can be far more changing than technologies.

When we talk about Zero Trust, we are not talking about any specific technology or layer in our technical stack - we are quite literally talking about every layer, including those we might not have yet, but must assume we will at some point in the future.

The big change in paradigm with its implicit required change in mindset, associated to this technical lack of specificity makes a shift to a Zero Trust security architecture a big challenge to any enterprise. It is, nevertheless, a challenge that must be tackled - the sooner, the better.

ZTA IN CORPORATE TERMS

It ends up boiling down to knowledge of what we have, vision of where we want to be, the strategic thinking to plan how we get there, and the perseverance to get there.

Needless to say, this is not a project to be implemented in 6 months or “an IT thing”, it is a corporation-wide subject that will take close collaboration between multiple departments and, as mentioned, won’t be resolved through an “out of the box” technical roll-out.

As a matter of fact, the NIST assumes a transition period (which might be long, depending on many factors, including the size of the business, commitment to infrastructure security, etc.), when the security architecture will be hybrid – both traditional security and ZTA will coexist.

Aligned with the above, the American DoD offers the following diagram for the assessment of the maturity of our Zero Trust architecture. The diagram clearly demonstrates that nowhere are we presumed to be able to simply “roll out Zero Trust”. Moving to ZT is long process, where an end-state ZTA will take time and commitment to be achieved.

PART 4.3

DESIGNING AND IMPLEMENTING ZERO TRUST

THE ‘THREE STEPS APPROACH’

A 3 STEPS APPROACH TO ZERO TRUST

There is a broadly recognized 3-step approach to Zero Trust: Visualization, Mitigation and Optimization? Where:

Visualization

- At this stage, the intention is to understand all the resources, their access points and the risks. Under NIST, this is an ongoing process because resources will change in terms of availability, risk, and importance.
- Key objectives are to see all entities (identities, workloads, endpoints) and understand any vulnerabilities or risks in order to see the attack path and key assets to defend

Mitigation

- We will be ready to detect and stop threats or mitigate the impact if a threat cannot be stopped immediately. The NIST calls for automation and orchestration because it is essential for real-time detection and response.
- Behavioral analytics should be incorporated to detect threats such as insiders. We also limit the impact of violations thanks to segmentation and the principles of least privilege.

Optimization

- The goal is to extend protection to all aspects of infrastructure and all resources, regardless of their location, without creating a poor user experience (which can lead to non-compliance and reduced productivity).
- Conditional Access is deployed based on risk, and it works through continuous verification without compromising a positive user experience.

PART 4.4

DESIGNING AND IMPLEMENTING ZERO TRUST

THE DOD APPROACH

THE DOD APPROACH TO ZERO TRUST

The American DoD reference architecture for Zero Trust is being used here, however, since it gives us very straight-forward advice on how to prepare for and reach ZTA:

1. **Define Mission Outcomes**: a Zero Trust design is derived from organization specific mission requirements and analysis that identify the critical protect surfaces - Data/Assets/Applications/Services (DAAS).
2. **Architect from the inside out:** first, focus on protecting the DAAS. Second, secure a path to access them.
3. **Define high level groups:** for users, devices, and applications/workloads.
4. **Determine who/what needs access:** to create and apply local security policies consistently across all environments (Local Area Network (LAN), Wide Area Network (WAN), Endpoint, Perimeter, Mobile, etc.).
5. **Inspect and log all traffic and events necessary to answer Commanders Critical Information Requirements (CCIRs) derived from mission analysis:** full visibility across all layers is required for analytics.

THE DOD APPROACH TO ZERO TRUST

The DoD presents us with 7 pilars for Zero Trust, and with a sequencial way to secure them. Those pilars are:

- **Z1: User**
- **Z2: Device**
- **Z3: Network/Environment**
- **Z4: Applications & Workload**
- **Z5: Data**
- **Z6: Visibility & Analytics**
- **Z7: Automation & Orchestration**

This is represented in the graphic that follows (an adaptation from the DoD's ZTA reference architecture publication).

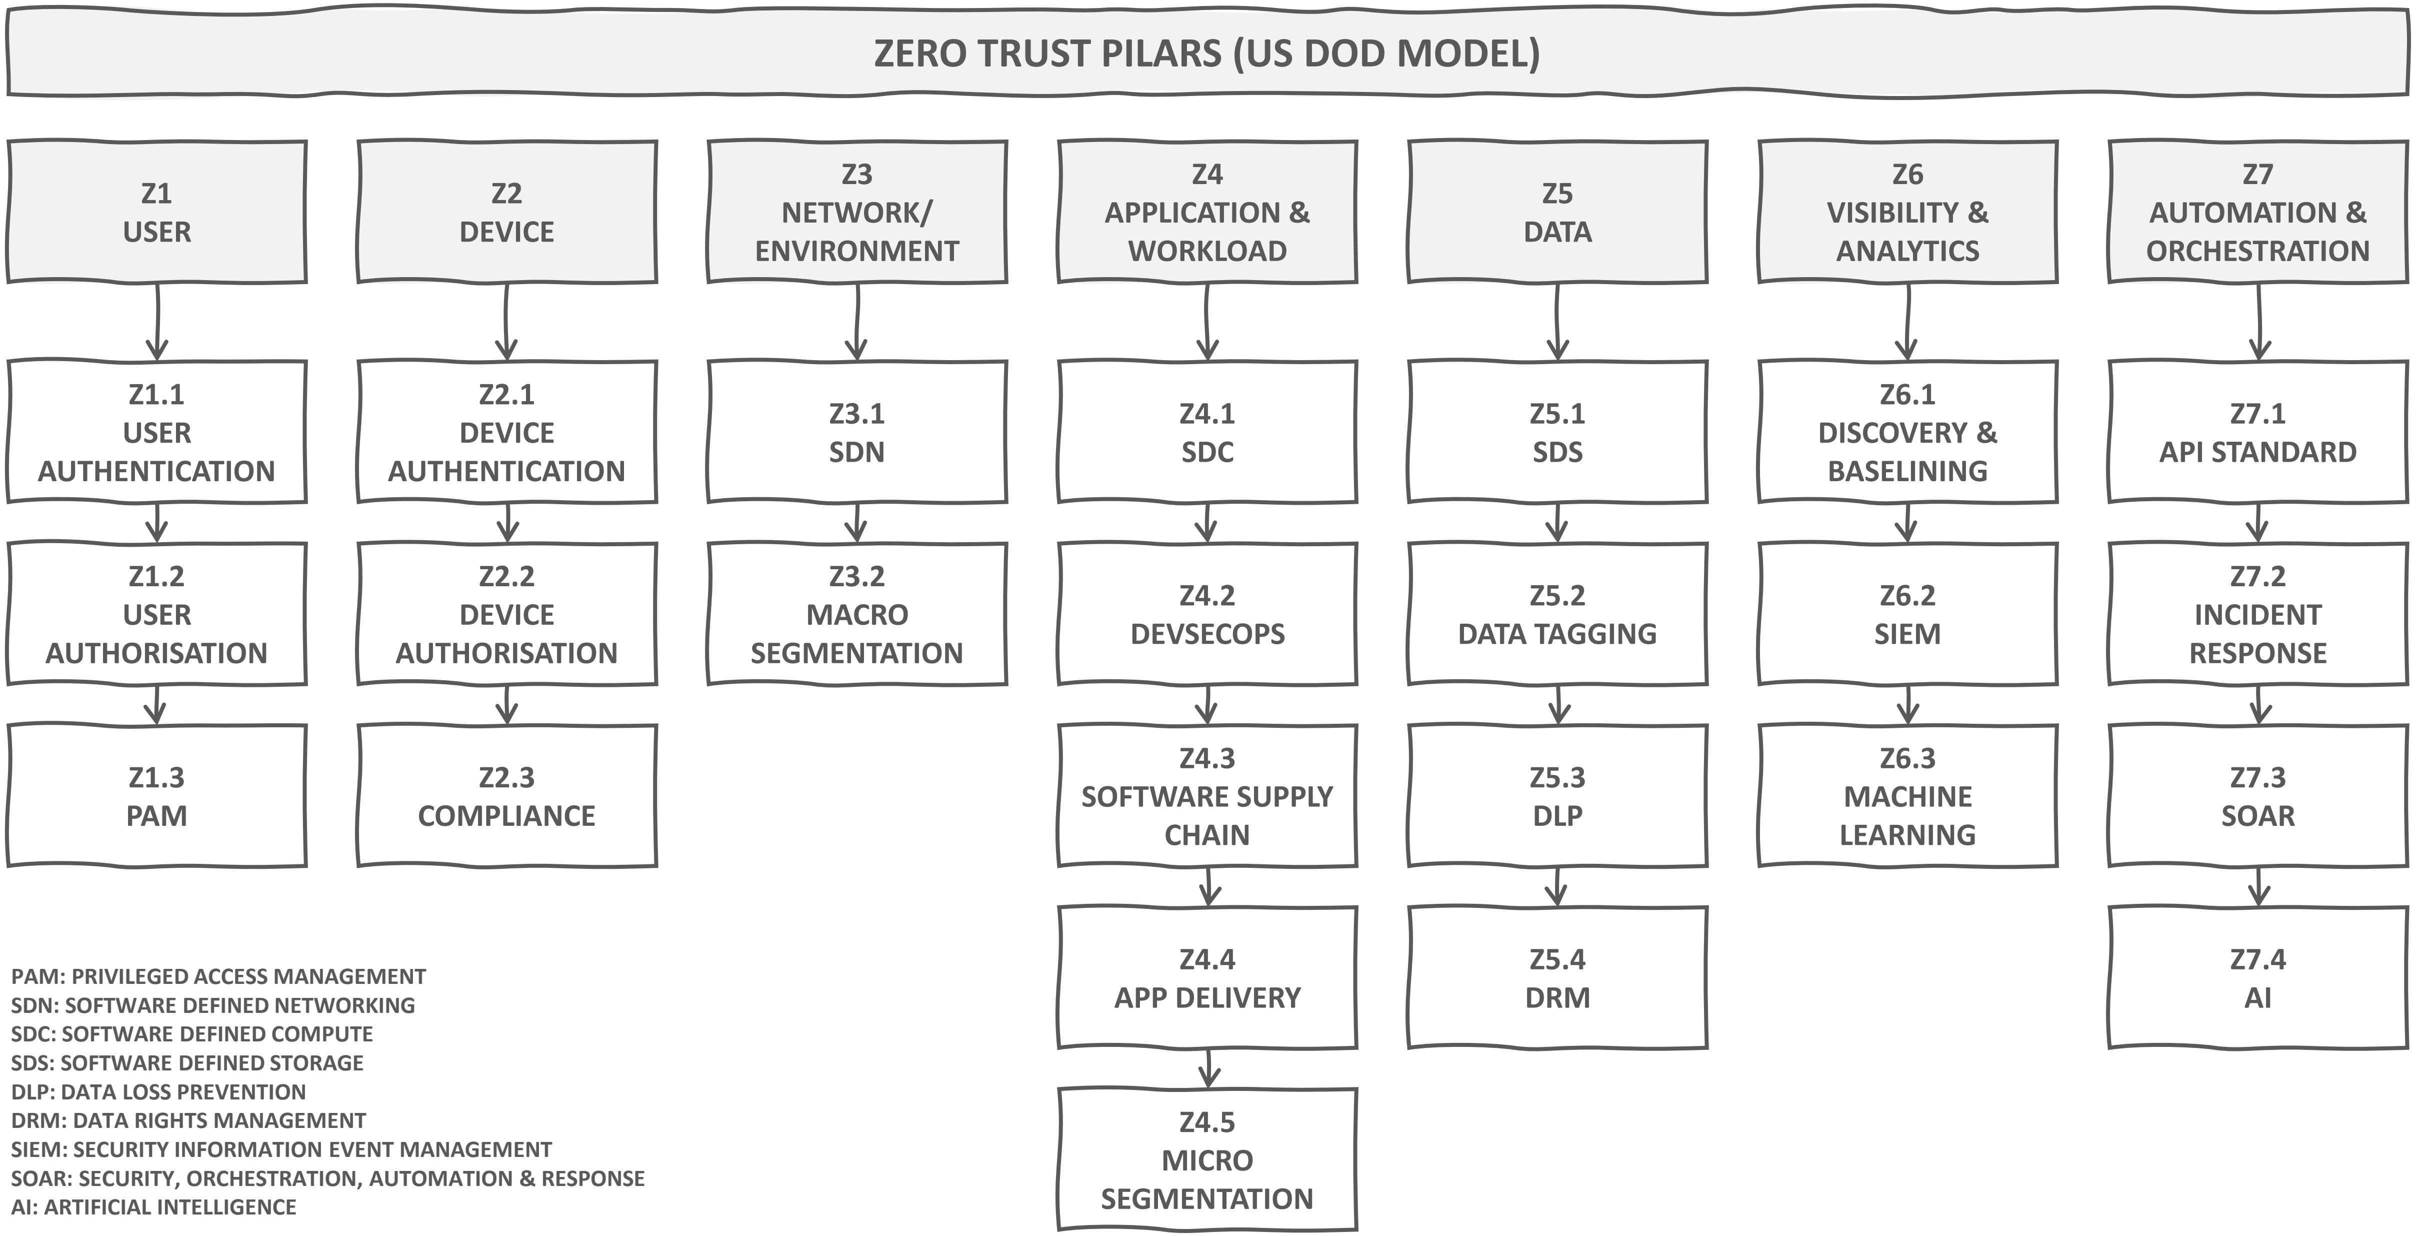

ZERO TRUST PILARS (US DOD MODEL)

- Z1 USER
  - Z1.1 USER AUTHENTICATION
  - Z1.2 USER AUTHORISATION
  - Z1.3 PAM
- Z2 DEVICE
  - Z2.1 DEVICE AUTHENTICATION
  - Z2.2 DEVICE AUTHORISATION
  - Z2.3 COMPLIANCE
- Z3 NETWORK/ ENVIRONMENT
  - Z3.1 SDN
  - Z3.2 MACRO SEGMENTATION
- Z4 APPLICATION & WORKLOAD
  - Z4.1 SDC
  - Z4.2 DEVSECOPS
  - Z4.3 SOFTWARE SUPPLY CHAIN
  - Z4.4 APP DELIVERY
  - Z4.5 MICRO SEGMENTATION
- Z5 DATA
  - Z5.1 SDS
  - Z5.2 DATA TAGGING
  - Z5.3 DLP
  - Z5.4 DRM
- Z6 VISIBILITY & ANALYTICS
  - Z6.1 DISCOVERY & BASELINING
  - Z6.2 SIEM
  - Z6.3 MACHINE LEARNING
- Z7 AUTOMATION & ORCHESTRATION
  - Z7.1 API STANDARD
  - Z7.2 INCIDENT RESPONSE
  - Z7.3 SOAR
  - Z7.4 AI

PAM: PRIVILEGED ACCESS MANAGEMENT
SDN: SOFTWARE DEFINED NETWORKING
SDC: SOFTWARE DEFINED COMPUTE
SDS: SOFTWARE DEFINED STORAGE
DLP: DATA LOSS PREVENTION
DRM: DATA RIGHTS MANAGEMENT
SIEM: SECURITY INFORMATION EVENT MANAGEMENT
SOAR: SECURITY, ORCHESTRATION, AUTOMATION & RESPONSE
AI: ARTIFICIAL INTELLIGENCE

ZERO TRUST

WHAT IT IS | WHY WE NEED IT | HOW WE GET THERE

PREPARE FOR ZERO TRUST

DISCOVERY

- ? IDENTIFY DAAS
- ? MAP DATA FLOWS
- ? INVENTORY USERS AND DEVICES
- ? IDENTIFY PRIVILEGED ACCOUNTS
- ? LOG NETWORK TRAFFIC

ASSESSMENT

- ? DETERMINE COMPLIANCE STATE LEVERAGING EXISTING HARDENING STANDARDS
- ? DETERMINE ACCOUNT PRIVILEGE LEVELS
- ? IDENTIFY IF EXISTING ENVIRONMENT/NETWORK SECURITY POLICIES ARE IMPLEMENTED IN THE LEAST PRIVILEGE MANNER

EVOLVING ZERO TRUST CAPABILITIES AND CONTROLS (DOD)

BASELINE

- ✓ DAAS IDENTIFIED
- ✓ ACCESS DAAS IS DETERMINED BY CYBERSEC POLICIES
- ✓ NETWORK SEGMENTATION WITH DENY ALL/ALLOW BY EXCEPTION
- ✓ MANAGED DEVICES COMPLIANT WITH IT SECURITY POLICIES
- ✓ PRINCIPLE OF LEAST PRIVILEGE IN PLACE
- ✓ MFA IN PLACE
- ✓ DATA CLASSIFICATION AND CRITICAL DATA TAGGING STARTED
- ✓ ENCRYPTION IN PLACE

INTERMEDIATE

- ✓ ACCESS BASED ON FINE-GRAINED USER AND DEVICE ATTRIBUTES (CYBERSEC POLICY ENHANCEMENT)
- ✓ MICROSEGMENTATION THROUGHOUT MOST OF THE NETWORK
- ✓ USER ID BASED ON ENTERPRISE FEDERATED SERVICES
- ✓ LEAST PRIVELEGE ENHANCEMENT WITH PAM
- ✓ BEGIN DEPLOYMENT OF DLP & DRM
- ✓ DATA IS CLASSIFIED AND TAGGED VIA FLOW ANALYSIS AND SOME AUTOMATION
- ✓ UEBA STARTS BEING USED TO DEVELOP BASELINE POLICY

ADVANCED

- ✓ CYBERSEC POLICIES DYNAMICALLY DETERMINE DAAS ACCESS? WITH ROBUST REAL-TIME ANALYTICS
- ✓ FULL MICROSEGMENTATION
- ✓ CONTINUOUS AND ADAPTATIVE AUTHENTICATION & AUTORISATION
- ✓ LEAST PRIVILEGE FULLY IMPLEMENTED
- ✓ MACHINE LEARNING TAGS AND CLASSIFIES DATA
- ✓ DLP & DRM FULLY IMPLEMENTED WITH DATA TAGS INCLUDED
- ✓ ADVANCED ANALYTICS USED FOR AUTOMATED AND ORCHESTRATED THREAT DETECTION

DAAS: DATA, ASSETS, APPLICATIONS, SERVICES
MFA: MULTI-FACTOR AUTHENTICATION

PAM: PRIVILEGED ACCESS MANAGEMENT
DLP: DATA LOSS PREVENTION

)RM: DATA RIGHTS MANAGEMENT
JEBA: USER AND ENTITY BEHAVIOUR ANALYTICS

PART 4.5

DESIGNING AND IMPLEMENTING ZERO TRUST

HANDS-ON APPROACH

HANDS-ON APPROACH TO ZERO TRUST

Based on the « roadmap » provided by the DoD's ZTA reference document, we will transform steps into questions and answers.

PREPARE FOR ZERO TRUST

DISCOVERY

? IDENTIFY DAAS
? MAP DATA FLOWS
? INVENTORY USERS AND DEVICES
? IDENTIFY PRIVILEGED ACCOUNTS
? LOG NETWORK TRAFFIC

ASSESSMENT

? DETERMINE COMPLIANCE STATE LEVERAGING EXISTING HARDENING STANDARDS
? DETERMINE ACCOUNT PRIVILEGE LEVELS
? IDENTIFY IF EXISTING ENVIRONMENT/NETWORK SECURITY POLICIES ARE IMPLEMENTED IN THE LEAST PRIVILEGE MANNER

? What data do we possess?
? Where is this data (and all copies thereof) stored?

✓ **Perform a full inventory of our data repositories. An initial « high-level » classification of this data can be done.**

? How do we manage the lifecycle of our data?
? Do we have any automated processes to manage data flows?

✓ **Find relevant documentation describing our data lifecycle and any data flows**

? What devices and appliances do we possess?
? Where are they located?
? What's their obsolescence state?

✓ **Perform a full inventory of devices and appliances, including licensing status and « end-of-life » labelling.**

? Who are our users?
? What devices do they use? (Are those devices provided by the business, BYOD, fully personal/unmanaged)?

✓ **Perform a full inventory of business-provided devices.**
✓ **Describe how the company manages BYOD (if applicable).**
✓ **Find out if fully unmanaged devices are allowed, and under which conditions.**

DAAS: DATA, ASSETS, APPLICATIONS, SERVICES
MFA: MULTI-FACTOR AUTHENTICATION

PAM: PRIVILEGED ACCESS MANAGEMENT
DLP: DATA LOSS PREVENTION

DRM: DATA RIGHTS MANAGEMENT
UEBA: USER AND ENTITY BEHAVIOUR ANALYTICS

HANDS-ON APPROACH TO ZERO TRUST

Preparation for Zero Trust (continued)

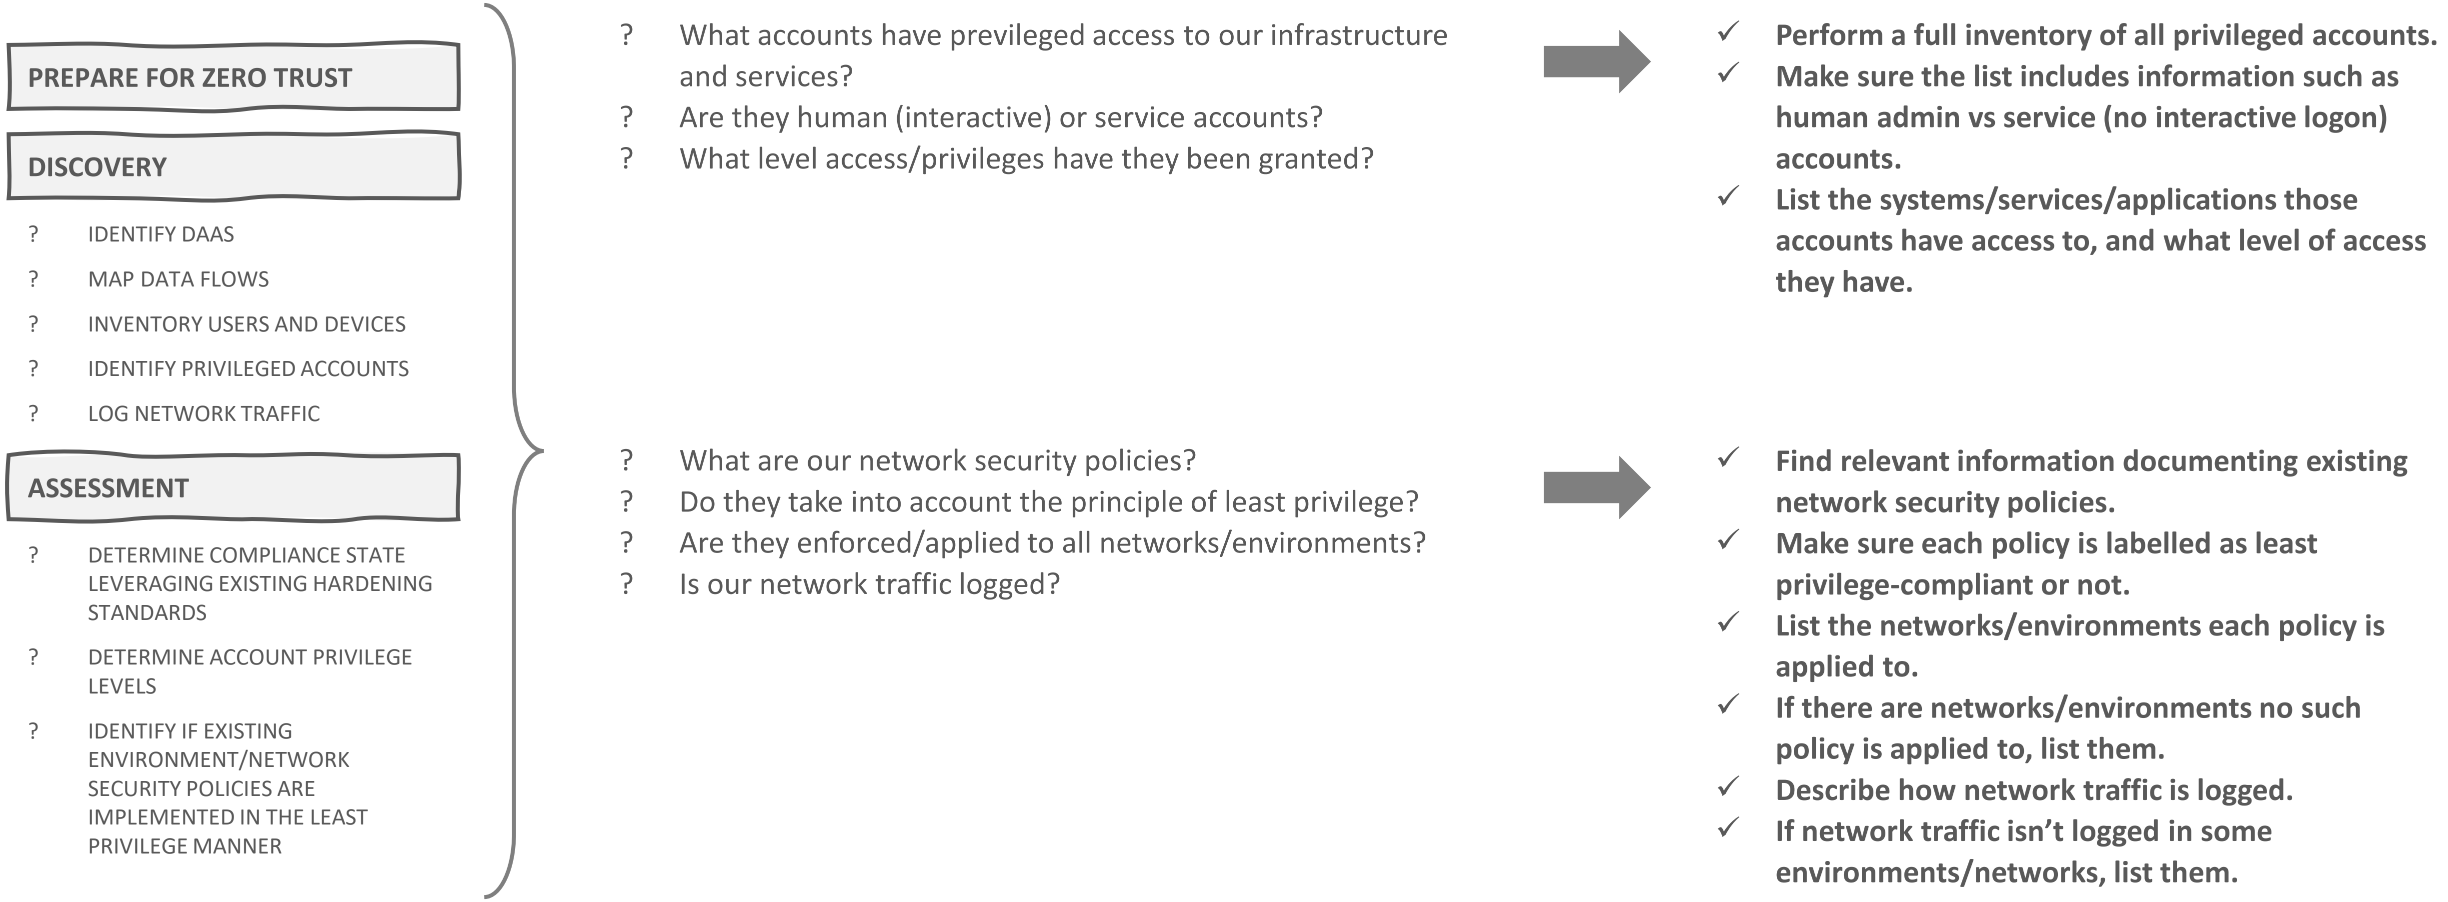

PREPARE FOR ZERO TRUST

DISCOVERY

- ? IDENTIFY DAAS
- ? MAP DATA FLOWS
- ? INVENTORY USERS AND DEVICES
- ? IDENTIFY PRIVILEGED ACCOUNTS
- ? LOG NETWORK TRAFFIC

ASSESSMENT

- ? DETERMINE COMPLIANCE STATE LEVERAGING EXISTING HARDENING STANDARDS
- ? DETERMINE ACCOUNT PRIVILEGE LEVELS
- ? IDENTIFY IF EXISTING ENVIRONMENT/NETWORK SECURITY POLICIES ARE IMPLEMENTED IN THE LEAST PRIVILEGE MANNER

- ? What accounts have previleged access to our infrastructure and services?
- ? Are they human (interactive) or service accounts?
- ? What level access/privileges have they been granted?

- ✓ **Perform a full inventory of all privileged accounts.**
- ✓ **Make sure the list includes information such as human admin vs service (no interactive logon) accounts.**
- ✓ **List the systems/services/applications those accounts have access to, and what level of access they have.**

- ? What are our network security policies?
- ? Do they take into account the principle of least privilege?
- ? Are they enforced/applied to all networks/environments?
- ? Is our network traffic logged?

- ✓ **Find relevant information documenting existing network security policies.**
- ✓ **Make sure each policy is labelled as least privilege-compliant or not.**
- ✓ **List the networks/environments each policy is applied to.**
- ✓ **If there are networks/environments no such policy is applied to, list them.**
- ✓ **Describe how network traffic is logged.**
- ✓ **If network traffic isn't logged in some environments/networks, list them.**

DAAS: DATA, ASSETS, APPLICATIONS, SERVICES
MFA: MULTI-FACTOR AUTHENTICATION

PAM: PRIVILEGED ACCESS MANAGEMENT
DLP: DATA LOSS PREVENTION

DRM: DATA RIGHTS MANAGEMENT
UEBA: USER AND ENTITY BEHAVIOUR ANALYTICS

HANDS-ON APPROACH TO ZERO TRUST

Reaching the baseline level

BASELINE

- ✓ DAAS IDENTIFIED
- ✓ ACCESS DAAS IS DETERMINED BY CYBERSEC POLICIES
- ✓ NETWORK SEGMENTATION WITH DENY ALL/ALLOW BY EXCEPTION
- ✓ MANAGED DEVICES COMPLIANT WITH IT SECURITY POLICIES
- ✓ PRINCIPLE OF LEAST PRIVILEGE IN PLACE
- ✓ MFA IN PLACE
- ✓ DATA CLASSIFICATION AND CRITICAL DATA TAGGING STARTED
- ✓ ENCRYPTION IN PLACE

A full inventory of all our data, assets, applications and services has been made.

- ✓ This (or these) inventory(ies) is (are) available for internal consultation.
- ✓ We should avoid Excel documents and privilege the use of a CMDB.
- ✓ Individuals must be tasked with keeping the information it contains up to date at all times.

Cybersec policies have been defined to dictate who/what has access to our DaaS, in which circumstances.

This includes:

- **Human-machine communications**
- **Machine-machine communications**
- **Interactive and non-interactive accounts**
- **Internal accounts and B2B or B2C access**
- **Internal and managed devices**

- ✓ The policies defined must always have in mind the principle of least privilege – just the right level of access is granted when needed, only for as long as needed.
- ✓ These policies will include, but will not be limited to:
  - RBAC
  - Password complexity
  - Access logging and tracing
  - PAM
  - C2C device policies
  - MFA required
  - Ecryption standards
  - Access only from certain devices/networks
  - Etc.
- ✓ Staff must be made aware of the contents of the cybersec policies, which must be centrally available for internal consultation.

DAAS: DATA, ASSETS, APPLICATIONS, SERVICES
MFA: MULTI-FACTOR AUTHENTICATION

PAM: PRIVILEGED ACCESS MANAGEMENT
DLP: DATA LOSS PREVENTION

DRM: DATA RIGHTS MANAGEMENT
UEBA: USER AND ENTITY BEHAVIOUR ANALYTICS

HANDS-ON APPROACH TO ZERO TRUST

Reaching the baseline level

BASELINE

- ✓ DAAS IDENTIFIED
- ✓ ACCESS DAAS IS DETERMINED BY CYBERSEC POLICIES
- ✓ NETWORK SEGMENTATION WITH DENY ALL/ALLOW BY EXCEPTION
- ✓ MANAGED DEVICES COMPLIANT WITH IT SECURITY POLICIES
- ✓ PRINCIPLE OF LEAST PRIVILEGE IN PLACE
- ✓ MFA IN PLACE
- ✓ DATA CLASSIFICATION AND CRITICAL DATA TAGGING STARTED
- ✓ ENCRYPTION IN PLACE

Our network is segmented and incoming traffic is denied by default.

- ✓ Incoming traffic to each segment will be denied unless a specific and justifiable rule is in place allowing traffic from specific sources, on specific ports, using specific protocols.
- ✓ Rules for M2M communications should be done in the narrowest way possible, no two services/applications that don’t have a real need to communicate should be allowed to.
- ✓ SDN is used whenever possible (reinforcing the above).
- ✓ NACLs and firewall rule-bases should be subject to automated monitoring for the identification and removal of obsolete rules.
- ✓ Traffic must be logged and logs must be inalterable and kept for a certain amount of time (to be defined according to cybersec policy).

Data classification is undergoing and critical data tagging his started

- ✓ Data classification and tagging should be automated for efficiency through maximum exploitation of metadata.
- ✓ Data access logs should start being collected to prepare for the automation of decision-making supported by UEBA.

DAAS: DATA, ASSETS, APPLICATIONS, SERVICES
MFA: MULTI-FACTOR AUTHENTICATION

PAM: PRIVILEGED ACCESS MANAGEMENT
DLP: DATA LOSS PREVENTION

DRM: DATA RIGHTS MANAGEMENT
UEBA: USER AND ENTITY BEHAVIOUR ANALYTICS

SOURCES

SOURCES

MAIN:

- **NIST 800-207** (download here)
- **Department of Defense (DOD) Zero Trust Reference Architecture** (download here)

OTHERS:

- Splunk - A Guide to Embracing a Zero Trust Security Model (download here)
- Fortinet - Build a Secure Remote Connection Solution for Today's Business (download here)
- Palo Alto - The State of Zero-trust Security Strategies (download here)
- https://www.cyberark.com/what-is/zero-trust/
- https://www.ssi.gouv.fr/agence/publication/le-modele-zero-trust/
- https://www.cloudflare.com/learning/security/glossary/what-is-zero-trust/
- http://www.actiac.org/documents/zero-trust-cybersecurity-current-trends
- https://www.securityweek.com/nists-zero-trust-taxonomy-introduces-components-threats-and-migration-routes#
- https://aws.amazon.com/fr/security/zero-trust/