TCG Trusted Network Communications

Federated TNC

Specification Version 1.0
Revision 27
May 18, 2009
Published

Contact:
admin@trustedcomputinggroup.org

TCG

TCG PUBLISHED

IWG TNC Document Roadmap

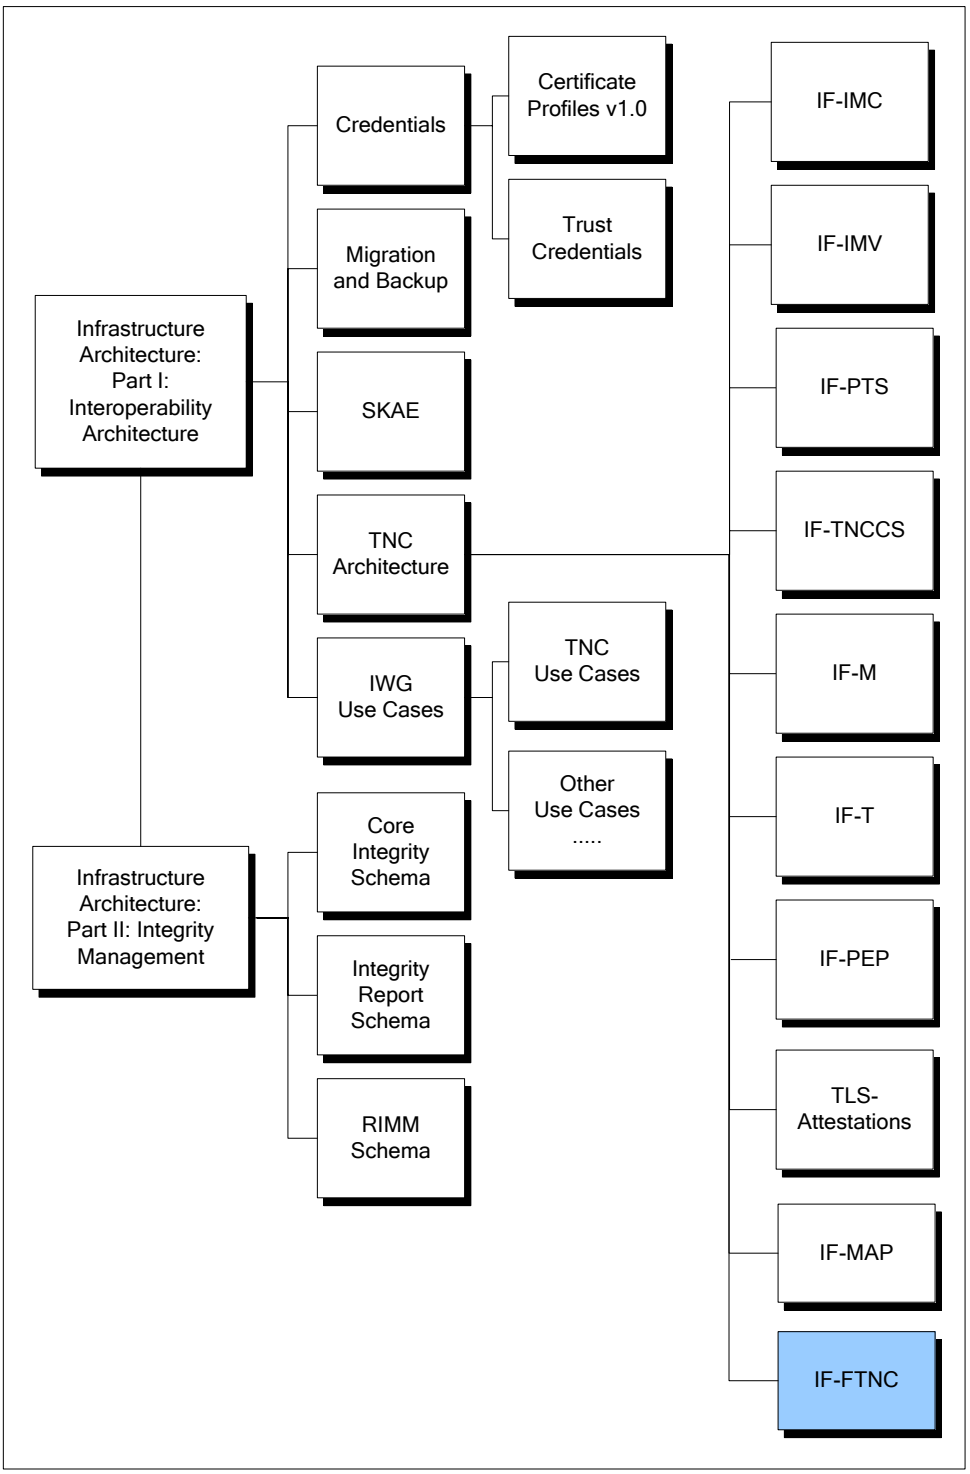

Table of Contents

Acknowledgements

This specification is dedicated to the memory of Manuel Sánchez Cuenca who, while working within the DAMe project at the University of Murcia before his premature passing, contributed towards the development of many of the concepts articulated within this specification, and early implementations that served as evidence of the feasibility and utility of the architecture.

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 and elsewhere; including the GÉANT2 project's JRA5 activity, the DAMe project of the Universities of Stuttgart and Murcia, and Internet2 MACE community.

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

Scott Kelly	**Aruba Networks**
Jeffery Dion	**Boeing**
Steven Venema	**Boeing**
Peter Wrobel	**CESG**
Mark Townsend	**Enterasys**
Sung Lee	**Fujitsu**
Mauricio Sanchez	**Hewlett-Packard**
Ren Lanfang	**Huawei**
Dr. Jiwei Wei	**Huawei**
Han Yin	**Huawei**
Stuart Bailey	**Infoblox**
Ravi Sahita (Co-editor)	**Intel**
Josh Howlett (Co-editor)	**JANET (UK)**
Steve Hanna (TNC co-chair)	**Juniper**
PJ Kirner	**Juniper**
Lisa Lorenzin	**Juniper**
Tom Price	**Lumeta**
Matt Webster	**Lumeta**
Paul Sangster (TNC co-chair)	**Symantec**
Brad Upson	**University of New Hampshire InterOperability Lab Lauren Giroux, US National Security Agency**

1 Introduction

1.1 Scope and Audience

The Trusted Network Communications Working Group (TNC) is defining an architecture that enables network operators to enforce policies regarding endpoint integrity when granting access to a network infrastructure. This specification is integral to the TNC's reference architecture. Specifically it defines the Federated TNC protocol (IF-FTNC), which enables communication of IF-M attributes, IF-TNCCS Access Recommendations and IF-MAP metadata from one security domain to another, as illustrated in Figure 1 below.

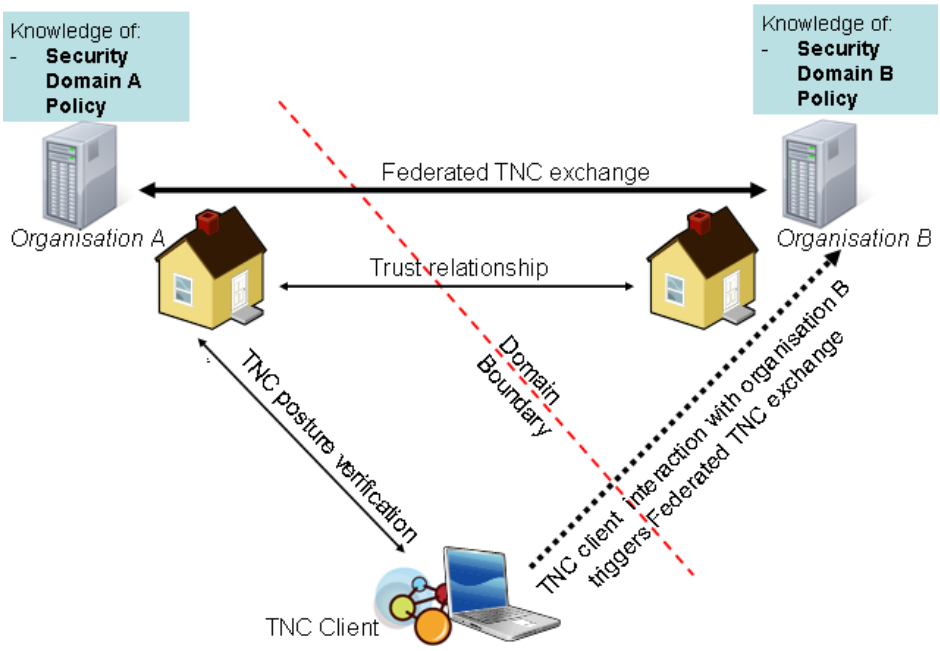

Figure 1

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 communicating integrity information. 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 IF-MAP [2], IF-TNCCS [3] and IF-M [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 [5]. 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 IF-FTNC

The Trusted Network Communications group has specified standards for assessing the security posture of an endpoint as it connects to the network. This assessment is performed by a TNC Server (TNCS) belonging to the same security domain as the endpoint (that is, there exists a direct trust relationship between the endpoint and the TNCS). This trust relationship is sufficient provided that the endpoint only accesses services within its own security domain. However it is undefined how the endpoint's posture should be assessed by a service within other security domains.

This specification defines how an endpoint's Security Posture Information (SPI) can be queried and supplied such that a security domain, other than the endpoint's own, can make authorization decisions controlling that endpoint's access to its networks and applications.

This specification defines SPI as any of: the IF-TNCCS Access Recommendation result; IF-M attributes reported by the endpoint during assessment; and IF-MAP metadata collected from other security devices.

Two or more systems that share a trust relationship permitting the exchange of information are sometimes said to be "federated". IF-FTNC uses these federated relationships to exchange endpoint SPI in a controlled and secure manner, thereby enabling the extension of the TNC architecture across security domains.

2.2 Conceptual Model

The IF-FTNC conceptual model consists of three main actors:

- The endpoint.
- The Asserting Security Domain (ASD), which has knowledge of the endpoint's SPI.
- The Relying Security Domain (RSD), which requires knowledge of the endpoint's SPI.

These parties, and their relationships, are illustrated in Figure 2 below.

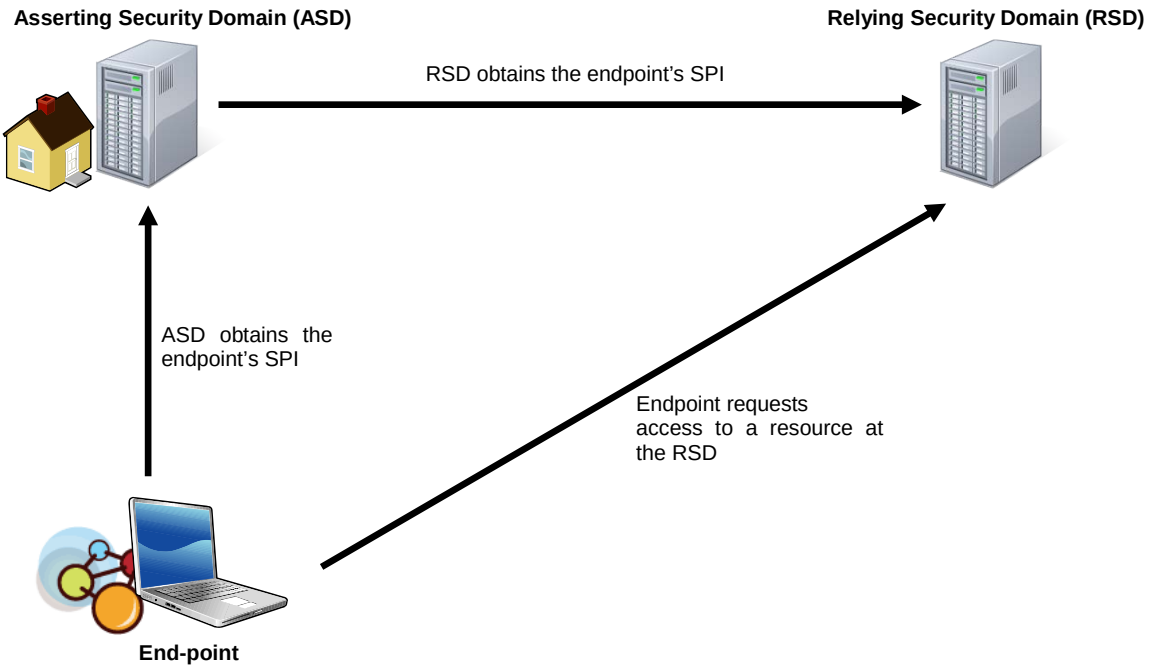

Figure 2

This specification realizes the conceptual model in two ways.

First, it defines how messages can be exchanged between the ASD and RSD to obtain and express an endpoint's SPI.

Secondly, it profiles the use of these messages within two supported use-cases:

- Roaming Assessment Profile: where the endpoint is requesting access to a network service operated by the RSD.
- Web Assessment Profile: where the endpoint is requesting access to a web resource operated by the RSD.

These profiles are discussed in the following sections. Future revisions of this specification may extend IF-FTNC to address other use cases.

2.2.1 Roaming Assessment Profile

Some organizations enter into roaming agreements with other organizations that enable roaming between their respective networks. The Roaming Assessment Profile is intended for use in these roaming network access scenarios. It enables the host network (in its role as the RSD) to request a roaming endpoint's SPI from the endpoint's 'home' network (in its role as the ASD).

Knowledge of the endpoint's SPI enables the host network to implement SPI-based authorization, allowing the local network administrator to assign an appropriate level of network access to the roaming endpoint. This could be used, for example, to move a roaming endpoint onto a remediation network if the endpoint is considered to be compromised, or at risk of compromise; or alternatively to move a roaming endpoint onto a privileged network with access to more sensitive corporate resources if the roaming endpoint is considered sufficiently trustworthy.

An example use of the Roaming Assessment Profile is illustrated in Figure 3 below.

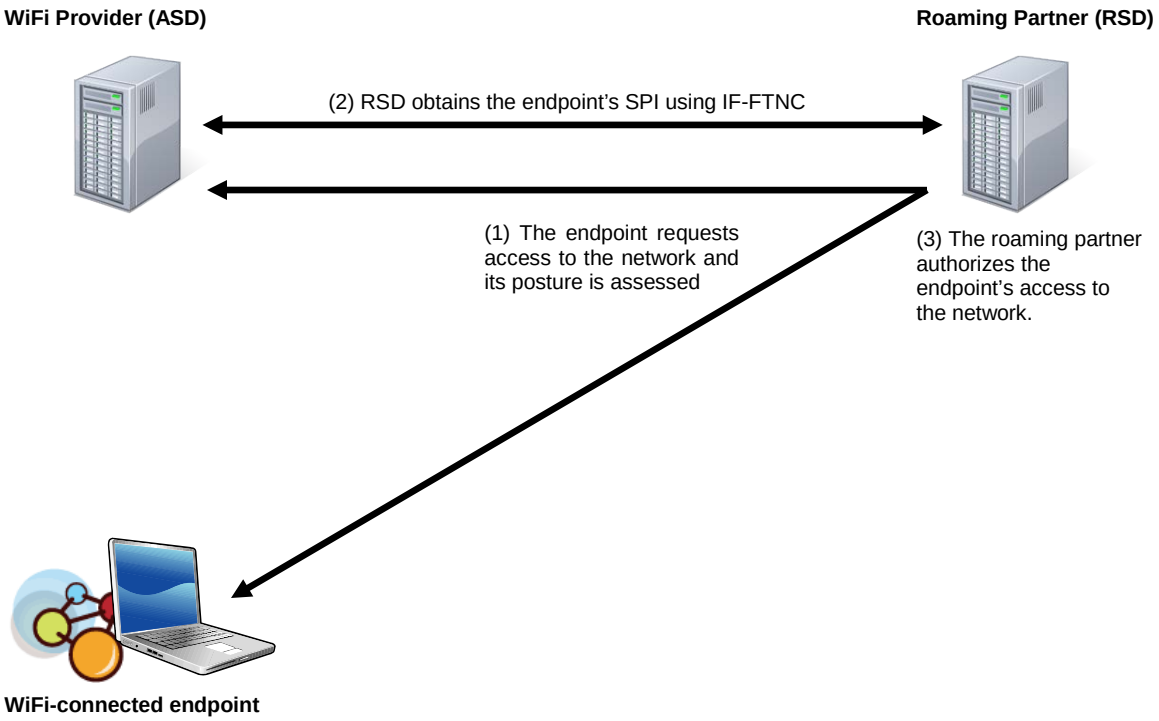

Figure 3

In this example, the WiFi Provider assesses the WiFi-connected endpoint's posture (step 1), using the IF-T EAP Bindings [7], as it connects to the Roaming Partner's network. The Roaming Partner acquires the endpoint's SPI by using IF-FTNC (step 2), enabling it to make a more informed authorization decision (step 3).

2.2.2 Web Assessment Profile

The TNC architecture is typically used to assess an endpoint's compliance with a network connection policy, and to enforce its requirements, at the time of connection to the network. Subsequent to the connection to the network, however, end users often interact with web resources that are provided by third parties (such as business partners, or software-as-a-service providers) who might have concerns about the security posture of the device being used by the end user. For example, the presence of malware could lead to the compromise of an end user's credentials.

Knowledge of the endpoint's SPI enables the web resource provider (in its role as the RSD) to request endpoint SPI from the endpoint's 'home' network (in its role as the ASD). This could be used, for example, to increase confidence that endpoints accessing the web resource are not infected with malware.

An example of the use of the Web Assessment Profile is illustrated in Figure 4 below.

Figure 4

In this example, the Employer assesses the endpoint's posture (step 1) when it connects to its home network. The Business Partner that operates the web resource acquires the endpoint's SPI by using IF-FTNC (step 2), allowing it to make a more informed authorization decision (step 3).

2.3 Overview

Figure 5 below illustrates the basic template for IF-FTNC.

The ASD operates a TNCS, which performs endpoint authentication and posture assessment, and a SAML attribute authority, which can respond to queries for attributes asserting properties of an endpoint and optionally also the user principal wielding that endpoint.

This specification supports the use of two separate SAML attribute authorities to assert endpoint and user principal attributes. However, in the interests of clarity, this configuration is not discussed explicitly in this overview and it is assumed that only a single SAML attribute authority is used.

The RSD operates a service provider, which is either a network or a web resource (for the Network and Web Assessment Profiles respectively) and a SAML requester that is used to obtain assertions describing properties of the endpoint and/or user principal that is attempting to gain access to that resource. In the interests of clarity, these are presented as a single entity (the service provider).

While the TNCS and SAML attribute authority are presented as distinct entities, it is possible to unite their functions within a single entity. Again, in the interests of clarity, this configuration is not discussed explicitly in this overview.

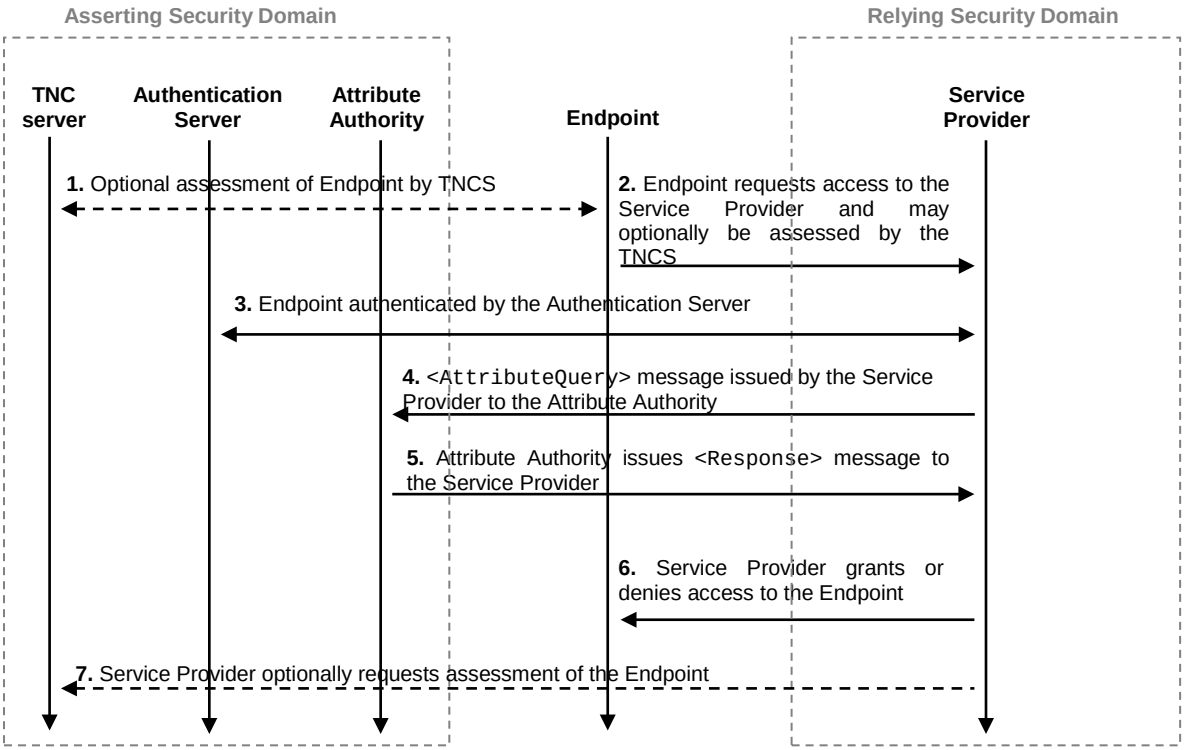

Figure 5

1. **Optional assessment of endpoint by TNCS**

   In step 1, the endpoint may be assessed by the TNCS to obtain more recent SPI for the endpoint.

   This step is optional because the endpoint may have been assessed previously by the TNCS before requesting access to the service provider, and therefore a fresh assessment may be redundant.

2. **Endpoint requests access to the Service Provider and may optionally be assessed by the TNCS**

   In step 2, the endpoint requests access to a service provider at the Relying Security Domain.

   The request method is specific to each profile. The endpoint may also be assessed during this request, if the request method supports the transport of the assessment between the endpoint and the TNCS.

   During this step, the service provider must also determine which ASD the endpoint is affiliated with; this is necessary for the following reasons:

   - To enable the endpoint to authenticate against the ASD.
   - To enable a possible assessment of the endpoint, if the profile in question supports it.
   - To enable the service provider to know where it should issue requests for the endpoint's SPI.

   This process is sometimes called 'discovery'; the discovery method is specific to each profile.

3. **Endpoint authenticated by the Authentication Server**

In step 3, the endpoint is authenticated by the authentication server at the Asserting Security Domain.

The method used to authenticate the endpoint is specific to each profile.

4. **<AttributeQuery> message issued by the Service Provider to the Attribute Authority**

   In step 4, the service provider issues an `<AttributeQuery>` message, requesting endpoint SPI, to the attribute authority.

   This request is performed using either the SAML 2.0 Assertion Query/Request Profile or the Shibboleth Attribute Exchange Profile.

5. **Attribute Authority issues <Response> message to the Service Provider**

   In step 5, the attribute authority issues a `<Response>` message to the service provider. The message may indicate an error or it will include an assertion containing the endpoint's SPI.

   This response is performed using either the SAML 2.0 Assertion Query/Request Profile or the Shibboleth Attribute Exchange Profile.

6. **Service Provider grants or denies access to the Endpoint**

   In step 6, having received the response from the attribute authority, the service provider may grant or deny access to the endpoint.

7. **Service Provider optionally requests assessment of the Endpoint**

   In step 7, the service provider may decide to request assessment of the endpoint, either immediately after authorization or at any time afterwards.

   Endpoint assessment can occur either as a consequence of the endpoint's SPI failing to satisfy the RSD's policy at time of authentication, or as a consequence of an event occurring while the endpoint is connected to the network. Some examples of events that might provoke assessment include:

   - The RSD becomes suspicious of the endpoint's behaviour.
   - The RSD receives a new policy requiring immediate action.
   - The endpoint notices a change in its local security posture.

   The method used to request assessment is specific to each profile.

   The Roaming Assessment Profile does not currently support the issuance of an assessment request by the service provider, but this is planned for future revisions of this specification.

2.4 Requirements

Here are the requirements that the IF-FTNC protocol must meet in order to successfully play its role in the TNC architecture.

- Meets the needs of the TNC architecture

  IF-FTNC must support all the functions and use cases described in the TNC architecture as they apply to the federation of TNC posture verification and all the use cases described in this document.

- Compatible with existing TNC protocols

  IF-FTNC must be compatible with existing TNC protocols for posture exchange (IF-TNCCS), metadata exchange (IF-MAP) and posture attribute exchange (IF-M)

- Transport Independent

  IF-FTNC must be capable of operating over any transport protocol, so long as this transport protocol meets the Security and other assumptions listed in Section 4. Half-duplex and full-duplex transports MUST both be supported.

- Security

  All security aspects of IF-FTNC must be fully analyzed and addressed to the greatest extent possible. At a minimum, IF-FTNC must support integrity-protection, authentication, and replay-protection for endpoint SPI. IF-FTNC should also support confidentiality of the security posture information and any other sensitive information. These security requirements may be provided by the transport protocol used by IF-FTNC.

- Scalability

  IF-FTNC must scale in terms of the ability to support:
  - The scalability requirements of the TNC protocols it federates
  - Number of connections requesting federated TNC data from IF-FTNC servers.
  - The amount of information that can be carried in the IF-FTNC exchange must scale.

- Extensibility

  IF-FTNC must permit extensibility allowing other IF-FTNC operations in the future. This must include the ability to extend the IF-FTNC data model to describe new IF-TNCCS Access Recommendations, IF-M attributes or IF-MAP metadata.

- Efficient

  IF-FTNC may delay network access until the endpoint is determined not to pose a security threat to the Relying Security Domain. To minimize user frustration, the IF-FTNC protocol MUST minimize delays and make communication as rapid and efficient as possible. Efficiency is also important when you consider that some network endpoints are small and low-powered and that some networks have high latency, high cost, or low bandwidth. Also, some transport protocols are half-duplex with a limited fragment size and require a full round trip per fragment.

- Vendor neutral

  The IF-FTNC protocol must be vendor neutral. However, it must include support for vendor-specific extensions.

- Internationalized

  If the IF-FTNC provides a way for the interacting Security Domains to exchange human readable strings then IF-FTNC must support exporting language preferences using standard message(s). In this case, any strings intended to be human readable must be adaptable so that they conform to the user's language preference.

- Fully interoperable

  IF-FTNC must be designed so that any Asserting and Relying Security Domain that comply with the IF-FTNC specification will interoperate (assuming that they support a common transport protocol and have a compatible set of federation policies).

- Flexibility

  IF-FTNC should be flexible, to enable one RSD to exchange assertion attributes with multiple ASDs. If the TNC client can be assessed by multiple ASDs, then the RSD should be able to negotiate which assertions it receives from the ASD to complete the IF-FTNC exchange.

- Privacy

  IF-FTNC must preserve user privacy. Security posture information may contain personal or domain-specific security data and must therefore be protected in a manner consistent with various laws, guidelines, and regulations. For example, user consent or policy configuration may be required before this data can be collected or released to another party or RSD.

- Discovery

  IF-FTNC must specify discovery mechanisms to enable RSDs to find the appropriate ASD for a given endpoint, so that the appropriate security posture information can be acquired. The discovery method specified must reuse existing schemes that are available in the specific protocol binding used for IF-FTNC.

2.5 Non-Requirements

There are certain requirements that the IF-FTNC protocol explicitly is not required to meet. This list may not be exhaustive.

- Failover

  IF-FTNC does not need to provide explicit support for failover mechanisms that allow an IF-FTNC server to gracefully recover from a failure.

2.6 Assumptions

Here are the assumptions that the IF-FTNC protocol makes about other components in the TNC architecture.

- Existence of a method for determining TNC client posture (for example, IF-TNCCS and IF-M)

- Existence of transport protocol that must provide Integrity, Mutual Authentication, Non-repudiation and should provide confidentiality between the Asserting and Relying Security Domains. The transport protocol must also provide reliable transport in the sense that failures can be noted by the sender. Either half duplex and full duplex transports are acceptable.

- The Relying Security Domain never performs a direct TNC assessment with the endpoint since it does not have a trust relationship with the endpoint; instead it relies on the Asserting Security Domain to provide it with information about the endpoint.

3 IF-FTNC Specification

3.1 Roaming Assessment Profile

In the scenario supported by the Roaming Assessment Profile, a roaming endpoint requests access to a network service provided by the RSD. The endpoint uses EAP to authenticate to the ASD, through the RSD, routed and transported over the RadSec or RADIUS protocols. The endpoint's posture may also be assessed during the EAP authentication. SAML name identifiers for the endpoint and/or user principals are established between the security domains, which the RSD uses to acquire SAML assertions, possibly including endpoint SPI, for these principals. The RSD uses these assertions to authorize the user and/or endpoint's access to the network.

3.1.1 Required Information

Identification: https://www.trustedcomputinggroup.org/XML/SAML/2008/ifftnc/1/profiles/roaming

Contact information: admin@trustedcomputinggroup.org

Description: Given below.

Updates: None.

3.1.2 Profile Overview

Figure 6 below shows the processing flow in the Roaming Assessment Profile. The following steps are described by the profile. Within an individual step, there may be one or more actual messages exchanged.

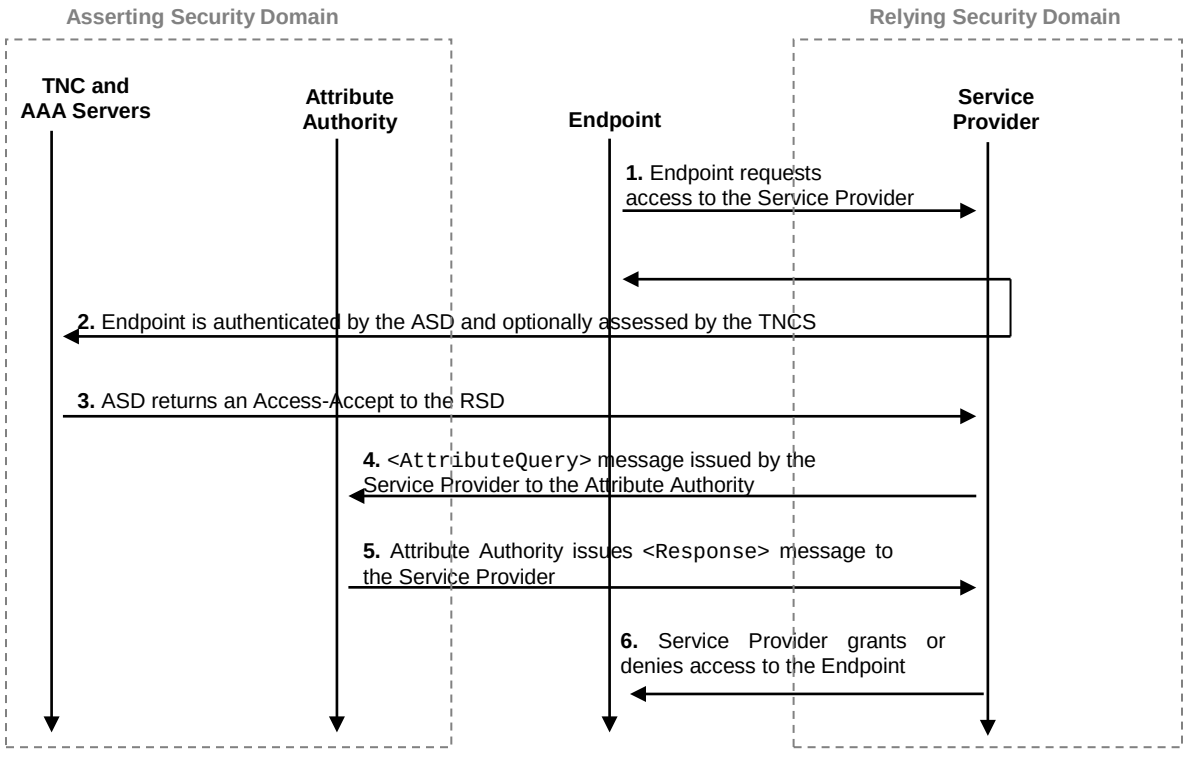

Figure 6

1. **Endpoint requests access to the Service Provider**

   In step 1, the endpoint requests access to the service provider.

2. **Endpoint is authenticated by the ASD and optionally assessed by the TNCS**

In step 2, the endpoint is authenticated by the ASD. The endpoint's posture may also be assessed by the ASD's TNCS to set fresh SPI for the endpoint.

3. **ASD returns an Access-Accept to the RSD**

   In step 3, if the endpoint authenticates successfully the ASD's RADIUS server should return a RADIUS Access-Accept packet to the service provider.

4. **<AttributeQuery> message issued by the Service Provider to the Attribute Authority**

   In step 4, the service provider may request SPI from the endpoint's SAML attribute authority.

5. **Attribute Authority issues <Response> message to the Service Provider**

   In step 5, the attribute authority issues a `<Response>` message to the service provider. The message may indicate an error or will include a SAML attribute assertion containing endpoint SPI.

6. **Service Provider grants or denies access to the Endpoint**

   In step 6, having received the response from the attribute authority, the service provider may grant or deny access to the endpoint.

3.1.3 Profile Description

3.1.3.1 Endpoint requests access to the Service Provider

The endpoint MUST initiate the profile by requesting access to the network operated by the service provider that is authenticated using EAP. It is RECOMMENDED that one of the EAP methods supported by [7] is used to permit posture assessment of the endpoint. The endpoint MUST identify itself by providing a UTF-8 encoded Network Access Identifier (NAI) [8] in its EAP-Identity/Response message.

The realm component of the NAI identifies the ASD with which the endpoint is affiliated, and which will authenticate it and optionally assess its posture (for example, if the NAI given by the endpoint is "user@example.com", the endpoint is affiliated with the security domain that is authoritative for the realm named "example.com").

The service provider MUST use the NAI to route the EAP packets towards the ASD. Implementations of this specification MUST support routing of unknown realms (sometimes called the 'default' realm).

If there is no realm given in NAI, or the realm is the same as the realm of the service provider, the service provider SHOULD assume that the endpoint is a local endpoint and terminate the processing of the Roaming Assessment Profile. The service provider MAY also attempt to authenticate and assess the posture of the endpoint.

Implementations of this specification MUST support RadSec [9], but RADIUS [10] MAY also be supported. Section 4.2.5 describes some of the security problems that may arise when using RADIUS proxies and which can be avoided by using RadSec.

The service provider MUST send single instances of the `IF-FTNC-RSD-Attribute-Requestor-EntityID` RADIUS attribute, giving the SAML entity identifier of the SAML requester that will subsequently request SPI for the endpoint if authentication succeeds, and the `IF-FTNC-Version` RADIUS attribute, giving the version of FTNC supported by the service provider (which MUST be set to a value of "1"), in every RADIUS Access-Request packet that it sends.

The `IF-FTNC-RSD-Attribute-Requestor-EntityID` and `IF-FTNC-Version` RADIUS attributes are defined in section 3.1.5.

3.1.3.2 Endpoint is authenticated by the ASD and optionally assessed by the TNCS

The ASD MUST attempt to authenticate the endpoint and MAY also attempt to assess its posture.

The ASD MUST check that the IF-FTNC-Version attribute is set to a value of "1" and MAY terminate the processing of the Roaming Assessment Profile if another value is found.

The IF-FTNC-RSD-Attribute-Requestor-EntityID RADIUS attribute (defined in section 3.1.3.4) informs the ASD which SAML requester it should expect to receive subsequent assertion requests and may be used for logging or other purposes.

3.1.3.3 ASD returns an Access-Accept to the RSD

If the endpoint is authenticated and authorized, the ASD MUST send a RADIUS Access-Accept packet to the service provider.

The ASD MUST decide if it wants to issue SAML assertions about the endpoint principal or the user principal wielding the endpoint at the instant of authentication.

If the ASD decides to issue SAML assertions expressing endpoint attributes (including, but not limited to, endpoint SPI) to the service provider, the RADIUS Access-Accept packet MUST contain the following attributes:

- A single instance of the IF-FTNC-Version RADIUS attribute, which MUST be set to a value of "1".
- One or more instances of the IF-FTNC-EIU-Name-Identifier RADIUS attribute giving the identity of the endpoint. Implementations MUST support the use of the MAC Address name identifier format (defined in section 3.4.1.1) expressed within
    - a SAML 2.0 <NameID> element
    - a SAML 2.0 <EncryptedID> element carrying the encrypted value of a SAML 2.0 <NameID> element

  Implementations MAY also use the SAML 1.1 <NameIdentifier> element.

  Implementations MAY support other name identifiers and formats.

  If the <EncryptedID> element is used, the name identifier MUST be encrypted using a previously established symmetric key. After encryption, the ciphertext MUST be placed in the <EncryptedData> element and the <EncryptedID> element MUST NOT contain an <EncryptedKey> element. The <EncryptedData> element MUST contain a <KeyInfo> element giving a <KeyName> element whose value MUST uniquely identify the RADIUS server.
- A single instance of the IF-FTNC-EIU-Attribute-Authority-EntityID RADIUS attribute, giving the SAML entity identifier of the endpoint's SAML attribute authority.

If the ASD decides to issue SAML assertions expressing attributes of the user principal wielding the endpoint to the service provider, the RADIUS Access-Accept packet MUST contain single instances of:

- A single instance of the IF-FTNC-Version RADIUS attribute, which MUST be set to a value of "1".
- One of more instances of the IF-FTNC-User-Name-Identifier RADIUS attribute giving the identity of the user. Implementations MUST support the use of the NAI name identifier format (defined in section 3.4.1.2) expressed within
    - a SAML 2.0 <NameID> element
    - a SAML 2.0 <EncryptedID> element carrying the encrypted value of a SAML 2.0 <NameID> element

  Implementations MAY also use the SAML 1.1 <NameIdentifier> element.

  Implementations MAY support other name identifiers and formats.

If the <EncryptedID> element is used, the name identifier MUST be encrypted using a previously established symmetric key. After encryption, the ciphertext MUST be placed in the <EncryptedData> element and the <EncryptedID> element MUST NOT contain an <EncryptedKey> element. The <EncryptedData> element MUST contain a <KeyInfo> element giving a <KeyName> element whose value MUST uniquely identify the RADIUS server.

- A single instance of the IF-FTNC-User-Attribute-Authority-EntityID RADIUS attribute, giving the SAML entity identifier of the user principal's SAML attribute authority.

The RSD MAY terminate the session if the IF-FTNC-Version RADIUS attribute is set to a value other than "1".

The IF-FTNC-EIU-Name-Identifier, IF-FTNC-EIU-Attribute-Authority-EntityID, IF-FTNC-User-Name-Identifier and IF-FTNC-User-Attribute-Authority-EntityID RADIUS attributes are defined in section 3.1.5.

If the RSD decides not to issue any assertions, it MUST not return these attributes in the RADIUS Access-Accept packet. If the service provider fails to receive these attributes as described, it MUST terminate the processing of this profile.

If the endpoint fails to authenticate, or is not authorized, the ASD MUST send a RADIUS Acccess-Reject packet to the service provider who MUST terminate the processing of this profile.

3.1.3.4 <AttributeQuery> message issued by the Service Provider to the Attribute Authority

The service provider MUST use either the SAML 2.0 Assertion Query/Request Profile [12] (which MUST be supported by implementations of this specification) or the Shibboleth Attribute Exchange Profile [13] (which MAY be supported by implementations of this specification) to obtain SAML assertions.

In both cases, the SAML requester MUST use the name identifier(s), given in the IF-FTNC-EIU-Name-Identifier and/or IF-FTNC-User-Name-Identifier attribute(s) obtained in the previous step, to name the subject of the query.

In the case of the SAML 2.0 Assertion Query/Request Profile, the SAML requester MUST issue the query as defined in section 6 of [12].

In the case of the Shibboleth Attribute Exchange Profile, the SAML requester MUST issue the query as defined in section 3 of [13].

The SAML requester MAY use the entity identifier(s), given in the IF-FTNC-EIU-Attribute-Authority-EntityID and/or IF-FTNC-User-Attribute-Authority-EntityID attribute(s), to determine the location of the appropriate SAML attribute authority(s) endpoint(s) and the supporting bindings; this information MAY be obtained from SAML 2.0 metadata [15].

If name and entity identifiers were obtained for both endpoint and user principals, the service provider MUST treat each attribute query and response separately.

3.1.3.5 Attribute Authority issues <Response> message to the Service Provider

The SAML attribute authority MUST process the query as defined in section 6 of [12], in the case of the SAML 2.0 Assertion Query/Request Profile, or section 3 of [13], in the case of the Shibboleth Attribute Exchange Profile.

3.1.3.6 Service Provider grants or denies access to the Endpoint

The service provider MUST process the <Response>(s) and <Assertions>(s) and grant or deny access to the network.

3.1.4 Use of Metadata

The use of metadata to described entities is RECOMMENDED.

If SAML metadata is used, the RSD's SAML requester MUST be described using the attribute query requester role descriptor type defined in [14].

3.1.5 RADIUS attributes

This section defines the RADIUS Vendor-Specific Attributes used by this profile.

A summary of the Vendor-Specific Attribute format, defined by [10], is shown below; the fields are transmitted from left to right.

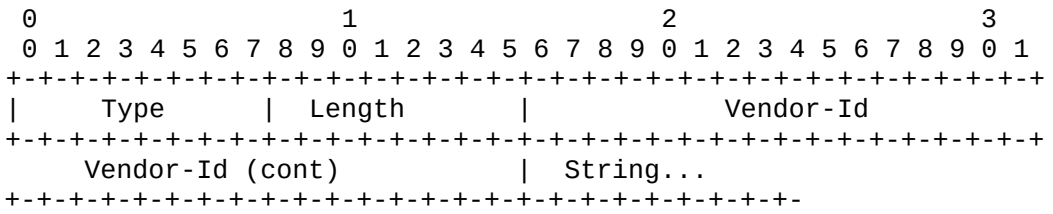

Type (8 bits)

26 for Vendor-Specific

Length (8 bits)

>= 7

Vendor-Id (32 bits)

The high-order octet is 0 and the low-order 3 octets are the TCG SMI PEN (0x005597).

String (variable length)

The string field is one or more octets giving one of the Vendor-Specific Attributes described in the following sections.

3.1.5.1 IF-FTNC-Version

This attribute indicates the version of IF-FTNC supported by this ASD or RSD. It MUST be sent in all Access-Request and Access-Accept packets and MUST be set to a value of "1". A summary of this attribute is shown below. The fields are transmitted from left to right.

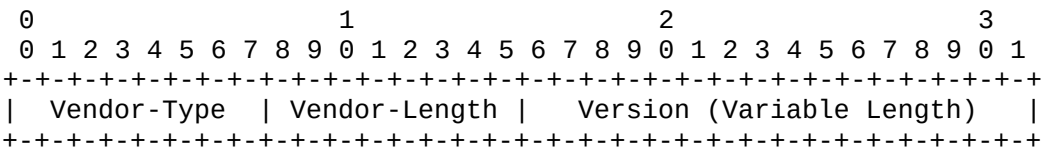

Vendor-Type (8 bits)

1

Vendor-Length (8 bits)

4

Version (16 bits)

This field indicates the version of IF-FTNC supported by this security domain and encoded using the RADIUS "text" type.

3.1.5.2 IF-FTNC-RSD-Attribute-Requestor-EntityID

This attribute gives the entity identifier of the RSD's SAML requester. It MUST only be used in the Access-Request packets sent to a RADIUS server. A summary of this attribute is shown below. The fields are transmitted from left to right.

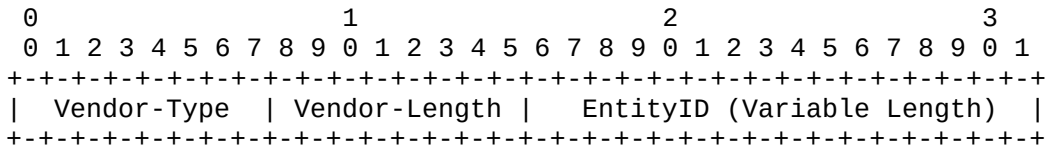

Vendor type (8 bits)

2

Vendor length (8 bits)

>= 3

EntityID (variable length)

The EntityID field is one or more octets giving the entity identifier of the service provider's SAML requester. This value MUST NOT exceed 245 characters. This is encoded as the "text" RADIUS type.

3.1.5.3 IF-FTNC-EIU-Attribute-Authority-EntityID

This attribute gives the entity identifier of the SAML attribute authority that can respond to attribute queries for the endpoint. It MUST only be used in Access-Accept packets. A summary of this attribute is shown below. The fields are transmitted from left to right.

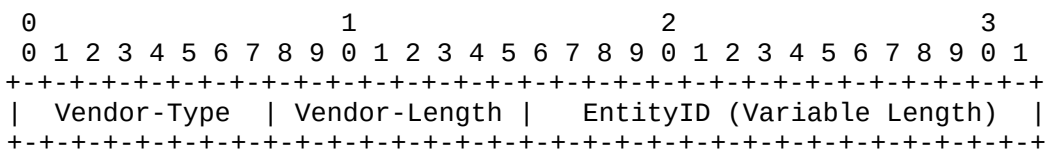

Vendor-Type (8 bits)

3

Vendor-Length (8 bits)

>= 3

EntityID (variable length)

The EntityID field is one or more octets giving the entity identifier of the endpoint's SAML attribute authority. This is encoded as the "text" RADIUS type.

3.1.5.4 IF-FTNC-User-Attribute-Authority-EntityID

This attribute gives the entity identifier of the SAML attribute authority that can respond to attribute queries for the user. It MUST only be used in Access-Accept packets. A summary of this attribute is shown below. The fields are transmitted from left to right.

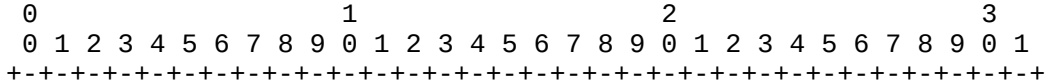

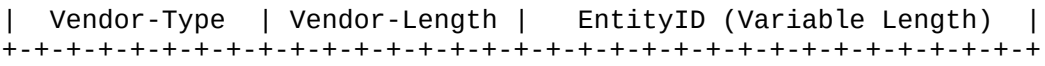

Vendor-Type (8 bits)

4

Vendor-Length (8 bits)

>= 3

EntityID (variable length)

The EntityID field is one or more octets giving the entity identifier of the user's SAML attribute authority. This value MUST NOT exceed 245 characters. This is encoded as the "text" RADIUS type.

3.1.5.5 IF-FTNC-EIU-Name-Identifier

This attribute gives a SAML name identifier for the endpoint. It MUST only be used in Access-Accept packets. A summary of this attribute is shown below. The fields are transmitted from left to right.

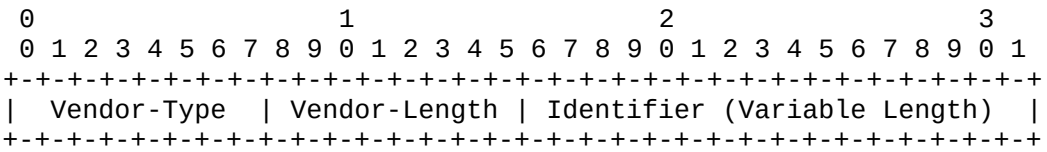

Vendor-Type (8 bits)

5

Vendor-Length (8 bits)

>= 3

Identifier (variable length)

This field gives a SAML name identifier for the endpoint. This is encoded as the "string" RADIUS type. If multiple `IF-FTNC-EIU-Name-Identifier` attributes are contained within an Access-Accept RADIUS packet they MUST be in order and they MUST be consecutive attributes. The attribute values are concatenated to recover the name identifier.

3.1.5.6 IF-FTNC-User-Name-Identifier

This attribute gives a SAML name identifier for the use. It MUST only be used in Access-Accept packets. A summary of this attribute is shown below. The fields are transmitted from left to right.

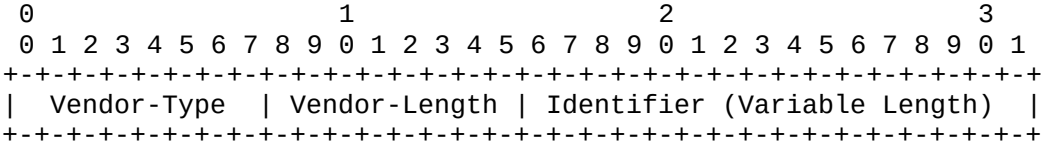

Vendor-Type (8 bits)

6

Vendor-Length (8 bits)

>= 3

Identifier (variable length)

This field gives a SAML name identifier for the user. This is encoded as the “string” RADIUS type. If multiple `IF-FTNC-User-Name-Identifier` attributes are contained within an Access-Accept RADIUS packet they MUST be in order and they MUST be consecutive attributes. The attribute values are concatenated to recover the name identifier.

3.2 Web Assessment Profile

In the scenario supported by the Web Assessment Profile, a user requests access to a web resource provided by the RSD. The user is authenticated using a Web Single Sign-On (SSO) profile of SAML to the ASD. A SAML name identifier for the endpoint principal is established between the security domains, which the RSD uses to acquire SAML assertions, possibly including endpoint SPI, for this principal. The RSD uses these assertions to authorize the user and/or endpoint’s access to the web resource.

3.2.1 Required Information

Identification: https://www.trustedcomputinggroup.org/XML/SAML/2008/ifftnc/1/profiles/web

Contact information: admin@trustedcomputinggroup.org

Description: Given below.

Updates: None.

3.2.2 Profile Overview

Figure 7 below shows the processing flow in the Web Assessment Profile. The following steps are described by the profile. Within an individual step, there may be one or more actual messages exchanged.

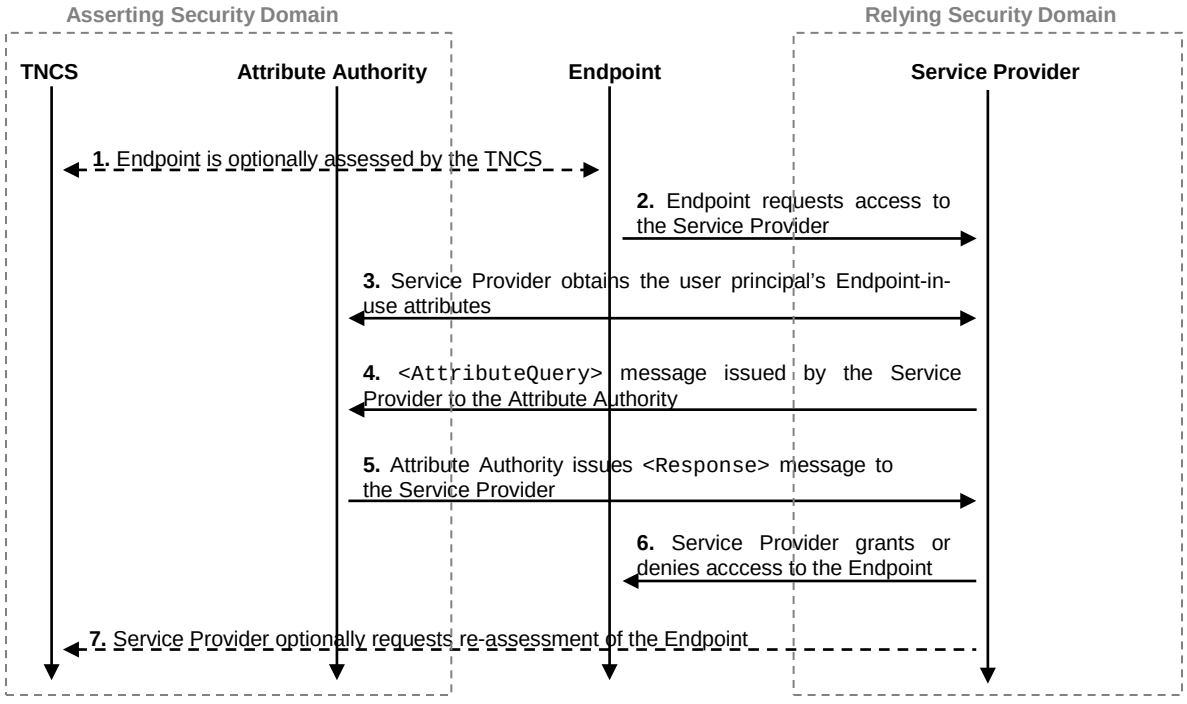

Figure 7

1. **Endpoint is optionally assessed by the TNCS.**

In step 1, the endpoint's posture may be assessed by the ASD's TNCS to set fresh SPI.

2. **Endpoint requests access to the Service Provider**

   In step 2, the endpoint requests access to the service provider.

3. **Service Provider obtains the user principal's Endpoint-in-use attributes**

   In step 3, the service provider obtains the user principal's Endpoint-in-use attributes using a Web SSO profile of SAML.

4. **<AttributeQuery> message issued by the Service Provider to the Attribute Authority**

   In step 4, the service provider requests endpoint SPI from the appropriate SAML attribute authority, using the information given in the Endpoint-in-use attributes.

5. **Attribute Authority issues <Response> to the Service Provider**

   In step 5, the SAML attribute authority issues a `<Response>` message to the service provider. The message may indicate an error or will include an attribute assertion containing the endpoint's SPI.

6. **Service Provider grants or denies access to the Endpoint**

   In step 6, having received the response from the attribute authority, the service provider may grant or deny access to the endpoint.

7. **Service Provider optionally requests assessment of the Endpoint**

   In step 7, the service provider may decide to request assessment of the endpoint, either immediately after authorization or some time afterwards.

3.2.3 Profile description

3.2.3.1 Endpoint is optionally assessed by the TNCS

The endpoint MAY be assessed by the TNCS at any time before requesting access to the service provider.

3.2.3.2 Endpoint requests access to the Service Provider

The user principal, wielding a browser, initiates any Web SSO Profile of SAML by requesting access to a secured resource at the service provider without an existing security context.

This does not require that the initial access request must be made at the service-provider. That would be undesirable because the SAML 1.1 Web SSO Profile [18] does not specify "SP-first" initiation, unlike the SAML 2.0 Web SSO Profile [12]. This step can therefore also be realised using the standard SAML 1.1 "IdP-first" initiation.

However, "SP-first" operation is often desirable and consequently SAML 1.1 implementations have adopted various conventions to provide SP-first operation, such as the Shibboleth Authentication Request Protocol [13].

Therefore, to facilitate interoperability, implementations of this specification MUST support the SAML 2.0 Web SSO Profile and MAY support the SAML 1.1 Web SSO Profile. If the SAML 1.1 Web SSO profile is implemented, it SHOULD support the Shibboleth Authentication Request Profile.

3.2.3.3 Service Provider obtains the user principal's Endpoint-in-use attributes

The service provider MUST use the Web SSO Profile initiated in the previous step to obtain the user principal's `tcg-tnc-ifftnc-eiu-name-identifier` and `tcg-tnc-ifftnc-eiu-SamlAaEntityId` attributes (these are defined in section 3.3.6.2) from the user principal's SAML attribute authority.

The ASD MUST decide if it wants to issue SAML assertions about the endpoint principal to the service provider. If the ASD decides not to issue any assertions, it MUST return these attributes with no values. If the service provider fails to obtain values for these attributes, it MUST terminate the processing of the profile.

3.2.3.4 <AttributeQuery> message issued by the Service Provider to the Attribute Authority

The service provider MUST use either the SAML 2.0 Assertion Query/Request Profile [12] (which MUST be supported by implementations of this specification) or the Shibboleth Attribute Exchange Profile [13] (which MAY be supported by implementations of this specification) to obtain SAML assertions.

In both cases, the SAML requester MUST use the name identifier, given in the `tcg-tnc-ifftnc-eiu-name-identifier` attribute obtained in the previous step, to name the subject of the query.

In the case of the SAML 2.0 Assertion Query/Request Profile, the SAML requester MUST issue the query as defined in section 6 of [12].

In the case of the Shibboleth Attribute Exchange Profile, the SAML requester MUST issue the query as defined in section 3 of [13].

The SAML requester MAY use the entity identifier, given by the `tcg-tnc-ifftnc-eiu-SamlAaEntityId` attribute, to determine the location of the appropriate SAML attribute authority's endpoint and the supported bindings; this information MAY be obtained from SAML 2.0 metadata [15].

3.2.3.5 Attribute Authority issues <Response> message to the Service Provider

The ASD's SAML attribute authority MUST process the query as defined in section 6 of [12], in the case of the SAML 2.0 Assertion Query/Request Profile, or section 3 of [13], in the case of the Shibboleth Attribute Exchange Profile.

3.2.3.6 Service Provider grants or denies access to the Endpoint

The service provider MAY process the `<Response>` and `<Assertions>`(s) and grant or deny access to the network.

3.2.3.7 Service Provider optionally requests assessment of the Endpoint

The service provider MAY request assessment of the endpoint immediately after authorization or at any time afterwards.

To request assessment of the endpoint, the service provider MUST redirect the user principal's browser to the ASD's FTNC service endpoint. The FTNC service endpoint is a web service, operated by the ASD, that provides services to RSDs.

The URI of this endpoint is constructed by appending the value of the `tcg-tnc-ifftnc-endpoint-FtncServiceEndpoint` attribute (given below as `<FtncServiceEndpoint>`) with the following parameters. This URI MUST be requested using the HTTP GET method.

```
<FtncServiceEndpoint>?request=assess&reason=<reason>&policy=<policyUri>&endpoint=<nameIdentifier>&sp=<spEntityId>&returnUri=<returnUri>
```

- `<request>`: this MUST take the value "assess".
- `<ftncServiceEndpoint>`: the endpoint's FTNC service endpoint; this is given by the `tcg-tnc-ifftnc-general-FtncServiceEndpoint` SAML attribute defined in section 3.3.6.2.
- `<reason>`: the reason given by the RSD for the assessment request. This MUST take one of the values specified in Table 1 below. This parameter is mandatory.

- <policyUri>: a reference to a policy, given as a URI, that names the policy controlling access to this resource. This parameter is optional.
- <nameIdentifier>: the name identifier associated with this endpoint, given by the tcg-tnc-ifftnc-eiu-name-identifier attribute. This parameter is mandatory.
- <spEntityId>: the entity identifier of the service provider. This parameter is mandatory.
- <returnUri>: the URI that the service provider suggests that the user principal's browser is redirected back to, following assessment. This URI MAY contain information that enables the service provider to recognise the endpoint when it returns (alternatively other state maintenance methods, such as cookies, MAY be used). This parameter is optional.

Table 1 below defines the reason code values.

Code	Reason description
0	**Reserved for experimental use.**
1	**No reason given by RSD.**
2	**The value of the OldestReceivedInformationAge attribute exceeds the RSD's policy.**
3	**The endpoint's SPI does not satisfy the RSD's policy.**
4	**The endpoint's behaviour is suspicious and suggests that the endpoint may not satisfy the RSD's policy.**

Table 1

The policyURI, endpoint, spEntityId and returnUri parameters MUST be encoded using the DEFLATE compression method [19] and then base64-encoded [20] and then URL-encoded [21].

The ASD, on receiving the request for this URI, MAY use the name identifier given in the endpoint parameter, or any other criteria, to determine if this is a valid endpoint before attempting to assess it. If the ASD does not assess the endpoint's posture, it SHOULD inform the user of this to reduce potential confusion and frustration.

The ASD MAY use any method it chooses to assess the endpoint; for example, the ASD might:

- Use IF-PEP [22] to forcibly disconnect the endpoint from the network and assess the endpoint when it next attempts to connect to the network.
- Use a layer 3 protocol to assess the endpoint.
- Use other proprietary web-based posture assessment systems.

After the endpoint has been assessed, the ASD MAY redirect the browser to the URI given by the ReturnUri parameter. The service provider MAY request new endpoint SPI by returning to step 3.2.3.4.

This request may not yield any assertions in the <Response> from the ASD if the endpoint's name identifier has expired or been cancelled. In this event, the service provider MAY attempt to obtain a new name identifier for the endpoint by returning to step 3.2.3.3.

The request for a new name identifier for the endpoint may also fail to yield any assertions in the <Response> from the ASD if the user principal's name identifier has expired or been cancelled. In this event, the service provider MAY attempt to obtain a new name identifier for the user principal by returning to step 3.2.3.2.

3.3 SAML Attribute Profile

3.3.1 Required Information

Identification: https://www.trustedcomputinggroup.org/XML/SAML/2008/ifftnc/1/profiles/attribute

Contact information: admin@trustedcomputinggroup.org

Description: Given below.

3.3.2 Profile Overview

This profile defines a common convention for the naming and representation of TNC IF-M attributes [4], IF-TNCCS Access Recommendation types [3] and IF-MAP metadata [2] when expressed as SAML 1.1 and 2.0 attributes.

3.3.3 Attribute Naming

The `Name` XML attribute is based on the OID assigned to the TNC object following the conventions described in Section 3.3.5 below.

SAML 1.1 does not specify any interoperable means of establishing the kind of name used, so the convention used within this profile is that the `AttributeNamespace` XML attribute in `<saml:Attribute>` elements MUST be set to:

```
urn:oasis:names:tc:SAML:2.0:attrname-format:uri
```

If the `FriendlyName` XML attribute is used then it MAY carry the name of the IF-M attribute name, IF-TNCCS Access Recommendation type or IF-MAP result-filter friendly name.

3.3.4 Attribute Name Comparison

Two `<Attribute>` elements refer to the same SAML attribute if and only if their `Name` XML attribute values are equal in the sense of [23]. The `FriendlyName` attribute plays no role in the comparison.

3.3.5 Attribute Encryption

SAML 2.0 provides support for encrypted attributes; implementations of this specification MAY encrypt attributes.

3.3.6 Attribute Values

The syntax rules defined by the SAML 2.0 X.500/LDAP attribute profile [12] are to be applied, with the caveat that the XML attribute named `Encoding` defined by that profile is NOT specified for use with this profile.

For SAML 1.1, the `<saml:AttributeValue>` element is also substituted for the `<saml2:AttributeValue>`.

To ensure uniqueness, each attribute type is assigned a name in the form of a URI. To construct attribute names, the URN `oid` namespace described in [23] is used.

This specification defines names attributes using object identifiers within the IF-FTNC sub-tree; this falls under the following hierarchy:

```
-- OID Hierarchy for TNC IF-FTNC

tcg OBJECT IDENTIFIER ::= { joint-iso-itu-t(2) international-organization(23) tcg(133) }

tcg-tnc OBJECT IDENTIFIER ::= { tcg 16 }

tcg-tnc-ifftnc OBJECT IDENTIFIER ::= { tcg-tnc 2 }
```

3.3.6.1 Posture attribute naming

OIDs that reference TNC posture attributes fall under the following hierarchy:

-- OID Hierarchy for TNC posture attributes

tcg-tnc-ifftnc-posture OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc 1 }

TNC posture attributes MUST be named within this hierarchy by sequentially suffixing the `tcg-tnc-ifftnc-posture` OID with the string values of the TNCCS-IF-M-Message's IF-M Vendor ID and IF-M Subtype IF-TNCCS Message Type fields and the string values of the IF-M attribute's Vendor Code and Attribute Type fields (in that order).

For example, the OID used to name the IF-M Port Filter attribute (which lists the set of network ports that are allowed or blocked on the endpoint) is:

`tcg-tnc-ifftnc-posture.5597.50.5597.50`

3.3.6.2 Endpoint-in-use Attributes

The Endpoint-in-use (EIU) attributes are attributes of the user principal that name the endpoint that the user is using, and its associated SAML attribute authority.

-- OID Hierarchy for TNC IF-FTNC Endpoint-in-use attributes

tcg-tnc-ifftnc-eiu OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc 2 }

tcg-tnc-ifftnc-eiu-NameIdentifier OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc-eiu 1 }

tcg-tnc-ifftnc-eiu-SamlAaEntityId OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc-eiu 2 }

tcg-tnc-ifftnc-eiu-NameIdentifier

This attribute takes a value giving the identity of the endpoint.

Implementations MUST support the use of MAC Address name identifier format (defined in section 3.4.1.1) within

- a SAML 2.0 `<NameID>` element
- a SAML 2.0 `<EncryptedID>` element carrying the encrypted value of a SAML 2.0 `<NameID>` element

Implementations MAY also use the SAML 1.1 `<NameIdentifier>` element.

Implementations MAY support other name identifiers and formats.

If the `<EncryptedID>` element is used, the name identifier MUST be encrypted using a previously established symmetric key. After encryption, the ciphertext MUST be placed in the `<EncryptedData>` element and the `<EncryptedID>` element MUST NOT contain an `<EncryptedKey>` element. The `<EncryptedData>` element MUST contain a `<KeyInfo>` element giving a `<KeyName>` element whose value MUST be the entity identifier of the issuing SAML attribute authority.

tcg-tnc-ifftnc-eiu-SamlAaEntityId

This attribute takes a value that MUST be the entity identifier of the endpoint's SAML attribute authority.

3.3.6.3 Endpoint Attributes

The Endpoint attributes describe a number of useful properties of an endpoint.

-- OID Hierarchy for TNC IF-FTNC Endpoint attributes

tcg-tnc-ifftnc-endpoint OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc 3 }

tcg-tnc-ifftnc-endpoint-FtncServiceEndpoint OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc-endpoint 1 }

tcg-tnc-ifftnc-endpoint-OldestReceivedInformationAge OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc-endpoint 2 }

tcg-tnc-ifftnc-endpoint-IftnccsAccessRecommendation OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc-endpoint 3 }

tcg-tnc-ifftnc-endpoint-FtncServiceEndpoint

This attribute takes a value that MUST be a URI giving the endpoint's FTNC service endpoint. This URI MUST use either the "http" or "https" scheme.

tcg-tnc-ifftnc-endpoint-OldestReceivedInformationAge

This attribute takes a value giving the instant, expressed as the number of seconds elapsed since the Unix epoch, when the endpoint's oldest collected security posture information was collected.

tcg-tnc-ifftnc-endpoint-IftnccsAccessRecommendation

This attribute takes a value giving the IF-TNCCS Access Recommendation type.

3.3.6.4 IF-MAP result-filters

OIDs that reference IF-MAP metadata result-filters fall under the following hierarchy:

-- OID Hierarchy for TNC IF-MAP metadata result-filters

tcg-tnc-ifftnc-ifmapResultFilters OBJECT IDENTIFIER ::= { tcg-tnc-ifftnc 4 }

IF-MAP result-filters MUST be named within this hierarchy by sequentially suffixing the `tcg-tnc-ifftnc` OID with a Private Enterprise Number and result-filter code (in that order).

This specification provides a set of standard TNC IF-MAP result-filters; these are laid out in Table 2 below. These filters are uniquely identified by the TCG SMI PEN (0x005597) and the result-filter code given. Non-standard result-filters may be defined by using the SMI and Filter code namespaces.

Code	Friendly name	Result-filter definition	Result-filter description
0	Experimental	This filter is reserved for experimental use.	**This filter is reserved for experimental use.**
1	**All**	**Treat as no result-filter attribute given (see section 3.8.2.6 of IF-MAP 1.0).**	**This filter is used to request all IF-MAP metadata on all identifiers and links.**

Table 2

For example, the OID used to name the 'All' result-filter is:
`tcg-tnc-ifftnc-ifmapResultFilters.5597.1`

3.4 Name Identifier Format Identifiers

The following identifiers MAY be used in the `Format` attribute of the SAML 2.0 `<NameID>` and SAML 1.1 `<NameIdentifier>` elements.

3.4.1.1 MAC Address

URI: https://www.trustedcomputinggroup.org/XML/SAML/2008/ifftnc/1/nameid-format/mac

Indicates that the content of the element is in the form of a MAC address expressed as six groups of two hexadecimal digits separated by hyphens (-) or colons (:), in transmission order.

3.4.1.2 Network Access Identifier

URI: https://www.trustedcomputinggroup.org/XML/SAML/2008/ifftnc/1/nameid-format/nai

Indicates that the content of the element is in the form of a Network Access Identifier as defined in [8].

4 Security Considerations

IF-FTNC defines and specifies standard interfaces for the exchange of endpoint SPI between security domains. This section provides a security analysis of IF-FTNC and the surrounding system as it relates to IF-FTNC. Three subsections define the trust model (which components are trusted to do what), the threat model (attacks that may be mounted on IF-FTNC and the system), and the countermeasures (ways to address or mitigate the threats previously identified).

4.1 Trust Model

The first step in analyzing the security of IF-FTNC is to describe the trust model, listing what each architectural component is trusted to do. The items listed here are assumptions but provisions are made in the Threat Model and Countermeasures sections for components that fail to perform as they were trusted to do.

4.1.1 Asserting Security Domain

The ASD is trusted to:

- Create or otherwise obtain accurate SPI.
- Protect the confidentiality and integrity of SPI.
- Convey SPI securely to authorized parties (e.g. Relying Security Domains).
- Resist attacks (including denial of service and attacks from endpoints, RSDs, etc.).

4.1.2 Relying Security Domain

The RSD is trusted to:

- Protect the confidentiality and integrity of SPI received from the ASD
- Verify the identity of the ASD and otherwise gauge the credibility of SPI
- Use SPI properly in making access control decisions
- Resist attacks (including denial of service and attacks from endpoints, ASDs, etc.)

4.1.3 Endpoint

The endpoint may be compromised or its user may be hostile. Therefore, the endpoint is not generally trusted. However, if the ASD gathers measurements from the endpoint and uses those measurements to create SPI then the endpoint is trusted to send accurate measurements. If the measurements come from the endpoint's TPM then it may not be necessary to trust the rest of the endpoint.

4.1.4 Network

The network used to transfer information between the ASD, RSD, and endpoint may include and depend on untrusted or actively hostile parties. Therefore, this network is not trusted.

4.1.5 RADIUS Proxies

When RADIUS proxies are employed in the Roaming Assessment Profile, they are trusted to transmit RADIUS packets to the next hop in the proxy chain without modifying or revealing any of the RADIUS attributes. Since these proxies are typically not part of the ASD or RSD's security domain, this trust may not be well founded. Therefore, this specification recommends that RadSec SHOULD be used instead of RADIUS proxies.

4.2 Threat Model

This section describes threats that can be undertaken against or in the context of IF-FTNC. These threats are broken down into several categories based on which party is mounting the attack.

4.2.1 Network Attacks

Many different network attacks can be mounted against IF-FTNC components. However, the effects of these attacks are limited because the network is not trusted. Nonetheless, here is a list of attacks that can be mounted. For the purposes of this list, the phrase “network traffic” should be taken to mean any messages or parts of messages sent among the endpoint, ASD, and RSD. Any of these attacks may be mounted by network elements, by parties who control network elements, and (in many cases) by parties who control network-attached devices.

- Network traffic may be passively monitored, gleaning information from any unencrypted traffic. Because IF-FTNC conveys information about endpoint security (SPI), it is essential to protect the confidentiality of this information and thus thwart this attack lest endpoint SPI fall into the wrong hands and be used to find and attack vulnerable endpoints.
- Even if all traffic is encrypted, some information can be gained by traffic analysis (volume, timing, source and destination addresses, etc.).
- Network traffic may be modified in transit.
- Previously transmitted network traffic may be replayed.
- New network traffic may be added.
- Network traffic may be blocked, perhaps selectively.
- A “Man In The Middle” (MITM) attack may be mounted where an attacker interposes itself between two communicating parties and poses as the other end to either party or impersonates the other end to either or both parties.
- Undesired network traffic may be sent in an effort to overload an architectural component, thus mounting a denial of service attack.

Clearly, strong protection must be provided against network attacks.

4.2.2 Endpoint Attacks

A compromised endpoint or one with a hostile user can mount several attacks. IF-FTNC is designed to withstand these attacks. However, it is still important to understand them and to understand the countermeasures that are employed so that the system designer and administrator can ensure that the countermeasures are properly used. Here is a list of attacks that the endpoint can mount on an IF-FTNC system.

- Measurements of endpoint posture can be modified or falsified. This could result in improper access if the ASD uses these measurements as the basis for generating SPI.
- An endpoint can mount denial of service or resource exhaustion attacks against the ASD or RSD by sending excess traffic. In principal, this attack is no different from the same attack mounted by any party on the network. However, the ASD and RSD must be careful to protect against these attacks even when dealing with endpoints.
- An endpoint can send improperly formatted messages to the ASD or RSD in an attempt to exploit bugs in the ASD or RSD code. If successful, this attack could result in failure or compromise of the ASD or RSD.
- An endpoint can gather information about the ASD’s or RSD’s policies (e.g. what is permitted). Such information can be used to exploit weaknesses in the policies.
- An endpoint can observe, modify, replay, block, or fabricate any message that passes through, originate from, or terminate at the endpoint. This is not a concern for the Roaming Assessment Profile since no messages pass through the endpoint for that profile. However, it is a concern for the Web Assessment Profile since SAML’s Web Browser SSO Profile sends messages through the endpoint.

- An endpoint can fail to properly comply with mandatory requirements contained in this specification or others. Such a failure could come about through something as simple as buggy code or an inattentive programmer. It could also happen because the endpoint is compromised or the user is hostile. We assume this worst case.
- An endpoint can engage in a “Man In The Middle” (MITM) attack between two communicating parties, such as between the ASD and the RSD. This is similar to the network MITM attack but can be more powerful if secure tunnels terminated at the endpoint are used. Such tunnels provide no protection against MITM attacks by the endpoint but do protect against the network MITM attack.

Because it is so easy to compromise an endpoint, IF-FTNC provides strong protection against endpoint attacks.

4.2.3 ASD Attacks

The Asserting Security Domain (ASD) is the source of an endpoint’s SPI. As such, it has a great deal of power. Therefore, it is essential to understand and protect against threats that can be mounted by the ASD.

First, let us consider the case of an unauthorized ASD (one that is not trusted by the RSD). All communications to and from the ASD are authenticated and integrity protected. So long as this authentication and integrity protection is properly designed, implemented, and enforced, an unauthorized ASD is no different from any other hostile network-connected entity. The threats that can be mounted by such a party are described in section 4.2.1. The difference in threats between that section and this one is substantial. Therefore, great care must be taken in providing authentication and integrity protection for communications to and from the ASD.

The following threats can be mounted by a compromised or hostile authorized ASD (one that is trusted by the RSD):

- Provide inaccurate SPI to the RSD or fail to gather and provide all SPI that should be provided. This could cause the RSD to grant access when it should not be granted or vice versa.
- Reveal SPI to unauthorized parties. This could help attackers target vulnerable endpoints. This threat is really not specific to IF-FTNC. It pertains to any party that has access to SPI such as a TNCS, IF-MAP server, or even most IF-MAP clients. Still, it’s worth highlighting here since it is a significant issue.
- Reveal or employ other information about user or endpoint activities. Unless countermeasures are taken, the ASD may know or be able to gather or infer a lot of information: which user in which role accessed which resources from which endpoint at which location and time. This information could be used to cause embarrassment, embezzle funds, or cause other injuries. Again, this attack is not specific to IF-FTNC but pertains to any authentication, authorization, or application server. Still, it is of concern. For more information about this attack, see the Privacy Considerations section.
- Mount denial of service or resource exhaustion attacks against the RSD by sending excess traffic. In principal, this attack is no different from the same attack mounted by any party on the network. However, the RSD must be careful to protect against these attacks even when dealing with an ASD.
- Send improperly formatted messages to the RSD in an attempt to exploit bugs in the RSD. If successful, this attack could result in failure or compromise of the RSD.
- Gather information about the RSD’s policies (e.g. what is permitted). Such information can be used to exploit weaknesses in the policies.

4.2.4 RSD Attacks

The Relying Security Domain (RSD) is fairly dependent on other components in the IF-FTNC system. One might think that the only attack it could mount would be to improperly grant or deny access to the systems and services that it protects. However, there are several other significant attacks enumerated below. For this reason, it is essential to protect against compromised or hostile RSDs.

Before going any further, let us consider the case of an unauthorized RSD (one that is not trusted by the ASD). All communications to and from the RSD are authenticated and integrity protected. So long as this authentication and integrity protection is properly designed, implemented, and enforced, an unauthorized RSD is no different from any other hostile network-connected entity. The threats that can be mounted by such a party are described in section 4.2.1. The difference in threats between that section and this one is substantial. Therefore, great care must be taken in providing authentication and integrity protection for communications to and from the RSD.

The following threats can be mounted by a compromised or hostile authorized RSD (one that is trusted by the ASD):

- Reveal SPI to unauthorized parties. This could help attackers target vulnerable endpoints. This threat is really not specific to IF-FTNC. It pertains to any party that has access to SPI such as an authentication server, TNCS, IF-MAP server, or even most IF-MAP clients. Still, it's worth highlighting here since it is a significant issue.
- Reveal or employ other information about user or endpoint activities. Unless countermeasures are taken, the RSD may know or be able to gather or infer a lot of information: which user in which role accessed which resources from which endpoint at which location and time. This information could be used to cause embarrassment, embezzle funds, or cause other injuries. Again, this attack is not specific to IF-FTNC but pertains to any authentication, authorization, or application server. Still, it is of concern. For more information about this attack, see the Privacy Considerations section.
- Mount denial of service or resource exhaustion attacks against the ASD or the endpoint by sending excess traffic. In principal, this attack is no different from the same attack mounted by any party on the network. However, the ASD and the endpoint must be careful to protect against these attacks even when dealing with an RSD.
- Send improperly formatted messages to the ASD or endpoint in an attempt to exploit bugs in the ASD or endpoint. If successful, this attack could result in failure or compromise of the ASD or endpoint.
- Improperly grant, deny, or restrict access to the service being protected by the RSD.

4.2.5 RADIUS Proxy Attacks

An untrustworthy RADIUS proxy can serve as an active Man In The Middle (MITM), reading and/or modifying most RADIUS attributes without detection. This includes all of the IF-FTNC RADIUS attributes described in section 3.1.5. Several attacks can be mounted:

- Insert, remove, or modify RADIUS attributes related to IF-FTNC. This may cause the RSD to grant an inappropriate level of access to the endpoint. This is not substantially different from the well-established RADIUS proxy attack of changing an Access-Accept to an Access-Reject or vice versa.
- Mount a denial of service attack on the ASD or RSD. Again, this is not a new attack but one that is always a concern when RADIUS proxies are concerned.

Of course, RADIUS proxies are not relevant to the Web Assessment Profile.

4.3 Countermeasures

This section lists the countermeasures present elsewhere in this specification and the threats against which they protect.

4.3.1 Countermeasures Against Network Attacks

Most network attacks are protected against by requiring encryption, authentication, integrity protection, and replay prevention for all communications between the endpoint, the ASD, and the RSD. These countermeasures protect against passive monitoring, modification in transit, replay, new traffic, selective blocking (except for truncation attacks), and MITM. Of course, these countermeasures must be properly designed, implemented, and enforced. If they are disabled or improperly configured, they will be ineffective. Also, authentication must be coupled with careful authorization to determine which parties are actually trusted.

Securing RADIUS traffic via RadSec is recommended but not required by this specification. If RadSec is employed, RADIUS traffic will be encrypted and authenticated between the RSD and the ASD and therefore properly protected against network attacks. If not, IPsec may be employed between RADIUS proxies but this does not protect against RADIUS proxy attacks since it is only hop-by-hop protection not end-to-end from the RSD to the ASD.

Securing SAML requests and responses (including encryption, authentication, integrity protection, and replay prevention) is also recommended but not required by this specification. This recommendation goes beyond the SAML 2.0 specification, which only recommends authentication. Encryption is recommended for IF-FTNC because the data conveyed in the SAML assertions is likely to be especially security-sensitive. Instead of simply conveying a user's identity, an assertion may convey information about the vulnerability of their endpoint. If attackers obtain this information, they can use it to target attacks on vulnerable endpoints.

No explicit protection against traffic analysis is provided in this specification. If desired, a countermeasure can be provided by maintaining a constant volume of sham traffic between all parties.

Protection against denial of service attacks is also out of the scope of this specification. Protection can be provided using conventional methods such as rate limiting.

4.3.2 Countermeasures Against Endpoint Attacks

The first defense against endpoint attacks is a complete distrust of the endpoint. The endpoint is the least protected element in the IF-FTNC system (with the possible exception of the network). A significant percentage of endpoints are currently compromised and this will probably always be the case. Therefore, IF-FTNC is designed so that the endpoint does not need to be trusted. All IF-FTNC components should take a similar approach. For example, RSDs and ASDs should carefully check any data they receive from endpoints since it may be designed to exploit vulnerabilities in their code. ASDs should support hardware measurements using a TPM and external scanning and monitoring technology so that more reliable SPI can be obtained.

The RSD and ASD should avoid sending data through the endpoint. For example, the Roaming Assessment Profile does not depend on endpoint-delivered data at all. In cases where this cannot be avoided, the ASD and RSD should ensure this data is protected against endpoint attacks by using authentication, confidentiality, integrity protection, replay prevention, and freshness detection. For example, the <Response> message sent by an ASD to an RSD during the Web SSO Profile MUST be signed if sent via the HTTP POST binding, as indicated in [12].

No protection is provided against endpoints gathering information about ASD and RSD policies. Performing access control without revealing information about policy is an area of active research.

4.3.3 Countermeasures Against ASD Attacks

The compromise of an authorized ASD is a serious matter. This section describes countermeasures that can be taken against various ASD attacks. However, it should be

emphasized that the compromise of an ASD is no more serious than the compromise of any shared authentication or authorization server, such as an LDAP server that is used for authentication.

Here is a list of countermeasures against ASD attacks. All of these should be employed. In fact, they are mainly good practices for managing any security-sensitive services.

- Prevent unauthorized parties from impersonating an authorized ASD by employing encryption, authentication, authorization, integrity protection, and replay prevention for all communications with the ASD. This ensures that only authorized ASDs are able to play the ASD role.
- Minimize the chance that an ASD server will be compromised. Harden ASD servers against attack and minimize them to reduce their attack surface. Force ASD servers to go through a regular TNC handshake (preferably with a TPM-based hardware health check) to verify their integrity. Manage ASD servers to minimize vulnerabilities in the underlying platform and in systems upon which the ASD depends. Monitor and limit traffic to and from the ASD with network security measures such as firewalls or intrusion detection systems. Carefully screen and monitor personnel with administrative access to detect problems as soon as possible. Do not use password-based authentication for administrators but instead use non-reusable credentials and multi-factor authentication. Deploy physical security measures to prevent physical attacks on ASDs.
- Plan ahead to reduce the impact of any ASD compromise. If possible, move users to multi-factor authentication and non-reusable credentials. This provides many other benefits as well: protection against endpoint compromise, preventing password reuse, etc. Ensure that devices connected to the ASD are hardened against attacks from the ASD. Improve endpoint security so that the ASD has few vulnerabilities that it can reveal.
- Detect any ASD compromise as quickly as possible. Monitor ASD behavior to detect unusual occurrences (such as a reboot or network traffic to unauthorized or atypical parties). Configure RSDs to log and/or notify administrators when peculiar ASD behavior is detected. RSDs should also check data sent from the ASD carefully to detect malformed data or denial of service attacks. Carefully investigate user reports of improper access (too generous or too strict) to determine the root cause. The investigator should not be the same person who administers the ASD or the servers upon which it depends. If a regular pattern of improper access emerges, conduct experiments and investigations to determine whether the problem is accidental or purposeful. To aid forensic investigation, maintain permanent read-only audit logs of security-relevant information pertaining to the ASD (especially administrative actions) and archived these logs in a safe location. The fact that an ASD often is comprised of several servers (e.g. RADIUS server, TNCS, and SAML Attribute Authority) can aid in detecting compromise since each server can monitor and log the behavior of the others. Furthermore, an ASD will generally be interacting with RSDs from other, independent organizations. This increases the likelihood that odd behavior will be noticed and decreases the likelihood of collusion between administrators.
- Respond promptly and properly to any ASD compromise. Isolate the compromised system and call in a forensic team to analyze its state. Perform a careful analysis of the impact of the compromise. Reissue any reusable credentials that may have been compromised. Investigate any endpoints and users that have received improper access and have terminate the improper access. However, recognize that the endpoints and users may have obtained such access without any knowledge that they were not supposed to receive it. In fact, they may not have realized that they had this access. Check logs to determine whether the ASD may have revealed security-sensitive or privacy-sensitive information to other parties and, if so, what the impact of this disclosure might be.

No protection is provided against ASDs gathering information about RSD policies. Performing access control without revealing information about policy is an area of active research. Furthermore, if an ASD is compromised this will not be the greatest concern.

4.3.4 Countermeasures Against RSD Attacks

As noted above, the compromise of an authorized RSD is less serious than the compromise of an authorized ASD. First, the RSD generally does not have any access to user credentials. Second, its access to SPI is limited to that which the ASD is willing to reveal across security domain boundaries. Generally, the impact is limited to the scope of the RSD's authority: the services to which it can grant access. Still, RSD attacks are serious.

The countermeasures against ASD attacks that are described in section 4.3.3 are largely relevant to RSD attacks also. Prevent unauthorized parties from impersonating an RSD. Minimize the chance that an RSD will be compromised. Reduce the impact of an RSD. Detect RSD compromise as soon as possible. And respond promptly to any RSD compromise. ASDs, endpoints, and service providers can all help in detecting RSD compromise.

One countermeasure that is particular to RSD attacks is to minimize the amount of SPI that the ASD reveals to the RSD. There are many other good reasons to minimize this, especially privacy and security considerations related to the reduced trust that the endpoint and ASD presumably have in the RSD since it is not part of their domain. Reducing the impact of RSD compromise is just one more reason to minimize SPI.

4.3.5 Countermeasures Against RADIUS Proxy Attacks

The best countermeasure against RADIUS proxy attacks is to use RadSec instead of RADIUS proxies. With RadSec, a secure TLS tunnel is established directly from the RSD to the ASD so no RADIUS proxies are used. Therefore, this specification recommends that RadSec be employed.

If RadSec cannot be employed, carefully manage RADIUS proxies to reduce the probability that one will be compromised. In addition, ASDs and RSD should authenticate and authorize SAML requests and responses.

5 Privacy Considerations

IF-FTNC involves transferring Security Posture Information (SPI) from one security domain to another. This SPI may include privacy-sensitive data and personally identifiable information such as a user identifier or IP address. IF-FTNC and implementations of IF-FTNC are obligated by ethics, policies, regulations, and laws to protect the privacy of such information in an appropriate manner. This section describes how such protection should be provided.

5.1 Local Configuration of Privacy Policy

Rules, regulations, laws, and user expectations vary with respect to privacy. In some environments, personal data cannot be transferred or even gathered. In other environments (generally with user consent), security and other considerations drive a need for personal data collection, storage, and transfer.

IF-FTNC cannot encompass the many changing rules and expectations for managing personal data. This is the responsibility of IF-FTNC system administrators, in consultation with management and users. Some organizations employ a central privacy management system that already provides some or all of the features described below, rendering them redundant. Other organizations have a privacy policy that does not require these features. Therefore, IF-FTNC implementations MUST permit local administrative configuration of privacy policy so that administrators can configure their systems to enforce the desired policy. All of the features described in this section should be subject to local administrative configuration.

5.2 Data Gathering and Storage

IF-FTNC implementations MUST allow system administrators to configure which personal data are gathered and how long they are stored. IF-FTNC implementations MUST provide opt-in and opt-out features that allow users to review the system's privacy policy (including a list of data to be gathered and a description of how this data will be used) and to indicate their consent to this policy or lack of consent.

The administrator should be permitted to enable or disable the opt-in and/or opt-out features to match local policy. These features may not be appropriate for a specific environment where users have already agreed to data gathering and storage. Furthermore, some environments may use an opt-in policy (where data is not gathered unless consent has been given) while others may use an opt-out policy (where data is gathered unless the user specifically requests otherwise).

5.3 Data Transfer

IF-FTNC implementations MUST allow system administrators to configure the parties to whom personal data may be transferred. IF-FTNC implementations MUST allow users to review this list of parties and receive notice of changes to the list.

6 References

[1] Trusted Computing Group, *TNC Architecture for Interoperability*, Specification Version 1.0, Revision 3, April 2005.

[2] Trusted Computing Group, *TNC IF-MAP Binding for SOAP*, Specification Version 1.0, Revision 25, April 2008.

[3] Trusted Computing Group, *TNC IF-TNCCS: TLV Binding*, Specification Version 2.0, Revision 10, January 2008.

[4] Trusted Computing Group, *TNC IF-M: TLV Binding Specification*, Specification Version 1.0, Revision 30, February 2008.

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

[6] B. Aboba et al, "Extensible Authentication Protocol (EAP)", Internet Engineering Task Force RFC 3748, June 2004.

[7] Trusted Computing Group, *TNC IF-T: Protocol Bindings for Tunneled EAP Methods*, Specification Version 1.1, Revision 10, May 2007.

[8] B. Aboba et al, "The Network Access Identifier", Internet Engineering Task Force RFC 4282, December 2005.

[9] S. Winter et al, "RadSec Version 2 - A Secure and Reliable Transport for the RADIUS Protocol", Internet Engineering Task Force Internet Draft draft-winter-RadSec-01, February 2008.

[10] C. Rigney et al, "Remote Authentication Dial In User Service (RADIUS)", Internet Engineering Task Force RFC 2865, June 2000.

[11] S. Cantor et al, "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS SSTC, March 2005.

[12] S. Cantor et al, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0". OASIS SSTC, March 2005.

[13] S. Cantor et al, "Shibboleth Architecture: Protocols and Profiles", Internet2 MACE, September 2005.

[14] T. Scavo et al, "Metadata Extension for SAML V2.0 and V1.1 Query Requesters", OASIS SSTC, November 2007.

[15] S. Cantor et al, "Metadata for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS SSTC, March 2005.

[16] S. Cantor et al, "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS SSTC, March 2005.

[17] E. Maler et al, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V1.1", OASIS SSTC, September 2003.

[18] E. Maler et al, "Bindings and Profiles for the OASIS Security Assertion Markup Language (SAML) V1.1", OASIS SSTC, September 2003.

[19] P. Deutsch, "DEFLATE Compressed Data Format Specification version 1.3", Internet Engineering Task Force RFC 1951, May 1996.

[20] N. Freed et al, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", Internet Engineering Task Force RFC 2045, November 1996.

[21] T. Berners-Lee et al, "Uniform Resource Locators (URL)", Internet Engineering Task Force RFC 1951, December 1994.

[22] Trusted Computing Group, *TNC IF-PEP: Protocol Bindings for RADIUS*, Specification Version 1.1, Revision 0.7, February 2007.

[23] M. Mealling, “A URN Namespace of Object Identifiers”, Internet Engineering Task Force RFC 3061, February 2001.