TCG Trusted Network Communications

Clientless Endpoint Support Profile

Specification Version 1.0
Revision 14
18 May 2009
Published

Contact:
      admin@trustedcomputinggroup.org

TCG PUBLISHED

TCG

IWG TNC Document Roadmap

```
                                    ┌──────────────────┐      ┌──────────────┐
                            ┌──────│ Certificate      │      │   IF-IMC     │
                    ┌───────────────┐│ Profiles v1.0    │      └──────────────┘
                    │  Credentials  │└──────────────────┘
                    └───────────────┘┌──────────────────┐      ┌──────────────┐
                            └──────│     Trust        │      │   IF-IMV     │
                                       │  Credentials     │      └──────────────┘
                    ┌───────────────┐ └──────────────────┘
                    │  Migration    │                            ┌──────────────┐
                    │  and Backup   │                            │   IF-PTS     │
  ┌──────────────┐  └───────────────┘                            └──────────────┘
  │ Infrastructure│
  │ Architecture: │ ┌───────────────┐                            ┌──────────────┐
  │   Part I:     │ │     SKAE      │                            │  IF-TNCCS    │
  │ Interoperability│└───────────────┘                           └──────────────┘
  │ Architecture │
  └──────────────┘ ┌───────────────┐                            ┌──────────────┐
                    │     TNC       │                            │    IF-M      │
                    │  Architecture │         ┌──────────────┐   └──────────────┘
                    └───────────────┘         │     TNC      │
                    ┌───────────────┐         │  Use Cases   │   ┌──────────────┐
                    │     IWG       │         └──────────────┘   │    IF-T      │
                    │  Use Cases    │         ┌──────────────┐   └──────────────┘
                    └───────────────┘         │    Other     │
                    ┌───────────────┐         │  Use Cases   │   ┌──────────────┐
                    │     Core      │         │    .....     │   │   IF-PEP     │
                    │   Integrity   │         └──────────────┘   └──────────────┘
  ┌──────────────┐  │    Schema     │
  │ Infrastructure│ └───────────────┘                            ┌──────────────┐
  │ Architecture: │ ┌───────────────┐                            │    TLS-      │
  │ Part II: Integrity│ │   Integrity   │                        │ Attestations │
  │ Management   │ │    Report     │                            └──────────────┘
  └──────────────┘ │    Schema     │
                    └───────────────┘                            ┌──────────────┐
                    ┌───────────────┐                            │   IF-MAP     │
                    │     RIMM      │                            └──────────────┘
                    │    Schema     │
                    └───────────────┘                            ┌──────────────┐
                                                                  │    CESP      │
                                                                  └──────────────┘
```

Acknowledgement

The TCG wishes to thank all those who contributed to this specification. This document builds on considerable work done in the various working groups in the TCG.

Table of Contents

1 Introduction

1.1 Scope and Audience

The Trusted Network Communications Work Group (TNC) defines an open solution architecture that enables network operators to enforce policies regarding endpoint integrity when granting access to a network infrastructure. Today's networks contain many "clientless endpoints", legacy devices that do not have a functional TNC client and therefore do not support integrity checking. In the absence of standards addressing clientless endpoints, every vendor may handle them in a different manner, negating the interoperability provided by TNC and causing situations where one vendor's Policy Enforcement Point (PEP) only works with that vendor's Policy Decision Point (PDP) because of different ways in which clientless endpoints are handled.

In addition to these compatibility problems, clientless endpoints represent a security risk in many environments because of the lack of identity and integrity information provided by the client. Mechanisms such as RADIUS-based MAC authentication and Link Layer Discovery Protocol (LLDP) can help classify clientless endpoints by function for authorization purposes; IF-MAP enables a standard way to achieve much greater security than previously available for clientless endpoints by incorporating information such as behavior, provisioning, etc. to infer an endpoint's integrity. The Clientless Endpoint Support Profile (CESP) specifies how PEPs, PDPs, and other TNC entities should handle clientless endpoints to ensure interoperability and enforce compliance in environments where endpoints lack a TNC client.

This specification is integral to the TCG Trusted Network Communications Work Group's (TNC) reference architecture. Specifically, this specification defines a Clientless Endpoint Support Profile, which is used to perform assessment and make access control decisions in the absence of a TNC client on the endpoint.

Architects, designers, developers, and technologists interested in the development, deployment, and interoperation of trusted systems will find this document necessary in providing a specific mechanism for assessing endpoints without a TNC client. Before reading this document any further, the reader should review and understand the TNC architecture as described in [1]. To understand this specification, the reader should review and understand the TNC IF-MAP Binding for SOAP [4].

1.2 Keywords
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [2]. This specification does not distinguish blocks of informative comments and normative requirements. Therefore, for the sake of clarity, note that lower case instances of must, should, etc. do not indicate normative requirements.

2 Background

2.1 Role of the CESP

Many environments contain IP-enabled endpoints that either cannot run a TNCC (e.g. printers, cameras, badge readers, etc.) or may not have a TNCC installed on them (e.g. guest laptops). This document defines a method for enforcing network security policies on clientless endpoints.

The Clientless Endpoint Support Profile offers a reasonable way to handle clientless endpoints. Network devices are designed to be flexible, providing many different types of policy and methods for applying policy; however, this flexibility may also lead to complexity and inability to interoperate. The CESP defines standard methods for handling clientless endpoints, which allows network devices from many vendors to interoperate and provide a predictable set of functionality. A vendor may offer functionality beyond what the CESP addresses; what the CESP provides is a standard set of functionality that users can rely upon across all devices, without limiting vendor creativity.

This solution addresses a variety of environments, including - but not limited to - a standalone switch or access point (AP) applying policy directly; a switch or AP that consults a policy server to determine what policy to apply (and what happens if that policy server is unavailable); a policy server that consults a Metadata Access Point (MAP) for available metadata to determine what policy to apply; and modification of an initial policy decision based on information that becomes available after the endpoint has connected to the network. In environments where a MAP is present, the MAP Server is used to aggregate information about clientless endpoints, and PDPs may use MAP data or static access control configuration to determine and apply permissible levels of access control.

2.1.1 Role of the CESP in the TNC Architecture

The Clientless Endpoint Support Profile utilizes a subset of the standard TNC components to provide a range of security measures for endpoints that, lacking a TNC client, would not otherwise be addressed by the TNC architecture. The PDP may consist of NAA, TNCS, and IMV(s), but only the NAA is utilized in this environment; likewise, the client may lack a TNCS and IMC(s) entirely, or they may be present but unavailable for interrogation. Also, the PDP and PEP may be separate, individual devices in the network (.e.g. a switch configured to consult a RADIUS server for MAC-based authentication) or may be combined in a single network device (e.g. a switch automatically assigning an unresponsive endpoint to a default guest VLAN).
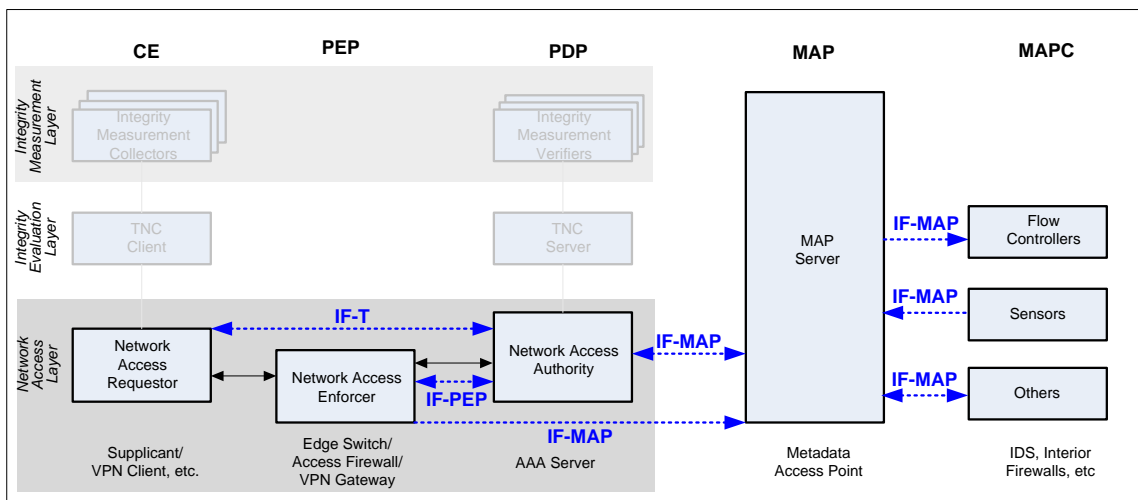

Figure 1 - Clientless Endpoint Support in the TNC Architecture

2.2 Supported Use Cases

Use cases that this version of the Clientless Endpoint Support Profile supports are as follows:

2.2.1   Independent Decision

A clientless endpoint attempts to obtain network access through a combined PEP/PDP with a local access control configuration (e.g., a switch with a designated guest VLAN, or a wireless AP with static access control policies). The combined PEP/PDP provisions access permission, denial, or isolation based upon its local configuration.

Figure 2 - Independent Decision

2.2.2   Consultative Decision

A clientless endpoint attempts to obtain network access through a PEP consuming policy from a PDP. The PDP may have access to a MAP; if so, it queries the MAP to look for information that will help it determine access privileges. The PDP then determines access permission, denial, or isolation.   It may also publish information to or subscribe to notifications from the MAP, if available.

Figure 3 - Consultative Decision

2.2.3   Modification of Decision

A PDP which has previously made an access control decision about a clientless endpoint needs to update that access control decision, possibly based upon a change of policy or because it receives updated information from a MAP. The PDP modifies its access control decision based on the new information.

2.3 Non-supported Use Cases

Several use cases, including but not limited to these, are not covered by this version of the Clientless Endpoint Support Profile:

• A Clientless Endpoint attempts to access a network environment in which no PEP is present.

2.4 Requirements

Here are the requirements that the Clientless Endpoint Support Profile must meet in order to successfully play its role in the TNC architecture.

• Meets the needs of the TNC architecture

The CESP must support all the functions and use cases described in the TNC Architecture as they apply to clientless endpoints.

- Efficient

  The TNC architecture delays network access until endpoint privileges are determined based on available local access control configuration or external metadata. To minimize user frustration, it is essential to minimize delays and make CESP decisions as rapid and efficient as possible.

- Extensible

  The CESP needs to expand over time as new features are added to the TNC architecture. The solution must allow new features to be added easily, providing for a smooth transition and allowing newer and older architectural components to continue to work together.

- Easy to use and implement

  The CESP should be easy for vendors to use and implement. It should allow them to enhance existing products to support the TNC architecture and integrate legacy code without requiring substantial changes. The solution should also make things easy for system administrators and end-users. Individual clientless endpoint access privileges should be determined automatically without requiring manual configuration for each endpoint or access attempt.

- Platform-independent

  Since network environments may contain many different types of clientless endpoints, the solution must function independently of the endpoint platform.

- Truly clientless

  The endpoint must not be expected to load or contain any special software or perform any functions beyond those expected of all network devices.

- Scalable

  The Clientless Endpoint Support Profile must be designed to scale to very large numbers of endpoints. This may require deployment of scalable services, but there should be no scaling limits in the protocol or design.

2.5 Non-Requirement

There are certain requirements that the Clientless Endpoint Support Profile explicitly is not required to meet. This list may not be exhaustive (complete).

- There should be no expectation that the Clientless Endpoint Support Profile will provide the same level of security provided for endpoints with clients. The goal is to provide improved security for clientless endpoints beyond the status quo before the introduction of this specification.

2.6 Assumptions

Here are the assumptions that the Clientless Endpoint Support Profile makes about other components in the TNC architecture.

- Existence of both a PEP and a PDP (either as standalone devices or as a combined PEP/PDP)

2.7 Objectives of the CESP

This section documents the objectives of the Clientless Endpoint Support Profile.

- Describe an infrastructure for the application of policy to devices which do not support or run a TNC Client. This application may be based on specific properties of the clientless devices, such as their MAC address or other metadata.

- Provide an overview of how to use static access control configuration or an external source of endpoint metadata to determine initial access permission, denial, or isolation for a clientless endpoint, and how to modify that determination based on updates to the endpoint metadata.

- Present examples that illustrate how to use compliant devices to apply access control in common use scenarios.

3   CESP Specification

3.1 Overview

The TNC Architecture clearly describes how to assess endpoint integrity and enforce compliance when a TNC Client is present on the endpoint; however, many environments have endpoints that either can't run a TNCC or may not have a TNCC installed on them. This specification provides a standard approach and enforcement mechanisms for deciding whether to allow an endpoint on the network, if it doesn't have a client that can report integrity.

This specification does not create a new binding or define any new syntax; rather, it specifies semantics of devices and required behavior involved in the decision, enforcement, and modification of access control applied to clientless endpoints. The goal of this specification is not to invent new technology, but to make existing technology easier to use, more consistent, and more reliably interoperable.

3.1.1 Components

The basic components of a clientless endpoint solution are: the endpoint itself, an enforcement point applying policy to the endpoint (a PEP), a policy server determining what policy is applied (a PDP), and, in some environments, a metadata clearinghouse (MAP) providing information that can inform a policy decision and other network devices contributing metadata (sensors).

3.1.1.1   Clientless Endpoint

At its most basic, a clientless endpoint is any endpoint that does not (or cannot) run a TNC client and provide verifiable identity and integrity data. That definition encompasses a multitude of real-world scenarios, each of which may lend itself to different access controls and enforcement methods.

When an endpoint attempts to connect to a network, it may offer a wide variety of information about itself, ranging from no information at all to complete identity and integrity information. For the purposes of this specification, we divide that spectrum of information into five classes of endpoints, the first four of which are clientless, and the fifth of which is not:
1. Completely unresponsive
    - Externally-observable MAC address, IP address, behavior
2. Link Layer Discovery Protocol - Media Endpoint Discovery (LLDP-MED)
    - Offers endpoint information such as capabilities, network policy, location, and inventory management. [6]
3. Unauthenticated endpoint
    - Offers invalid credentials (unknown user, failed authentication)
4. Authenticated endpoint with unverifiable integrity
    - Offers valid credentials, but no integrity information
5. Authenticated endpoint with TNC Client - verifiable identity and integrity
    - Offers valid credentials and verifiable integrity information

Case #5, where an endpoint offers verifiable identity and integrity information, is thoroughly addressed by other specifications in the TNC architecture. The Explicit Solution for Clientless Endpoint focuses on the first four cases, where the endpoint is completely unresponsive, offers only LLDP-MED information, is unauthenticated, or offers valid identity but no integrity data.

Examples of a completely unresponsive endpoint include, but are not limited to: a printer, badge reader, or IP-enabled security camera with no 802.1X or LLDP-MED support; a laptop with a wireless client, but no 802.1X supplicant.

Examples of an LLDP-enabled endpoint include, but are not limited to: a media gateway, conference bridge, or media server; an IP phone or a PC running a VoIP softphone.

Examples of an unauthenticated endpoint include, but are not limited to: a laptop running an 802.1X supplicant providing credentials for its native environment, when brought into another environment; or an 802.1X-capable endpoint connecting to a non-802.1X-enabled network device.

Examples of an authenticated endpoint with unverifiable integrity include: an endpoint running an 802.1X supplicant but no TNCC; a laptop running a TNC stack which is configured not to share information with the network; or an endpoint logging into a Windows domain in an environment where an inline PEP is capable of intercepting the Windows login.

To refer back to the TNC architecture, a clientless endpoint - an endpoint without a TNC client - may have no client whatsoever, or it may run a NAR (802.1X supplicant, IPsec VPN client); however, if those components are providing information that is invalid or unverifiable in the environment to which they are trying to authenticate, they offer no more useful information than an endpoint without any client at all.   From an access control policy perspective, all clientless endpoints are similar in that access control decisions cannot be made based solely on identity and/or integrity data volunteered by the endpoint; such decisions must be made based on externally-observable information such as MAC address, IP address, LLDP-MED information, and/or metadata such as behavior or registration.

3.1.1.2   PDP and PEP

A PDP makes policy decisions and provisions them to a PEP over IF-PEP.   A PEP consumes policy decisions from a PDP via IF-PEP and enforces those decisions.   For more detailed descriptions of PDP and PEP, see the TNC Architecture for Interoperability [1].

3.1.1.3   Combined PEP/PDP

A combined PEP/PDP is a single network device performing both PEP and PDP functions. The device both makes policy decisions locally (PDP function) and enforces them (PEP function). Examples include, but are not limited to, a switch with a default VLAN for endpoints that do not respond to EAP messages or that fail authentication, or a firewall with existing static access control policies.

3.1.1.4   Sensors

A sensor publishes information to a MAP via IF-MAP. For more detailed descriptions of Sensors, see the TNC Architecture for Interoperability [1] and the TNC IF-MAP Binding for SOAP [4].

3.1.1.5   MAP

A MAP is a metadata access point - a clearinghouse for metadata, i.e. observable information in a network. Metadata could be information about flows in the network, or information about a specific endpoint that has connected to the network. For a more detailed description of a MAP, see the TNC Architecture for Interoperability [1] and the TNC IF-MAP Binding for SOAP [4].

3.1.1.6   Relationships between components

All of the components above work together to provide access control for clientless endpoints based on externally-observable information. In addition, a single network device can perform the functions of multiple components simultaneously.

One common scenario is guest connectivity in a conference room. In the most basic environment, a guest connects their laptop (the clientless endpoint) to an 802.1X-capable conference room switch (the combined PEP/PDP). The switch attempts to start EAP negotiation, but receives no response from the laptop. After the EAP timeout period, the switch applies its default access policy and assigns the endpoint's port to a guest VLAN, which has access only to the Internet.

On the other end of the spectrum, a more sophisticated environment might include a printer (the clientless endpoint), a switch capable of both 802.1X and MAC authentication (the PEP), a policy server (the PDP), a MAP, a DHCP server (a sensor), and an IPS (also a sensor).

When the printer is rebooted, it brings up its network connection; the switch attempts to start EAP negotiation, but receives no response from the printer. The switch then sends the MAC address of

the printer to the policy server, which looks it up in a database of unmanaged endpoints; the database could reside locally on the policy server, or be provided by a third-party directory service. The database identifies the MAC address as belonging to a printer, so the policy server sends attributes back to the switch to place the printer on the printer VLAN. It then publishes the printer MAC address and role designation as a printer to the MAP, and it subscribes to the MAP for notifications on that endpoint.   The DHCP server allocates an IP address to the printer and publishes the IP address assignment to the MAP, linking the IP address to the printer's MAC address.

A week later, an attacker disconnects the printer and connects his malicious endpoint to the network using the printer's MAC address. The entire scenario above is repeated, and the attacker ends up in the printer VLAN, where he starts probing the network around him. The IPS detects non-printer traffic coming from the IP address of the printer and publishes an update to the MAP for that IP address. The MAP notifies the policy server of the updated information, and the policy server instructs the switch to quarantine or disconnect the misbehaving endpoint.

3.1.2 Approach

3.1.2.1   General Approach

Control of a clientless endpoint may never be as secure as that of an endpoint with a TNC Client, since a TNCC allows interrogation of the endpoint to determine its integrity as well as the identity of the user requesting access. However, with the appropriate tools, information gathered externally to a clientless endpoint - both from the endpoint itself and from other aspects of the network - can be brought to bear upon access control decisions for the clientless endpoint, both at and after connection to the network.

A MAP allows correlation of network traffic, management, and security data; the PDP can leverage that data when provisioning access to a clientless endpoint, as in the scenario above. Because of the breadth and volume of information that a MAP can make available to the decision-making process, the MAP-enabled environment provides the greatest confidence in the security of access control for clientless endpoints; however, not all environments contain a MAP. In the absence of this powerful clearinghouse of centralized metadata, other available security and information-gathering mechanisms may be used to offer lesser, but still valuable, degrees of protection.

When a clientless endpoint connects to the network, the access control device it encounters may be a standalone PEP, or a combined PEP/PDP. The standalone PEP might be a switch or VPN concentrator without any local access control configuration that must consult a PDP to obtain policy. The combined PEP/PDP might be a firewall that applies static ACLs to network traffic, or a switch that assigns an endpoint to a VLAN based on static guest VLAN configuration, or an AP assigning an endpoint to the default VLAN provisioned on the SSID with which it associated.

In all of the above cases, the network device makes the initial access control decision in one of two ways. The combined PEP/PDP makes an independent access control decision, using static local configuration; the standalone PEP makes an access control decision based on consultation with a PDP or a MAP respectively.

3.1.2.2   Independent Decisions

One form of an independent decision is a static local configuration, such as a default guest VLAN, applied to endpoints that do not offer identity or integrity information; local configuration could also include a local data store such as a MAC database or IP address database used to determine appropriate provisioning for the clientless endpoint. In these cases, the PEP may apply additional enforcement mechanisms, such as DHCP snooping and/or dynamic ARP inspection, to increase the level of security associated with the endpoint.

Another instance of an independent access control decision is when a PEP that normally consults a PDP for access control policy loses connectivity with its PDP.   One real-world example is

distributed access control environments, where the PDP is a centralized RADIUS server and the PEP is an 802.1X-capable network device at a remote location. If the WAN link or connectivity between the remote location and central site is disrupted, the PEP may experience a temporary disconnect. With many current implementations, authenticated endpoints retain access until a re-authentication is required by a timeout or reconnection; endpoints attempting to authenticate will be denied access due to the inability of the PEP to reach the PDP.

In that scenario, link downtime leads to extensive local access disruption. This can be avoided by a fallback local access control configuration that is enforced only in cases when the RADIUS server is unreachable.

A third instance of an independent access control decision is the use of LLDP by the PEP to determine access based upon information received from an LLDP agent on the clientless endpoint. In this case, the PEP may assign a VLAN, apply a filter, or start forwarding traffic in accordance with the default port configuration, based on LLDP attributes received from the endpoint.

3.1.2.3   Consultative Decision

A consultative decision occurs when a PEP consults a PDP, which makes an access control decision and provisions it to the PEP, which consumes and enforces that decision. One familiar instance of a consultative decision is an 802.1X negotiation resulting in a RADIUS lookup; another is MAC-based authentication, where an endpoint MAC address is sent to a PDP for comparison against a database of MAC addresses.

In both the above cases, the PDP may also act as the AAA data store, consulting an internal database of credentials or MAC addresses, or it may consult an external AAA data store such as an Active Directory server, LDAP database of MAC addresses, or separate guest access management or endpoint profiling solution.

In addition to validating the authentication credential or endpoint identifier, the PDP may also consult a MAP to determine whether any additional metadata – such as behavior, results of a vulnerability scan, DHCP allocation, etc. – is available to inform the access control decision. For an initial connection to the network, this metadata may not yet have been collected; however, for an access control decision on a connected endpoint, e.g. when the endpoint requests access to privileged resources, metadata may be useful in determining whether to allow access.

3.1.2.4   Modification of Decision

Once the initial access control decision has been made and enforced, circumstances may change, requiring a re-evaluation of permitted access. For example, in an environment with sensors publishing information to a MAP, the PDP may make an initial access control decision for a clientless endpoint, and then subscribe to the MAP for information about that endpoint. If the sensors detect inappropriate or unauthorized activity and publish that information to the MAP, the MAP will notify the PDP, resulting in a re-evaluation of the endpoint's access privileges and probably restriction or termination of the endpoint's access.

3.1.2.5   Data Availability

The various types of enforcement devices apply policy to the differing types of endpoints in predictable ways, depending upon the amount and type of data available for the decision. There may be no data at all, externally observable data, or data provided by the endpoint itself. The more data available to inform an access control decision, the greater the relative security of the clientless endpoint solution. Here are some examples of enforcement mechanisms that may be applied, based on how much data is available:

	Independent Decision	**Consultative Decision**	
Data Used	**Combined PEP/PDP - local default or**	**Separate PEP and**	**Separate PEP and**

	local data	PDP	PDP plus MAP
No data	PEP/PDP provisions default access based on static local configuration	N/A	N/A
LLDP-MED data	PEP/PDP provisions appropriate access based on local authorization data	N/A	N/A
MAC address	PEP/PDP provisions appropriate access based on local authorization data	PDP provisions appropriate access based on local or external authorization data	PDP provisions appropriate access based on local or external authorization data and metadata
IP address	PEP/PDP provisions appropriate access based on local authorization data	PDP provisions appropriate access based on local or external authorization data	PDP provisions appropriate access based on local or external authorization data and metadata
Identity	N/A	PDP provisions appropriate access based on local or external authentication data	PDP provisions appropriate access based on local or external authentication data and metadata
Identity plus integrity	N/A	PDP provisions appropriate access based on local or external authentication data	PDP provisions appropriate access based on local or external authentication data and metadata

Table 1 - Enforcement Options by Data Availability

3.1.2.6   Relative Security Assessment

The level of security provided by the Clientless Endpoint Support Profile depends upon the amount and type of externally-observable information available for consideration when making the access decision. If no information is available, and the default policy is enforced, then a minimal level of security is provided.

If the access control decision is based upon a single, or static, piece of externally-observable information, the environment becomes incrementally more secure, since access can be provisioned to the endpoints based on their function or device identity. Authorizing a phone based on LLDP-MED information enables a switch to assign the endpoint to a VoIP VLAN; authorizing a printer based upon its MAC address enables the policy server to instruct the switch to assign the endpoint to a printer VLAN.

When a MAP is added to the environment, the scope of the access control solution is considerably increased; network security policies can be enforced based on multiple dynamic pieces of information. IF-MAP can be used to gather and centralize information about a device's behavior and status from a variety of network resources, ranging from DHCP servers and vulnerability scanners to IDSes and traffic logs. This metadata can be used to express device classification, and PDPs can consume IF-MAP data to inform policy decisions. For example, correlating a DHCP server's MAC-to-IP address mapping with an IDS's detection of unauthorized traffic enables active, intelligent policy enforcement; an endpoint that was initially provisioned access by MAC address, via RADIUS-based MAC authentication, may have its access privileges modified or revoked based

on unauthorized behavior associated with its IP address. A clientless endpoint environment that leverages a MAP to inform its policy enforcement provides the highest level of security addressed by the CESP.

3.2 Enforcement Mechanisms

For a CESP-compliant solution, the PEP and PDP MUST provide a predictable mechanism for allowing clientless endpoints to receive a set of services. Various mechanisms exist in the market today, including: port-based authorization, local authorization (by MAC or IP address), protocol-based authorization (LLDP), RADIUS-based MAC authentication, and 802.1X authentication without a functioning TNC client. These mechanisms offer varying levels of security:

- o 802.1X authentication provides credential-based authentication with optional integrity verification and offers the most specific intelligence for access authorization and provisioning.

- o RADIUS-based MAC authentication provides device-based authentication and is typically used for authorizing printers and other utility devices. It is easily spoofed, but can leverage external information (such as a MAC database generated by endpoint profiling and monitoring, and/or metadata provided by a MAP) to detect and mitigate such attacks.

- o LLDP offers a mechanism for the clientless endpoint to provide information about itself to the PEP and to receive configuration data from the PEP (VLAN, ACL, QoS). LLDP is subject to same vulnerability as MAC authentication, as the information may be spoofed to gain access.

- o Local authorization (often called local administrative overrides or static MAC / IP bypass) is a mechanism that allows for a MAC- or IP-address-specific rule, resident on the PEP, to target individual devices and provision access; it is often used for quarantine functions. Like RADIUS-based MAC authentication, it is easily spoofed, but without the advantage of verification against an external database or MAP.

- o A default access policy (local configuration of the network port or SSID to a default VLAN or filter) specifies port behavior that provides provisioning of a default "guest access" network. It offers the least security and is commonly used to allow endpoints that may not have a functioning 802.1X supplicant or LLDP agent access to a restricted network, either to gain Internet access or to access a remediation server.

Note that the term "port" is used generically throughout to refer to both wired ports and wireless SSIDs.

The PEP and PDP devices MUST support the following enforcement mechanisms: 802.1X authentication, RADIUS-based MAC authentication, and default access policy. The PEP SHOULD support LLDP-MED to provide local device configuration and network authorization of media-enabled endpoints. The PEP MAY support local authorization based upon MAC or IP address tables.

3.2.1 Precedence of Enforcement Mechanisms

If a PEP or PDP implements multiple authorization mechanisms, the PEP or PDP MUST implement a precedence scheme that determines which authorization mechanism will determine the end-point's resulting authorization. The precedence scheme MAY be configurable, allowing a network administrator to adjust the order in which the options for authorizing a clientless endpoint take precedence. The PEP or PDP MUST be capable of implementing the following precedence order, which SHOULD be the default authorization precedence.

In order of highest to lowest precedence:

• Local Authorization	Device identity
• 802.1X	User identity, device identity

• RADIUS-based MAC	Device identity
• LLDP	Device identity, device class information
• Default Access Policy	Location only

Table 2 - Authorization Precedence

The following diagram illustrates the default CESP authorization precedence:
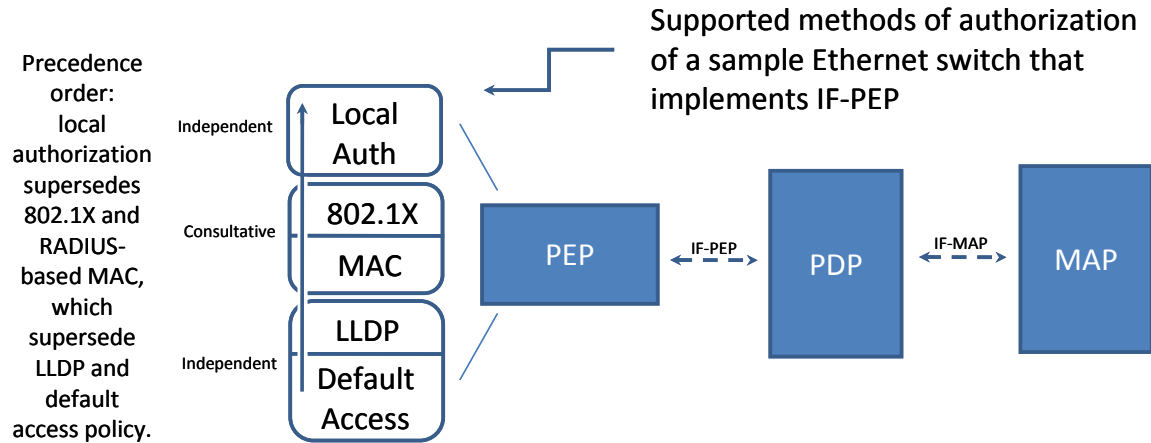

Figure 4 - Authorization Precedence

Not all PEP or PDPs will support the entire spectrum of authorization functions.

From a security standpoint, it is preferable to prioritize enforcement mechanisms based on more-specific information (such as 802.1X, which employs information about a user, and RADIUS-based MAC authentication, which employs information about an individual endpoint) over those based on less-specific information (such as LLDP, which employs information about a class of endpoints; local authorization, which employs information about individual endpoints or a wildcarded category of endpoints; or default access policy, which employs no information about the endpoint). Local authorization is the exception to that rule, due to the purposes for which it is most commonly utilized (local whitelist / blacklist).

3.2.2 Enforcement Mechanisms for Independent Decisions

3.2.2.1   Default Access Policy

The default access policy is the policy of last resort, generally used to assign limited access to unresponsive and/or unrecognized endpoints. For example, visitors may lack an 802.1X supplicant or fail 802.1X authentication, and the MAC address of their endpoint may not be recognized on the host network; a default access policy can assign them to a guest VLAN, perhaps restricted to Internet-only access, or apply a restrictive access control list (ACL) when all other enforcement mechanisms are unavailable or have failed.

In a case where no authentication has succeeded, either because no authentication was offered or because all available authentication failed, an L2 switch or wireless AP (acting as a combined PEP/PDP) MUST provide the following enforcement options:

- Automatic assignment of unresponsive endpoint to a locally-configured VLAN after timeout

- Ability for administrator to choose which VLAN on a per-port or per-SSID basis

and SHOULD provide:

- Automatic application of locally-configured ACLs to unresponsive endpoint after timeout

3.2.2.2   Temporary Disconnect

The temporary disconnect scenario occurs when a standalone PEP, consulting a PDP via RADIUS to receive access control decisions, loses access to the PDP from which it is consuming policy. This scenario is often found in distributed access control environments, where a centralized PDP is used by geographically dispersed PEPs which communicate with it remotely (via WAN links, IPsec tunnels, MPLS, or other connections), when the underlying network connection is interrupted.

In cases where the RADIUS server is identified as disconnected, the PEP MUST provide the following enforcement options:

- Automatic assignment of endpoint to a locally-configured VLAN

- Ability for administrator to choose which VLAN on a per-port or per-SSID basis

and SHOULD provide:

- Automatic application of locally-configured ACLs to endpoint

The PEP MUST provide a configurable method of identifying the RADIUS server as disconnected. The exact method can vary. Typical methods are to have a configurable maximum number of seconds from the last valid communication with the RADIUS server or to have a configurable maximum number of consecutive timeouts in attempting to communicate with the RADIUS server...

In addition, the PEP SHOULD provide a method of identifying when the RADIUS server has become reachable and forcing re-authentication for endpoints permitted access under the temporary disconnect policy. Again, the exact method may vary, but care should be taken to avoid flooding the RADIUS server when connectivity is re-established. One possible method is to set a reduced authentication timeout value on endpoints authenticated via the temporary disconnect mechanism, so that they will frequently re-attempt authentication. If the RADIUS server is still unavailable when an endpoint attempts to re-authenticate, they will be authorized by the temporary disconnect mechanism and receive the short authentication timeout value again; if connectivity to the RADIUS server has been restored, they will authenticate as usual and receive the normal authentication timeout value.

3.2.2.3   Local Authorization

Local authorization allows a network administrator to selectively provision specific, locally-configured access policies based on endpoint MAC address or IP address.

Local authorization based on endpoint MAC address allows for provisioning of a pre-defined VLAN or ACL to endpoints lacking a supplicant or known not to be able to participate successfully in 802.1X. For example, this allows an administrator to assign a voice VLAN, or apply an ACL that permits only voice traffic, to a VoIP handsets based on its MAC OUI; the handset can function on the network without waiting for timeout of an 802.1X negotiation attempt. This feature also allows the network administrator to provision blacklists of undesirable MAC addresses (such as prohibited endpoints or the MAC address 00-11-22-33-44-55, which is commonly used as a placeholder example in documentation).

Another option is local authorization based on endpoint IP address, a method where the PEP extracts observable information (the IP address of the endpoint) and checks to see whether the endpoint has a registered or unregistered IP address. This allows a network administrator to provision blacklists of undesirable IP addresses (bogons) or to permit specific IP addresses used for network monitoring.

A combined PEP/PDP MAY be pre-provisioned with a list of MAC addresses, IP addresses, or both. If local authorization is implemented, the PEP/PDP MUST indicate accept or reject and SHOULD state authorization level. Authorization level for accepted endpoints MAY be defined as a VLAN or ACL.

3.2.2.4   LLDP-MED

LLDP [6] is a link-layer protocol that transmits advertisements containing device information, device capabilities and media specific configuration information periodically to neighbors attached to the

same network. The LLDP agent operates only in an advertising mode, and hence does not support any means for soliciting information or keeping state between two LLDP entities. The LLDP agent advertises information over LLDP Data Units (LLDPDUs) and records the information received from other agents in IEEE-defined MIB modules. [6]

LLDP-MED [9] is an enhancement to LLDP to support the automatic configuration of resources for media-enabled devices providing "plug and play" networking. A layer 2 network device acting as a combined PEP/PDP may use LLDP-MED as a mechanism to assign a VLAN or ACL to an endpoint or start forwarding traffic using the default port settings (VLAN or ACL). By monitoring LLDP for specific attributes, the PEP/PDP may determine that a media device such as an IP telephone (VoIP) or IP security camera has connected. The PEP/PDP MAY use that information to authorize the endpoint to a pre-determined VLAN or ACL, and MAY maintain a table of LLDP device types and access authorization to provide unique access controls or VLAN attributes for specific device types (e.g., mapping VoIP devices to a phone VLAN and IP video camera devices to a camera VLAN).

LLDP defines the following mandatory TLVs. All compliant LLDPDUs MUST contain at a minimum the following four mandated TLVs in the following order:

- Chassis ID TLV (Type = 1)

- Port ID TLV (Type = 2)

- Time To Live TLV (Type = 3)

- End of LLDPDU TLV (Type = 0)

The LLDP-MED specification defines the following set of TIA Organizationally Specific TLVs:

- LLDP-MED Capabilities TLV (OUI = 00-12-BB, Subtype = 1)

- Network Policy TLV (OUI = 00-12-BB, Subtype = 2)

- Location Identification TLV (OUI = 00-12-BB, Subtype = 3)

- Extended Power-via-MDI TLV (OUI = 00-12-BB, Subtype = 4)

- Inventory - Hardware Revision TLV (OUI = 00-12-BB, Subtype = 5)

- Inventory - Firmware Revision TLV (OUI = 00-12-BB, Subtype = 6)

- Inventory - Software Revision TLV (OUI = 00-12-BB, Subtype = 7)

- Inventory - Serial Number TLV (OUI = 00-12-BB, Subtype = 8)

- Inventory - Manufacturer Name TLV (OUI = 00-12-BB, Subtype = 9)

- Inventory - Model Name TLV (OUI = 00-12-BB, Subtype = 10)

- Inventory - Asset ID TLV (OUI = 00-12-BB, Subtype = 11)

CESP does not require a combined PEP/PDP to support LLDP. However, if a Layer 2 PEP/PDP implements LLDP, it MUST support LLDP-MED and MUST also support reporting of the LLDP and LLDP-MED TLVs via SNMP or syslog to a device capable to proxy this information to a MAP (in addition to the reporting requirements specified in section 3.2.2.5).

An example of this is a VoIP handset attaching to an Ethernet switch acting as a combined PEP/PDP. The switch is configured to support LLDP and LLDP-MED and maintain a LLDP profile for VoIP handsets. This profile configures the network port with VLAN and QoS settings to support "plug and play" operation of the VoIP handset. The CESP-compliant switch provisions access for the VoIP handset and generates an informational message that is communicated to a MAP via a proxy, which may be a separate PDP or a network management application.

For CESP compliance, the combined PEP/PDP MUST report the following TLVs when using LLDP as a mechanism to provision access policy:

- Chassis ID TLV (Type = 1)

- Port ID TLV (Type = 2)

- Time To Live TLV (Type = 3)

- End of LLDPDU TLV (Type = 0)

- LLDP-MED Capabilities TLV (OUI = 00-12-BB, Subtype = 1)

- Network Policy TLV (OUI = 00-12-BB, Subtype = 2)

- Location Identification TLV (OUI = 00-12-BB, Subtype = 3)

- Extended Power-via-MDI TLV (OUI = 00-12-BB, Subtype = 4)

and SHOULD report the remaining TLVs:

- Inventory - Hardware Revision TLV (OUI = 00-12-BB, Subtype = 5)

- Inventory - Firmware Revision TLV (OUI = 00-12-BB, Subtype = 6)

- Inventory - Software Revision TLV (OUI = 00-12-BB, Subtype = 7)

- Inventory - Serial Number TLV (OUI = 00-12-BB, Subtype = 8)

- Inventory - Manufacturer Name TLV (OUI = 00-12-BB, Subtype = 9)

- Inventory - Model Name TLV (OUI = 00-12-BB, Subtype = 10)

- Inventory - Asset ID TLV (OUI = 00-12-BB, Subtype = 11)

```
⊞ Frame 1 (268 bytes on wire, 268 bytes captured)
⊞ Ethernet II, Src: HewlettP_57:ca:7f (00:13:21:57:ca:7f), Dst: LLDP_Multicast (01:80:c2:00:00:0e)
⊟ Link Layer Discovery Protocol
  ⊞ Chassis Subtype = MAC address
  ⊞ Port Subtype = Locally assigned
  ⊞ Time To Live = 120 sec
  ⊞ Port Description = 1
  ⊞ System Name = ProCurve Switch 2600-8-PWR
  ⊞ System Description = ProCurve J8762A Switch 2600-8-PWR, revision H.08.89, ROM H.08.5X (/sw/code/build/fish(ts_08_5))
  ⊞ Capabilities
  ⊞ Management Address
  ⊞ IEEE 802.3 - MAC/PHY Configuration/Status
  ⊞ TIA - Media Capabilities
  ⊞ TIA - Network Policy
  ⊞ TIA - Location Identification
  ⊞ TIA - Extended Power-via-MDI
  ⊞ End of LLDPDU
```

Figure 5 - Sample Packet Capture of LLDP-MED Packet With Required TLVs.

3.2.2.5   Reporting Requirements for Independent Decisions

In cases where an independent decision has been made, information about that decision may have significance for access control decisions in other areas of the network. To make that information available, the combined PEP/PDP device MUST provide:

- Ability to report the authorization of an endpoint with the following attributes:

  o MAC address of client

  o Time of authorization

  o Identity of device provisioning access (IP address of combined PEP/PDP)

  o Location of endpoint (Interface or SSID to which endpoint is connected)

  o Authorization level granted (e.g. guest VLAN, ACL, and/or default port assignment)

- Ability to communicate that information using an accepted standard such as SNMP or syslog, to allow for a proxy to be implemented for publication of the information to a MAP. (The proxy could be a management application which receives syslog or SNMP trap information from the combined PEP/PDP and uses a MAP Client to publish the authorization information to the MAP.)

  o If SNMP or syslog is used, the vendor MUST publish the format of SNMP/syslog messages to enable third-party developers to adapt these messages for centralized reporting via IF-MAP.

and SHOULD provide:

- Ability to publish location and authorization attributes to a MAP directly (i.e., by implementing a MAP Client and acting as a Sensor)

3.2.3 Enforcement Mechanisms for Consultative Decisions

3.2.3.1 RADIUS-Based MAC Authentication

MAC authentication is not true authentication, but authorization of endpoints based on their MAC address, and is primarily used in 802.1X environments to handle endpoints without an 802.1X supplicant (such as printers or networked security cameras). The PEP extracts observable information - the MAC address of the endpoint - and sends it to the PDP, which determines whether it is a registered or unregistered MAC address. The PDP consults a local or external pre-provisioned list of MAC addresses to determine authorization of the endpoint. The PDP MUST indicate accept or reject and SHOULD state authorization level. The format of the authorization MUST be compliant with current TNC IF-PEP specification and SHOULD provide Accept, Accept with VLAN ID, Accept with filter-id, or Reject.

Within RADIUS, if a user authentication fails (the identity is unknown, or the credential is invalid), the RADIUS server sends an access-reject. [5] For MAC authentication via RADIUS, an unknown or/ unregistered MAC will therefore trigger an access-reject. The PEP SHOULD deny access upon receiving an Access-Reject. In some environments, it may be desirable to allow endpoints with unregistered MAC addresses limited access to the network to enable discovery of new MAC addresses. In these cases, the PDP may be configured with a wildcard at the end of its MAC database to accept all unregistered MAC addresses and, if desired, place them on a restricted VLAN so that MAC registration or learning may be employed.

The PEP MUST use RADIUS authentication as the mechanism to authorize a particular MAC address (endpoint), and MUST be configured to present the endpoint MAC address to the PDP as an identity credential for validation.

The market currently offers multiple options for format of the MAC address presented by the PEP to the PDP during RADIUS-based MAC authentication:

- Endpoint MAC Address, Byte-Dash Format (RFC 3580)
  Example: 00-11-22-33-44-55 (upper case only)

- Endpoint MAC Address, Colon Format
  Example: 00:11:22:33:44:55

- Endpoint MAC Address, Unseparated Format
  Example: 001122334455

- Endpoint MAC Address, Dot Separated format
  Examples: 0011.2233.4455 or 001122.334455

A PDP MUST accept any of the formats presented for authentication to assure interoperability with legacy installations. A PEP MUST send only one format to a PDP for authorization, and it MUST be one of the above formats. PEP manufacturers SHOULD implement the BYTE-DASH format standardized in RFC -3580, since this format is used for the Called-Station-ID and Calling-Station-ID, to avoid using two different MAC-address formats in the same frame...

In addition to the MAC Address format convention, the various PEPs on the market today use different formats for creating and appending a password to the MAC address – creating an identity credential for authentication and authorization. Some current implementations may use the MAC address as both the identity and the password; other implementations use the MAC address as the identity and a pre-shared secret known to the RADIUS server and the network device as the password. Current options in common use are:

- Username: MAC Address / Password: [empty]
  Example: 00-11-22-33-44-55 / []

- Username: MAC Address / Password: MAC Address
  Example: 00-11-22-33-44-55 / 00-11-22-33-44-55

- Username: MAC Address / Password: [shared secret]
  Example: 00-11-22-33-44-55 / [secret]

A PDP MUST accept any of these formats for compatibility. A PEP SHOULD send the MAC Address / MAC Address format and SHOULD use MS-CHAP-v2 as the authentication protocol.

3.2.3.2   IEEE 802.1X

The IEEE 802.1X standard defines port-based network access control. PEPs MUST support IEEE 802.1X [7] and SHOULD fully support the RADIUS usage guidelines set forth in RFC 3580 [8].

A common 802.1X-related scenario that affects clientless endpoints is authentication failure - when an 802.1X supplicant presents credentials that are invalid in the network environment to which it is requesting access. While the cause may be as simple as an incorrectly typed username or password, the specific failure mode affecting clientless endpoints is an 802.1X supplicant configured for one environment which is then brought into another 802.1X environment that does not share a trust relationship with the endpoint. A common example is a traveling user, with a laptop running an 802.1X supplicant configured to provide credentials for their employer, who plugs their laptop into the 802.1X-enabled conference room switch at another company. The IEEE 802.1X standard permits optional default behavior for the PEP to fail open or allow connection with a default policy upon receiving an access-reject / EAP-failure from the PDP.

In cases where the endpoint has failed 802.1X authentication, the PEP MUST provide the following enforcement options, which MUST be configurable separately from the default local policy:

- Automatic assignment of endpoint to a locally-configured VLAN

- Ability for administrator to choose which VLAN on a per-port or per-SSID basis

and SHOULD provide:

- Automatic application of locally-configured ACLs to endpoint

In addition, the PEP SHOULD provide the alternate ability to fall back through the other enforcement mechanisms in order of precedence.

3.2.4 Enforcement Mechanisms for Modification of Decision

3.2.4.1   RADIUS Change of Authorization (CoA)

RADIUS CoA is a mechanism by which dynamic changes can be made to the access control policy applied to an endpoint, even after the endpoint has connected to the network and undergone an initial authorization. CoA allows quarantine or disconnection of an endpoint which transitions from an authorized state to an unauthorized state while connected to the network; it also provides the ability to expand the access privileges of an endpoint after initial connection if additional information about the endpoint is available (e.g. MAP metadata).

PEPs using RADIUS-based enforcement mechanisms (MAC auth, 802.1X) MUST be fully compliant with RFC 5176.

4 Security Considerations

The purpose of the Clientless Endpoint Support Profile is to determine and manage network access for clientless endpoints in accordance with policies established by the network administrator, using information about those endpoints gathered in a variety of ways. The fact that the endpoints are not configured with TNC client software makes this problem more challenging than the typical TNC environment. However, even a conventional TNC environment is subject to the lying endpoint problem so the threats and problems are not so different.

Securing clientless endpoints and the networks to which they connect is a challenging problem. The techniques described in this document improve the security of this challenging situation. However, not all the security threats relevant to the situation are completely addressed by the techniques described here. Like any other security system, the Clientless Endpoint Support Profile cannot promise perfect security. Still, it provides better security than would otherwise be available.

This section enumerates the threats present in a network with clientless endpoints, the countermeasures that are provided by the techniques described in this document, and the issues still remaining after these countermeasures are applied.

4.1 Threat Model

A variety of attacks can be mounted against a CESP environment.

4.1.1 Falsified Endpoint Information

As described in section 3.1.2.5, there are many differences in the amount of information available regarding clientless endpoints: no information; MAC address; device information obtained via LLDP or device profiling; identity information obtained via a captive portal, 802.1X supplicant, or other mechanism; etc. If this information is incorrect, then improper access may be granted to the endpoint. This section describes ways that endpoint information can be falsified.

- MAC spoofing

  MAC authentication has vulnerabilities, particularly since a MAC address is easy to observe in network traffic and easy to change on certain endpoints. One attack is to unplug a printer, connect a laptop, and have the laptop use the printer's MAC address as its own in order to gain network connectivity. There are variations on an attack model with MAC spoofing and a sample of MAC attacks is presented here.

    o Unauthorized endpoint spoofs authorized MAC address to gain access

    o Unauthorized endpoint spoofs authorized MAC address then acts badly to get authorized endpoint with that MAC address banned

- LLDP spoofing

  An endpoint can send any LLDP information that it wants. If this information is used to determine network access, an endpoint may be able to gain inappropriate access by sending false information via LLDP.

- Identity spoofing

  If identity information is used to determine network access (as with a captive portal) and an attacker can obtain identity credentials (via a keystroke logger, rogue AP, or other mechanism), the attacker can use those credentials to gain unauthorized access.

- Refusal to supply endpoint information

  Since the CESP is designed to handle endpoints that do not have a client and therefore cannot supply much or any endpoint information, a hostile endpoint can simply refuse to supply information about itself. Network administrators should consider this possibility when establishing policies. An unresponsive endpoint may be quite intelligent but simply choosing to not supply information about itself.

4.1.2 Attacks on the MAP

Some CESP models depend on consulting a MAP. The IF-MAP specification contains a complete discussion of the threat model for IF-MAP and for the MAP. Those threats also apply when MAP is used in conjunction with the CESP but they will not be repeated here since they are no different with the CESP than they are in the non-CESP case.

4.1.3 Attacks on the PEP

The PEP is a critical part of the CESP. Attacks against the PEP and against IF-PEP are covered by the Security Considerations sections of the TNC Architecture and the IF-PEP specification. However, the CESP does provide several new attack opportunities.

- MAC Address Attacks on PEP

  Some endpoints can modify their MAC addresses. An attacker could flood the PEP with different MAC addresses. If the PEP is not robust against this attack, it could cause unpredictable failures that might grant unauthorized access to the endpoint. More likely, it could fill tables on the PEP that would result in other endpoints being denied access.

- LLDP Attacks on PEP

  If the PEP's implementation of LLDP is not robust, a malicious endpoint could send malformed LLDP packets that could result in compromise of the PEP, unauthorized access, or degradation or denial of service to other endpoints.

4.1.4 Attacks on the PDP

Attacks against the PDP are covered by the Security Considerations sections of the TNC Architecture. The CESP only presents one new attack opportunity.

- MAC Address Attacks on PDP

  Some endpoints can modify their MAC addresses. An attacker could cycle through many different MAC addresses, resulting in a flood of authentication requests being sent to the PDP. If the PDP is not robust against this attack, it could cause unpredictable failures that might grant unauthorized access to the endpoint or degraded access to others.

4.2 Countermeasures

This section describes countermeasures against the threats described in section 4.1.

4.2.1 Countermeasures Against Falsified Endpoint Information

A variety of countermeasures can be employed against falsified endpoint information. Some of these countermeasures are specific to one kind of endpoint information (e.g. spoofed MAC addresses). Others are more general.

- MAC locking

  PEP or PDP locks MAC address to a particular port, rejecting MAC authentication that comes from other ports. This thwarts MAC spoofing when a MAC address is already in use. To improve mobility and manageability, a MAC locking table entry can be automatically created when a clientless endpoint connects and automatically removed when the endpoint disconnects or after a period of inactivity. Instead of MAC locking (which can block a legitimate device instead of the spoofing device), a duplicate MAC address can result in limited access and close monitoring for both endpoints and in further investigation to track down the spoofing device.

  In cases where duplicate MAC addresses appear on separate switches, IF-MAP provides additional protection against MAC spoofing by associating locality information with a given MAC address. When a MAC address appears on the network, its device classification may be confirmed by sensor. If two devices on the network are using the same MAC address, they may

be differentiated by including locality information (the switch and port to which they are connected) in the MAP database.

- MAP correlation

    The behavior of clientless endpoints should be monitored by Network Behavior Anomaly Detection, Intrusion Detection Systems, Device Characterization systems, etc. Behavior which is not consistent with the endpoint information supplied can be reported to the MAP, which can result in quarantine and other consequences.

- Limited access

    The network access granted to a clientless endpoint should be limited to the access necessary for that endpoint (based on the endpoint information supplied). For example, telephones can be placed on a phone VLAN and printers on a printer VLAN (or restricted with filters or firewall rules). This will reduce the impact of falsified endpoint information, including the refusal to supply endpoint information.

- Identity reputation

    To reduce the impact of identity spoofing, the authentication server can monitor authentication for odd patterns such as the same endpoint being simultaneously authenticated from two locations at once. Suspect behavior can result in the revocation of credentials.

4.2.2 Countermeasures Against Attacks on the MAP

The IF-MAP specification describes many good countermeasures against attacks on the MAP. These will not be duplicated here.

4.2.3 Countermeasures Against Attacks on the PEP

The Security Considerations sections of the TNC Architecture and the IF-PEP specifications describe countermeasures against attacks on the PEP. These will not be duplicated here. However, this section does describe specific countermeasures against the attacks specific to the CESP.

- Limit MAC addresses per port

    To reduce the impact of MAC address flooding attacks, PEPs should limit the number of MAC addresses that they will accept on a port or employ other similarly effective mechanisms. For example, they may provide ways to block MAC authentication on problematic ports or send problem alerts via syslog or other similar mechanisms.

- LLDP robustness

    PEPs should be robust against malformed LLDP packets.

4.2.4 Countermeasures Against Attacks on the PDP

The Security Considerations section of the TNC Architecture describes countermeasures against attacks on the PDP. These will not be duplicated here. However, this section does describe specific countermeasures against the attacks specific to the CESP.

- Scalable PDP

    The PDP should be designed to scale well in the face of a heavy authentication load. This is especially important because MAC address authentication allows a single endpoint to easily trigger a large number of simultaneous authentications. A PDP may wish to throttle MAC

authentications so that, under heavy load, each PEP receives a fair share of attention. That will prevent one endpoint on one PEP from locking out other endpoints using other PEPs.

5 CESP Examples

The Clientless Endpoint Support Profile describes enforcement mechanisms that can be applied to a variety of scenarios:

Scenario	Description	Use	Decision Type	
			Independent	Consultative
Completely unresponsive endpoint	Systems added to network without transmitting observable traffic	Printers, electronic door locks, HVAC systems…	Default Access Policy Local Auth	RADIUS-Based MAC Auth
Endpoint with LLDP-MED	Media endpoint protocol used to communicate information on device type & capabilities	PEP uses information to provision network settings for endpoint media devices such as VoIP phones and IP surveillance systems	LLDP	
Unauthenticated endpoint - fails authentication (802.1X)	Endpoint configured with 802.1X client that fails authentication	Guest Networking for systems / users; remediate internal user that may have misconfigured supplicant	Default Access Policy	RADIUS-Based MAC
Unauthenticated endpoint - fails authentication (MAC)	Endpoint not found in consultative MAC authentication database	Guest Networking for systems / users; remediate internal users with unregistered MAC addresses	Default Access Policy	RADIUS-Based MAC Auth
Authenticated endpoint with unverifiable integrity (802.1X)	Supplicant-enabled endpoints that don't support a TNC client or have TNC client error (misconfigured, disabled, information restricted, etc.)	Verified user or machine identity; remediate internal users' endpoints to correct TNC Client error	N/A	802.1X Auth
Authenticated endpoint with unverifiable integrity (MAC)	Endpoints without an 802.1X supplicant or TNC Client that have registered MAC addresses with policy settings in authentication database	Printers, VoIP telephones, other devices lacking 802.1X supplicant	N/A	RADIUS-Based MAC Auth

Table 3 - Clientless Endpoint Example Scenarios

5.1 Consultative Decision - RADIUS-Based MAC Authentication

Description: A printer is connected to the network and requires provisioning to provide services.

Implementation: Using the RADIUS-Based MAC Authentication enforcement mechanism, the MAC address for the printer must be entered in the Network Access Authority (NAA) database prior to the first connection, with the entry matching the MAC address as the username, and VLAN or Filter-ID attributes associated for forwarding.

The PEP uses IF-PEP to authenticate the printer, using the MAC address learned on the interface, in the BYTE-DASH format in a RADIUS request to the PDP. The PDP consults the NAA database (either locally or remote) and responds with a VLAN or Filter-ID to the PEP.

The PEP starts forwarding traffic using the information contained in the RADIUS-Accept e.g., VLAN or Filter-ID. The printer is then properly connected to the network.

5.2 Independent Decision - Default Access Policy

Description: A company guest connects a mobile device, such as a laptop, to a network for which they do not have domain credentials. The mobile device supports 802.1X and has credentials and a configuration for their home network. The guest would like to access the Internet, and the host company would like to prevent connection to internal networks.

Implementation: The guest mobile device makes a connection to either an open (non-encrypted) wireless LAN or a physical connection to the wired LAN. The mobile device responds to any 802.1X challenge, but is not able to supply proper credentials (wrong domain) or is not properly configured with the correct EAP type (client supports PEAP on a TLS network).

The PEP is configured to request 802.1X from connecting systems, and also with a default port access configuration equaling a guest VLAN or guest forwarding policy (implementations vary by PEP vendor).

The PEP forwards the credentials supplied by the mobile device to the PDP, and the PDP issues a RADIUS-Reject message to the PEP. The PEP then starts forwarding traffic based on the static access control configuration, e.g. guest VLAN.

The guest mobile device is able to access services provisioned by the network administrator: for example, access to the internet but not internal resources.

6 References

6.1 Normative References

[1]   Trusted Computing Group, *TNC Architecture for Interoperability*, Specification Version 1.4, May 2009.

[2]   Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", Internet Engineering Task Force RFC 2119, March 1997.

[3]   Chiba, M., Dommety, G., Eklund, M., Mitton, D., and Aboba, B., "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", RFC 5176, January 2008.

6.2 Informative References

[4]   Trusted Computing Group, *TNC IF-MAP Binding for SOAP*, Specification Version 1.1, May 2009.

[5]   Rigney, C., Rubens, A., Simpson, W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[6]   IEEE Standard 802.1AB-2005, *Station and Media Access Control Connectivity Discovery*, 802.1AB™- May 2005.

[7]   IEEE Standard 802.1X-2004, *Port-Based Network Access Control,* IEEE Std 802.1X™-2004, November 2004.

[8]   Congdon, P., Aboba, B., Smith, A., Zorn, G., and Roese, J., "IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines", RFC 3580, September 2003.

[9]   TIA Standard TIA-1057, *Link Layer Discovery Protocol for Media Endpoint Devices,* ANSI/TIA-1057-2006, April 2006.

7 Glossary

When used in the CESP, the following terms are defined as below. Please also see the TNC Architecture [1] for broader TNC related terminology.

Authenticated Endpoint: Endpoint offering verifiable identity credential when connecting to a network.

Clientless Endpoint: Endpoint that does not run a TNC Client, or does not make TNC Client data available.

Combined PEP/PDP: A single component (often a network device) performing both PEP and PDP functions. The component both makes policy decisions locally (PDP function) and enforces them (PEP function).

Consultative Decision: An access control decision in which a policy server, acting as a standalone PDP, determines what access control policy to apply to an endpoint and provisions that policy to a network device, acting as a standalone PEP. The network device consumes the access control policy from the PDP and enforces it by applying the access control decision to the endpoint.

Independent Decision: An access control decision made by a network device acting as a combined PEP/PDP, which locally determines and enforces access control policy.

Local default-based decision: An access control decision made by a combined PEP/PDP based on default configuration of a port or SSID, such as a defined "guest" VLAN.

Local data-based decision: An access control decision made by a combined PEP/PDP based on locally-stored data, such as MAC address or IP address tables.

 Modification of Decision: A change to an access control decision resulting in revision of the access control policy applied to an endpoint after its initial connection to the network.

Unauthenticated Endpoint: Endpoint offering unverifiable or no identity credential when connecting to a network.

Unresponsive Endpoint: Endpoint offering no data when connecting to a network.

8   Annex A: Selection of Authorization Method

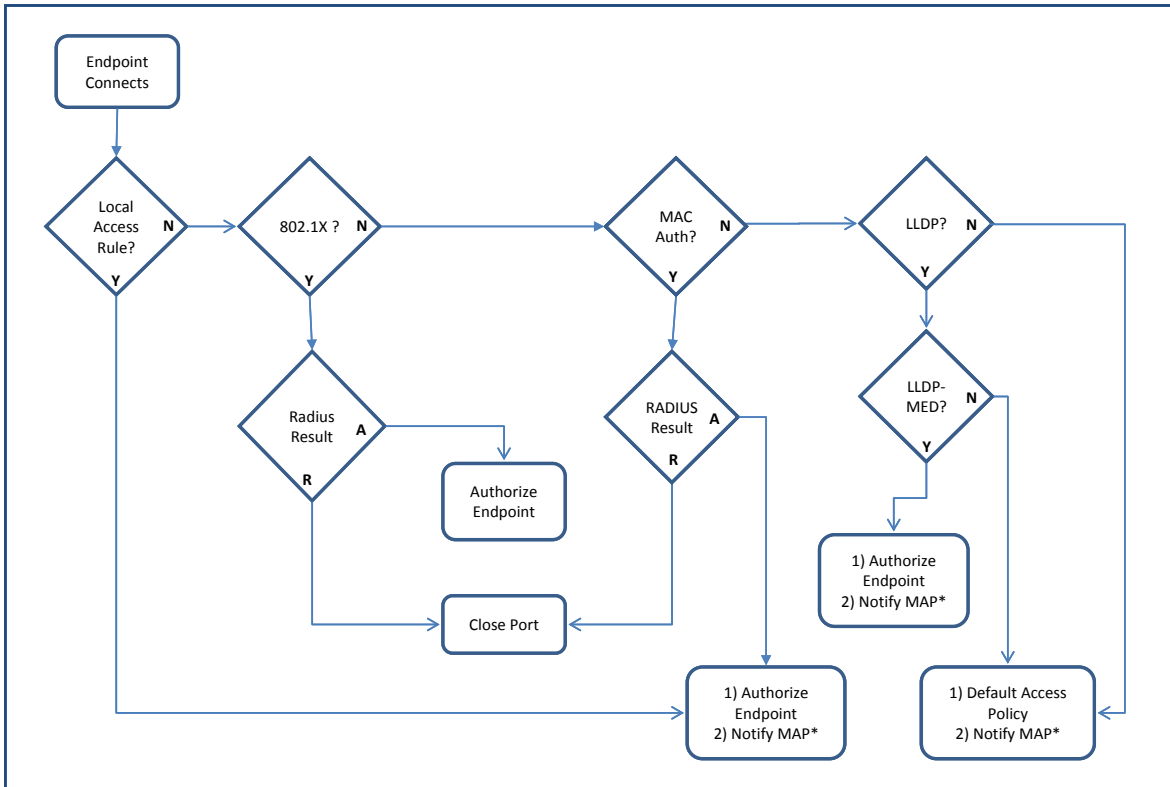The following diagram illustrates how a CESP-compliant PEP selects an authorization method when a clientless endpoint connects:

Figure 6 - Selection of Authorization Method