TCG Trusted Network Communications
TNC IF-PEP: Protocol Bindings for RADIUS

Specification Version 1.1
Revision 0.8
5 February 2007
Published

Contact:
admin@trustedcomputinggroup.com

TCG

TCG PUBLISHED

Disclaimer

IWG TNC Document Roadmap

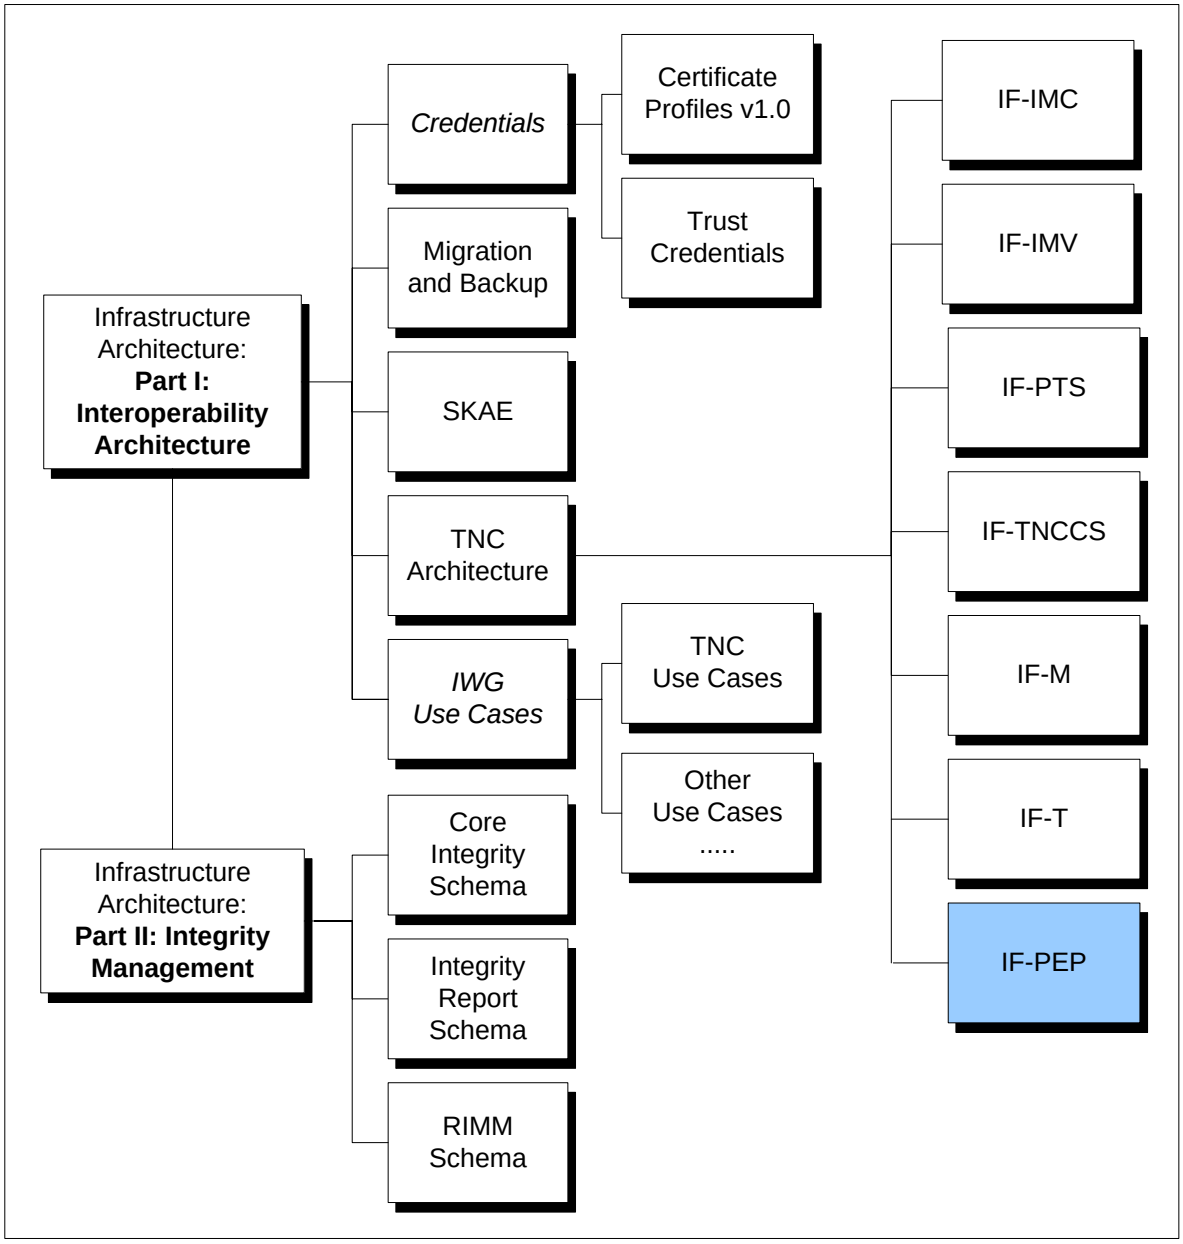

Table of Contents

Acknowledgements

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

Special thanks to the members of the TNC contributing to this document:

Aman Garg	**3Com**	
Bipin Mistry	**3Com**	
Mahalingam Mani	**Avaya**	
Hidenobu Ito	**Fujitsu Limited**	
Sung Lee	**Fujitsu Limited**	
Kazuaki Nimura	**Fujitsu Limited**	
Boris Balacheff	**Hewlett-Packard**	
Mauricio Sanchez	**Hewlett-Packard**	
Diana Arroyo	**IBM**	
Lee Terrell	**IBM**	
Stuart Bailey	**Infoblox**	
Ravi Sahita	**Intel Corporation**	
Ned Smith	**Intel Corporation**	
Chris Trytten	**iPass**	
Barbara Nelson	**iPass**	
Steve Hanna (TNC co-chair)	**Juniper Networks, Inc.**	
Alex Romanyuk	**Meetinghouse Data Communications**	
Gene Chang	**Meetinghouse Data Communications**	
John Vollbrecht	**Meetinghouse Data Communications**	
Sandilya Garimella	**Motorola**	
Jeff Six	**National Security Agency**	
Joseph Tardo	**Nevis Networks**	
Pasi Eronen	**Nokia Corporation**	
Meenakshi Kaushik	**Nortel Networks**	
Thomas Hardjono	**SignaCert, Inc.**	
Bryan Kingsford	**Symantec Corporation**	
Paul Sangster (TNC co-chair)	**Symantec Corporation**	
Rod Murchison	**Vernier Networks**	
Scott Cochrane	**Wave Systems**	
Greg Kazmierczak	**Wave Systems**	

1 Introduction

1.1 Scope and Audience

The Trusted Network Communications Work Group (TNC) is defining an open solution architecture that enables network operators to enforce policies regarding endpoint integrity when granting access to a network infrastructure. Part of the TNC architecture is IF-PEP, a standard interface between the Policy Decision Point (PDP) and the Policy Enforcement Point (PEP). This document defines and specifies IF-PEP using RADIUS (Remote Authentication Dial In User Service) [7].

Architects, designers, developers and technologists who wish to implement, use, or understand IF-PEP should read this document carefully. Before reading this document any further, the reader should review and understand the TNC architecture as described in [8].

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 [4].

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 Purpose of IF-PEP

This document describes and specifies IF-PEP protocol bindings for RADIUS. For the purposes of this document, this document will refer to IF-PEP for RADIUS as just IF-PEP. Future versions of this specification may describe additional protocol bindings with other protocols, such as Diameter [15] or SNMP [16]. IF-PEP allows communication between the PDP's Network Access Authority (NAA) and the Policy Enforcement Point (PEP).

IF-PEP is used to send network access decisions from the NAA to the PEP so the PEP can enforce those decisions on an endpoint's network traffic. By the time IF-PEP is used, the integrity state of the endpoint has already been established, the NAA has made a network access decision, and the PEP is ready to enforce the decision under direction of the NAA.

The network access decision should be considered an event that necessitates some action by the PEP, whether to allow unimpeded access (e.g. for an endpoint that complies with policy) or isolation (e.g. for an endpoint that does not comply with policy). Isolation in TNC may consist of no access or limited access. In particular for non-compliant endpoints, there is a clear need to describe the steps and interaction involving the AR, PEP, and PDP that must be undertaken to dynamically alter the PEP's enforcement: first, enforcing isolation while the endpoint performs remediation, and later removing isolation.

This document describes a standard set of isolation techniques available for performing endpoint isolation and specifies the RADIUS protocol bindings and requirements for both an NAA and PEP relating to each isolation technique. Isolation techniques are presented by their relevant network layer and a use case for each is described. For each isolation technique, a mapping to RADIUS is specified by describing the RADIUS messages or attributes that are to be utilized. Other documents may later be written to describe how other isolation techniques may be implemented.

Architects, designers, developers and technologists who wish to implement and use IF-PEP MUST abide to this document.

2.1.1 Supported Use Cases

Use cases that this version of IF-PEP supports are as follows:

- An **access requestor** (AR) is attempting network access through an 802.1X-enabled wired Ethernet (IEEE802.3) switch **policy enforcement point** (PEP) that is controlled by a RADIUS-based **policy decision point** (PDP).
- An **access requestor** (AR) is attempting network access through an 802.1X-enabled wireless (IEEE802.11) access point based **policy enforcement point** (PEP) that is controlled by a RADIUS-based **policy decision point** (PDP).
- An **access requestor** (AR) that is 802.1Q VLAN tag capable is attempting network access through an 802.1Q VLAN tag capable **policy enforcement point** (PEP) that is controlled by a RADIUS-based **policy decision point** (PDP). The PDP uses VLAN-based isolation employing 802.1Q tagged VLANs.
- An **access requestor** (AR) is attempting network access through an IPsec VPN server **policy enforcement point** (PEP) that is controlled by a RADIUS-based **policy decision point** (PDP).
- A **policy decision point** (PDP) may need to update the access policy enforced by the **policy enforcement point** (PEP) without having the **access requestor** (AR) lose network connectivity.

- A **policy decision point** (PDP) may at different times enforce different integrity compliance requirements on an **access requestor** (AR) that require changes to the access policy enforced by the **policy enforcement point** (PEP).
- An **access requestor** (AR) is attempting network access through a non-802.1X or non-IPsec **policy enforcement point** (PEP) that is controlled by a RADIUS-based **policy decision point** (PDP).

2.1.2 Non-supported Use Cases

Several use cases, but not limited to these, not covered is this version of IF-PEP are as follows:

- An **access requestor** (AR) is attempting network access through a non-network based **policy enforcement point** (PEP) that is controlled by a non-RADIUS based **policy decision point** (PDP). Descriptions for the various forms of PEP embodiments can be found in Annex A.
    - A non-network based PEP consists of an enforcement entity that is a physical and/or logical component of either the AR or the PDP. Additional discussion about the various embodiments for TNC policy enforcement points can be found in Annex A.
    - Furthermore, this use case means that only RADIUS is supported. Other protocols, such as SNMP, are not supported.
- An **access requestor** may connect one network with a protected network (e.g., site to site VPN).
- An **access requestor** (AR) is granted access to a single device based on compliance.
    - In this use case, the single device the AR is accessing implements both PEP and PDP functionality in addition to its nominal endpoint role (e.g. network attached printer, file server, etc.).

2.2 Requirements

The following are the requirements which IF-PEP must meet in order to successfully play its role in the TNC architecture. These are stated as general requirements, with specific requirements called out as appropriate.

a. Meets the needs of the TNC architecture

IF-PEP must support all the functions and use cases described in the TNC architecture as they apply to the relationship between the NAA and the PEP.

Specific requirements include:

- IF-PEP should support various methods of isolation for endpoints that fail integrity verification.
- IF-PEP must be compatible and usable with network access technologies supporting the TNC architecture, especially 802.1X networks and IPsec VPNs.
- IF-PEP must be compatible and usable with message transport technologies supporting the TNC architecture, such as tunnel EAP methods.
- IF-PEP must be compatible and usable with authentication server technologies supporting the TNC architecture, namely RADIUS.

b. Secure

The communications between a PEP and NAA must be protected. A PEP and NAA must provide its own security mechanisms as suggested in the Security Considerations section.

Specific requirements include:

- Communication between a PEP and NAA server must be authenticated and its integrity maintained.
- Communication between a PEP and NAA server should be confidential.

c. Extensible

IF-PEP will need to expand over time as new features and supported network, message, and authentication technologies are added to the TNC architecture. IF-PEP must allow new features to be added easily, providing for a smooth transition and allowing newer and older architectural components to work together.

d. Easy to use and implement

IF-PEP should be easy for PEP and NAA 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. IF-PEP should also make things easy for system administrators and end-users. Components of the TNC architecture should plug together automatically without requiring extensive manual configuration.

e. Substantially network media independent

IF-PEP must not be solely limited to supporting just Ethernet networks. It must be applicable for other network types.

2.3 Assumptions

Here are the assumptions that this document makes about components in the TNC architecture.

- Network access decision
  The network access decision is about limiting network access in some way. The network decision is derived from the TNCS' recommendation – *allow, deny,* or *isolate*.

- NAA and RADIUS
  The NAA of a TNC PDP is a RADIUS server.

- Control of PEP
  The PEP is under the control of the NAA.

- A PEP cannot change an NAA *deny* network decision to either *allow* or *isolate*.
  A PEP can reject an NAA's *allow* or *isolate* network decision, if it is incapable of fully enforcing the decision because of, but not limited to, mis-configuration, PEP resource depletion, or violation of local PEP security policy.

2.4 Out of Scope

While IF-PEP does allow passage of network access policy between the NAA and the PEP, it does not facilitate all configuration or management tasks/processes between both. Several noteworthy items that have been identified as out of scope for this version of IF-PEP include:

- The NAA and PEP are each in different administrative domains
  Discussion on how to handle an NAA and PEP that are each under the control of a different administrative entity is out of scope.

- Control of multiple PEPs concurrently
  Discussion on how an NAA can control multiple PEPs concurrently for a given AR access request is out of scope. If multiple PEPs exist in an environment, the network decision is not communicated to them until the AR attempt to gain access through one of those PEPs.

- Reconciliation of PEP and NAA capabilities
  Discussion on how to reconcile enforcement capabilities or perform capability discovery is out of scope. The administrator is responsible for reconciling capability differences between the NAA and PEP during solution deployment. That is, he or she is responsible for establishing that an NAA can/will instruct the PEP to perform some type of isolation in a meaningful manner.

2.5 Features Provided by IF-PEP

This section documents the features provided by IF-PEP.

2.5.1 Endpoint Isolation

Within the TNC architecture, a PEP serves the vital role of enforcing network access decisions and in particular for non-compliant endpoints, it serves to isolate an endpoint's access to just those network resources necessary for remediation or basic network access. IF-PEP describes various isolation techniques that differ in the network layer (L2, L3, etc.) they affect.

2.5.2 Network Access Decision Transport

The NAA needs to communicate to the PEP the network access decision to be enforced for a particular endpoint. IF-PEP describes the usage of RADIUS messages and Attribute Value Pairs (AVPs) as a means for sending the network access decision from the NAA to the PEP.

2.5.3 Support of Remediation and Handshake Retry

The initial network access decision may later be modified. For example, endpoints that are found to be in non-compliancy may return into compliancy, thereby requiring a PEP to no longer enforce an isolation policy. For such cases, the TNC architecture describes the need for handshake retries to establish a new integrity decision. IF-IMV [10] and IF-IMC [9] describe the process of integrity revalidation. IF-PEP describes the role of the PEP as part of the handshake retry.

3 Isolation Techniques and Use Cases

This section describes the TNC standard set of isolation techniques available for performing endpoint isolation. For each isolation technique, a use case is presented that describes a nominal use in a TNC solution.

It should be noted that an implementer is free to conceive and deploy additional isolation techniques (i.e. vendor specific), but MUST support at least one of the three isolation techniques described herein to be considered TNC compliant.

3.1 Binary-based Isolation

Binary-based isolation provides an all or none network access proposition. In terms endpoint isolation, binary access is the simplest technique. Either the endpoint is allowed onto the network or it is completely blocked from the network.

The use case for binary isolation is to grant network access to compliant endpoints while blocking network access to non-compliant endpoints. While binary isolation is the most straightforward isolation technique, it does not allow network based remediation; it can prohibit access to those with invalid setups, but does not help the network operator provide remediation.

3.2 VLAN-based Isolation

In Ethernet environments supporting IEEE 802.1D[17] and IEEE 802.1Q[18], virtual LANs (VLANs) allow administrators to logically compartmentalize their networks into distinct logical networks while still maintaining the same physical connection topology. Originally VLANs were positioned as a design tool to improve network performance and scalability by facilitating the creation of multiple layer-2 broadcast domains. However, the benefits of VLANs have been extended into the network access domain by enabling segregation of network traffic based on security requirements and data sensitivity.

The use case for data link layer isolation via VLANs is to compartmentalize endpoints based on their endpoint integrity state. It is expected that there may be different levels of restricted access, and therefore different VLAN policies could be used. For example, VLANs can be used for containment of non-compliant endpoints in an administrator-defined “Fix-Up” VLAN. Usually this VLAN has limited connectivity to just those network resources fundamental to the remediation process. This could mean only a DHCP, DNS and remediation server would be present on the VLAN. Meanwhile, endpoints that are integrity compliant could be assigned “full-access” VLAN access.

3.3 Filter-Based Isolation

Filter-based enforcement refers to setting up filter rules (e.g. ACLs - Access Control Lists) in the PEP such that when a user initiates traffic, the PEP examines the set of rules associated with the service granted to the user. These rules determine what traffic is allowed to proceed through the PEP and what traffic will be filtered (i.e. blocked).

The filtering action usually either permits or denies traffic from traversing a PEP. The PEP will match an incoming packet’s network header portions to its permit or deny rule list and determine whether to allow the packet to continue or be blocked. Packets that are allowed to continue are usually unaltered.

The use case for filter-based isolation is to filter traffic based on their endpoint integrity state. Filter-based isolation provides a more granular method for controlling endpoint network access than the data link layer isolation method. For example, non-compliant endpoints can be isolated from one another when they reside on the same IP subnet by allowing them to only communicate with remediation servers, but not amongst themselves. The endpoints are constrained to communicating to those just those IP addresses and TCP ports needed for remediation.

It is expected that there may be different levels of restricted access, and therefore different filters can be used. As an example, endpoints that are non-compliant might be assigned the “Fix-Up” filter, which could limit access to the IP addresses and TCP ports of just the DNS and remediation server(s). Meanwhile, endpoints that are integrity compliant could be assigned the “Full-Access” filter, which would impose no IP address or TCP port access restrictions on the endpoint.

4 IF-PEP for Network-Based PEPs

The TNC architecture describes the role of the PEP as the entity that enforces the decisions of the NAA regarding network access. In many instances, TNC documentation uses network devices (e.g. VPN, 802.1X) as typical embodiments of PEPs. However, the TNC architecture does not limit PEPs to just network devices, but rather considers PEPs as logical entities that can take on a number of different embodiments, as described in Annex A. This version of IF-PEP, however, is limited in scope to treatment of just network-based PEPs.

4.1 Model

The diagram below shows, in the upper-half, the TNC architecture and, in the lower-half, a sample representation and relationship between physical entities.

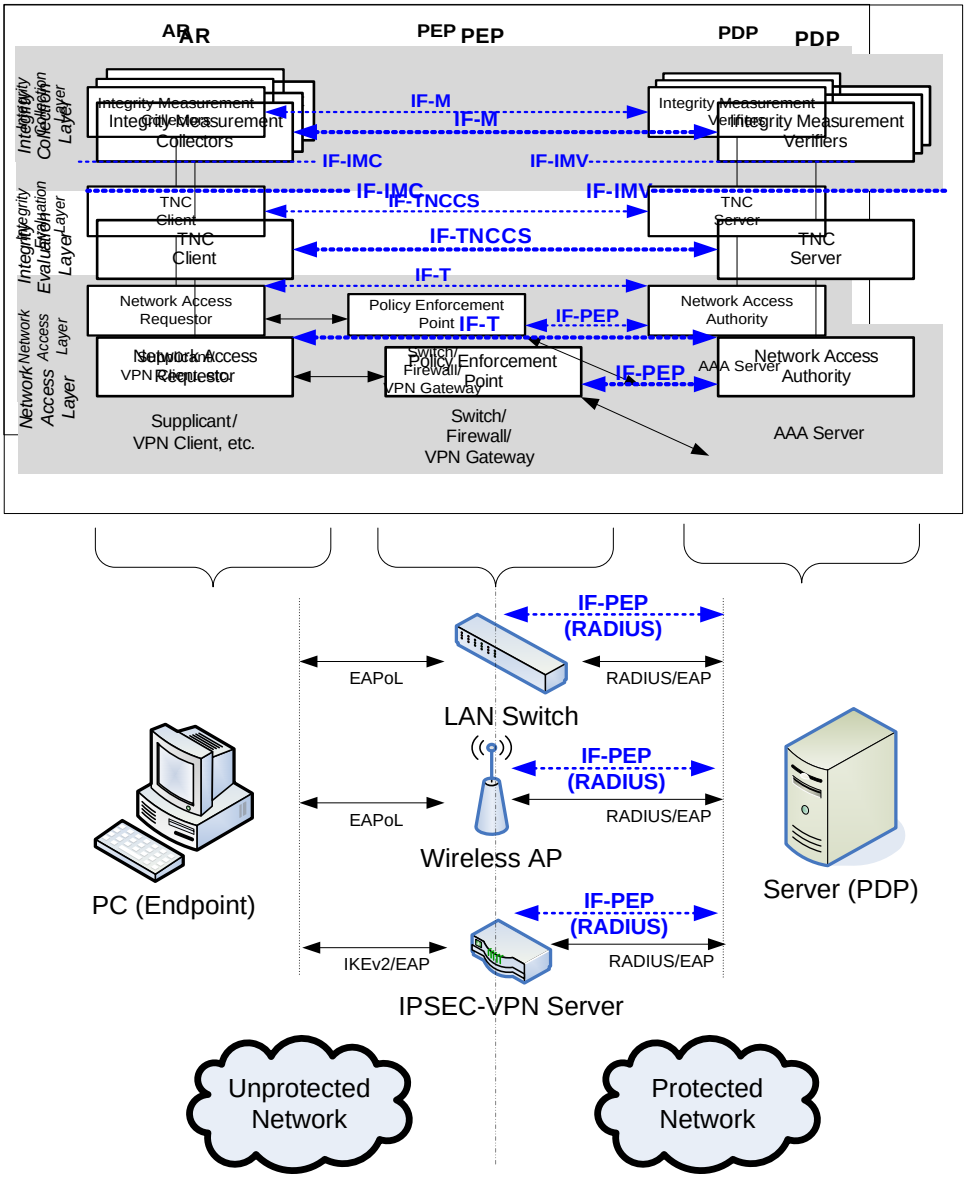

Figure 1 - Network-Based PEP in TNC Architecture

In the diagram, the PC endpoint is trying to acquire network access either over a LAN switch, wireless access point, or IPsec VPN server. The LAN switch, wireless access point, and IPsec VPN server are all examples of network-based PEPs who are under the control by the NAA in the PDP via IF-PEP. In all instances IF-PEP uses RADIUS as a transport protocol foundation.

4.2 NAA and PEP Requirements related to TNC

The following are the requirements which an NAA and network-based PEP must meet to function in the TNC architecture.

- Both an NAA and PEP MUST implement at least one of the isolation techniques described in section 3 and SHOULD implement two or more for greater deployment flexibility. A list of the isolation techniques follows:
    - Binary- based
    - VLAN-based
    - Filter-based
- If an NAA and PEP intend to support VLAN-based isolation, the NAA and PEP MUST support untagged (native) VLAN assignment and enforcement, respectively. Both the NAA and PEP SHOULD support tagged VLAN assignment and enforcement, respectively.
- Both an NAA and PEP SHOULD allow one or more isolation techniques to be concurrently active.

  For example, VLAN-based isolation coupled with filter-based isolation allows for deployment scenario whereby multiple endpoints on the same isolation VLAN cannot communicate with each other directly. Direct communication between endpoints is prohibited by the traffic filter.
- An NAA and PEP MAY implement additional vendor-specific isolation techniques that both the NAA and PEP have agreed upon. However, the requirement that both the NAA and PEP MUST implement at least one of the three isolation techniques described in section 3 remains.
- Both an NAA and PEP MUST allow dynamic access policy update. This update SHOULD be performed without affecting existing network connectivity for the endpoint.

  The TNC architecture describes the use case for integrity revalidation. Depending on the capabilities and limitations of various TNC components, it is desirable to perform revalidation without interrupting network connectivity.

5 RADIUS as IF-PEP Transport Protocol

This section specifies how RADIUS is used as a network access decision transport protocol between a TNC PEP and NAA. This section elaborates on:

- How RADIUS is to be used to meet TNC PEP requirements.
- RADIUS-specific requirements imposed on the NAA and PEP.

5.1 Why RADIUS?

Since the TNC architecture is part of the authentication, authorization, and accounting (AAA) architecture, it is only natural that some of the same protocols from contemporary AAA implementations be pressed into service within the TNC framework. A key AAA protocol is RADIUS.

RADIUS is already the popular protocol for providing authentication, authorization, and accounting (AAA) services in many varied environments, ranging from enterprise to ISP operator environments, and many diverse deployments, spanning from 802.1X, IPsec-VPN, to SSL-VPN.

This document described how IF-PEP can be implemented using RADIUS. This allows compatibility with the huge installed base of network devices that already support RADIUS, easing the deployment of TNC-based solutions.

5.2 Relevant RFCs

RADIUS is defined by a body of IETF RFC (Request For Comment) documents, not all of which are relevant for IF-PEP. The table below enumerates the RFCs relevant to IF-PEP.

RFC Number	**Title**
RFC2865	**Remote Authentication Dial In User Service**
RFC2868	**RADIUS Attributes for Tunnel Protocol Support**
RFC3576	**Dynamic Authorization for RADIUS**
RFC3579	**RADIUS support for Extensible Authentication Protocol (EAP)**
RFC3580	**IEEE802.1X RADIUS Usage Guidelines**
RFC4675	**RADIUS attributes for Virtual LAN and Priority Support**

5.3 Isolation Techniques Mapped to RADIUS Attributes

5.3.1 Binary Isolation

No RADIUS attributes are necessary to implement binary access control. RADIUS already natively supports binary access control via the ACCESS-ACCEPT and ACCESS-REJECT messages per RFC2865 [7]. These two messages comprise the only two outcomes to a RADIUS access request. The NAA (i.e. RADIUS server) responds back to the PEP with an ACCESS-ACCEPT when the endpoint should be given network access and with an ACCESS-REJECT when the endpoint should be given no network access. For example, if the endpoint is integrity compliant, the NAA may send the PEP an ACCESS-ACCEPT to allow the endpoint onto the network. If the endpoint is integrity non-compliant, the NAA sends the PEP an ACCESS-REJECT to block all network access by the endpoint.

5.3.2 VLAN-based Isolation

For PEPs that implement IEEE 802.1D and IEEE 802.1Q bridging functionality, VLAN-based isolation provides a layer 2 based method for isolating an endpoint. Isolation of endpoints using VLANs can be achieved by having the PEP support a number of different VLANs. For example, the TNC *allow* and *isolate* recommendations can be mapped to separate VLANs, a "Full-Access" and "Fix-UP" VLAN respectively.

VLAN aware PEPs enforce them based on either a tag, using IEEE 802.1Q, or if not tagged (i.e. untagged), then based on some characteristic associated with the traffic, such as the ingress interface (e.g. physical port, wireless SSID, etc.).

RADIUS provides various methods for assigning both untagged and tagged VLANs as described in the following two sections.

5.3.2.1 Untagged VLANs

Two alternatives exist for assigning policy for untagged VLANs with RADIUS. The first method is by using tunnel attributes defined in RFC2868 and following the guidelines specified in RFC 3580 [6]. Used together both of these RFCs can be used to assign the ingress untagged VLAN ID (commonly known as the PVID) and the egress untagged VLAN ID.

The table below shows the specific tunnel attributes out of RFC2868 that are used in assignment of the untagged VLAN:

Attribute	**Value**
Tunnel-Type	**VLAN (13)**
Tunnel-Medium	**802**
Tunnel-Private-Group-ID	**VLANID**[1]

Table 1 RFC2868 attributes for PVID VLAN Assignment

The second method to assign untagged VLANs is by using the VLAN attributes defined in RFC4675 [19]. The VLAN attributes described by this RFC allow assignment of one or more egress untagged VLAN(s) and allow manipulation of the per-port Ingress Filter variable defined in IEEE 802.1Q.

The table below shows the specific VLAN attributes out of RFC 4675 that are relevant for egress VLAN assignment and Ingress Filter manipulation.

Attribute
Egress-VLANID
Egress-VLAN-Name
Ingress-Filters

Table 2 RFC4675 attributes for VLAN Assignment

[1] The value VLANID is the desired 802.1Q-2003 VID value (1-4096)

5.3.2.2 Tagged VLANs

Only one option exists for assigning policy for tagged VLANs with RADIUS, namely the same attributed listed in Table 2 in the previous section (Egress-VLANID, Egress-VLAN-Name and Ingress-Filter attributes). Besides being capable of assigning policy for untagged VLANs, they can also be used to assign policy for egress tagged VLANs.

5.3.3 Filter-based Isolation

Filter-based isolation refers to setting up filter rules in the PEP and assigning those rules to an endpoint such that when the endpoint initiates traffic, the PEP examines the set of rules associated with the endpoint.

The RADIUS attribute that allows the assignment of filter rules is the Filter-ID attribute, which is defined in RFC2865. The Filter-ID attribute allows an NAA to assign named filters to a PEP for enforcement upon an endpoint's traffic.

Isolation of endpoints using filters can be achieved by having the PEP support a number of different filters. For example, the TNC *allow* and *isolate* recommendations can be mapped to separated filters, a "Full-Access" and "Fix-Up" filters respectively.

5.4 Remediation and Handshake Retry Mapped to RADIUS

The TNC architecture describes the use case for integrity revalidation. Depending on the capabilities and limitations of various TNC components, it is desirable to perform revalidation without interrupting network connectivity. A PEP that allows dynamic changes of active access policy without connectivity interruption is an extremely desirable trait.

For example, a non-compliant endpoint that is on a restricted access network may have remediated itself and desires a new integrity handshake. While the initial integrity handshake resulted in the endpoint being identified as non-compliant, the subsequent integrity handshake may result in the endpoint being identified as compliant. As such, the isolation policy being enforced by the PEP needs to be changed from one of isolation to one providing greater access.

The RADIUS messages that allows dynamic policy changes on a PEP to be made are defined in RFC3576[5]. RFC3576 describes two new RADIUS messages, Disconnect and Change-of-Authorization (CoA), which can be sent arbitrarily from the RADIUS server to a PEP to perform a policy change. The disconnect message causes an immediate termination of a session and the CoA message can be used to update access policy by including a new set of VLAN and/or filtering attributes.

5.5 NAA and PEP Requirements related to RADIUS

The following are the RADIUS requirements which an NAA and PEP must meet to comply with IF-PEP.

- An NAA and PEP MUST support RADIUS as defined in RFC2865.
- An NAA and PEP SHOULD support RADIUS as defined in RFC2866.
- If an NAA and PEP intend to support VLAN-based isolation that consists in assignment of endpoint to an untagged VLAN, they both MUST support the attributes listed in Table 1. Additionally:
    - Both NAA and PEP MUST support RFC3580, section 3.31 usage guidelines.
    - An NAA MUST use all Table 1 attributes as part of RADIUS Access-Accept or CoA-req message.
    - Both the NAA and PEP MUST support the 'Tag' field in the Tunnel-Private-Group-ID attribute.
    - An NAA MUST always set the 'Tag' field in the Tunnel-Private-Group-ID to zero (0x00).

    - Both the NAA and PEP MUST interpret the Tunnel-Private-Group-ID as a string encoding of the integer VLAN ID value between 1 and 4094 (12-bit value).

- If an NAA and PEP intend to support VLAN-based isolation that consists in assignment of endpoint to an untagged VLAN, they both SHOULD support the attributes listed in Table 2. Additionally:
    - Both NAA and PEP MUST follow the usage rules for attributes in Table 2 as described in RFC4675.

- If an NAA and PEP intend to support VLAN-based isolation that consists in assignment of endpoint to tagged VLAN(s), they both MUST support the attributes listed in Table 2. Additionally:
    - Both NAA and PEP MUST follow the usage rules for attributes in Table 2 as described in RFC4675.

- If an NAA and PEP intend to support the filter-based isolation method, both MUST support the Filter-ID attribute as defined in RFC2865.

- If an NAA and PEP are implementing an isolation technique not described in section 3.1, they MAY use other standard and vendor-specific RADIUS attributes as necessary. However, the requirement that both the NAA and PEP support at least one of the isolation techniques in section 3.1, and thereby some set of mapped RADIUS messages or attributes as described in section 5.3, remains.

- If an NAA and PEP intend to support dynamic policy changes, as described in section 5.4, both the NAA and PEP MUST support RFC3576.

6 Security Considerations

6.1 Threat Analysis

IF-PEP is vulnerable to a number of threats associated with the PEP and NAA as identified in the TNC threat analysis [13]. Moreover, since this document describes the use of RADIUS, it is vulnerable to all of the threats associated with RADIUS. Several of the most noteworthy threats are outlined below, but a complete discussion of threats is found the TNC threat analysis, as well as RFC2607[2], RFC3162[3], RFC3579[1], and RFC3580.

6.1.1 Rogue NAA

A rogue NAA may be able to misuse IF-PEP using RADIUS in the following ways:

- Send invalid messages to PEP, leading to PEP crashes or compromise, excessive PEP resource consumption, or lack of connectivity (denial of service) for endpoint.
- Mount a man-in-middle (MITM) attack between a legitimate PEP and NAA to compromise the remediation process and further subvert an endpoint. Specific threats include:
    - Attribute editing – Attributes removed or added with malicious intent.
    - Connection hijacking – Subversion of the communication between the PEP and NAA.
    - Eavesdropping – Acquisition of confidential data, identities, and integrity information.
    - Offline attack – Dictionary attack against enciphered data.
    - Replay attack – Previous (legitimate) protocol exchanges replayed.
- Provide incorrect access policy to PEP, causing compliant endpoint to be rejected or non-compliant endpoints to be let on the network.
- Use vendor specific extensions to IF-PEP to perform other attacks.

6.1.2 Tampering

If an NAA or PEP lack strong configuration security controls, an attacker may be able to tamper with the configuration of cryptographic keys or credentials and lead to:

- Disclosure of information within messages exchanged between NAA and PEP.
- Association of a legitimate PEP with a rogue NAA, or vice versa.

6.1.3 Threats beyond IF-PEP

IF-PEP is part of the larger TNC architecture. Successful attacks against other parts of the TNC architecture will generally result in negative effects for PEPs, NAAs, and the system as a whole. See the Security Considerations section of the TNC Architecture or the TNC threat analysis document for an analysis of considerations that pertain to other parts of the TNC architecture.

6.2 Suggested Remedies

A number of suggested security remedies exist for IF-PEP using RADIUS, as follows:

- Mutual authentication and integrity preservation between PEP and NAA by having both PEP and PDP:
    - MUST support usage of non-obvious RADIUS secrets as described in RFC3579 section 4.3.3.
        - MUST support RADIUS secrets of at least 16 octets.
    - MUST support Message-Authenticator attribute as described in RFC3579, section 4.3.2.
- Enciphering of access policy transport protocol by having both PEP and PDP:

    - SHOULD use IPsec/IKE as described in RFC3579, section 4.2, and RFC3580, section 5.1.

- Replay protection by having both PEP and PDP:
    - MUST support Event-Timestamp attribute as described in RFC3576, section 5.4

Additionally both the NAA and PEP SHOULD support authentication of administrators before allowing changes to configuration.

7 Use Case Walkthrough

This section provides an informative (non-binding) walkthrough of the typical TNC use case showing how IF-PEP supports the use case for an 802.1X-capable Ethernet switch acting as a PEP.

The text describing IF-PEP usage is in **bold.** Sequence diagrams that illustrate the main parts of this walkthrough are included at the end of this section.

7.1 Configuration

1. The IT administrator configures any addressing and security information needed for server-side components (PEP, NAA, TNCS, and IMVs) to securely contact each other. The manner in which the PEP, NAA, and TNCS find each other is not specified.
2. The IT administrator configures policies in the NAA, TNCS, and IMVs for what sorts of user authentication, platform authentication, and integrity checks are required when.
3. For an 802.1X Ethernet switch: The IT administrator configures any VLAN or named access policy (for use with RADIUS 'filter-id' attribute) information to be used for isolation purposes.

7.2 Network Connect

1. The NAR attempts to connect to a network protected by a PEP, thus triggering an Integrity Check Handshake.
2. The PEP sends a network access request to the NAA.
3. The NAA performs user authentication of the NAR. User authentication may also involve having the NAR authenticate the NAA. Based on the NAA's policy, the user identity established through this process may be used to make immediate access decisions (like deny). **[IF-PEP] The immediate rejection consists of a RADIUS ACCESS-REJECT message.** If the decision is to deny the user, the use case ends here.
4. The TNCC and TNCS exchange integrity measurements via IF-TNCCS and IF-T [11]. The outcome is a TNCS Action Recommendation that is that is passed to the NAA.
5. The NAA sends its network access decision response to the PEP. **[IF-PEP] The access decision consists of a RADIUS ACCESS-ACCEPT message coupled with attributes (e.g. Filter-ID, Tunnel-ID, etc.) expressing the access policy.**
6. The PEP implements the network access decision response. During this process, the NAR is typically informed of the decision.
7. If the IMCs need to perform remediation, they perform that remediation. Then they continue with Handshake Retry. If no remediation was needed, the use case ends here.

7.3 Handshake Retry

1. Either because on IMC has completed remediation or because the TNCS desires to recheck the security state of the AR periodically, there may be a need to retry the Integrity Handshake.
2. Depending on limitations of the NAR, the NAA, and the PEP, the TNCC may need to disconnect from the network and reconnect to retry the Integrity Check Handshake. **[IF-PEP] If the PEP does not support dynamic policy changes without network interruption, then the TNCC is forced to disconnect from the network.** In that case (especially if the previous handshake resulted in full access), it may decide to skip the handshake retry. If the retry is skipped, the use case ends here.

3. The integrity handshake is redone as in step 4 of the Network Connection section above. Note that IF-TNCSS does not yet support handshake retry without network disconnect.
4. **[IF-PEP] The NAA sends an update to the network access decision via a RADIUS CoA message and attribute set expressing the access policy.**

7.4 Sequence Diagram for Network Connect

The following sequence diagram illustrates the Network Connect use case, as described in section 7.2.

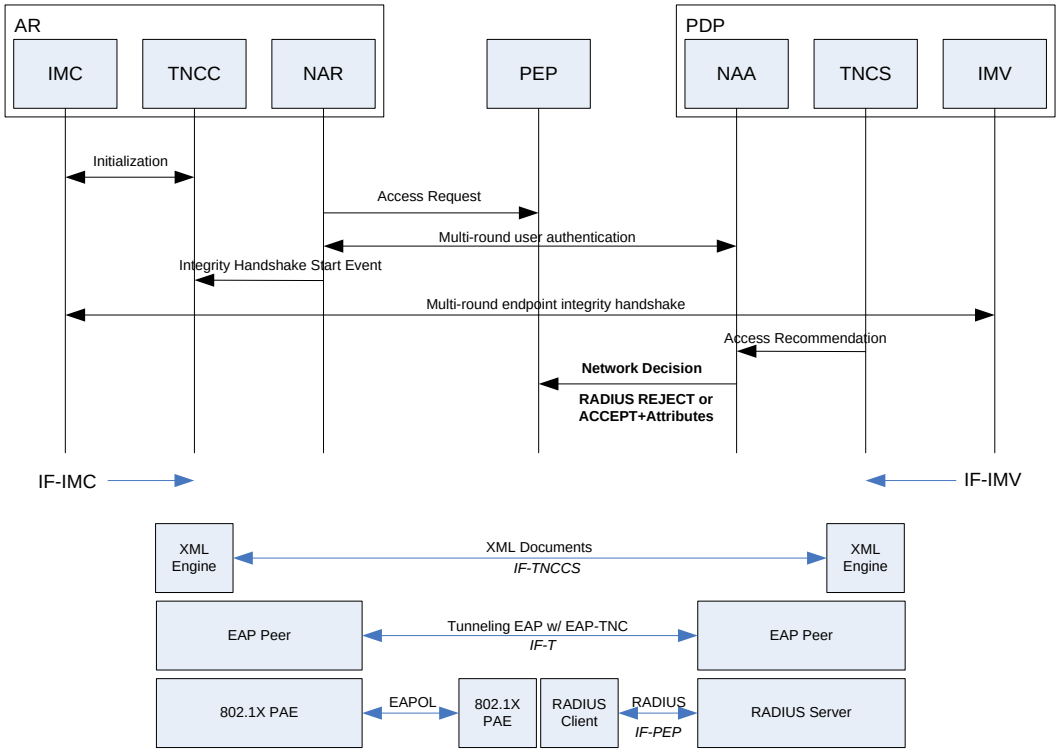

Figure 2 - IF-PEP Network Connect Sequence Diagram

7.5 Sequence Diagram for Handshake Retry

The following sequence diagram illustrates the Handshake Retry use case, as described in section 7.3.

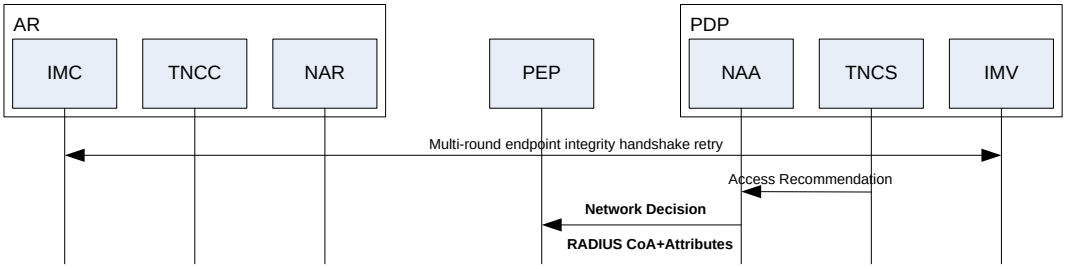

Figure 3 - IF-PEP Handshake Retry Sequence Diagram

8 References

[1] Aboba, B. and P. Calhoun, "RADIUS Support for Extensible Authentication Protocol (EAP)", RFC 3579, September 2003.

[2] Aboba, B. and J. Vollbrecht, "Proxy Chaining and Policy Implementation in Roaming", RFC 2607, June 1999.

[3] Aboba, B., Zorn, G. and D. Mitton, "RADIUS and IPv6", RFC 3162, August 2001.

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

[5] Chiba, M., et. al., "Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)", RFC 3576, July 2003.

[6] Congdon, P., et. al., “IETF 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines”, RFC3580, September 2003.

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

[8] Trusted Computing Group, TNC *Architecture for Interoperability*, Specification Version 1.1, May 2006.

[9] Trusted Computing Group, TNC *IF-IMC*, Specification Version 1.1, May 2006.

[10] Trusted Computing Group, TNC *IF-IMV*, Specification Version 1.1, May 2006.

[11] Trusted Computing Group, TNC *IF-T*, Specification Version 1.0, May 2006.

[12] Trusted Computing Group, TNC *IF-TNCCS*, Specification Version 1.0, May 2006.

[13] Trusted Computing Group, TNC Threat Analysis, December 2005, work in progress.

[14] Zorn, G., et al., "RADIUS Attributes for Tunnel Protocol Support", RFC 2868, June 2000.

[15] Colhoun, P., et al., “Diameter Base Protocol”, RFC 3588, September 2003.

[16] Case. J., et al., “A Simple Network Management Protocol (SNMP)”, May 1990.

[17] IEEE Standards for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges, IEEE Std 802.1D-2004, June 2004.

[18] IEEE Standards for Local and Metropolitan Area Networks: Draft Standard for Virtual Bridged Local Area Networks, P802.1Q-2003, January 2003.

[19] Congdon, P., Sanchez, M., and Aboba, B., “RADIUS Attributes for Virtual LAN and Priority Support”, RFC 4675, September 2006.

9 Annex A: PEP Embodiment Spectrum

The TNC architecture describes the role of the PEP as the entity that enforces the decisions of the PDP regarding network access. In many instances, TNC documentation uses network devices (e.g. VPN, 802.1X) as typical embodiments of PEPs. However, the TNC architecture fundamentally does not limit PEPs to just network devices, but rather considers PEPs as logical entities that can take on any number of physical embodiments.

This section describes four physically conceivable PEP embodiments and shows how IF-PEP maps in for each of these embodiments. It should be noted that the IF-PEP protocol bindings for RADIUS detailed in this document generally apply only to network-based PEPs. Future versions of this document may take the other embodiments into consideration.

9.1 PEP Types

9.1.1 Network-based PEP

A network-based PEP consists of an enforcement entity that is physically and logically separate from either the AR or the PDP. A network-based PEP is likely to be the most common PEP embodiment and is widely used in examples described in TNC documentation. In fact, this version of IF-PEP is limited in scope to treatment of just network-based PEPs.

The diagram below shows, in the upper-half, the TNC architecture and, in the lower-half, a sample representation and relationship between physical entities.

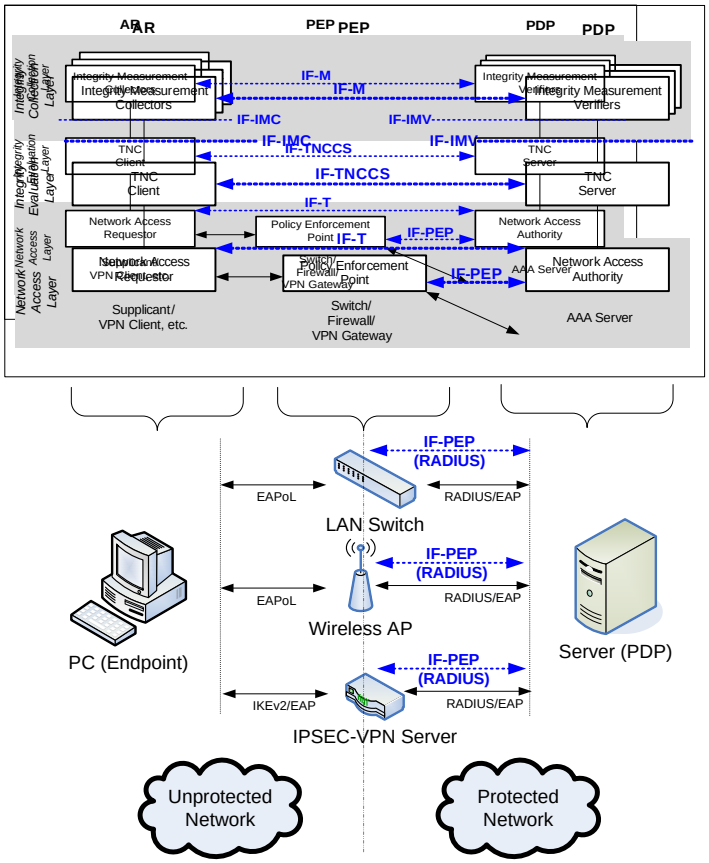

Figure 4 : Network-Based PEP in TNC Architecture

Other examples of network-based PEPs include (non-exhaustive list):

- SSL-VPN
- 802.16 “Wi-MAX” wireless broadband
- Network firewall
- Network router

9.1.2 Endpoint-based PEP

An endpoint-based PEP consists of an enforcement entity that is physically part of the AR, but logically separate from the AR. An endpoint-based PEP does not require the network to enforce any isolation policy as this function is self contained within the AR. The AR is capable of enforcing an isolation policy on itself.

The diagram below shows, in the upper-half, the TNC architecture and, in the lower-half, a sample representation and relationship between physical entities. The PEP is shown as being part of the endpoint under control by the NAA in the PDP via IF-PEP. It should be noted that this version of IF-PEP does not consider how RADIUS protocol bindings can be used, if at all, with endpoint-based PEPs. A future version of this document may take endpoint-based PEPs into consideration in a more complete fashion.

The network is shown as a cloud given its relevance only as a communication medium and without any capability or need to enforce isolation policy.

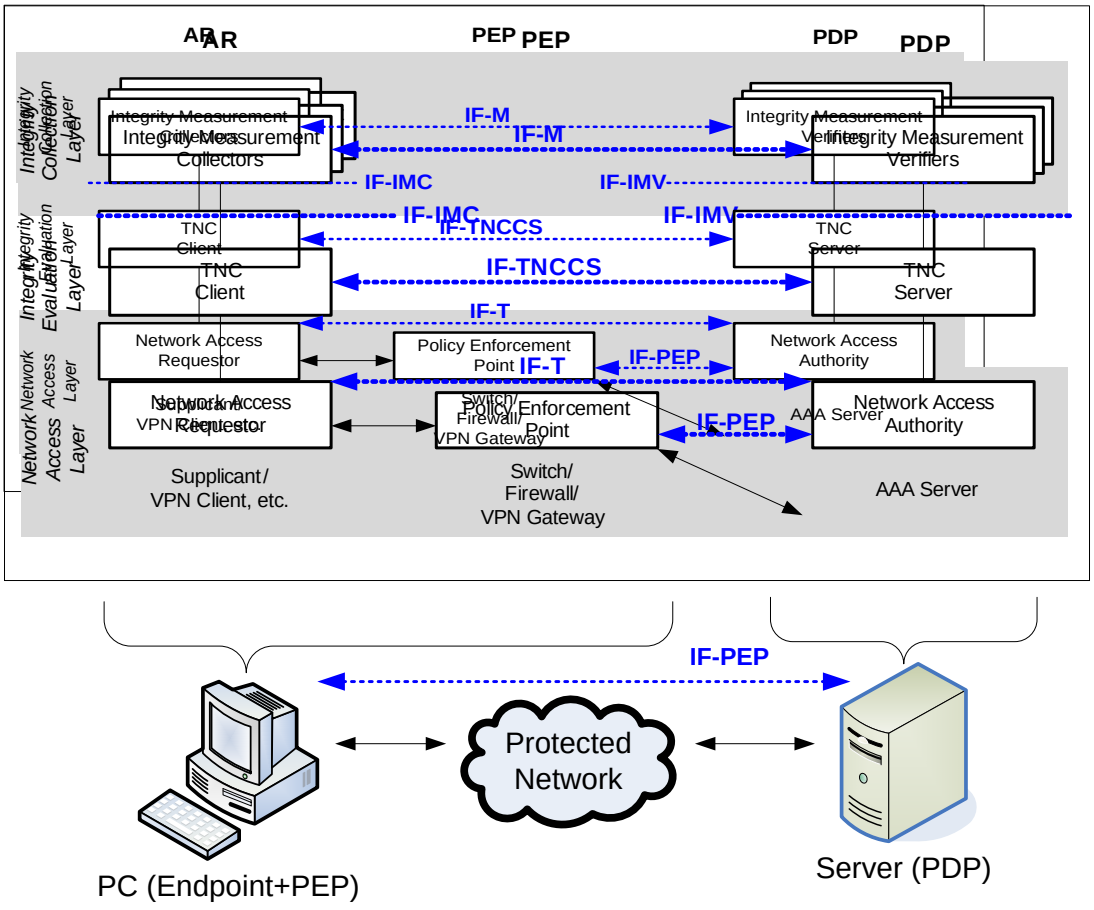

Figure 5 : Endpoint-Based PEP in TNC Architecture

A common example of an endpoint-based PEP is the class of applications known as “personal firewalls” or “managed host firewalls.”

9.1.3 Server-based PEP

A server-based PEP consists of an enforcement entity that is physically part of the PDP, but logically separate from the PDP. A server-based PEP does not require the network to enforce any isolation policy as this function is contained within the PDP.

The diagram below shows, in the upper-half, the TNC architecture and, in the lower-half, a sample representation and relationship between physical entities. The PEP is shown as being part of the PDP under control by the NAA in the PDP via IF-PEP. It should be noted that this version of IF-PEP does not consider how RADIUS protocol bindings can be used, if at all, with server-based PEPs. A future version of this document may take server-based PEPs into consideration in a more complete fashion.

The network is shown as a cloud given its relevance only as a communication medium and without any capability or need to enforce isolation policy.

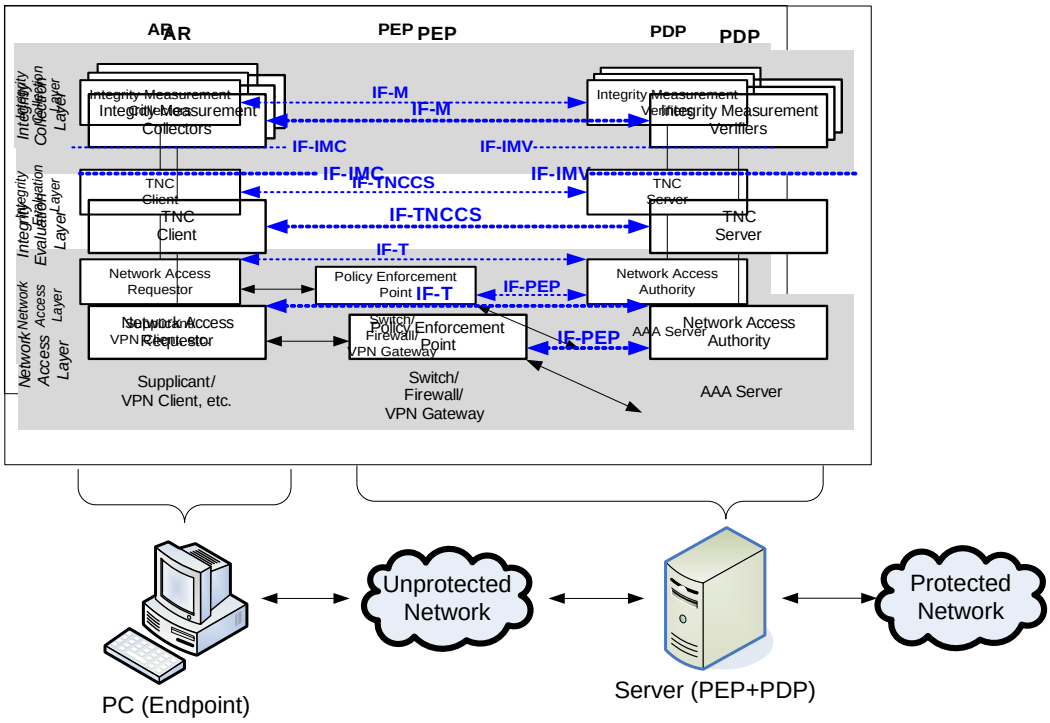

Figure 6 : Server-Based PEP in TNC Architecture

Where as one of IF-PEP’s features is access policy transport, this feature is not likely of great use for a server-based PEP, given the form IF-PEP is currently defined. One can assume that a server-based PEP will likely come from the same vendor that provides the NAA. As such, a vendor specific API between the NAA and PEP functionality will suffice and there is no need for IF-PEP’s standardized access policy transport feature.

9.1.4 Fully Integrated Endpoint

The fully integrated endpoint is a derivative of the endpoint-based PEP and consists of an endpoint that has collapsed all the TNC logical entities into the physical embodiment. Both the PEP and PDP are part of the same physical AR (i.e. the endpoint). As with the endpoint-based

PEP, the fully integrated endpoint does not require the network to enforce any isolation policy as this function is contained within the PDP.

The diagram below shows, in the upper-half, the TNC architecture and, in the lower-half, a sample fully integrated endpoint. The PEP is shown as being part of the endpoint along with the rest of the TNC logical entities. The network is shown as a cloud given its relevance only as a communication medium and without any capability or need to enforce isolation policy.

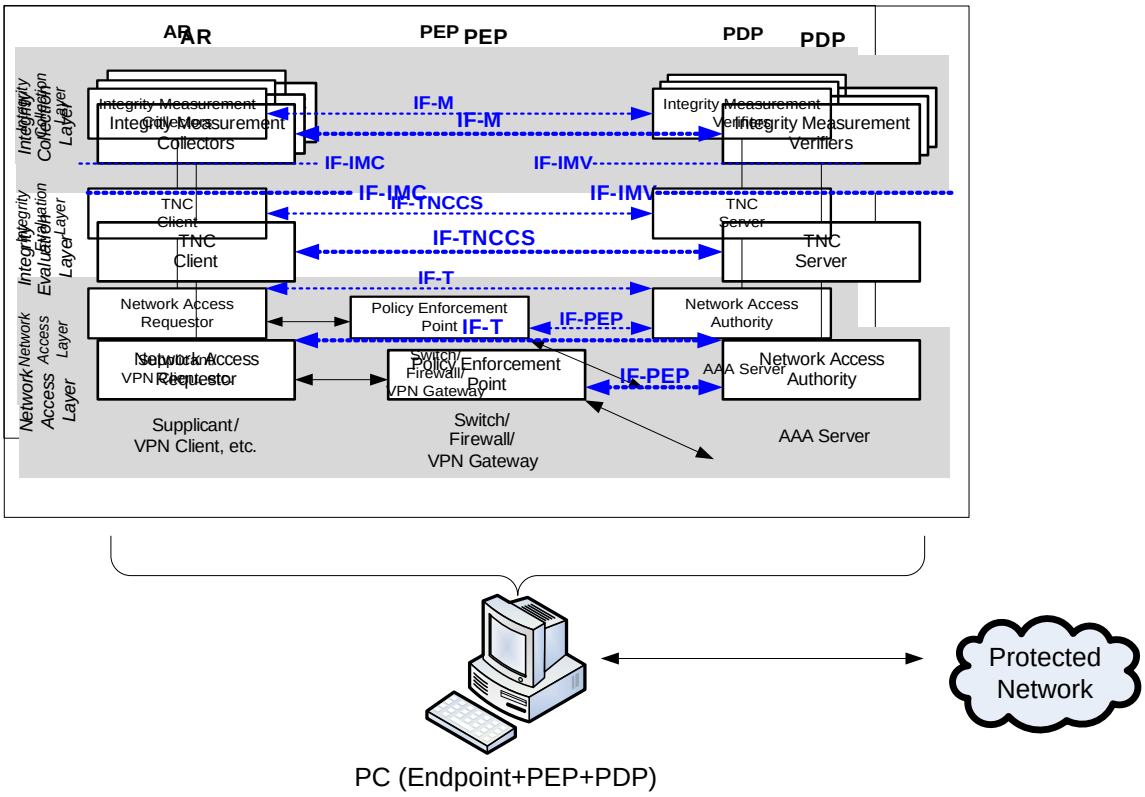

Figure 7 : Fully Integrated Endpoint in TNC Architecture

Whereas one of IF-PEP's features is access policy transport, this feature is not likely of great use for a fully integrated endpoint. One can assume that a fully integrated endpoint will likely come from just one vendor. As such, a vendor specific API between the NAA and PEP functionality will suffice and there is no need for IF-PEP's standardized access policy transport feature.