ITUPublications

International Telecommunication Union
Standardization Sector

ITU-T Technical Report

(03/2024)

TR.zt-acp

Guidelines for zero-trust based access control platform in telecommunication networks

Technical Report ITU-T TR.zt-acp

Guidelines for zero-trust based access control platform in telecommunication networks

Summary

Technical Report ITU-T TR.zt-acp specifies a framework to solve the risk of excessive trust, using unknown equipment, privilege abuse, and exposing resources in the process of resource access in telecommunication networks. Based on the access control platform, the framework utilizes key advantages of zero-trust technologies, realizes enhanced identity management, continuous trust evaluation, dynamic access control and continuous real-time audit.

Keywords

Access control, zero trust.

Note

This is an informative ITU-T publication. Mandatory provisions, such as those found in ITU-T Recommendations, are outside the scope of this publication. This publication should only be referenced bibliographically in ITU-T Recommendations.

Table of Contents

Technical Report ITU-T TR.zt-acp

Guidelines for zero-trust based access control platform in telecommunication networks

1 Scope

This Technical Report provides guidelines for zero-trust (ZT) based access control platform (ACP) in telecommunication networks. This Technical Report covers the following topics:

– Security challenges of resource access process in telecommunication networks;
– Security requirements to settle these challenges;
– Framework of zero-trust based access control platform in telecommunication networks;
– Workflow of zero-trust based access control platform in telecommunication networks.

2 References

None.

3 Definitions

3.1 Terms defined elsewhere

This Technical Report uses the following terms defined elsewhere:

3.1.1 access control [b-ITU-T X.1252]: A procedure by which an administrator can restrict access to resources, facilities, services, or information based on pre-established rules and specific rights or authority associated with the requesting party.

3.1.2 authentication [b-ITU-T X.1252]: Formalized process of verification that, if successful, results in an authenticated identity for an entity.

3.1.3 authorization [b-ITU-T X.1252]: The granting of rights and, based on these rights, the granting of access.

3.1.4 bastion host [b-ISO/IEC 27033-1]: Specific host that is used to intercept packets entering or leaving a network and the system that any outsider must normally connect with to access a service or a system that lies within an organization's firewall.

3.1.5 entity [b-ITU-T X.1252]: Something that has a separate and distinct existence and which can be identified in context.

3.1.6 identity [b-ITU-T X.1257]: Set of attributes related to an entity.

3.1.7 resource [b-ISO/IEC 2382]: Any element of a data processing system needed to perform required operations.

NOTE – Example: Storage devices, input-output units, one or more processing units, data, files, and programs.

3.1.8 user [b-ITU-T X.1252]: Any entity that makes use of a resource, e.g., system, equipment, terminal, process, application, or corporate network.

3.1.9 zero trust [b-NIST SP 800-207]: A collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised.

3.1.10 zero trust architecture [b-NIST SP 800-207]: An enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies.

3.2 Terms defined in this Technical Report

This Technical Report defines the following terms:

3.2.1 access control platform: A set of functions such as account management, asset management, authentication management, authorization management, and audit management, which together provides the capabilities of the access control in the zero-trust architecture.

3.2.2 gateway: Point of connection between networks, or between subgroups within networks, intended to provide a necessary protocol translation, or to establish a security tunnel with the agent, or to protect a network according to a given security policy.

3.2.3 implicit trust zone: An area where all the entities are trusted to at least the level of the last access control gateway.

3.2.4 security capability: A capability that relates to a particular security feature.

4 Abbreviations and acronyms

This Technical Report uses the following abbreviations and acronyms:

ACP	Access Control Platform
CDM	Continuous Diagnostics and Mitigation
EDR	Endpoint Detection and Response
ID	Identifier
MTLS	Mutual Transport Layer Security
SPA	Single Packet Authorization
TIS	Threat Intelligence Service
VPN	Virtual Private Network
ZT	Zero Trust
ZTA	Zero-trust Architecture
ZT-ACP	Zero-trust Access Control Platform

5 Conventions

None.

6 Overview

In the telecommunication network, the access control platform (ACP) provides functions of account management, asset management, authentication management, authorization management, and audit management. It provides security capabilities of access control to network and device administrators of the telecommunication operators in telecommunication networks. The ACP framework, which comprises control layer and data layer, is shown in Figure 1.

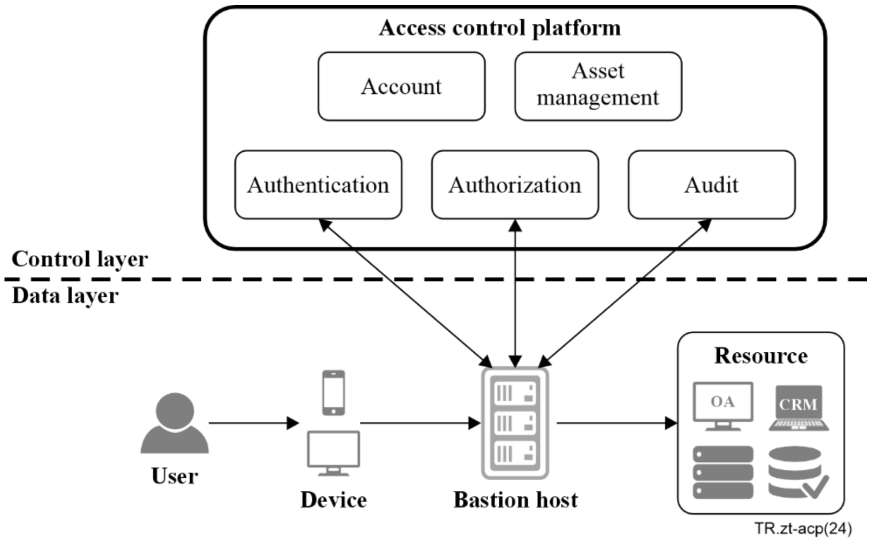

Figure 1 – Framework of access control platform in telecommunication networks

When a user (i.e., network and device administrator) needs to access the resources in the telecommunication network, the user needs to request to access the resources via the bastion host. The bastion host forwards the user's access request to the authentication module in ACP. The authorization module authorizes the user to access the target resources based on security policies. During the access process, the bastion host needs to collect and forward the logs to the audit module.

7 Security challenges and requirements of the access control platform in telecommunication networks

7.1 Security challenges

With the introduction of new technologies such as cloud computing and the increasing applications of remote access, the new security challenges arise as follows:

– Excessive trust

In the traditional perimeter-based model, location determines the degree of trust. The users in the implicit trust zone are trusted by default, which gives users excessive access rights resulting in excessive trust.

– Usage of unknown devices

During authentication, the user and the device are not bounded. Unknown devices may be used by users to access the resources which will bring security threats. The user's credentials may be stolen and used by attackers to access the resources. The authentication which only relies on the user's identity is not enough to guarantee security.

– Privilege abuse

During the access process, the devices, resources and network environment may become insecure, and the user may have abnormal behaviours. The ACP cannot collect the risk information on time, and it cannot change the user's access right dynamically. The audit module may help to discover the security threat after the event.

– Exposing resource

With the introduction of cloud environments into telecommunication networks, the application of remote access increases significantly. The virtual private network (VPN) device that provides remote access is frequently exposed to high-risk vulnerabilities. The open resource access mode of "connect first, authenticate later" provided by the VPN maximizes the exposure of intranet resources to the attacker.

Therefore, single identity authentication, static access control, and post-audit provided by the traditional access control platform cannot fully cope with the new security challenges of resource access process in telecommunication networks.

7.2 Security requirements

The security requirements of the access control platform are as follows:

– All users, service and device need to be identified. An identity can represent a user (a human), service (software process) or device. Each of them needs to be uniquely identifiable in a zero-trust architecture (ZTA).

– All the requests need to be authenticated and authorized.

– All communication needs to be secured to ensure confidentiality, integrity and availability. All of the exchanged information needs to provide source authentication.

– The device and the user which initiates the request need to be treated as one entity during the accessing process. The threat related to the users (e.g., missing user credential) and devices (e.g., usage of unknown devices) can be decreased.

– The risk information needs to be collected continuously and the trust level is evaluated based on the risk information. The authorization needs to be adjusted dynamically based on the evaluation results, besides the traditional static access control methods, such as role-based or attribute-based access control.

– The agent and the controller need to provide pre-authentication, which can help to reduce the network exposure.

– The access to individual resources is granted on a per-session basis. Trust in the requester is evaluated before the access is granted.

– The entire enterprise's local network is not considered an implicit trust zone.

8 Framework of the zero-trust based access control platform in telecommunication networks

8.1 Brief introduction of zero trust

Zero trust is a cyber-security paradigm focused on resource protection and the premise that trust is never granted implicitly but must be continually evaluated. Zero-trust architecture is an end-to-end approach to enterprise resource and data security that encompasses identity, credentials, access management, operations, endpoints, hosting environments and the interconnecting infrastructure [b-NIST SP 800-207]. The introduction of zero trust into the access control platform can meet the security requirements in clause 7.2. The core zero-trust logical components can be found in Appendix I.

8.2 Reference framework

Based on the traditional ACP in Figure 1, zero-trust technology could be used to reconstruct the telecommunication operator's network operation and maintenance system, strengthen the ability of authentication and audit, and provide new security capabilities (e.g., risk data collection, continuous trust evaluation, dynamic access control). The reference framework of zero-trust based access control platform (ZT-ACP) is shown in Figure 2. In Figure 2, new components introduced for the reference framework of zero-trust based access control platform in telecommunication networks are indicated by the gray-shaded boxes.

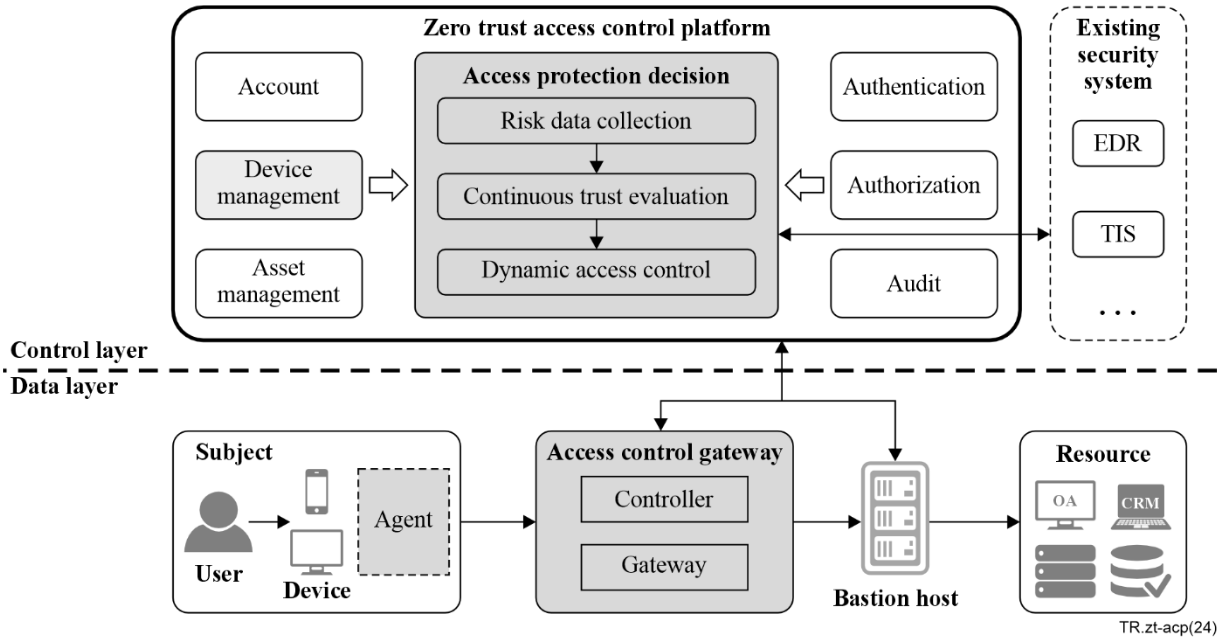

Figure 2 – Reference framework of zero-trust based access control platform in telecommunication networks

8.3 Components

8.3.1 Agent

The agent provides the following functions:

– It collects and reports the device information to the device management module in the zero-trust access control platform (ZT-ACP).

– It collects and reports the risk data to the risk data collection module in the ZT-ACP.

– It routes all the resource access requests to the access control gateway.

8.3.2 Access control gateway

8.3.2.1 Controller

The access control gateway is composed of a controller and a gateway. The controller has the following functions:

– The controller authenticates the subject and establishes a security tunnel between the gateway and the agent before the subject accesses the resources.

– The controller receives the dynamic access control decisions from the dynamic access control module, interprets the decisions and passes the decisions to the gateway.

– The controller records and forwards the real-time logs to the audit module. The logs include access control logs (e.g., exception logs of the access control gateway, user access request initiation logs, and the dynamic decision logs) and access behaviour logs (e.g., normal and abnormal behaviour logs about user access through the access control gateway).

8.3.2.2 Gateway

The gateway performs dynamic decisions and establishes a security tunnel with the agent.

8.3.3 Bastion host

The bastion host receives and implements the dynamic access control decisions from the dynamic access control module.

8.3.4 Device management

Device management module has the following functions:

– Device identity registration management. The device management module collects device information (e.g., device fingerprinting) via the agent, and assigns a unique device identifier (ID) to the device and registers the device.

– Device identity verification. It verifies and authenticates devices by manual intervention of administrators. The verified devices could be registered in the ZT-ACP. The identities of the device and the associated user are bounded to be an authenticated subject. A device with a unique device identifier can be bounded to multiple users. The authenticated subjects will be stored in the device management module as a trusted list.

– Device management. Provide the capability to manage devices with various kinds of operation systems (such as Windows, Mac, iOS, Android, etc.), and allow administrators to register, audit, and deregister the device ID.

8.3.5 Access protection decision

8.3.5.1 Risk data collection

The risk data collection module collects the user related information, device related information, and network environment information. The user related information includes the user account, abnormal behaviour, etc. The device related information includes hardware signatures, operation system version, vulnerability risk level, etc. The device related information can be obtained from existing security systems, such as an endpoint detection and response (EDR), and threat intelligence service (TIS). The network environment information includes protocols, credentials, etc. The collected risk data are used for continuous trust evaluation.

8.3.5.2 Continuous trust evaluation

The continuous trust evaluation module analyses the logs and risk data, evaluates the risk of the access process, and ranks the risk level. The trust evaluation algorithm is the process used by the policy engine to ultimately grant or deny access to a resource. The risk assessment level is used for dynamic access control.

8.3.5.3 Dynamic access control

It makes dynamic access control decisions based on the risk assessment level. The dynamic access control decisions will be transmitted to the access control gateway or the bastion host.

9 Workflow of the zero-trust based access control platform in telecommunication networks

9.1 Pre-authentication

Before a user initiates an access request to the ZT-ACP, a secure connection between the agent and the gateway will be established. The agent information needs to be registered and recorded in the device management module before it connects to the gateway.

1) The agent sends a single packet authorization (SPA) request to the controller. The request contains the identities of devices and users (IP+Port, device ID, user ID, etc.), usually called subject identity.

2) The controller extracts and verifies the subject ID in the SPA request to ensure it is in the trusted list of the device management module.

3) After the successful verification, the controller sets the firewall policy for the gateway based on the subject ID. There will be no response and no action if the verification fails.

4) The agent connects the gateway and establishes a mutual transport layer security (MTLS) connection.

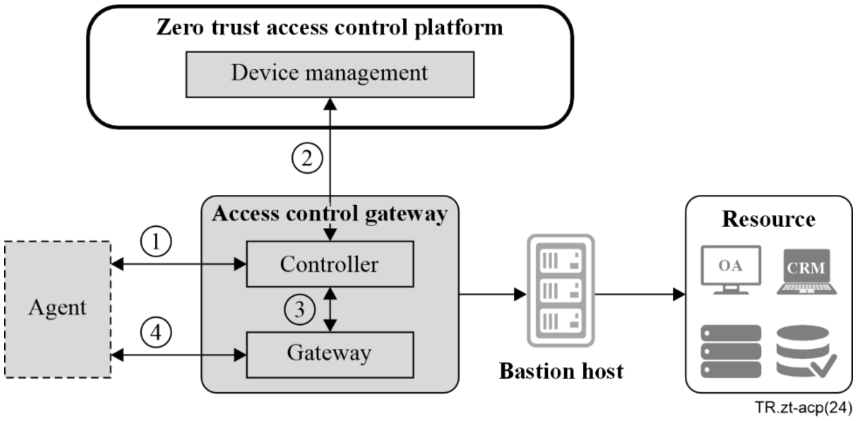

Figure 3 – Workflow of pre-authentication

9.2 Logging in to the ZT-ACP

The agent needs to login to the ZT-ACP before it accesses the target resources. In this process, the security of the agent is assessed dynamically.

1) The agent sends an access request to the ZT-ACP via the controller. The request contains the identities of the device and the user (i.e., the subject).
2) The subject is authenticated by the authentication module in the ZT-ACP.
3) The access protection decision module evaluates the trust level of the subject, generates and sends an access policy message to the authentication module based on the evaluation. The access policy message indicates the result to access the ZT-ACP. For example, the result could be granted, denied, or granted after re-authentication.
4) The authentication module forwards the authentication result, the ZT-ACP information and the access policy to the controller.
5) The controller forwards the authentication result and the access policy to the agent. The ZT-ACP information will be forwarded to the gateway by the controller if the access is granted.
6) The agent logins to the ZT-ACP through the gateway.

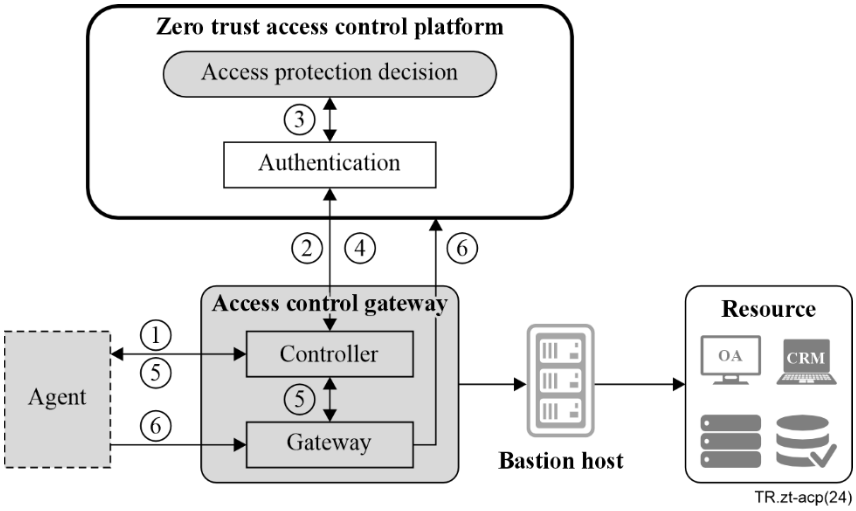

Figure 4 – Workflow of login to the ZT-ACP

9.3 Accessing the application resource

The agent selects the application resource or system resource to be accessed when it logs in the ZT-ACP. The workflow to access the application resource is as follows:

1) The agent sends an application resource access request to the access protection decision module.
2) The access protection decision module evaluates the trust level of the subject, generates and sends an access policy message, a gateway routing control message, and an application resource message to the controller.
3) The controller forwards the access policy message to the agent.
4) The controller forwards the gateway routing control message and the application resource message to the gateway.
5) The agent establishes an encrypted tunnel with the gateway. The gateway configures the route from the agent to the application resource. Then the user can login and access the application resource.
6) The access protection decision module evaluates the trust level of the access process continually. When the trust level changes, it updates and sends the access policy to the gateway.

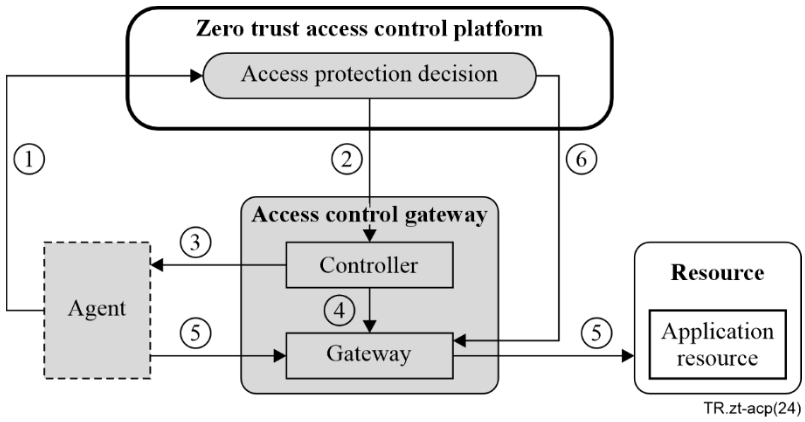

Figure 5 – Workflow of accessing the application resource

9.4 Accessing the system resource

The workflow to access the system resource is as follows:

1) The agent sends a system resource access request to the access protection decision module.
2) The access protection decision module evaluates the trust level of the subject, generates and sends an access policy message, a gateway routing control message and a bastion host message to the controller.
3) The controller forwards the access policy message to the agent.
4) The controller forwards the gateway routing control message and the bastion host message to the gateway.
5) The agent establishes an encrypted tunnel with the gateway. The gateway configures the route from the agent to the bastion host, and the bastion host establishes a connection with the system resource. Then the user can login and access the system resource.
6) The access protection decision module evaluates the trust level of the access process continually. When the trust level changes, it updates and sends the access policy to the gateway.

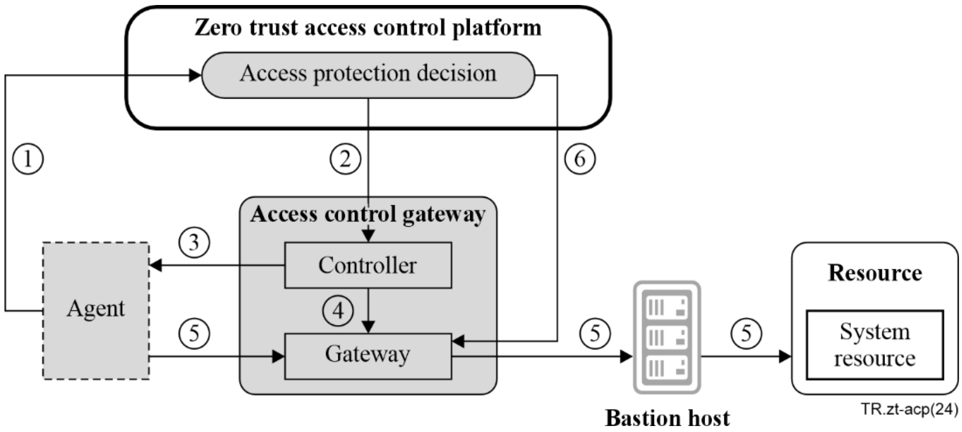

Figure 6 – Workflow of accessing the system resource

Appendix I

The core zero-trust logical components

The core zero trust logical components in Figure I.1 shows the basic relationship between the components and their interactions [b-NIST SP 800-207].

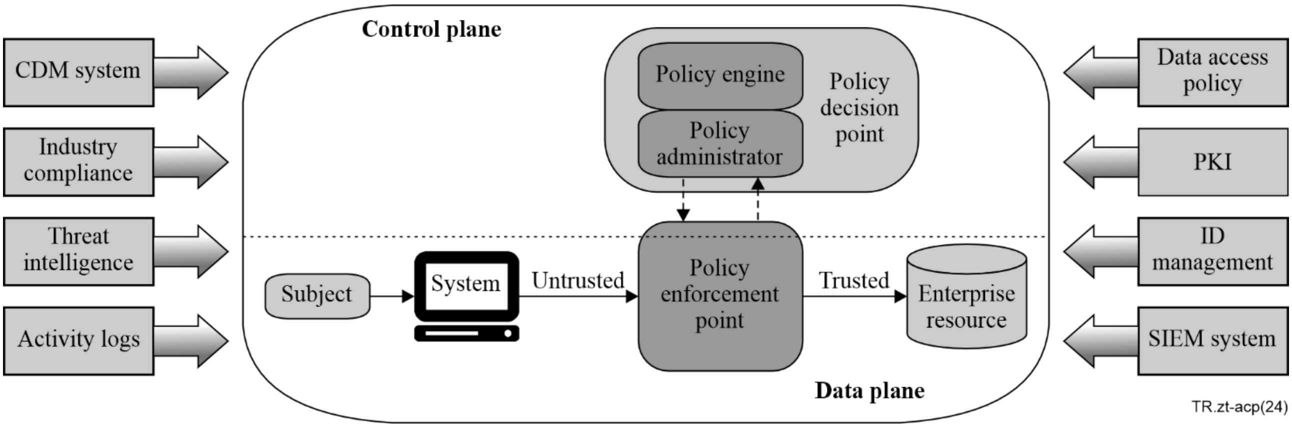

Figure I.1 – Core zero trust logical components

The components are described as follows:

– Policy engine (PE): This component is responsible for the ultimate decision to grant access to a resource for a given subject. The PE uses the enterprise policy as well as the input from external sources (e.g., continuous diagnostics and mitigation (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.

– Policy administrator (PA): 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 the PEP to shut down the connection.

– 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 a 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.

Bibliography

[b-ITU-T X.1252] Recommendation ITU-T X.1252 (2021), *Baseline identity management terms and definitions*.

[b-ITU-T X.1257] Recommendation ITU-T X.1257 (2016), *Identity and access management taxonomy*.

[b-ISO/IEC 2382] ISO/IEC 2382:2015, *Information technology – Vocabulary*.
<https://www.iso.org/standard/63598.html>

[b-ISO/IEC 27033-1] ISO/IEC 27033-1:2015, *Information technology – Security techniques – Network security – Part 1: Overview and concepts*.
<https://www.iso.org/standard/63461.html#:~:text=%2D%20introduces%20how%20to%20achieve%20good,addresses%20the%20issues%20associated%20with>

[b-NIST SP 800-207] National Institute of Standards and Technology Special Publication 800-207, *Zero Trust Architecture*.
<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>