

# **TCG Trusted Network Connect**

## **TNC IF-MAP Binding for SOAP**

**Specification Version 2.0**

**Revision 47**

**November 11, 2011**

**Published**

**Contact:**

[admin@trustedcomputinggroup.org](mailto:admin@trustedcomputinggroup.org)

### **TCG Published**

Copyright © TCG 2011

**TCG**

Copyright © 2011 Trusted Computing Group, Incorporated.

**Disclaimer**

THIS SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Without limitation, TCG disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification and to the implementation of this specification, and TCG disclaims all liability for cost of procurement of substitute goods or services, lost profits, loss of use, loss of data or any incidental, consequential, direct, indirect, or special damages, whether under contract, tort, warranty or otherwise, arising in any way out of use or reliance upon this specification or any information herein.

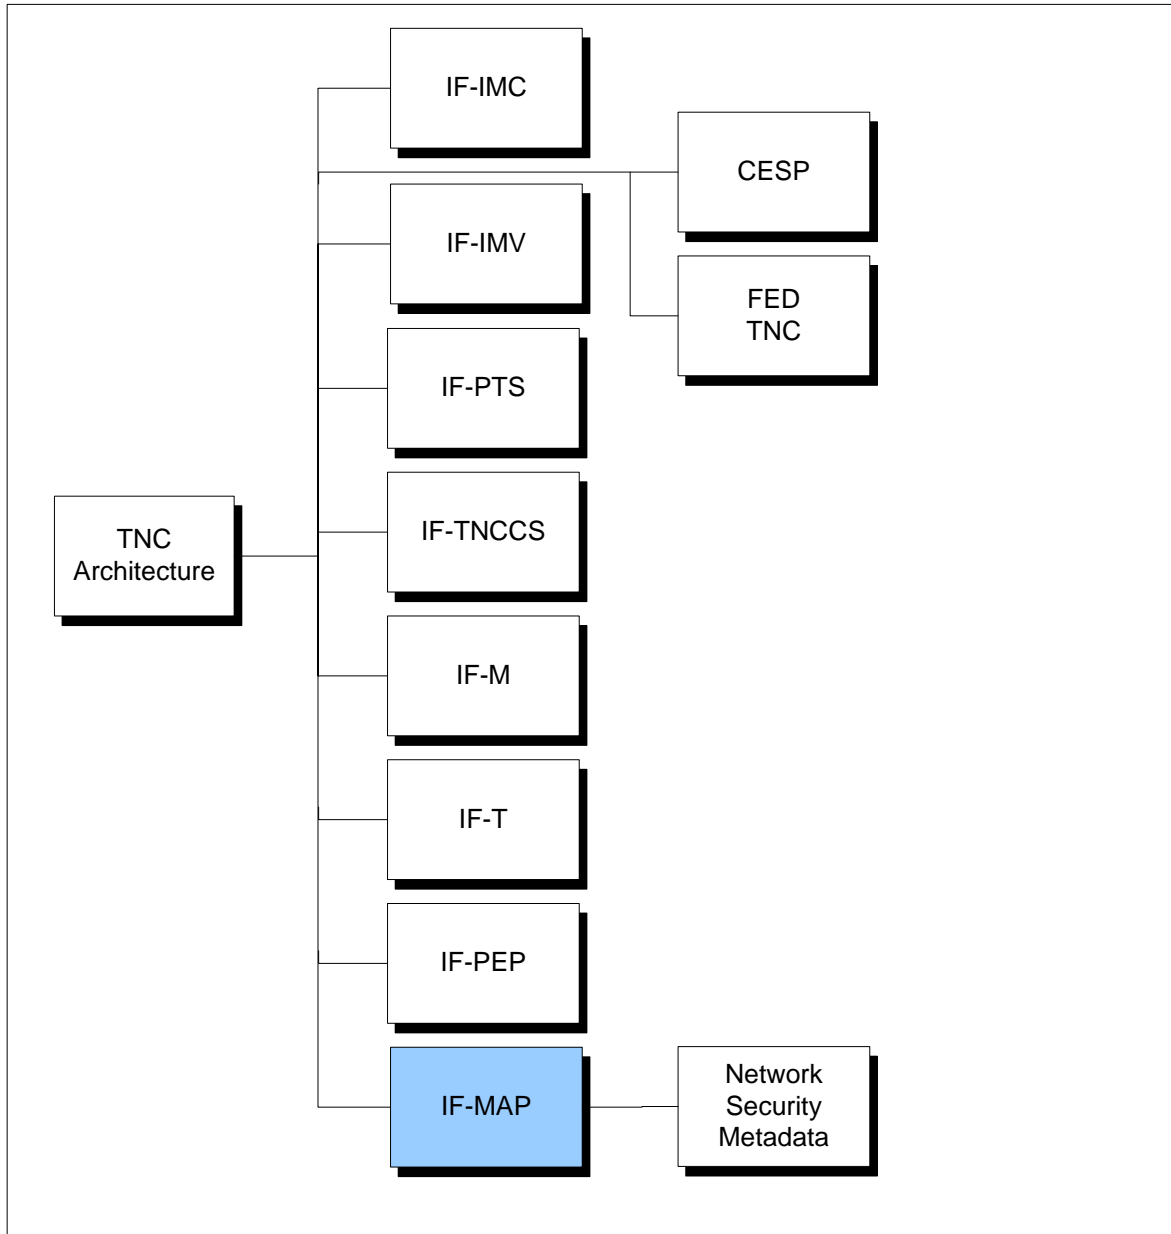
No license, express or implied, by estoppel or otherwise, to any TCG or TCG member intellectual property rights is granted herein.

**Except that a license is hereby granted by TCG to copy and reproduce this specification for internal use only.**

Contact the Trusted Computing Group at [www.trustedcomputinggroup.org](http://www.trustedcomputinggroup.org) for information on specification licensing through membership agreements.

Any marks and brands contained herein are the property of their respective owners.

## IWG TNC Document Roadmap



## Acknowledgements

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

Scott Kelly	<b>Aruba Networks</b>
Amit Agarwal	<b>Avaya</b>
Mahalingam Mani	<b>Avaya</b>
Craig Dupler	<b>Boeing Corporation</b>
David Mattes	<b>Boeing Corporation</b>
Steven Venema (Co-Editor)	<b>Boeing Corporation</b>
Eric Byres (Invited Expert)	<b>Byres Security</b>
Mark Townsend	<b>Enterasys</b>
Michael McDaniels	<b>Extreme Networks</b>
Ingo Bente	<b>Fachhochschule Hannover</b>
Josef von Helden	<b>Fachhochschule Hannover</b>
Arne Welzel	<b>Fachhochschule Hannover</b>
Hidenobu Ito	<b>Fujitsu Limited</b>
Seigo Kotani	<b>Fujitsu Limited</b>
Houcheng Lee	<b>Fujitsu Limited</b>
Sung Lee	<b>Fujitsu Limited</b>
Graeme Proudler	<b>Hewlett-Packard</b>
Mauricio Sanchez	<b>Hewlett-Packard</b>
Andreas Steffen	<b>Hochschule für Technik Rapperswil</b>
Han Yin	<b>Huawei Technologies</b>
Diana Arroyo	<b>IBM</b>
Guha Prasad Venkataraman	<b>IBM</b>
Sean Convery	<b>Identity Engines</b>
Chris Hessing	<b>Identity Engines</b>
Morteza Ansari	<b>Infoblox</b>
Stuart Bailey (Co-Editor)	<b>Infoblox</b>
Andrew Benton	<b>Infoblox</b>
Navin Boddu	<b>Infoblox</b>
Tom Clark	<b>Infoblox</b>
Peter Lee	<b>Infoblox</b>
Rod Murchison	<b>Infoblox</b>
Ivan Pulleyn	<b>Infoblox</b>
David Vigier	<b>Infoblox</b>
Rena Yang	<b>Infoblox</b>
Ravi Sahita	<b>Intel Corporation</b>
Ned Smith	<b>Intel Corporation</b>
Josh Howlett	<b>JANET(UK)</b>
Yan Avlasov	<b>Juniper Networks</b>
Roger Chickering (Co-Editor)	<b>Juniper Networks</b>
Charles Goldberg	<b>Juniper Networks</b>
Steve Hanna (TNC co-chair)	<b>Juniper Networks</b>
Clifford Kahn	<b>Juniper Networks</b>
PJ Kirner	<b>Juniper Networks</b>
Lisa Lorenzin (Co-Editor)	<b>Juniper Networks</b>
John Jerrim	<b>Lancope</b>
Mark Labbancz	<b>Lumeta</b>
Matt Webster	<b>Lumeta</b>

Bill Nemec	<b>Lumeta</b>
Eric Fitzgerald	<b>Microsoft</b>
Ryan Hurst	<b>Microsoft</b>
Sandilya Garimella	<b>Motorola</b>
Meenakshi Kaushik	<b>Nortel</b>
Paul Sangster (TNC co-chair)	<b>Symantec Corporation</b>
Ted Fornoles	<b>Trapeze Networks</b>
Matthew Gast	<b>Trapeze Networks</b>
Tim McCarthy	<b>Trapeze Networks</b>
Jeffrey Peden	<b>Trapeze Networks</b>
Brian Wangerian	<b>Trapeze Networks</b>
Brad Upson	<b>UNH InterOperability Lab</b>
Mike Boyle	<b>US National Security Agency</b>
Lauren Giroux	<b>US National Security Agency</b>
Chris Salter	<b>US National Security Agency</b>
Gloria Serrao	<b>US National Security Agency</b>
Thomas Hardjono	<b>Wave Systems</b>
Greg Kazmierczak	<b>Wave Systems</b>

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

David Mattes	Boeing
Steve Venema	Boeing
Peter Lee	Infoblox
David Vigier	Infoblox
Steve Hanna	Juniper
Clifford Kahn	Juniper
Lisa Lorenzin	Juniper
Mark Labbancz	Lumeta
Matt Webster	Lumeta
Matthew Gast	Trapeze
Jeffrey Peden	Trapeze
Gloria Serrao	US National Security Agency
Mike Boyle	US National Security Agency

## Table of Contents

<b>1</b>	<b>Introduction .....</b>	<b>8</b>
1.1	Scope and Audience .....	8
1.2	Keywords.....	8
1.3	Overview of Changes from Version 1.1 .....	8
1.3.1	Separation into Two Specifications .....	8
1.3.2	Changes to Identifiers, Operations, and Protocol .....	8
1.4	Overview of Changes from Version 2.0r36 .....	9
<b>2</b>	<b>Background .....</b>	<b>10</b>
2.1	MAP Servers and Clients .....	10
2.2	Operational Scope of IF-MAP .....	10
2.3	Supported Use Cases .....	10
2.4	Requirements .....	10
2.5	IF-MAP in TNC Architecture.....	11
2.6	Data Model.....	12
2.6.1	Identifier.....	13
2.6.2	Link.....	13
2.6.3	Metadata .....	14
<b>3</b>	<b>IF-MAP Interface.....</b>	<b>15</b>
3.1	String Encoding .....	15
3.2	Identifiers.....	15
3.2.1	access-request.....	15
3.2.2	device .....	15
3.2.3	identity .....	16
3.2.4	ip-address.....	18
3.2.5	mac-address.....	18
3.3	Metadata .....	19
3.3.1	ifmap-publisher-id.....	19
3.3.2	ifmap-timestamp.....	20
3.3.3	ifmap-cardinality .....	20
3.3.4	Extension by Adding Attributes .....	21
3.3.5	Lifetime of Metadata.....	21
3.3.6	Metadata Size .....	21
3.3.7	IF-MAP Standard Metadata types.....	21
3.4	Filters.....	21
3.5	XML Validation .....	23
3.6	Responses .....	23
3.6.1	Error Codes in responses .....	24
3.6.2	Error Strings in responses.....	25
3.6.3	Logging of Errors in Responses.....	25
3.7	Requests .....	25
3.7.1	publish .....	26
3.7.2	search.....	28
3.7.3	Search Results .....	32
3.7.4	subscribe .....	33
3.7.5	poll.....	35
3.7.6	purgePublisher .....	39
3.8	Schema Versioning .....	39
3.9	Vendor-specific Metadata .....	40
3.10	Atomicity.....	40
<b>4</b>	<b>SOAP Binding and Session Management.....</b>	<b>41</b>
4.1	Client-Server Communication Model .....	41
4.1.1	Sessions.....	41
4.2	SOAP Transport.....	42
4.3	Session ID.....	43

4.4	Session Renewal .....	45
4.5	ARC error handling .....	45
4.6	WSDL and Example Code .....	46
<b>5</b>	<b>Recommendations for Backward Compatibility .....</b>	<b>47</b>
5.1	Translating identifiers .....	47
5.2	Translating Metadata .....	47
5.3	Searches and Subscriptions .....	47
5.4	notify .....	48
5.5	1.1 Clients and IF-MAP Metadata for Network Security .....	48
5.6	Vendor-specific Metadata .....	48
5.7	2.0 Clients and IF-MAP 1.1 Metadata .....	48
<b>6</b>	<b>Security Considerations .....</b>	<b>49</b>
6.1	Trust Model .....	49
6.1.1	MAP Server .....	49
6.1.2	MAP Clients .....	49
6.1.3	Network .....	49
6.2	Threat Model .....	50
6.2.1	Network Attacks .....	50
6.2.2	MAP Clients .....	50
6.2.3	MAP Servers .....	51
6.3	Countermeasures .....	52
6.3.1	Securing the IF-MAP Protocol .....	52
6.3.2	Securing MAP Clients .....	52
6.3.3	Securing MAP Servers .....	53
6.4	Summary .....	53
<b>7</b>	<b>Privacy Considerations .....</b>	<b>55</b>
7.1	identity Identifier .....	55
7.2	mac-address Identifier .....	55
7.3	ip-address Identifier .....	55
<b>8</b>	<b>References .....</b>	<b>56</b>
<b>9</b>	<b>Basic Example .....</b>	<b>58</b>
9.1	Webcam Conferencing .....	58
<b>10</b>	<b>IF-MAP Schema .....</b>	<b>62</b>
10.1	Identifier Types, Requests and Responses .....	62

# 1 Introduction

## 1.1 Scope and Audience

The Trusted Network Connect Working Group (TNC-WG) has defined an open solution architecture that enables network operators to control access to a network. Part of the TNC architecture is IF-MAP, a standard interface between the Metadata Access Point and other elements of the TNC architecture. This document defines and specifies IF-MAP.

Architects, designers, developers and technologists who wish to implement, use, or understand IF-MAP should read this document carefully. Before reading this document any further, the reader should review and understand the TNC architecture as described in [1].

## 1.2 Keywords

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

## 1.3 Overview of Changes from Version 1.1

IF-MAP 2.0 (referring to this document) includes many changes, as outlined in this section. Some of these changes have backwards compatibility implications, as described in section 5.

### 1.3.1 Separation into Two Specifications

The concepts in version 1.1 are now contained in two specifications: 1) "IF-MAP Binding for SOAP Version 2.0" which specifies the identifiers, operations, and all aspects of the base publish/search/subscribe protocol and 2) "IF-MAP Metadata for Network Security Version 1.0" which specifies the use cases, metadata and semantics for IF-MAP enabled network security related systems to interoperate in a standard way to provide comprehensive network security. Progress toward two goals are met by separating the base protocol from the use case specific standard metadata. First, changes to either the base protocol or use case specific metadata have a high likelihood of independence which enables more responsive standards in both categories. Second, an independent base protocol reflects the anticipation of a wide variety of use cases and metadata to be developed and satisfied by the publish/search/subscribe primitives provided in IF-MAP 2.0 and any future versions.

### 1.3.2 Changes to Identifiers, Operations, and Protocol

- Schema is WSI Compliant to be better suited to existing SOAP tool chains
- New namespace for operations and identifiers
- New HIP/HIT subtype of Identity
- New "notify" operation for non-persistent updates
- Application level sessions between clients and servers
- Terminal matches for search
- Data lifetime may optionally be constrained by session lifetime
- Poll results consist of changes since the last poll result rather than full search results.
- Attributes on metadata are extensible to enable the base schema to be revised without affecting metadata schema
- Clarified required metadata attributes



- Clarified subscription updates and poll queue behavior
- Changed prefix for vendor-defined values from DNS name to SMI PEN
- Removed trusted-platform-module identity type

## 1.4 Overview of Changes from Version 2.0r36

IF-MAP 2.0r47 (this document) contains the following changes from the initial published version, IF-MAP 2.0r36:

- In section 3.3.3, clarified handling of cardinality of vendor-specific metadata on individual or separate identifiers. Added norms about the validity and effect of the ifmap-cardinality attribute.
- In section 3.3.6, clarified how to calculate maximum length of metadata.
- In section 3.7.2.5, clarified how to calculate maximum size of search results.
- In section 3.7.5, clarified independence of metadata from encoding in processing pollResults; clarified the desired behavior when a client sends a pollRequest and has no pending subscriptions. Noted that section 3.10 affects handling of pollResults.
- In section 3.7.6, clarified what a MAP Server must do upon receipt of a purgePublisher request.
- In section 3.10, clarified what to return when multiple publish / delete operations occur in a single publish request vs. in separate publish requests.
- In section 4.1.1, clarified what happens when a client with an existing SSRC opens a new connection.
- Created section 5.7 to clarify how to handle device identifiers with an aik-name in a mixed IF-MAP 1.1 and IF-MAP 2.0 environment.

## 2 Background

### 2.1 MAP Servers and Clients

A MAP (Metadata Access Point) is a TNC element providing the MAP Server function, which stores state information about devices, users, and flows in a network. This information includes registered address bindings, authentication status, endpoint policy compliance status, endpoint behavior, and authorization status. For example, the user joe has authenticated through an 802.1X switch using an endpoint with MAC address 00:11:22:33:44:55. An endpoint assessment has revealed that the endpoint has the proper anti-virus software installed and enabled, as required by policy. The endpoint has subsequently been assigned IPv4 address 192.0.2.4, and has been engaged in instant messaging (IM) traffic with the corporate IM server 192.0.2.69.

MAP Clients may publish information to a MAP, search the information in a MAP, and subscribe to notifications from a MAP when information stored in the server changes. A single MAP Client may publish, search, and subscribe; however, many MAP Clients are solely a publisher or a subscriber. For example, a TNC Server publishes information about the policy compliance of an endpoint and a Flow Controller (such as a layer 3 firewall) subscribes to notification of changes to this information. When the TNC Server detects that the endpoint is no longer policy compliant, the TNC Server updates the information in the MAP Server. The MAP Server notifies the Flow Controller. The Flow Controller blocks access to the network by the newly non-compliant device. In this example, both the TNC Server and the Flow Controller are MAP Clients. The TNC Server is a publisher. The Flow Controller is a subscriber.

IF-MAP is the protocol used for communication between MAP Clients and Servers.

### 2.2 Operational Scope of IF-MAP

A MAP allows elements in the TNC architecture to share and correlate stateful runtime metadata. This data *augments* other sources of data for security related decision-making. Searches and subscriptions using IF-MAP return data that *nominally* reflects recent metadata values and relationships as reported by MAP Clients. A MAP server cannot guarantee that the information it dispenses is accurate. MAP clients control the accuracy of the data. Validation of proper MAP Client behavior for a specific use case (e.g. correctly reporting a de-provisioning operation via IF-MAP) is out of the scope of this specification. No global transactional guarantees are provided for IF-MAP 2.0 (e.g. ordering of publish requests). IF-MAP does not provide historical information.

### 2.3 Supported Use Cases

Use cases that this version of IF-MAP supports:

- A MAP Client, such as a PDP, Sensor, or PEP, publishes metadata to a MAP Server.
- A MAP Client, such as a PDP or Flow Controller, searches a MAP Server for metadata associated with an endpoint.
- A MAP Client, such as a PDP or Flow Controller, subscribes to notifications from a MAP Server about changes in metadata for an endpoint.

### 2.4 Requirements

The following are the requirements that IF-MAP 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.

#### 1. Meets the needs of the TNC architecture

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

Specific requirements include:

- IF-MAP must support both synchronous response and asynchronous notification queries.
- IF-MAP must support frequent updates to metadata. While directory protocols like LDAP are optimized for infrequent updates and frequent reads, IF-MAP is required to have strong support for frequent updates and reads.

## 2. Secure

- Communication between MAP Clients and Servers MUST be authenticated and integrity-protected against unauthorized modifications en route.
- Communication between MAP Clients and Servers MUST provide confidentiality against unauthorized disclosure.
- Communication between MAP Clients and Servers MUST NOT be susceptible to replay attacks.

## 3. Extensible

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

## 4. Easy to use and implement

IF-MAP should be easy for MAP Client and Server 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.

## 5. Unambiguous

There should be clarity and lack of ambiguity for identification of specific entities (ARs, users, etc.) for which metadata exists and which are interacting with the MAP Server. For example users, devices, ARs and all other instances of TNC elements should be uniquely identifiable within an IF-MAP implementation.

## 6. Scalable and Efficient

IF-MAP is intended to be used for interfacing thousands of networking devices to an IF-MAP service in a large organization in which thousands of updates occur per second. It is expected that within a few years, MAP Servers will be required to serve millions of networking devices. Therefore, the IF-MAP specification should not place implementation burdens on clients or servers that would prevent scaling to meet the demands of a large organization.

## 2.5 IF-MAP in TNC Architecture

As described in the TNC Architecture specification [1], the TNC architecture includes:

- A MAP Client role which includes Flow Controllers and Sensors. A Flow Controller is an enforcer that is not required to communicate directly with a PDP. A Sensor shares metadata about the network.
- A Metadata Access Point role which coordinates metadata exchange.

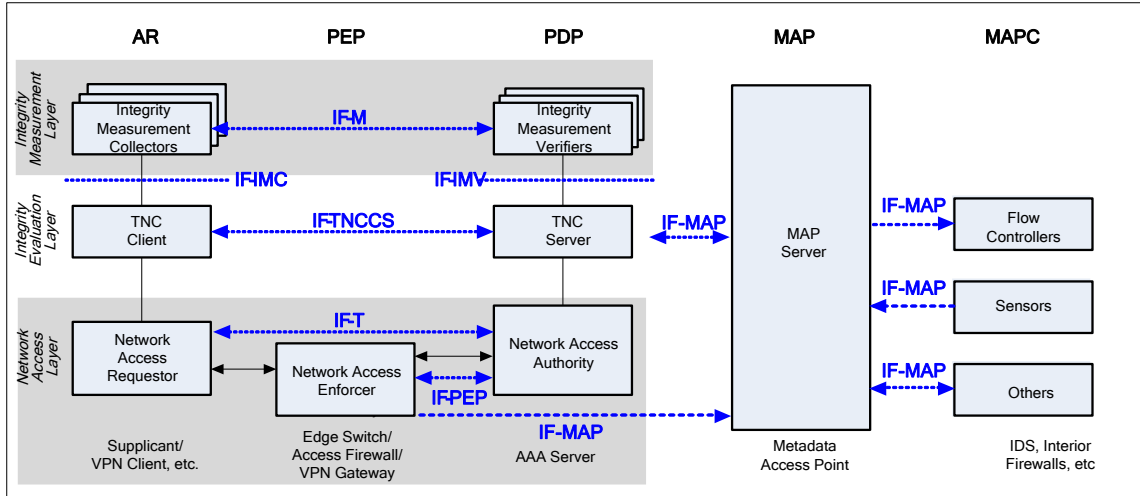


Figure 1

The Metadata Access Point role of the TNC architecture is performed by an element that has a MAP Server function. All network elements that use IF-MAP to access a MAP Server (PEPs, Flow Controllers, Sensors, PDPs, and any other elements) are MAP Clients. All MAP Clients must be authenticated and authorized to use the MAP Server.

## 2.6 Data Model

In IF-MAP there are two types of data: identifier and metadata. There is one type of relationship: link. All IF-MAP operations and data types are represented as XML documents.




Figure 2

In Figure 2, identifiers are represented by ovals, metadata is represented by rectangles, and links are represented by lines connecting identifiers.

### 2.6.1 Identifier

IF-MAP specifies an **identifier** as a single, globally unique value within a space of values described by an identifier type specified in the IF-MAP XML schema. An identifier or set of identifiers is required for all IF-MAP metadata operations. The schemas for all types of identifiers are defined by this document in order to provide a unified namespace for metadata.

For example, the IPAddrType identifier type schema element defines an identifier space consisting of all possible IP addresses.

All identifiers in an identifier space implicitly exist at all times and are always legal for any operation within the limits of a client's authorization policy. In other words, an identifier does not need to be explicitly created before being used (and in fact there are no operations in IF-MAP for creating and destroying identifiers).

The publish operation associates metadata with identifiers or links between identifiers. The search and subscribe operations use an identifier as the starting point for a query. See section 3.7.

### 2.6.2 Link

IF-MAP specifies a **link** as an unnamed, bi-directional binding relationship between two **identifiers**. For example, a DHCP server might create a link between a mac-address identifier and an ip-address identifier.

### 2.6.3 Metadata

In IF-MAP, **metadata** is represented as typed values which are well described by schema. Each instance of metadata in a MAP Server is associated with a particular **identifier** or **link**. There are two types of metadata:

- Standard Metadata
- Vendor-specific Metadata

For purposes of extensibility, the schema for standard metadata is defined in supporting specifications, such as TNC IF-MAP Metadata for Network Security[15].

Vendor-specific metadata is used to supplement the use cases defined by the IF-MAP Metadata specifications. All MAP Clients and Servers **MUST** support both standard and vendor-specific metadata. A MAP Client **MUST** ignore vendor-specific metadata that it does not understand. A MAP Server **MUST** allow storage and retrieval of vendor-specific metadata. All MAP Server implementations **MAY** support XML Schema validation as described by [3] and [4]. All MAP Clients **MAY** likewise validate documents using the same mechanisms and **MUST NOT** send metadata that does not comply with the relevant schema. No provisions are made in this specification for uploading XML Schemas.

## 3 IF-MAP Interface

### 3.1 String Encoding

MAP Clients and MAP Servers MUST exchange string fields in UTF-8.

### 3.2 Identifiers

Identifiers provide a unified namespace that can be used to refer to specific metadata items. There are several types of identifiers. All MAP Server and Client implementations MUST support the entire set of identifiers.

There are no explicit limits on the size of an identifier. MAP Clients and Servers MUST support identifiers up to 1000 bytes in length, and SHOULD support longer identifiers. If a MAP Server receives an identifier from a client that exceeds its maximum identifier length, the MAP Server SHOULD respond with an IdentifierTooLong error (see section 3.6.1). If a MAP Client receives an identifier from a server that exceeds its maximum identifier length, the MAP Client SHOULD treat the operation as having failed and log an administrator-viewable message.

#### 3.2.1 access-request

An access-request identifier represents a request for access to a network by a logical endpoint. Multiple access-request identifiers for the same endpoint may be stored in the MAP database, such as when a multi-homed endpoint requests access to multiple networks.

IF-MAP specifies an AccessRequestType consisting of administrative domain string and name string. A MAP Server MUST process two **access-request** identifiers as equivalent if and only if ALL corresponding fields in the two **access-request** identifiers are equivalent.

When utilizing the AccessRequestType specified in IF-MAP, a MAP Client MUST NOT specify an administrative domain. This field is retained from previous versions of IF-MAP for backward compatibility reasons. A MAP Server MUST process the administrative domain field as CASE SENSITIVE. A MAP Server MUST process as equivalent (i.e. belonging to the same administrative domain) administrative domains which are specified as an empty string or unspecified.

MAP Clients MUST choose the value of the name attribute in such a way that the odds of having two logically different access-request identifiers with the same name are negligible. The following strategies all satisfy this requirement:

- The name attribute takes the form "ifmap-publisher-id:UID" where ifmap-publisher-id refers to a specific MAP Client and UID may be a simple ordinal value
- The name attribute is a UUID as described in RFC 4122 [16]
- The name attribute is based on a cryptographically strong random number of at least 128 bits

```
<xsd:complexType name="AccessRequestType">
  <xsd:attribute name="administrative-domain" type="xsd:string"/>
  <xsd:attribute name="name" type="xsd:string" use="required"/>
</xsd:complexType>
```

#### 3.2.2 device

A device identifier represents a physical or virtual asset which is attempting to gain entry to a network or a PDP or other authenticating element.

IF-MAP specifies a *DeviceType* consisting of either an aik-name string or a name string. Because aik-names are not guaranteed to be globally unique, that field is deprecated. MAP Clients **MUST NOT** create a device identifier with an aik-name.

A MAP Server **MUST** process two **device** identifiers as equivalent if and only if the corresponding names in the two **device** identifiers are equivalent in both type (i.e. name or aik-name) and value.

The name attribute **MUST** satisfy the same uniqueness requirements as the name attribute of an access-request identifier.

For multiple virtual devices operating on the same physical device, each virtual device **SHOULD** have its own device identifier in IF-MAP.

```
<xsd:complexType name="DeviceType">
  <xsd:choice>
    <xsd:element name="aik-name" type="xsd:string"/>
    <xsd:element name="name" type="xsd:string"/>
  </xsd:choice>
</xsd:complexType>
```

### 3.2.3 identity

An identity identifier represents an end-user, device, application or other logical or physical entity. A single organization may have several administrative domains which provision identities. Identities may take several different syntactic forms.

IF-MAP specifies an *IdentityType* consisting of administrative domain string, name string, type enumeration, and other-type-definition string. A MAP Server **MUST** process two identities as equivalent if and only if **ALL** corresponding fields in the two identities are equivalent.

When utilizing the *IdentityType* specified in IF-MAP, a MAP Client **MAY** specify an administrative domain. A MAP Server **MUST** process the administrative domain field as **CASE SENSITIVE**. It is up to MAP Clients to perform canonicalization of the administrative domain field if it is derived from a case insensitive external source. A MAP Server **MUST** process as equivalent (i.e. belonging to the same administrative domain) administrative domains which are specified as an empty string or unspecified.

A MAP Client **MUST** specify a name field consisting of a non-empty string. A MAP Client **MUST** specify a type in order to provide a hint to MAP Clients how to parse the name field. A MAP Server **MUST** process the name field according to the syntax specified by the type enumeration. A MAP Server **MUST** process two syntactically equal name fields as distinct if they are associated with distinct types. For example an identity with an unspecified administrative domain, type=username, and name="foo.bar" **MUST** be considered distinct from an identity with an unspecified administrative domain, type=dns-name name="foo.bar" since the types are distinct.

Furthermore, the uniqueness of the name field is dependent on the syntax of the type enumeration. Because the exact syntax and equivalence (i.e. matching) rules for each of the identity types may vary widely from organization to organization, the choice is left to the MAP Server implementation. MAP Servers **SHOULD** provide both pre-configured and configurable options for both syntax and equivalence rules of all the various identity types.

The following identity types are supported:

Identity Type	Usage
aik-name	TCG AIK Name [20]
distinguished-name	An X.500 Distinguished Name [21]
dns-name	A name from the Domain Name System [22]



email-address	An email address, e.g. sally@example.com [23]
kerberos-principal	A Kerberos unique identity [24]
username	A user's login name
sip-uri	A SIP identifier [25]
tel-uri	A telephone number URI [26]
hip-hit	<p>Host Identity Tag (HIT) from the Host Identity Protocol. This is typically a 128-bit hash of a public key with some of the upper bits masked out for other purposes. [27] HIT MUST be expressed as x:x:x:x:x:x:x, where the 'x's are the lowercase hexadecimal values of the eight 16-bit pieces of the HIT.</p> <p>Examples:</p> <p style="padding-left: 40px;">2001:10:7654:3210:fedc:ba98:7654:3210</p> <p style="padding-left: 40px;">2001:10:0:0:8:800:200c:417a</p> <p>No leading zeros are allowed except that the number 0 is represented by a single 0 character.</p>
other	<p>If a MAP Client specifies identity type as "other", the client MUST specify a non-empty string for the other-type-definition field. The other-type-definition field's value MUST take one of two forms:</p> <ol style="list-style-type: none"> <li>1. "Vendor-ID:Name": A vendor-defined type. Vendor-ID is an SMI Private Enterprise Number [19] owned by the vendor, and Name is the type name</li> <li>2. "Name": A TCG-defined type. A TCG-defined type may be specified in a future version of IF-MAP or in a supplement to IF-MAP.</li> </ol>

```

<xsd:complexType name="IdentityType">
  <xsd:attribute name="administrative-domain" type="xsd:string"/>
  <xsd:attribute name="name" type="xsd:string" use="required"/>
  <xsd:attribute name="type" use="required">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="aik-name"/>
        <xsd:enumeration value="distinguished-name"/>
        <xsd:enumeration value="dns-name"/>
        <xsd:enumeration value="email-address"/>
        <xsd:enumeration value="kerberos-principal"/>
        <xsd:enumeration value="username"/>
        <xsd:enumeration value="sip-uri"/>
        <xsd:enumeration value="tel-uri"/>
        <xsd:enumeration value="hip-hit"/>
        <xsd:enumeration value="other"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
  <xsd:attribute name="other-type-definition" type="xsd:string"/>
</xsd:complexType>

```

### 3.2.4 ip-address

An ip-address identifier represents a single IP address. Due to the fact that many networks are deployed using overlapping IP address spaces, an additional "administrative-domain" qualifier can optionally be added to an ip-address identifier in order to uniquely identify that address. The administrative-domain is a string whose format is organizationally defined.

Both IPv4 and IPv6 address identifiers are supported. These two classes of addresses are differentiated by use of the "type" attribute, which can have either the value "IPv4" or "IPv6". In the absence of a "type" attribute, "IPv4" is assumed.

IP addresses MUST be canonicalized by MAP Clients. MAP Servers MUST reject IP addresses which are not in canonical form.

The canonical form of an IPv4 address is dot-decimal notation (i.e. dotted quad notation) consisting of four dot-separated decimal numbers between 0 and 255. No leading 0s are allowed except that the number 0 is represented by a single 0 character.

IPv4 address => octet "." octet "." octet "." octet

octet => 0..255

For the purposes of this specification, the canonical form of an IPv6 address is x:x:x:x:x:x:x:x, where the 'x's are the lowercase hexadecimal values of the eight 16-bit pieces of the address.

Examples:

2001:db8:7654:3210:fedc:ba98:7654:3210

2001:db8:0:0:8:800:200c:417a

No leading zeros are allowed except that the number 0 is represented by a single 0 character.

A MAP Client MAY specify an administrative domain. A MAP Server MUST process the administrative domain field as CASE SENSITIVE. A MAP Server MUST process as equivalent (i.e. belonging to the same administrative domain) administrative domains which are specified as an empty string or unspecified.

```
<xsd:complexType name="IPAddressType">
  <xsd:attribute name="administrative-domain" type="xsd:string"/>
  <xsd:attribute name="value" type="xsd:string" use="required"/>
  <xsd:attribute name="type">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="IPv4"/>
        <xsd:enumeration value="IPv6"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:complexType>
```

### 3.2.5 mac-address

A mac-address identifier represents an Ethernet MAC address. Due to the prevalence of virtual MAC addresses, an optional administrative-domain qualifier can be added to a MAC address in order to uniquely identify the address. The administrative-domain is a string whose format is organizationally defined.

MAP Clients and Servers MUST specify the MAC address as six groups of two lowercase hexadecimal digits, separated by colons (:) in transmission order, e.g. 01:23:45:67:89:ab.

A MAP Client MAY specify an administrative domain. A MAP Server MUST process the administrative domain field as CASE SENSITIVE. A MAP Server MUST process as equivalent

(i.e. belonging to the same administrative domain) administrative domains which are specified as an empty string or unspecified.

```
<xsd:complexType name="MACAddressType">
  <xsd:attribute name="administrative-domain" type="xsd:string"/>
  <xsd:attribute name="value" type="xsd:string" use="required"/>
</xsd:complexType>
```

### 3.3 Metadata

Metadata takes the form of any XML element that is allowed to have the attributes in the metadataAttributes attributeGroup.

```
<xsd:attributeGroup name="metadataAttributes">
  <xsd:attribute name="ifmap-publisher-id"/>
  <xsd:attribute name="ifmap-timestamp" type="xsd:dateTime"/>
  <xsd:anyAttribute/>
</xsd:attributeGroup>
```

ifmap-publisher-id and ifmap-timestamp are special attributes (known as “operational attributes”) that MAP Servers add to stored metadata. MAP Clients MUST NOT specify operational attributes in publish requests (section 3.7.1). MAP Servers MUST include operational attributes in search responses (section 3.7.2).

The ifmap- prefix in metadata attribute names is reserved for use by this specification and its successors. Metadata elements MUST NOT include attributes that begin with the ifmap- prefix other than attributes specified in this document or its successors. Any unrecognized attributes beginning with the ifmap- prefix MUST be ignored by MAP Clients and MAP Servers.

#### 3.3.1 ifmap-publisher-id

ifmap-publisher-id is a unique value assigned by a MAP Server and associated with a specific MAP Client which performs a publish operation.

The ifmap-publisher-id is assigned by a MAP Server to identify a particular MAP Client. A MAP Server MUST NOT assign the same ifmap-publisher-id to multiple different MAP Clients and MUST consistently use the same ifmap-publisher-id for a particular MAP Client. The ifmap-publisher-id MUST NOT be a function of specific credentials. A MAP Client MUST be able to change credentials (e.g. new cert after expiration) and continue to be assigned the same ifmap-publisher-id. The ifmap-publisher-id MUST NOT be a function of time or of any connection specific information; it must remain consistent across time, and across multiple connections.

The MAP Server MAY include other server specific information when generating an ifmap-publisher-id, which means the same MAP Client might be represented by two different ifmap-publisher-ids on two different MAP Servers.

These rules require a MAP Server to maintain configuration information about each client that might publish data. The configuration information includes the client’s ifmap-publisher-id. When a client connects, the MAP Server determines which client configuration to use to retrieve the ifmap-publisher-id. The criteria a MAP Server uses to map from a particular SOAP connection to a client configuration is implementation dependent.

### 3.3.2 ifmap-timestamp

ifmap-timestamp is the time, as understood by the MAP Server, of the completion of an IF-MAP publish operation. The granularity of ifmap-timestamp is one second. An ifmap-timestamp without a timezone component **MUST** be interpreted as UTC time.

### 3.3.3 ifmap-cardinality

A metadata type may be either singleValue or multiValue. The MAP Server or Client determines whether a metadata type is singleValue or multiValue using the ifmap-cardinality attribute.

All metadata type schemas **MUST** include either the singleValueMetadataAttributes attributeGroup or the multiValueMetadataAttributes attributeGroup (per the box below). Inclusion of one or the other of these attributeGroups ensures that all instances of a particular metadata type will have consistent ifmap-cardinality.

Every metadata item **MUST** include a valid ifmap-cardinality attribute (per the box below). ("Metadata item" means a child of a <metadata> element.) If a metadata item in a request lacks a valid ifmap-cardinality attribute, the MAP Server **MUST** return an InvalidMetadata errorResult. If a metadata item in an update request specifies different if-map cardinality from an instance of the same metadata type *already associated with the identifier or link specified in the update request* (in violation of the ifmap-cardinality requirement in the preceding paragraph), the MAP Server **MUST** return an InvalidMetadata errorResult. The MAP Server is not required to consider cardinality of instances of the same metadata type on identifiers or links other than the one specified in the update request.

Whenever there is an errorResult, the request **MUST** have no other effect. For the balance of this section, we assume that there is no errorResult.

If a metadata item in an update request specifies singleValue, the MAP Server **MUST** replace any previous instance of that metadata type associated with the same identifier or link with the new metadata.

If a metadata item in an update request specifies multiValue, the MAP Server **MUST** append the metadata to any existing metadata on the identifier or link even if it duplicates existing metadata on that identifier or link.

```
<xsd:attributeGroup name="singleValueMetadataAttributes">
  <xsd:attributeGroup ref="metadataAttributes"/>
  <xsd:attribute name="ifmap-cardinality" use="required">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="singleValue"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:attributeGroup>

<xsd:attributeGroup name="multiValueMetadataAttributes">
  <xsd:attributeGroup ref="metadataAttributes"/>
  <xsd:attribute name="ifmap-cardinality" use="required">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="multiValue"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:attributeGroup>
```

Note that `singleValueMetadataAttributes` and `multiValueMetadataAttributes` both include the `metadataAttributes attributeGroup`.

### 3.3.4 Extension by Adding Attributes

Subsequent versions of this specification or other TNC specifications may define new XML attributes for IF-MAP elements. To anticipate this, the schema places `<anyAttribute/>` on every element's declaration.

### 3.3.5 Lifetime of Metadata

Liveness is essential for much metadata: if metadata becomes stale it can create security exposures or other bad effects. Normally, the publishing client is responsible for keeping such metadata current. But if the client's session ends, the MAP server can automatically delete such metadata.

A publisher may specify whether metadata should be deleted at session end by attaching the XML attribute `lifetime` to the update request. The values are `lifetime="session"` and `lifetime="forever"`. "session" is the default.

If an element was published with `lifetime="session"` and the client session ends, either due to inactivity (see section 4.1.1) or at the client's request, the MAP server **MUST** delete the metadata. This deletion **MUST** be completed before the publishing client is allowed to create another session.

The lifetime attribute is meaningful only in update requests. A client **SHOULD** not use it in a notify request and a MAP server **MUST** ignore it if it appears there.

A MAP server **MUST** retain all metadata from an update request until (a) it is explicitly deleted by a client, via a `delete` or `purgePublisher` request, (b) it is deleted at session end, as discussed above, or (c) the server deletes all data that has the same `ifmap-publisher-id`. This may happen due to hardware failure or administrator action, for example.

For uses of this attribute, see the IF-MAP Metadata for Network Security specification [15].

### 3.3.6 Metadata Size

There are no explicit limits on the size of metadata. MAP Clients and Servers **MUST** support metadata at least 100000 bytes in length, and **SHOULD** support longer metadata. Specifically, this is a measure of the size of an individual metadata element within a metadata tag, in wire format. If a MAP Server receives metadata from a MAP Client that exceeds its maximum metadata length, the MAP Server **SHOULD** respond with a `MetadataTooLong` error (see section 3.6.1). If a MAP Client receives metadata from a MAP Server that exceeds its maximum metadata length, the MAP Client **SHOULD** treat the operation as having failed and log an administrator-viewable message.

### 3.3.7 IF-MAP Standard Metadata types

IF-MAP Standard Metadata types are defined in the TNC IF-MAP Metadata specifications listed in section 2.6.3.

## 3.4 Filters

An IF-MAP request may include filters to specify elements to which the request applies. A MAP Server **MUST** consider a filter consisting of the empty string as a request to match nothing.

```
<xsd:simpleType name="FilterType">
  <xsd:restriction base="xsd:string"/>
</xsd:simpleType>
```

All `FilterType` values MUST be constructed with the IF-MAP filter syntax, which is defined by the following grammar:

```

Filter           ::= ElemOrExpr
ElemOrExpr      ::= ElemExpr ( "or" ElemExpr )*
ElemExpr       ::= ( ( QName ) ( "[" PredOrExpr "]" )? )
                | ( "[" PredOrExpr "]" )
PredOrExpr     ::= PredAndExpr ( "or" PredAndExpr )*
PredAndExpr    ::= PrimaryPredExpr ( "and" PrimaryPredExpr )*
PrimaryPredExpr ::= PredExpr | ParenPredExpr
ParenPredExpr  ::= "(" PredOrExpr ")"
PredExpr       ::= Selector Operation Value
Selector       ::= ( ( ElemSelector ) ( "/" AttributeSelector )? )
                | AttributeSelector
ElemSelector   ::= QName ( "/" QName )*
AttributeSelector ::= "@" QName
Operation      ::= "=" | "!=" | "<" | ">" | "<=" | ">="
Value         ::= Literal
Literal       ::= NumericLiteral | StringLiteral
NumericLiteral ::= IntegerLiteral | DecimalLiteral | DoubleLiteral
IntegerLiteral ::= Digits
DecimalLiteral ::= ( "." Digits ) | ( Digits "." [0-9]* )
DoubleLiteral  ::= ( ( "." Digits )
                | ( Digits ( "." [0-9]* )? ) ) [eE] [+-]? Digits
StringLiteral  ::= ( "'" (EscapeQuot | [^"])* "'" )
                | ( '"' (EscapeApos | [^'])* '"' )
EscapeQuot    ::= '\"'
EscapeApos    ::= '\''
Digits       ::= [0-9]+
QName        ::= [http://www.w3.org/TR/REC-xml-names/#NT-QName]

```

A filter consists of a set of `ElemExprs`, which are expressions for selecting elements. `ElemExprs` may be combined using “or” and parentheses. `ElemExprs` may not be combined using “and” since such combinations would typically result in filters that match no elements.

An `ElemExpr` may select elements based on element name, attributes and subelements, or both. To select elements by name, an `ElemExpr` with a `QName` is used. To select elements by attributes and subelements, a set of `PredExprs` (for “predicate expressions”) inside square brackets is used. `PredExprs` may be combined using “and”, “or”, and parentheses.

A `QName` is the name of an XML element or attribute. A `QName` may optionally contain a prefix followed by a “.” character.

A `PredExpr` consists of a `Selector`, an `Operation`, and a `Value`. A `Selector` is a path expression that specifies matching subelements and/or attributes. Subelements are matched against `QNames`, and attributes are matched against `QNames` preceded by “@” prefixes.

Subelements and attributes may be matched against `Values` using the `Operations` “=”, “!=”, “<”, “>”, “<=”, and “>=”. If both the `Selector` and the `Value` are numeric, then numeric comparisons are used. Otherwise, case-sensitive string comparisons are used.

Given the `xmlns` attribute `xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-METADATA/2"`, here are some example IF-MAP filters:

- `meta:role`
- `meta:vlan[administrative-domain="Main Campus" and name>=30 and name<=152]`
- `meta:role[name="sales"] or meta:vlan[name=42]`

- meta:event[confidence > 50 and significance="critical"]
- [@ifmap-publisher-id="my publisher id"]

Given a vendor-specific metadata schema and xmlns attribute xmlns:vend="http://example.com/IFMAP/metadata-schema/1"

- vend:ike-policy[@gateway="1.2.3.4" and phase1/@identity="joe"]

In contexts where an IF-MAP filter is used, the filter is applied separately to each individual piece of metadata. Conceptually, an ElemExpr matches a top-level node in an XML document consisting of a single piece of metadata. The result of applying a filter to a piece of metadata is either "match" or "don't match".

The namespace prefixes in a filter MUST be declared in the XML document containing the filter, and the declaration of the namespace prefix MUST apply to the scope of the filter. A QName in a filter with no namespace prefix refers to the default namespace of the XML document it is contained in.

### 3.5 XML Validation

MAP Servers and Clients have the ability to assert that IF-MAP XML message documents adhere to their specified XML schema. A MAP Server MAY perform XML validation on non-metadata (i.e. operations and identifiers) and on metadata. MAP Servers and Clients SHOULD inform each other about whether transmitted XML has been validated.

If a MAP Server does validate XML, any MAP Clients it communicates with MAY validate XML that they receive.

In order to allow MAP Clients and Servers to communicate about whether or not XML has been validated, IF-MAP requests and responses include the validationAttributes attributeGroup.

```
<xsd:attributeGroup name="validationAttributes">
  <xsd:attribute name="validation" use="optional">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="None"/>
        <xsd:enumeration value="BaseOnly"/>
        <xsd:enumeration value="MetadataOnly"/>
        <xsd:enumeration value="All"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:attributeGroup>
```

### 3.6 Responses

IF-MAP involves the client sending requests to the server, and the server sending responses back to the client. A MAP Server MUST respond to every MAP Client request. If a MAP Client request cannot be processed, the MAP Server MUST respond with an errorResult element and appropriate error code and message or with a lower-level error response (e.g. SOAP). Responses to search and poll requests include additional results which are described in detail in section 3.7.3.

```
<xsd:complexType name="ResponseType">
  <xsd:choice>
    <xsd:element name="errorResult" type="ErrorResultType"/>
    <xsd:element name="pollResult" type="PollResultType"/>
    <xsd:element name="searchResult" type="SearchResultType"/>
  </xsd:choice>
</xsd:complexType>
```

```

<xsd:element name="subscribeReceived"/>
<xsd:element name="publishReceived"/>
<xsd:element name="purgePublisherReceived"/>
<xsd:element name="newSessionResult"
  type="SessionResultType"/>
<xsd:element name="renewSessionResult"/>
<xsd:element name="endSessionResult"/>
</xsd:choice>
<xsd:attributeGroup ref="validationAttributes"/>
</xsd:complexType>

```

### 3.6.1 Error Codes in responses

A response from a MAP Server may indicate an error by including an `errorResult` element:

```

<xsd:complexType name="ErrorResultType">
  <xsd:sequence>
    <xsd:element name="errorString" type="xsd:string"/>
  </xsd:sequence>
  <xsd:attribute name="errorCode" use="required">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="AccessDenied"/>
        <xsd:enumeration value="Failure"/>
        <xsd:enumeration value="InvalidIdentifier"/>
        <xsd:enumeration value="InvalidIdentifierType"/>
        <xsd:enumeration value="IdentifierTooLong"/>
        <xsd:enumeration value="InvalidMetadata"/>
        <xsd:enumeration value="InvalidSchemaVersion"/>
        <xsd:enumeration value="InvalidSessionID"/>
        <xsd:enumeration value="MetadataTooLong"/>
        <xsd:enumeration value="SearchResultsTooBig"/>
        <xsd:enumeration value="PollResultsTooBig"/>
        <xsd:enumeration value="SystemError"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
  <xsd:attribute name="name"/>
</xsd:complexType>

```

The meanings of the error codes are:

Code	Meaning
InvalidIdentifier	Syntax of an identifier in the request is invalid. For example, an ip-address identifier of 1.2.3.A would result in an InvalidIdentifier error.
InvalidIdentifierType	An unrecognized identifier type was specified within an identifier element in the request.
IdentifierTooLong	The MAP Client specified an identifier that exceeds the maximum identifier size supported by the server.
InvalidMetadata	Invalid metadata value in the request. This can happen if the specified metadata does not



	match the schema associated with its type.
InvalidSchemaVersion	The MAP Server was unable to process the metadata specified in the request because the MAP Client and MAP Server do not agree on the schema version.
InvalidSessionID	The MAP Client specified an invalid session-id.
MetadataTooLong	The MAP Client specified metadata that exceeds the maximum identifier size supported by the server.
AccessDenied	The MAP Client issued an unauthorized request. Which requests are allowed from a particular client depends on configuration settings of the MAP Server which are beyond the scope of this document.
SearchResultsTooBig	A search or subscribe command generates too much data.
PollResultsTooBig	Updates that match the client's subscription have generated too much data
SystemError	The MAP Server was unable to successfully process a request because of a system error on the server
Failure	The MAP Server was unable to successfully process a request for an unspecified reason.

The name attribute in an errorResult element is used for poll results to indicate the name of the subscription that caused an error (see section 3.7.5).

### 3.6.2 Error Strings in responses

If present, the errorString in an errorResult element contains a human readable string further describing the error that occurred. MAP Servers SHOULD include a good errorString to aid in debugging.

### 3.6.3 Logging of Errors in Responses

A MAP Client SHOULD log errors so that administrators can trace the causes of IF-MAP errors. Log messages SHOULD include the errorCode as well as the errorString if present.

## 3.7 Requests

MAP Clients interact with MAP Servers by sending requests.

All requests except for newSession MUST include a session-id attribute, which is specified in the IF-MAP schema using the sessionAttributes attributeGroup.

```
<xsd:attributeGroup name="sessionAttributes">
  <xsd:attribute name="session-id" type="xsd:string"
    use="required"/>
</xsd:attributeGroup>
```

### 3.7.1 publish

A publish request may create, modify, or delete metadata associated with one or more identifiers or links. A publish request may also be used to tell the MAP server to notify subscribers about the existence of transitory, non-persistent metadata. Publish requests are for publishing metadata and not creating or destroying identifiers or links. Identifiers and links are never explicitly created or destroyed. Conceptually, all identifiers exist at all times (see section 2.6.1) and links between identifiers have no meaning without having metadata attached to them. A successful metadata publish MUST result in a publishReceived message. Otherwise, the entire publish request MUST fail without effect and the response MUST contain an errorResult element with an errorCode attribute indicating the cause of the failure.

There are three subtypes of publish: update, notify and delete. A MAP Server MUST process valid publish messages which contain any combination of the three publish subtypes.

```
<xsd:complexType name="UpdateType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
        type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
    <xsd:element name="metadata" type="MetadataListType"
      minOccurs="1" maxOccurs="1"/>
  </xsd:sequence>
  <xsd:attribute name="lifetime" default="session">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="session"/>
        <xsd:enumeration value="forever"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:complexType>

<xsd:complexType name="DeleteType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
        type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attribute name="filter" type="FilterType" use="optional"/>
</xsd:complexType>

<xsd:complexType name="PublishRequestType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="unbounded">
      <xsd:element name="update" type="UpdateType"/>
      <xsd:element name="notify" type="UpdateType"/>
      <xsd:element name="delete" type="DeleteType"/>
    </xsd:choice>
  </xsd:sequence>
</xsd:complexType>
```

```
</xsd:choice>
</xsd:sequence>
<xsd:attributeGroup ref="sessionAttributes"/>
<xsd:attributeGroup ref="validationAttributes"/>
</xsd:complexType>
```

### 3.7.1.1 update

update adds or changes metadata associated with identifiers and links. An update request with ifmap-cardinality="singleValue" MUST replace the previous value for that metadata type if it exists on the MAP Server. An update request with ifmap-cardinality="multiValue" MUST append the new value to a list of values for that metadata type on the MAP Server.

For example, a PDP publishes a list of roles to a user:

```
<?xml version="1.0"?>
<ifmap:publish session-id="111"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
  METADATA/2">
  <update>
    <identity name="joe" type="username"/>
    <metadata>
      <meta:role ifmap-cardinality="multiValue">
        <name>Guest</name>
      </meta:role>
      <meta:role ifmap-cardinality="multiValue">
        <name>Contractor</name>
      </meta:role>
    </metadata>
  </update>
</ifmap:publish>
```

### 3.7.1.2 notify

notify tells the MAP server to notify subscribers with metadata that does not persist in the database. The format of a notify operation is the same as the format of an update operation. One common use of notify is with event metadata, to notify other MAP clients about events without storing events in the MAP server.

Metadata published with notify MUST be queued for delivery to all clients with subscriptions whose searches match the published data at the time the data is published. Metadata published with notify MUST be delivered to clients inside notifyResult elements within pollResult elements (see section 3.7.5). Metadata published with notify MUST NOT be delivered to clients in response to a search request.

For example, a Sensor publishes event metadata using notify:

```
<ifmap:publish session-id="111"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
  METADATA/2">
  <notify>
    <ip-address value="192.0.2.11" type="IPv4"/>
    <metadata>
      <meta:event ifmap-cardinality="multiValue">
        <name>attack</name>
        <discovered-time>2009-05-30T13:10:11</discovered-time>
      </meta:event>
    </metadata>
  </notify>
</ifmap:publish>
```

```
<discoverer-id>1273</discoverer-id>
<magnitude>45</magnitude>
<confidence>100</confidence>
<significance>important</significance>
<type>worm infection</type>
</meta:event>
</metadata>
</notify>
</ifmap:publish>
```

### 3.7.1.3 delete

delete removes metadata from identifiers and links; identifiers and links are never explicitly deleted. A MAP Server MUST delete metadata specified by a valid delete message.

If a delete request has no filter, a MAP server MUST delete all metadata associated with the identifier or link. If a delete request has a filter that is the empty string, the MAP server MUST delete nothing.

For example, when an IF-MAP enabled DHCP server revokes the lease on an IP address it deletes the associated metadata from the link between the ip-address identifier and the mac-address identifier.

```
<?xml version="1.0"?>
<ifmap:publish session-id="111"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <delete
    filter='meta:ip-mac[@ifmap-publisher-id="222"]'>
    <ip-address value="192.0.2.11"/>
    <mac-address value="00:11:22:33:44:55"/>
  </delete>
</ifmap:publish>
```

### 3.7.1.4 Multiple operations in a single publish request

When a publish request contains multiple update and/or delete elements which operate on the same identifiers or links, the result of the publish request MUST be consistent with the update and delete elements having been applied to the IF-MAP database in the order in which they are specified by the client.

An IF-MAP Client MUST publish multiple operations in a single publish request when the operations are atomically related. An IF-MAP Client MUST refrain from publishing multiple operations in a single publish request for all reasons other than an achieving atomic database semantics, for example, but not limited to, any attempts to increase efficiency of communication with the IF-MAP Server.

A MAP Client SHOULD refrain from publishing very frequent updates to the same metadata. If a MAP Client is observing rapidly changing behavior in the network that needs to be reflected in MAP, the MAP Client SHOULD throttle its publish requests for that metadata to a reasonable rate such as once per second.

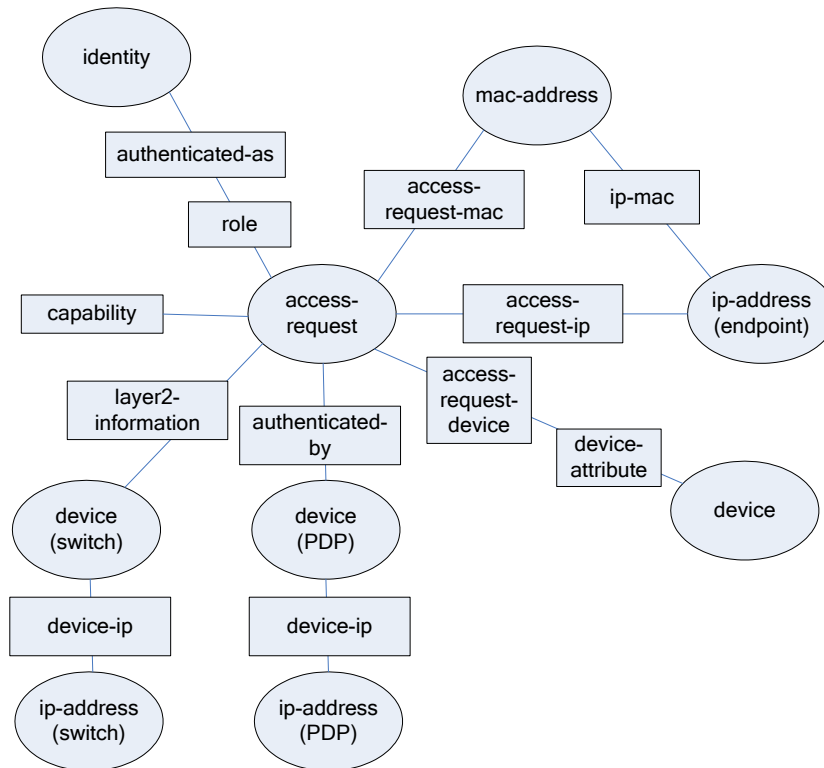
See section 3.10 for further discussion of the atomicity of multiple operations in a single publish request.

## 3.7.2 search

A search request retrieves metadata associated with an identifier instance and any linked identifiers.

### 3.7.2.1 The Retrieval Model

The retrieval model is that of an ever-widening search of an arbitrarily connected graph by examination and rule matching of metadata on identifiers and links. The results are bounded by reachability by following valid link types, specified maximum result depth, specified maximum result size, and terminal identifier types.



**Figure 3**

It may be helpful to consider the MAP Server state as a graph where searches find connected sub-graphs starting at a particular node. Figure 3 depicts identifiers as ovals, links as lines, and metadata as squares.

The following six values parameterize searches:

#### 3.7.2.2 identifier

identifier specifies the starting place of the search.

#### 3.7.2.3 match-links

match-links specifies the criteria for positive matching for including metadata from any link visited in the search. match-links also specifies the criteria for including linked identifiers in the search.

#### 3.7.2.4 max-depth

max-depth specifies the maximum distance of any included identifiers. Distance is measured by number of links away from the starting identifier.

#### 3.7.2.5 max-size

max-size specifies the maximum size of the results. Specifically, this is a measure of the size of the searchResult element or pollResult element in its entirety, in wire format. This excludes the SOAP envelope and result wrapper; effectively, it includes everything under the ifmap:response element.

### 3.7.2.6 result-filter

The filter specifies any further rules for deleting data from the results. If there is no result-filter attribute, all metadata on all identifiers and links that match the search is returned to the client. If an empty result-filter attribute is specified, the identifiers and links that match the search are returned to the client with no metadata.

### 3.7.2.7 terminal-identifier-type

The terminal-identifier-type specifies identifier types that when encountered terminate a search. This is used when the depth of a search might otherwise cause the search to gather metadata associated with more links or identifiers than desired. Metadata associated with a terminal identifier type is included in search results if it matches result-filter. Syntactically, terminal-identifier-type's value is a comma-separated list of one or more of the following strings:

- access-request
- identity
- ip-address
- mac-address
- device

### 3.7.2.8 The Search Algorithm

MAP Servers MUST provide results equivalent, with respect to the identifiers, links, and metadata, to a search which would result from the following algorithm.

Some terms and conditions:

A search is specified in IF-MAP syntax for SearchRequestType (see below).

The record contains the results and is returned to the MAP Client.

Current depth is defined as the number of links between the current identifier and the starting identifier.

A sub-result contains partial results.

Algorithm:

1. Run recursive search subroutine below with the search identifier and depth zero.
2. Apply "result-filter" to the results, deleting any unmatched metadata from the results.
3. Give the results to the client or return an error (and no results) indicating max-size was reached.

Recursive Search Subroutine:

1. Start at the current identifier, with empty current results, at the current depth.
2. Add any metadata on the current identifier to current results.
3. If the current identifier's type is contained within "terminal-identifier-type", return current results.
4. If the current depth is greater than or equal to "max-depth", return current results.
5. For all links associated with the current identifier, add all metadata on the links that match the "match-links" filter to current results. Remember the set of links that matches the "match-links" filter for use in step #6 below.
6. For all identifiers associated with the current identifier via links remembered in step 5, start the recursive search subroutine (step 1) with a depth equal to the current depth plus one. Add the return value of each recursive search subroutine to current results.

## 7. Return current results.

If a MAP Client does not specify max-depth, the MAP Server MUST process the search with a max-depth of zero.

MAP Servers MUST support size constraints up to and including 100KB<sup>1</sup>. If a MAP Client does not specify max-size, the MAP Server MUST process the search with a max-size of 100KB. If a MAP Client specifies a max-size that exceeds what the MAP Server can support, the MAP Server MUST enforce its own maximum size constraints.

A MAP Server MUST return an error (and no partial results) if result exceeds max-size. In other words, the MAP Client will get back either full results or an error. The table below summarizes the possible combinations concerning max-size requests.

Client's Requested "max-size" value	Server's maximum supported result size	Size of search result	Server's response
Unspecified	ANY (i.e. >= 100KB)	<=100KB	Result
Unspecified	ANY (i.e. >= 100KB)	>100KB	Error - result too large
specified (i.e. > 0)	< client's requested "max-size"	<= server's maximum supported result size	Result
Specified (i.e. > 0)	< client's requested "max-size"	> server's maximum supported result size	Error – result too large
specified (i.e. > 0)	>= client's requested "max-size"	<= client's requested "max-size"	Result
specified (i.e. > 0)	>= client's requested "max-size"	> client's requested "max-size"	Error – result too large

```

<xsd:complexType name="SearchType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="1">
      <xsd:element name="access-request"
        type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attribute name="match-links" type="FilterType"/>
  <xsd:attribute name="max-depth" type="xsd:unsignedInt"/>
  <xsd:attribute name="max-size" type="xsd:unsignedInt"/>
  <xsd:attribute name="result-filter" type="FilterType"/>
  <xsd:attribute name="terminal-identifier-type"
    type="xsd:string"/>
</xsd:complexType>

```

<sup>1</sup>A limit is specified so that MAP Server implementations may protect themselves from resource exhaustion due to unreasonable searches. 100KB is large enough to return results for reasonable searches.

SearchType is used by SearchRequestType as well as SubscribeRequestType (see section 3.7.4.1).

```
<xsd:complexType name="SearchRequestType">
  <xsd:complexContent>
    <xsd:extension base="SearchType">
      <xsd:attributeGroup ref="sessionAttributes"/>
      <xsd:attributeGroup ref="validationAttributes"/>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

For example, when a Flow Controller detects a new flow from a previously unseen IP address, it searches a MAP Server for the list of capabilities assigned to the corresponding access-request to make enforcement decision about this flow:

```
<?xml version="1.0"?>
<ifmap:search session-id="123"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2"
  match-links="meta:access-request-ip or meta:ip-mac or
  meta:access-request-mac"
  max-depth="3" result-filter="meta:capability or
  meta:device-attribute or meta:roles"
  terminal-identifier-type="identity,device">
  <ip-address value="192.0.2.11" type="IPv4"/>
</ifmap:search>
```

### 3.7.3 Search Results

A successful search MUST result in a response message containing a searchResult element comprised of identifiers and links along with their associated metadata. The resulting XML document contains a resultItem element for each node and each edge in the metadata graph that matches the query. The XML structure itself does not directly reflect the structure of the metadata graph. Metadata appearing in a searchResult is filtered by the result-filter attribute in the search or subscription corresponding to the searchResult. If the result-filter filters out all metadata associated with an identifier or link, the identifier or link MUST still be included in the searchResult even though no metadata is associated with the identifier or link in the search result. If no result-filter is present in a search or subscription, ALL metadata that matches the search MUST be returned.

Since all identifiers for a given identifier type are always valid to search, the MAP Server MUST never return an "identifier not found" error when searching for an identifier. In this case, the MAP Server MUST return the identifier with no metadata or links attached to it.

Each search result may contain a name attribute. The name attribute is used for poll results (see section 3.7.5).

```
<xsd:complexType name="ResultItemType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
        type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
    </xsd:choice>
  </xsd:sequence>
</xsd:complexType>
```



```

    <xsd:element name="device" type="DeviceType"/>
  </xsd:choice>
  <xsd:element name="metadata" type="MetadataListType"
    minOccurs="0" maxOccurs="1"/>
</xsd:sequence>
</xsd:complexType>

<xsd:complexType name="SearchResultType">
  <xsd:sequence>
    <xsd:element name="resultItem" type="ResultItemType"
      minOccurs="0" maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="name"/>
</xsd:complexType>

```

For example, a MAP Server may respond to the “search” request from the example above with the following message:

```

<?xml version="1.0"?>
<ifmap:response
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
  METADATA/2">
  <searchResult>
    <resultItem>
      <ip-address value="192.0.2.11" type="IPv4"/>
    </resultItem>
    <resultItem>
      <access-request name="123"/>
      <metadata>
        <meta:capability ifmap-cardinality="multiValue">
          <name>finance-server-access</name>
        </meta:capability>
      </metadata>
    </resultItem>>
    <resultItem>
      <ip-address value="192.0.2.11" type="IPv4"/>
      <access-request name="123"/>
    </resultItem>
  </searchResult>
</ifmap:response>

```

### 3.7.4 subscribe

A MAP Client uses subscribe requests to manage its subscription to searches which may be polled on a MAP Server.

A subscription is a list of SearchType items. The subscription list is consulted by the MAP Server when determining what clients need to be notified about the results of a publish request.

A MAP Server MUST maintain only one subscription list per connected MAP Client.

A subscribeRequest contains one or more update or delete elements used to manage the subscription list.

A MAP Server MUST respond to a valid subscribe message with a subscribeReceived message. If the subscribeRequest is not valid, the MAP Server MUST respond with an appropriate errorResult.

When a MAP Client initially connects to a MAP Server, the MAP Server MUST delete any previous subscriptions corresponding to the MAP Client. In other words, subscription lists are only valid for a single MAP Client session.

```

<xsd:complexType name="DeleteSearchRequestType">
  <xsd:attribute name="name" type="xsd:string" use="required"/>
</xsd:complexType>

<xsd:complexType name="SubscribeRequestType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="unbounded">
      <xsd:element name="update">
        <xsd:complexType>
          <xsd:complexContent>
            <xsd:extension base="SearchType">
              <xsd:attribute name="name" type="xsd:string"
                use="required"/>
            </xsd:extension>
          </xsd:complexContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="delete" type="DeleteSearchRequestType"/>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attributeGroup ref="sessionAttributes"/>
  <xsd:attributeGroup ref="validationAttributes"/>
</xsd:complexType>

```

### 3.7.4.1 update

update adds or changes subscribed searches. Each subscribed search is identified by a name attribute. The value of the name attribute is generated and managed by the MAP Client. To add a new search request to a subscription, the MAP Client specifies a new name attribute on the update element. To modify an existing search request, the MAP Client specifies the name attribute of an existing search. The value of the name attribute MUST be a string between 1 and 20 characters in length, such as a unique integer value for each search the client subscribes to. The client MUST specify the name attribute in every update element of a subscriptionRequest.

For example, a Flow Controller subscribes to a MAP Server for changes in the list of roles assigned to an IP address via access-request and identity to make enforcement decisions about this flow:

```

<?xml version="1.0"?>
<ifmap:subscribe session-id="111"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
  METADATA/2">
  <update name="35" result-filter="meta:role"
    match-links="meta:access-request-ip or
      meta:access-request-mac or meta:ip-mac or
      meta:access-request-identity"
    max-depth="3">
    <ip-address value="192.0.2.11" type="IPv4"/>
  </update>

```

```
</ifmap:subscribe>
```

### 3.7.4.2 delete

delete removes the subscribed search associated with the specified name.

For example, a Flow Controller deletes its subscription to a MAP Server for changes in the list of roles assigned to an AR:

```
<?xml version="1.0"?>
<ifmap:subscribe session-id="111"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
  METADATA/2">
  <delete name="35"/>
</ifmap:subscribe>
```

### 3.7.5 poll

A poll request is sent by a MAP Client to a MAP Server to request notification of metadata changes based on the MAP Client's subscription. Metadata may be changed by publish requests (section 3.7.1), purgePublisher requests (section 3.7.6), or automatic deletion of metadata by the MAP Server (section 3.3.5).

```
<xsd:complexType name="PollRequestType">
  <xsd:attributeGroup ref="validationAttributes"/>
  <xsd:attributeGroup ref="sessionAttributes"/>
</xsd:complexType>
```

A MAP Server MUST respond to a valid poll with a pollResult, endSessionResult (see section 4), or errorResult message. If the poll is not valid, the MAP Server MUST respond with an appropriate errorResult. A valid poll may result in an error (most likely SearchResultsTooBig or PollResultsTooBig). In this case the MAP Server MUST respond with a pollResult message containing an errorResult. If a server responds to a poll with an errorResult, all of the client's subscriptions are automatically invalidated and MUST be removed by the server. After receiving an errorResult in response to a poll, a client will need to issue new subscribe requests in order to receive more poll results.

In response to a poll request, the MAP Server checks to see if any search results are available for subscriptions the client has made. If any results are available, the MAP Server MUST send a pollResult response containing search results. If no results are available, including the case where there are no active subscriptions for this client, the MAP Server MUST NOT send a response to that poll. At some future time when metadata changes occur that match client subscriptions, the MAP Server MUST send a pollResult response containing search results.

The first time a pollResult contains search results for a new subscription, the search results MUST consist of the complete set of identifiers, links, and metadata for the subscription as specified in section 3.7.2.1. The complete search results are sent to the MAP client in a searchResult element. Subsequent pollResults (known as "delta pollResults") MUST contain updates in updateResult elements as metadata is added and deletes in deleteResult elements as metadata is removed. The updateResult and deleteResult elements returned by the server MUST reflect ALL of the updates which have occurred since the last poll which affect the client's subscriptions. The server MUST NOT compress search results by removing updateResult and deleteResult elements which cancel each other out. The server MAY return the accumulated result elements in multiple pollResult messages. The server accomplishes this by returning a subset of the accumulated results to the client in response to a single poll request. In response to the next poll request the server returns another subset, continuing in this fashion until all pollResults have been sent. Note that atomic changes require special handling; see section 3.10 for details.

A metadata change may introduce or remove a link between two subgraphs that match a subscription. When this happens after an initial searchResult message has been sent, the MAP Server informs the MAP Client about the metadata that was added or removed to the subscription using updateResult and deleteResult elements.

The server MUST return a separate searchResult, updateResult, deleteResult or errorResult element for each subscription that has changes. Each searchResult, updateResult, deleteResult and errorResult element in a pollResult response MUST include a name attribute which identifies the subscription corresponding to the result.

The server MUST NOT return searchResult, updateResult or deleteResult elements for a subscription that has no changes associated with it.

pollResult is also used to send metadata to MAP clients that a MAP client publishes using the notify publish subtype. When a MAP client publishes using notify, the MAP server MUST add the published metadata to the poll results for each subscriber whose subscription matches the published metadata. For each such client that has an active poll request, the MAP server SHOULD return the accumulated search results to the client as soon as possible. For a client that does not have an active poll request, the MAP server MUST retain metadata published using notify until the client issues a poll request or the client's session ends. Metadata published using notify MUST be sent to a subscribing MAP client in a notifyResult element.

notifyResult is only used to return metadata that was published using notify, even if other metadata matches the result-filter in the subscription. A MAP Client receiving a notifyResult MUST NOT assume that metadata missing from a notifyResult has been deleted.

When processing pollResults, a MAP client SHOULD consider a result as the combination of metadata, identifier, result type (notify, update, or delete), and subscription name, but not the manner in which the results are encoded in the XML. When generating delta pollResults, a MAP server MUST ensure that metadata elements are delivered in the same order they are received, and a MAP client SHOULD maintain the order of metadata elements as delivered in the XML when processing these results.

Because the server might choose to group the results differently, it SHOULD NOT be assumed that there is a 1:1 relationship between the number of notify requests and actual notifyResult responses. For example, the following two results are functionally equal:

```
<ifmap:response>
  <pollResult>
    <notifyResult name="sub1">
      <resultItem>
        <identity type="username" name="id1"></identity>
        <metadata>
          <meta:event ifmap-cardinality="multiValue"
            ifmap-publisher-id="x"
            ifmap-timestamp="xxx">...event1...</meta:event>
        </metadata>
      </resultItem>
    </notifyResult>
    <notifyResult name="sub1">
      <resultItem>
        <identity type="username" name="id1"></identity>
        <metadata>
          <meta:event ifmap-cardinality="multiValue"
            ifmap-publisher-id="x"
            ifmap-timestamp="xxx">...event2...</meta:event>
        </metadata>
      </resultItem>
    </notifyResult>
  </pollResult>
</ifmap:response>
```

```

</pollResult>
</ifmap:response>

```

```

<ifmap:response>
  <pollResult>
    <notifyResult name="sub1">
      <resultItem>
        <identity type="username" name="id1"></identity>
        <metadata>
          <meta:event ifmap-cardinality="multiValue"
            ifmap-publisher-id="x"
            ifmap-timestamp="xxx">...event1...</meta:event>
          <meta:event ifmap-cardinality="multiValue"
            ifmap-publisher-id="x"
            ifmap-timestamp="xxx">...event2...</meta:event>
        </metadata>
      </resultItem>
    </notifyResult>
  </pollResult>
</ifmap:response>

```

A MAP Server MUST buffer at least 5,000,000 bytes of poll results for each MAP Client (see section 4.3). If the size of a MAP Client's poll results exceeds the MAP Server's limit, the MAP Server MUST indicate this to the MAP client by responding to a poll request with an `errorResult` containing an `errorCode` of `PollResultsTooBig`. In this situation, the MAP Server MAY opt to first return a subset of the accumulated `notifyResult` elements before responding to a subsequent poll request with an `errorResult` response.

```

<xsd:complexType name="PollResultType">
  <xsd:sequence>
    <xsd:choice minOccurs="0" maxOccurs="unbounded">
      <xsd:element name="searchResult" type="SearchResultType"/>
      <xsd:element name="updateResult" type="SearchResultType"/>
      <xsd:element name="deleteResult" type="SearchResultType"/>
      <xsd:element name="notifyResult" type="SearchResultType"/>
      <xsd:element name="errorResult" type="ErrorResultType"/>
    </xsd:choice>
  </xsd:sequence>
</xsd:complexType>

```

The figure below illustrates subscription and polling between the MAP Client and MAP Server.

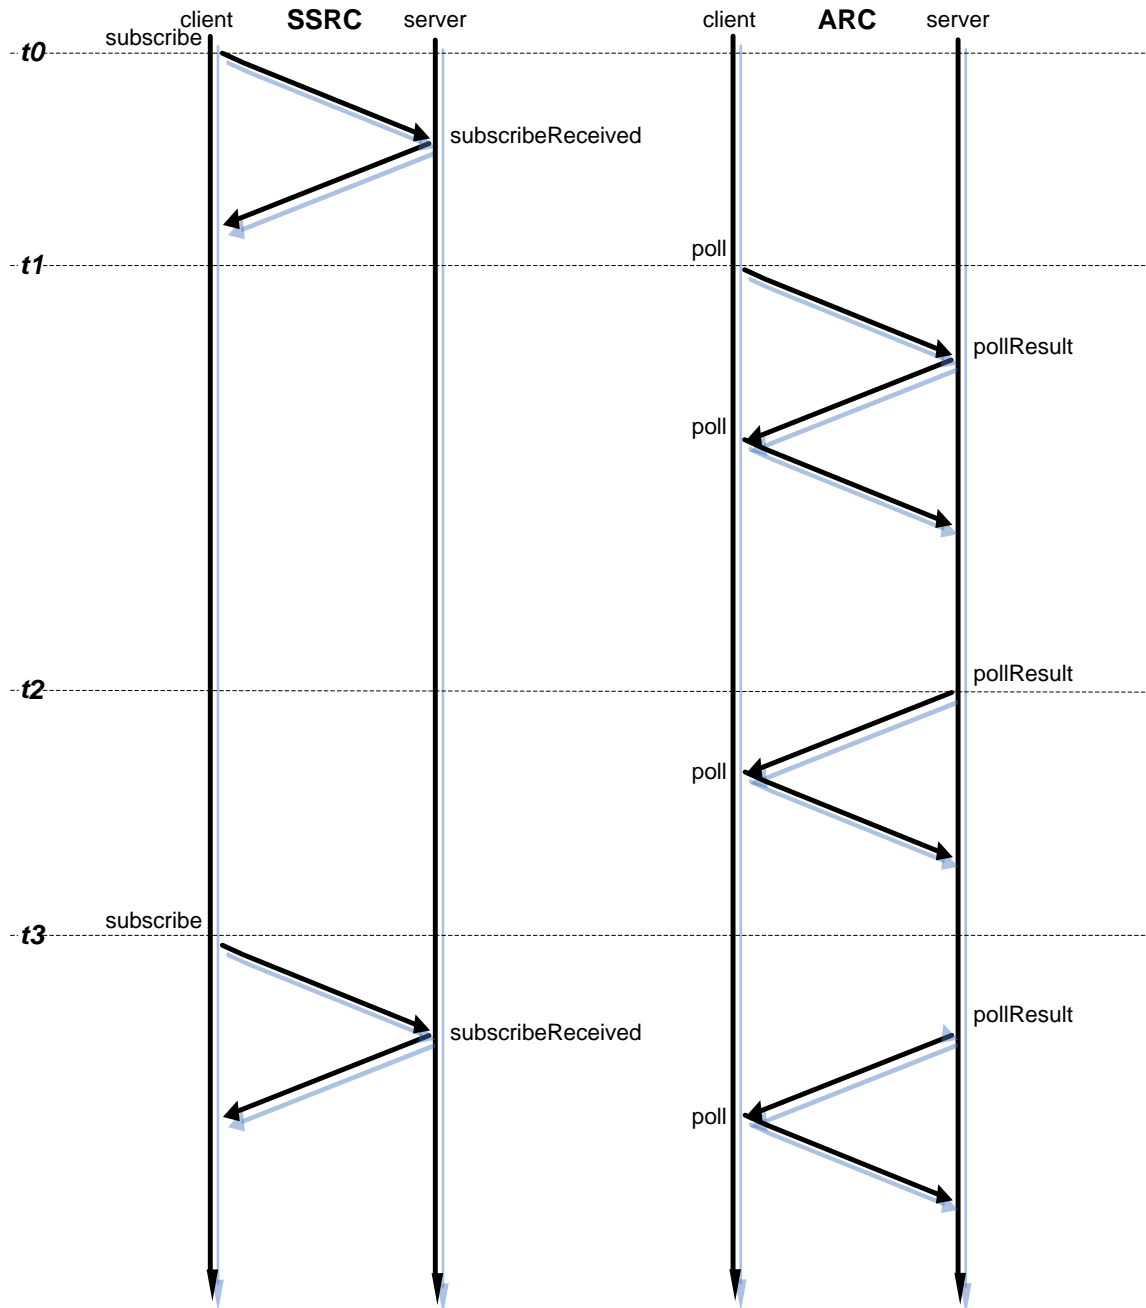


Figure 4

The figure shows two channels between the client and the server (see section 4 for a full explanation of SSRC and ARC). On the SSRC channel, the client sends subscribe requests. On the ARC channel, the client sends poll requests. In the figure, the vertical arrows represent elapsed time starting at the top. At  $t_0$ , the client sends a subscribe request, and the server responds immediately with a subscribeReceived response. At  $t_1$ , the client sends a poll request. Since the client has already subscribed and the subscription search matches some results, the server immediately replies with a pollResult response containing the search results. Upon receiving the pollResult response, the client issues another poll request. When the server receives the second poll request, the server does not respond right away because there are no changes to the search results since the last poll.

At t2, another MAP Client makes changes to the IF-MAP metadata that match the subscription. The server sends a pollResult response containing the new search results. This pollResult is in response to the last poll request the client issued. When the client receives the pollResult response it issues another poll request. Again, the server does not respond right away because there are no more changes to the search results.

At t3, the client issues a subscribe request on the SSRC channel that alters the search results for the subscription. On the SSRC channel, the server sends a subscribeReceived response. At the same time, the server sends a pollResult response on the ARC channel containing the new search results. When the client receives the pollResult response, it sends another poll request so the client can be notified of further changes.

A race condition is inherent in updating a subscribed search. There is no mechanism to synchronize management of a subscribed search with delivery of results for that subscription. A MAP client may receive results that pertain to the subscription as it existed prior to an update after having received a subscribedReceived confirmation from the MAP server for the update in question. This can happen because poll results are delivered to the MAP client on a different channel than the one used for subscription management requests.

To avoid this race condition, a MAP client may delete the old subscription and create a new subscription with a new name. That way, the MAP client can distinguish between results that pertain to the old subscription and results that pertain to the new subscription.

A MAP Client that issues subscribe requests MUST create an ARC and issue poll requests. After receiving each pollResult response from a MAP Server, a MAP Client SHOULD immediately send another poll request in order to minimize the amount of poll result data the MAP Server accumulates on the MAP Client's behalf.

### 3.7.6 purgePublisher

A purgePublisher request is sent by a MAP Client to ask that the MAP Server remove all metadata associated with a particular publisher. The purgePublisher request is typically used by a MAP Client to purge its own data after a power cycle or system reset if the client has no persistent knowledge of metadata it published prior to the reset. A MAP Server MAY forbid a MAP Client to use the purgePublisher request to remove data published by a different MAP Client, in which case the MAP Server MUST respond with an AccessDenied error.

```
<xsd:complexType name="PurgePublisherRequestType">
  <xsd:attribute name="ifmap-publisher-id" type="xsd:string"/>
  <xsd:attributeGroup ref="sessionAttributes"/>
</xsd:complexType>
```

A MAP Server MUST either respond to a purgePublisher request with a purgePublisherReceived message and remove all data with that publisher ID, or return an errorResult message and not delete any data.

## 3.8 Schema Versioning

Schema version agreement between MAP Servers and MAP Clients is based on XML namespaces. XML namespaces included in IF-MAP message documents SHOULD be compared by the MAP Server to determine compatibility of MAP Client requests. A MAP Server MUST include the XML namespaces associated with the newest schema versions in any message documents. If a MAP Server is unable to process a request due to a schema version mismatch it SHOULD return an InvalidSchemaVersion error.

### 3.9 Vendor-specific Metadata

IF-MAP standard metadata schema SHOULD NOT be extended directly with vendor-specific extensions. Separate vendor-specific metadata schema MAY be defined. Vendor-specific metadata schema MAY rely on IF-MAP standard schemas. MAP Servers MUST support operations using vendor-specific metadata which is defined by an XML schema.

ifmap-publisher-id and standard metadata names can be used to uniquely associate instances of vendor-specific metadata which add more detail to instances of standard metadata. For example, an IDS acting as a Sensor might update a MAP Server with an IF-MAP standard **event** and update the MAP Server with a vendor-specific event which contains the name of the IF-MAP standard **event** and more vendor-specific detail. A Flow Controller which understands the vendor-specific event can combine the standard event and vendor-specific event by matching on ifmap-publisher-id and **event** name.

MAP Clients MUST ignore unrecognized vendor-specific metadata returned by searches and subscriptions. This enables MAP Clients that use vendor-specific metadata to coexist with MAP Clients that do not use vendor-specific metadata.

### 3.10 Atomicity

In this section, we consider three cases (publish, purgePublisher, and automatic metadata deletion), each of which constitutes a “change”:

- Publish requests that contain multiple update, notify and/or delete requests have special behavior. The Server MUST ensure that they appear atomic.
- The Server MUST ensure that purgePublisher requests appear atomic.
- When (per section 3.3.5) the Server deletes all metadata that came from a particular client and had lifetime=”session”, the Server MUST ensure that this bulk deletion appears atomic.

Where “atomic” means:

- No interim results are visible to Clients. In particular, a pollResult or searchResult will reflect the data in a state either before or after the entire change. A MAP Client would never be able to tell the order of the operations within a multi-part publish request.
- Changes are all-or-nothing. Either a change fails and leaves the data unchanged, or the change succeeds in entirety. This holds even in case of a failure.

This requirement for special handling of atomic changes supersedes the requirement in section 3.7.5 that “[t]he updateResult and deleteResult elements returned by the server MUST reflect ALL of the updates which have occurred since the last poll which affect the client’s subscriptions.”

Per section 3.7.5, if a MAP Client sends a publish request containing a delete operation, followed by another publish request containing an update operation, polling MAP Clients should see the results of both operations.

However, if a MAP Client sends a single publish request containing both a delete operation and an update operation, polling MAP Clients should see only the end result due to the atomicity of the publish request; the MAP Server MUST send a pollResult containing the end state after the delete and update operations in the single publish request.



## 4 SOAP Binding and Session Management

As a half-duplex web services transport protocol for XML payloads which is easy to set up for synchronous and asynchronous modes of operation, SOAP is a good match for the needs of IF-MAP.

All connections over which MAP Clients and Servers communicate are initiated by MAP Clients. All IF-MAP operations are initiated by the MAP Client. A MAP Server may respond to an operation from a client either synchronously or asynchronously.

A MAP Client initiates a connection using an https URI. A MAP Client **MUST** be capable of communicating with a MAP Server using any valid URI as specified in [11] and [17]. In particular, a MAP Client **MUST** be capable of communicating with a MAP Server using a tcp port specified in an https URI.

### 4.1 Client-Server Communication Model

MAP Clients and Servers communicate within the context of a session.

#### 4.1.1 Sessions

An IF-MAP session consists of one synchronous send-receive channel (SSRC) and no more than one optional asynchronous receive channel (ARC); a MAP Client **MUST** open exactly one SSRC and no more than one ARC. Each session is uniquely identified by a session id. A session lasts as long as the client actively uses it and doesn't attempt to establish a new session. A client keeps a session alive either by keeping a TCP connection associated with the SSRC or the ARC open or by periodically sending a renewSession request to the server. If a MAP client has an existing SSRC connection and sends an SSRC request over a different connection, the MAP server **MUST** associate the new connection with the SSRC and **MUST** close the old connection (if possible) or respond to new requests on the old connection with an AccessDenied errorResult.

An IF-MAP server **SHOULD** end a session after a period of client inactivity which **MUST** be no shorter than three minutes. The length of the period of client inactivity that results in session termination **MAY** be configurable on an IF-MAP server.

##### 4.1.1.1 Synchronous Send-Receive Channel (SSRC)

The SSRC is a SOAP/HTTPS channel over which the following IF-MAP messages **MAY** be communicated: newSession, renewSession, endSession, publish, search, subscribe, purgePublisher, response. The SSRC **MUST NOT** be used for poll requests. After opening a session, the MAP Client sends a series of requests on the SSRC, and the MAP Server synchronously responds to each request over the same channel.

Valid request -response pairs are:

- newSession -> newSessionResult | errorResult
- publish -> publishReceived | errorResult
- subscribe -> subscribeReceived | errorResult
- search -> searchResult | errorResult
- purgePublisher -> purgePublisherReceived | errorResult
- renewSession -> renewSessionResult | errorResult
- endSession -> endSessionResult | errorResult

A MAP Client **SHOULD** timeout if the MAP Server does not respond in a timely fashion to a request.

If after sending a publish request the client does not get a response from the server for any reason (timeout, connection closed, client shutting down, etc) the client should generally

reconnect when it is able to and should generally repeat the publish. However, repeating may result in the publish happening twice. If the publish includes an update of a multi-valued piece of metadata, this would result in a duplicate value. A client SHOULD prevent this duplication from happening. To prevent it, when republishing the client may prepend a delete request that would delete the identical piece of metadata with the identical ifmap-publisher-id. For example:

```
<ifmap:publish session-id="222">
  <delete
    filter="meta:capability[@ifmap-publisher-id='111']">
    <access-request name="111:42"/>
  </delete>
  <update>
    <access-request name="111:42"/>
    <metadata>
      <meta:capability ifmap-cardinality="multiValue">
        <name>unrestricted-access</name>
      </meta:capability>
    </metadata>
  </update>
</ifmap:publish>
```

#### 4.1.1.2 Asynchronous Receive Channel (ARC)

The ARC is an optional channel intended to allow a MAP Client to asynchronously receive the results of its current subscription by using the poll request. The following IF-MAP messages MAY be communicated over the ARC: poll and response. Other messages MUST NOT be communicated over the ARC. The response element sent in response to a poll request MUST contain either an appropriate pollResult element corresponding to the MAP Client's subscription, an errorResult element indicating the cause of an error, or an endSessionResult indicating that the session ended.

A MAP Client MUST have a valid session with a MAP Server before opening an ARC. A MAP Client MUST NOT send any IF-MAP request other than "poll" on an ARC. A "poll" request MUST contain a session-id attribute referring to a valid session. A MAP Client MUST have no more than one ARC active at a time, meaning that a MAP Client MUST NOT have more than one outstanding "poll" request for a particular session.

## 4.2 SOAP Transport

All implementations of the IF-MAP protocol MUST support Simple Object Access Protocol (SOAP) v. 1.2 as defined in [7]. All IF-MAP messages are encapsulated inside SOAP bodies which in turn are inside SOAP envelopes. HTTP compression for responses, if used, MUST be negotiated according to HTTP 1.1 [17]. A MAP client MAY compress requests according to HTTP 1.1, where the compression algorithm used is indicated using the Content-Encoding HTTP header. MAP servers MUST accept requests that have been compressed using gzip and identity transformations.

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:publish session-id="128738734">
      ...
    </ifmap:publish>
  </env:Body>
```

```
</env:Envelope>
```

### 4.3 Session ID

All IF-MAP messages are associated with a session-id. The session-id is a value chosen by the MAP Server. The content of the session-id is not meaningful to a MAP Client, nor is the method used to derive a session-id defined by this specification. A session-id is a string that matches NMTOKEN, as defined in the XML 1.0 specification. A session-id may be up to 128 characters in length. The purpose of the session-id is to enable the server to share state across multiple incoming TCP connections from a single client.

When a MAP Client first connects to a MAP Server, the MAP Client requests a new session-id by sending a SOAP request containing a “newSession” element in the SOAP body:

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:newSession max-poll-result-size="10000000"/>
  </env:Body>
</env:Envelope>
```

max-poll-result-size is an optional attribute that indicates to the MAP Server the amount of buffer space the client would like to have allocated to hold poll results. A MAP Server MUST support buffer sizes of at least 5,000,000 bytes, and MAY support larger sizes.

If the MAP Server can create a session for the client, the MAP Server’s response MUST contain a “newSessionResult” element that has a “session-id” attribute that specifies the MAP Client’s session-id along with an “ifmap-publisher-id” attribute that the MAP Client can use to recognize metadata that it published by examining operational attributes. If the MAP Client included a max-poll-result-size attribute in its newSession request, the MAP Server MUST include a max-poll-result-size in its response indicating the actual amount of buffer space available for poll results for the client:

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:response>
      <newSessionResult session-id="222" ifmap-publisher-id="111"
        max-poll-result-size="7500000"/>
    </ifmap:response>
  </env:Body>
</env:Envelope>
```

If a MAP Client sends more than one SOAP request containing a “newSession” element in the SOAP body, the MAP Server MUST respond by ending the previous session and starting a new session. The server’s response MUST contain a “newSessionResult” element, and any state associated with the old session MUST be discarded. The new session MAY use the same

session-id or allocate a new one. If an ARC is associated with the old session, the server MUST send an endSessionResult on the ARC.

A MAP Client associates an ARC channel with the same session-id as its SSRC channel. It does this by sending a SOAP request containing a “poll” element in the SOAP body.

Each subsequent request (including a “poll” request on an ARC) MUST contain a “session-id” attribute in the top level element of the SOAP body, specifying the session-id assigned to the connection by the MAP Server:

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:search session-id="222">
      ...
    </ifmap:search>
  </env:Body>
</env:Envelope>
```

If the MAP Client specifies an invalid session-id, the MAP Server MUST indicate an InvalidSessionID errorResult in its response. A MAP Server MUST NOT accept a request from a MAP Client unless the session-id is valid for that client. A MAP Server MUST NOT permit one MAP Client to use a session that is created by another MAP Client.

To explicitly terminate a session with a MAP server, a MAP Client MAY send an endSession request:

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:endSession session-id="222"/>
  </env:Body>
</env:Envelope>
```

If the session is valid, The MAP Server MUST respond with an endSessionResult response.

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:response>
      <endSessionResult/>
    </ifmap:response>
  </env:Body>
</env:Envelope>
```

When a session ends for any reason, and there is an outstanding poll request on the ARC, the MAP Server MUST send an endSessionResult to the MAP Client on the ARC.

If a MAP Server receives a message containing a SOAP body containing a poll element that specifies a session which already has an ARC with an outstanding poll request, the MAP Server MUST:

- end the session
- respond to the poll request on the older ARC with an endSessionResult
- respond to the poll request on the newer ARC with an errorResult response with an errorCode of InvalidSessionID

## 4.4 Session Renewal

In order to keep an IF-MAP session from timing out, a MAP Client MUST either keep the underlying TCP connection associated with the SSRC open, or send periodic renewSession requests to the MAP Server. The MAP Client should send renewSession requests somewhat more frequently than the session timeout on the MAP Server. Since the minimum session timeout is 180 seconds, a MAP Client that doesn't keep the SSRC open might send renewSession requests every 150 seconds unless configured otherwise.

Like other IF-MAP requests, renewSession request MUST specify a valid session-id.

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:renewSession session-id="222"/>
  </env:Body>
</env:Envelope>
```

If the session-id is valid, the server MUST respond with a renewSessionResult element. Otherwise, the server MUST respond with an errorResult element, specifying an InvalidSessionID errorCode.

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:meta="http://www.trustedcomputinggroup.org/2010/IFMAP-
METADATA/2">
  <env:Body>
    <ifmap:response>
      <renewSessionResult/>
    </ifmap:response>
  </env:Body>
</env:Envelope>
```

## 4.5 ARC error handling

If a MAP Server or MAP Client detects an error at the TCP or SSL layer of an ARC, the MAP Server or Client MUST end the session. The reason is that after a transport error on the ARC the MAP Server does not know whether the MAP Client received the last pollResult sent by the MAP Server, and the MAP Client does not know whether it missed any pollResult data.

After the MAP Client realizes that the session has ended, either because the MAP Client ended it explicitly or because an IF-MAP request resulted in an errorResult with an InvalidSessionID errorCode, the MAP Client SHOULD establish a new session and reestablish any subscriptions it is interested in. This enables the MAP Client and MAP Server to resynchronize pollResults.

## 4.6 WSDL and Example Code

An IF-MAP WSDL file and code examples utilizing various SOAP toolchains can be found on the TCG web site: <http://trustedcomputinggroup.org/> by searching for IF-MAP WSDL.

## 5 Recommendations for Backward Compatibility

In an environment where both IF-MAP 1.1 and IF-MAP 2.0 clients are deployed, an IF-MAP proxy or server that supports both IF-MAP 1.1 and IF-MAP 2.0 may want to expose metadata published by 1.1 clients to 2.0 clients and vice versa. If such backward compatibility is desired, both identifiers and metadata SHOULD be translated as described in this section. In the sections that follow, the entity that translates metadata is referred to as the IF-MAP Translator.

### 5.1 Translating identifiers

The IF-MAP 2.0 identifiers are nearly identical to the IF-MAP 1.1 identifiers with the exception of IF-MAP 2.0 identity identifiers having a type of “hip-hit” and IF-MAP 1.1 identity identifiers having a type of “trusted-platform-module”. An IF-MAP 2.0 identity identifier with a type of “hip-hit” SHOULD be represented to an IF-MAP 1.1 client as an identity identifier with a type of “other” and an other-type-definition of “hip-hit”. An IF-MAP 1.1 identity identifier with a type of “trusted-platform-module” SHOULD be represented to an IF-MAP 2.0 client as an identity identifier with a type of “other” and an other-type-definition of “trusted-platform-module”.

### 5.2 Translating Metadata

Metadata from the IF-MAP 1.1 specification published by 1.1 clients SHOULD be represented in the IF-MAP Metadata for Network Security 1.0 namespace in order to expose it seamlessly to IF-MAP 2.0 clients.

Some metadata from the IF-MAP Metadata for Network Security 1.0 namespace published by 2.0 clients may be represented as metadata in the IF-MAP 1.1 namespace. IF-MAP Metadata for Network Security 1.0 specifies a superset of the metadata in IF-MAP 1.1, so metadata that has no equivalent in the IF-MAP 1.1 specification cannot be represented as metadata in the IF-MAP 1.1 namespace. However, this metadata may be visible to IF-MAP 1.1 clients in the IF-MAP Metadata for Network Security 1.0 namespace.

### 5.3 Searches and Subscriptions

An IF-MAP Translator SHOULD translate metadata and identifiers between the two metadata namespaces when returning search results and poll results to an IF-MAP Client. An IF-MAP 2.0 Client SHOULD only specify the new metadata namespace in publish requests and in result filters, and SHOULD rely on the IF-MAP Translator to translate identifiers and metadata from the older schemas in order to interoperate with 1.1 clients.

When calculating results for a search or subscription with no result filter, an IF-MAP Translator SHOULD include metadata in both the IF-MAP 1.1 and IF-MAP Metadata for Network Security 1.0 namespaces. Metadata that has representations in both namespaces is returned twice, once in each namespace.

In IF-MAP 1.1, a negative max-depth was used to specify an unbounded search. IF-MAP 2.0 does not support unbounded searches. When a search request with a negative max-depth is received from an IF-MAP 1.1 client, an IF-MAP translator SHOULD substitute a very large positive number such as 4294967295.

IF-MAP 2.0 uses an unsigned integer for max-depth, while IF-MAP 1.1 used a signed integer. When a search request with a max-depth that exceeds the range of a signed 32 bit integer is received from an IF-MAP 2.0 client, an IF-MAP translator SHOULD translate max-depth to -1, turning the search into an unbounded search.

The specification of the search algorithm changed between IF-MAP 1.1 and IF-MAP 2.0. The algorithm as specified in IF-MAP 2.0 is the intended behavior for search even in IF-MAP 1.1; the IF-MAP 1.1 specification did not properly describe the search algorithm.

## 5.4 notify

An IF-MAP Translator SHOULD expose metadata published by the IF-MAP 2.0 notify operation to subscriptions made by 1.1 clients. Because IF-MAP 1.1 does not specify a notifyResult element in pollResult, the IF-MAP Translator should simulate notifyResult by sending an IF-MAP 1.1 pollResult that contains the metadata published by notify, followed a short time later by another IF-MAP 1.1 pollResult that does not contain the data published by notify. The recommended time between the two pollResults is 30 seconds, and IF-MAP Translators MAY enable administrators to configure different values for this time.

An IF-MAP metadata specification may require that notify be used to publish metadata that was originally specified by IF-MAP 1.1. For example, IF-MAP Metadata for Network Security 1.0 requires that event metadata be published using notify. When an IF-MAP 1.1 client publishes such metadata using the IF-MAP 1.1 metadata schema, an IF-MAP translator SHOULD use notifyResult instead of updateResult in pollResult messages sent to IF-MAP 2.0 clients. When the IF-MAP 1.1 deletes such metadata, the IF-MAP Translator SHOULD NOT send a deleteResult to IF-MAP 2.0 clients.

## 5.5 1.1 Clients and IF-MAP Metadata for Network Security

IF-MAP 1.1 Clients MAY use the new elements from the IF-MAP Metadata for Network Security namespace in order to take advantage of the new features before IF-MAP 2.0 servers are generally available. Such clients SHOULD only use the newly introduced elements from the new schema. Any elements that exist in both the IF-MAP 1.1 metadata schema and the IF-MAP Metadata for Network Security schema SHOULD be published by IF-MAP 1.1 clients using the IF-MAP 1.1 namespace. This enables interoperability with other IF-MAP 1.1 clients, as well as interoperability with IF-MAP 2.0 clients that follow the recommendations outlined in section 5.3.

## 5.6 Vendor-specific Metadata

An IF-MAP Translator SHOULD NOT perform any translation on vendor-specific metadata without coordinating the translation with the vendor that defined the vendor-specific metadata.

## 5.7 2.0 Clients and IF-MAP 1.1 Metadata

Section 3.2.2 states that “Because aik-names are not guaranteed to be globally unique, that field is deprecated. MAP Clients MUST NOT create a device identifier with an aik-name.” However, IF-MAP 1.1 permits publishing a device identifier with an aik-name. In practice:

- An IF-MAP 1.1 MAP Client is permitted to publish a device identifier with an aik-name.
- An IF-MAP 2.0 MAP Client is not permitted to publish a device identifier with an aik-name.
- An IF-MAP 2.0 MAP Server should return an InvalidIdentifier errorResult to an IF-MAP 2.0 Client that publishes a device identifier with an aik-name.
- An IF-MAP 2.0 MAP Client can create a search that returns a device identifier with an aik-name that was published by an IF-MAP 1.1 MAP Client.



## 6 Security Considerations

A MAP serves as a metadata clearing house for MAP Clients such as PEPs, PDPs, Flow Controllers, and Sensors, using a publish-subscribe-search model of information exchange and lookup. By increasing the ability of MAP Clients to learn about and respond to security-relevant events and data, IF-MAP can improve the timeliness and utility of the security system. However, this integrated security system can also be exploited by attackers if they can compromise it. Therefore, strong security protections for IF-MAP are essential.

This section provides a security analysis of the IF-MAP protocol and the architectural elements that employ it, specifically with respect to their use of this protocol. Three subsections define the trust model (which elements are trusted to do what), the threat model (attacks that may be mounted on the system), and the countermeasures (ways to address or mitigate the threats previously identified).

### 6.1 Trust Model

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

#### 6.1.1 MAP Server

The MAP Server is trusted to:

- Store data and protect the integrity of this data throughout its lifecycle
- Perform service requests in a timely and accurate manner
- Create and maintain accurate operational attributes
- Resist attacks (including denial of service and other attacks from MAP Clients)
- Only reveal data to and accept service requests from authorized parties

The MAP Server is not expected (trusted) to:

- Verify the truth (correctness) of data

The MAP Server MAY validate data against schema but is not required to do so.

#### 6.1.2 MAP Clients

Authorized MAP Clients are trusted to:

- Preserve the confidentiality of sensitive data retrieved from the MAP Server
- Ensure the accuracy of data in the MAP Server database, by avoiding database corruption and inaccurate data
- Avoid placing too much data on the MAP Server
- Avoid creating too many links on the MAP Server
- Avoid creating too many subscriptions on the MAP Server
- Not delete valuable data from the MAP Server

#### 6.1.3 Network

The network used to carry IF-MAP messages is trusted to:

- Perform best effort delivery of network traffic

The network used to carry IF-MAP messages is not expected (trusted) to:

- Provide confidentiality or integrity protection for messages sent over it
- Provide timely or reliable service

## 6.2 Threat Model

To secure the IF-MAP protocol and the architectural elements that implement it, this section identifies the attacks that can be mounted against the protocol and elements.

### 6.2.1 Network Attacks

A variety of attacks can be mounted using the network. For the purposes of this subsection the phrase “network traffic” should be taken to mean messages and/or parts of messages. 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
- Even if all traffic is encrypted, valuable 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

### 6.2.2 MAP Clients

An unauthorized MAP Client (one which is not recognized by the MAP Server or is recognized but not authorized to perform any actions) cannot mount any attacks other than those listed in the Network Attacks section above.

An authorized MAP Client, on the other hand, can mount many attacks. These attacks might occur because the MAP Client is controlled by a malicious, careless, or incompetent party (whether because its owner is malicious, careless, or incompetent or because the MAP Client has been compromised and is now controlled by a party other than its owner); because the MAP Client is running malicious software; because the MAP Client is running buggy software (which may fail in a state that floods the network with traffic); or because the MAP Client has been configured improperly. From a security standpoint, it generally makes no difference why an attack is initiated. The same countermeasures can be employed in any case.

Here is a list of attacks that may be mounted by an authorized MAP Client:

- Incorrectly create, delete, or modify metadata, perhaps causing network access to be incorrectly blocked or allowed
- Cause many false alarms or otherwise overload the MAP Server or other elements in the network security system (including human administrators) leading to a denial of service or disabling parts of the network security system
- Omit important actions (such as posting incriminating data), resulting in incorrect access

- Use confidential information obtained from the MAP Server to enable further attacks (such as using endpoint health check results to exploit vulnerable endpoints)
- Upload metadata crafted to exploit vulnerabilities in the MAP Server or in other MAP Clients, with a goal of compromising those systems
- Issue a search request or set up a subscription that matches an enormous result, leading to resource exhaustion on the MAP Server and/or the network
- Establish an ARC channel using another client's session-id

Dependencies of or vulnerabilities of authorized MAP Clients may be exploited to effect these attacks. Another way to effect these attacks is to gain the ability to impersonate a MAP Client (through theft of the MAP Client's identity credentials or through other means).

Even a clock skew between the MAP Client and MAP Server can cause problems if the MAP Client assumes that old metadata should be ignored.

### 6.2.3 MAP Servers

An unauthorized MAP Server (one which is not trusted by MAP Clients) cannot mount any attacks other than those listed in the Network Attacks section above.

An authorized MAP Server can mount many attacks. Similar to the MAP Client case described above, these attacks might occur because the MAP Server is controlled by a malicious, careless, or incompetent party (either a MAP Server administrator or an attacker who has seized control of the MAP Server). They might also occur because the MAP Server is running malicious software, because the MAP Server is running buggy software (which may fail in a state that corrupts data or floods the network with traffic), or because the MAP Server has been configured improperly.

All of the attacks listed for MAP Clients above can be mounted by the MAP Server. Detection of these attacks will be more difficult since the MAP Server can create false operational attributes and/or logs that imply some other party created any bad data.

Additional MAP Server attacks may include:

- Expose different database state to different MAP Clients to mislead investigators or cause inconsistent behavior
- Mount an even more effective denial of service attack than a single MAP Client could
- Send results to a MAP Client that claim to have been validated as schema compliant by the server but are not
- Leverage control of the MAP Server to attack other systems (e.g. attack other MAP Servers employed for availability that may be vulnerable to attacks from peer MAP Servers or use privilege escalation to gain control of the machine where the MAP Server is running)
- Obtain and cache MAP Client credentials so they can be used to impersonate MAP Clients even after a breach of the MAP Server is repaired
- Obtain and cache MAP Server administrator credentials so they can be used to regain control of the MAP Server after the breach of the MAP Server is repaired

Dependencies of or vulnerabilities of the MAP Server may be exploited to obtain control of the MAP Server and effect these attacks.

## 6.3 Countermeasures

### 6.3.1 Securing the IF-MAP Protocol

To address network attacks, the IF-MAP binding for SOAP described in this document requires that the IF-MAP protocol **MUST** be carried over TLS ([8], [10], or [18]) as described in [11]. The MAP Client **MUST** verify the MAP Server's certificate and determine whether the MAP Server is trusted by this MAP Client before completing the TLS handshake. The MAP Server **MUST** authenticate the MAP Client either using mutual certificate-based authentication in the TLS handshake or using Basic Authentication as described in [12]. All MAP Servers and MAP Clients **MUST** implement both mutual certificate-based authentication and Basic Authentication. The selection of which client authentication technique to use in any particular deployment is left to the administrator. Since Basic Authentication has many security disadvantages (especially the transmission of reusable client passwords to the server), it **SHOULD** only be used when absolutely necessary. SOAP intermediaries **MUST NOT** be used.

Per the HTTP specification [17], when basic authentication is in use, a MAP server **MAY** respond to any request that lacks credentials with HTTP code 401. A client **MAY** avoid this code by submitting basic auth credentials with every request. If it does not do so, a client **MUST** respond to this code by resubmitting the same request with credentials (unless the client is shutting down).

Upon successful authentication, the trusted client entities **MUST** be verified for authorization to serve the MAP Client role. The means of authorization is out of scope of this specification.

These protocol security measures provide protection against all the network attacks listed in section 6.2.1 except denial of service attacks. If protection against these denial of service attacks is desired, ingress filtering [13], rate limiting per source IP address, and other denial of service mitigation measures [14] may be employed.

### 6.3.2 Securing MAP Clients

MAP Clients (such as branch office firewalls) may be deployed in locations that are susceptible to physical attacks<sup>2</sup>. Physical security measures may be taken to avoid compromise of MAP Clients, but these may not always be practical or completely effective. An alternative measure is to configure the MAP Server to provide read-only access for such systems. MAP Servers **MUST** allow the administrator to configure read-only access for MAP Clients. The MAP Server **SHOULD** also include a full authorization model so that individual clients may be configured to have only the privileges that they need. The MAP Server **MAY** provide functional templates so that the administrator can configure a specific client as a DHCP server and authorize only the operations and metadata types needed by a DHCP server to be permitted for that client. These techniques can reduce the negative impacts of a compromised MAP Client without diminishing the utility of the overall system.

To handle attacks within the bounds of this authorization model, the MAP Server **MAY** also include rate limits and alerts for unusual MAP Client behavior. MAP Servers **SHOULD** make it easy to revoke a MAP Client's authorization when necessary. Another way to detect attacks from MAP Clients is to create fake entries in the IF-MAP database (honeytokens) which normal MAP Clients will not attempt to access. The MAP Server **SHOULD** include auditable logs of client activities.

To avoid content-based attacks, the MAP Server **MAY** validate metadata posted by MAP Clients. However, MAP Servers and MAP Clients **SHOULD** also be robust against malformed data. This is especially important for vendor-specific metadata, which the MAP Server may not be able to validate.

---

<sup>2</sup> Example is a WLAN access point that may have to be placed strategically for radio coverage but in physically ill-secured locations.

To avoid compromise of MAP Clients, MAP Clients SHOULD be hardened against attack and minimized to reduce their attack surface. They MAY include a Trusted Platform Module (TPM) and go through a TNC handshake to verify the integrity of the MAP Client. They should be well managed to minimize vulnerabilities in the underlying platform and in systems upon which the MAP Client depends. Personnel with administrative access should be carefully screened and monitored to detect problems as soon as possible.

To detect clock skew relative to the MAP Server, MAP Clients MAY occasionally add a small amount of harmless metadata and check the operational attributes to see how the server's concept of time differs from the client's. MAP Clients and Servers SHOULD be configurable to synchronize their clocks with an NTP server.

### 6.3.3 Securing MAP Servers

Because of the serious consequences of MAP Server compromise, MAP Servers SHOULD be especially well hardened against attack and minimized to reduce their attack surface. They SHOULD include a Trusted Platform Module (TPM) and go through a regular TNC handshake to verify the integrity of the MAP Server. They should be well managed to minimize vulnerabilities in the underlying platform and in systems upon which the MAP Server depends. Network security measures such as firewalls or intrusion detection systems may be used to monitor and limit traffic to and from the MAP Server. Personnel with administrative access should be carefully screened and monitored to detect problems as soon as possible. Administrators should not use password-based authentication but should instead use non-reusable credentials and multi-factor authentication (where available). Physical security measures SHOULD be employed to prevent physical attacks on MAP Servers.

To ease detection of MAP Server compromise should it occur, MAP Server behavior should be monitored to detect unusual behavior (such as a reboot, a large increase in traffic, or different views of the database for different clients). MAP Clients should log and/or notify administrators when peculiar MAP Server behavior is detected. MAP Clients should also check data sent from the MAP Server carefully to detect malformed data. To aid forensic investigation, permanent read-only audit logs of security-relevant information (especially administrative actions) should be maintained. If MAP Server compromise is detected, a careful analysis should be performed of the impact of this compromise. Any reusable credentials that may have been compromised should be reissued.

#### 6.3.3.1 Limit on search result size

A MAP Server MAY have a limit to the amount of data it is willing to return in search or subscription results (see section 3.7.2.8). This mitigates the threat of a MAP Client causing resource exhaustion by issuing a search or subscription that leads to an enormous result.

#### 6.3.3.2 Cryptographically random session-id and authentication checks for ARC

A MAP Server SHOULD ensure that the client establishing an ARC is the same client as the client that established the corresponding SSRC. The MAP Server SHOULD employ both of the following strategies:

1. session-ids SHOULD be cryptographically random
2. The HTTPS transport for the SSRC and the ARC SHOULD be authenticated using the same credentials. SSL session resumption MAY be used to establish the ARC based on the SSRC SSL session.

## 6.4 Summary

IF-MAP's considerable value as a clearing-house for security-sensitive data exchange distribution also makes the protocol and the network security elements that implement it a target for attack. Therefore, strong security has been included as a basic design principle within the IF-MAP design process.

The IF-MAP protocol provides strong protection against a variety of different attacks. In the event that a MAP Client or MAP Server is compromised, the effects of this compromise have been reduced and limited with the recommended role-based authorization model and other provisions, and best practices for managing and protecting IF-MAP systems have been described. Taken together, these measures should provide protection commensurate with the threat to IF-MAP systems thus ensuring that they fulfill their promise as a network security clearing-house.

## 7 Privacy Considerations

MAP Clients may publish information about endpoint health, network access, events (which may include information about what services an endpoint is accessing), roles and capabilities, and the identity of the end user operating the endpoint. Any of this published information may be queried by other MAP Clients and could potentially be used to correlate network activity to a particular end user.

Dynamic and static information published to a MAP Server, ostensibly for purposes of correlation by Flow Controllers for intrusion detection, could be misused by a broader set of MAP Clients which hitherto have been performing specific roles with strict well-defined separation of duties.

Care should be taken by deployers of IF-MAP to ensure that the information published by MAP Clients does not violate agreements with end users or local and regional laws and regulations. This can be accomplished either by configuring MAP Clients to not publish certain information or by restricting access to sensitive data to trusted MAP Clients<sup>3</sup>.

### 7.1 identity Identifier

The identity identifier may include specific information about an end user's identity, enabling MAP Clients to determine how a particular end user is accessing the network.

### 7.2 mac-address Identifier

It may be possible to determine the identity of an end user by correlation with the MAC address of an endpoint. For example, an employee may be issued a laptop and a company database may store the MAC address of the laptop along with information that identifies the employee. If an association between MAC address and end user is known, then a MAP Client could determine how a particular end user is accessing the network by querying for the known MAC address.

### 7.3 ip-address Identifier

If an endpoint has a static IP address or a dynamic IP address with a very long lease, it may be possible over time to make an association between a particular IP address and a particular end user. In this case, a MAP Client could determine how a particular end user is accessing the network by querying for the known IP address.

---

<sup>3</sup>A MAP Server implementation may provide an authorization model which protects data published by one MAP Client from being visible to another MAP Client. The specifics of such an authorization model are outside the scope of this specification.

## 8 References

- [1] Trusted Computing Group, *TNC Architecture for Interoperability*, Revision 1.4, May 2009
- [2] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, RFC 2119, Best Practices, March 1997, IETF
- [3] W3C, *Extensible Markup Language (XML) 1.1 (Second Edition)*, September 2006
- [4] W3C, *XML Schema Part 0: Primer Second Edition*, October 2007
- [5] W3C, *XML Path Language (XPath) 2.0*, January 2007
- [6] C. Rigney, S. Willens, A. Rubens, W. Simpson, *Remote Authentication Dial In User Service (RADIUS)*, RFC2865, Standards Track, June 2000, IETF
- [7] W3C, *SOAP Version 1.2 Part 1: Messaging Framework (Second Edition)*, April 2007
- [8] T. Dierks, C. Allen, *The TLS Protocol Version 1.0*, RFC 2246, Standards Track, January 1999, IETF
- [9] R. Hinden, S. Deering, *IP Version 6 Addressing Architecture*, RFC 4291, Standards Track, February 2006, IETF
- [10] T. Dierks, E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.1*, RFC 4346, Standards Track, April 2006, IETF
- [11] E. Rescorla, *HTTP Over TLS*, RFC 2818, Informational, May 2000, IETF
- [12] J. Franks, P. Hallam-Baker, J. Hostetler, S. Lawrence, P. Leach, A. Luotonen, L. Stewart, *HTTP Authentication: Basic and Digest Access Authentication*, RFC 2617, Standards Track, June 1999, IETF
- [13] P. Ferguson, D. Senie, *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*, RFC 2827, Best Current Practices, May 2000, IETF
- [14] M. Handley, Ed., E. Rescorla, Ed., *Internet Denial-of-Service Considerations*, RFC 4732, Informational, November 2006, IETF
- [15] Trusted Computing Group, *TNC IF-MAP Metadata for Network Security*, Revision 1.0, May 2010
- [16] P. Leach, M. Mealling, R. Salz, *A Universally Unique Identifier (UUID) URN Namespace*, RFC 4122, July 2005, IETF
- [17] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, *Hypertext Transport Protocol – HTTP/1.1*, RFC 2616, June 1999, IETF
- [18] T. Dierks, E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.2*, RFC 5246, August 2006, IETF
- [19] K. McCloghrie, D. Perkins, J. Schoenwaelder, *Structure of Management Information Version 2 (SMIv2)*, April 1999, IETF
- [20] Trusted Computing Group, *TPM Main Specification Level 2 Version 1.2*, Revision 103, June 2007
- [21] ITU-T, *X.500 Directory Specification*, August 2005
- [22] P. Mockapetris, *Domain names - Implementation and Specification*, RFC 1035, November 1987, IETF
- [23] P. Resnick, *Internet Message Format*, RFC 5322, October 2008, IETF



- [24] Neuman, C., Yu, T., Hartman, S., and K. Raeburn, *The Kerberos Network Authentication Service (V5)*, RFC 4120, July 2005, IETF
- [25] M. J. Handley, H. Schulzrinne, E. M. Schooler, J. Rosenberg, *SIP: Session Initiation Protocol*, RFC 3261, March 1999, IETF
- [26] H. Schulzrinne, *The tel URI for Telephone Numbers*, RFC 3966, December 2004, IETF
- [27] R. Moskowitz, P. Nikander, *Host Identity Protocol (HIP) Architecture*, RFC 4423, May 2006, IETF

## 9 Basic Example

The following example is intended to demonstrate simple publish and search operations of IF-MAP, basic coordination between IF-MAP clients, and the use of vendor specific metadata. This example is not intended to be comprehensive or address a well thought out use-case. For detailed examples and plausible use-cases one should refer to a use-case driven IF-MAP metadata standard such as [15].

### 9.1 Webcam Conferencing

A video conferencing application uses IF-MAP to store and retrieve information about the webcam capabilities of computers. The application publishes metadata defined by the following XML schema:

```
<?xml version="1.0" ?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns="urn:example.com:webcam"
  targetNamespace="urn:example.com:webcam">

  <!-- webcam-capabilities is attached to a device identifier and
    describes the capabilities of a computer's webcam -->
  <xsd:element name="webcam-capabilities">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="enabled" type="xsd:int"/>
        <xsd:element name="video-format" type="xsd:string"/>
      </xsd:sequence>
      <xsd:attributeGroup
        ref="ifmap:singleValueMetadataAttributes"/>
    </xsd:complexType>
  </xsd:element>

  <!-- webcam-user is attached to the link between a device
    identifier and an identity identifier, and is used to
    associate a user with a computer that has a webcam -->
  <xsd:element name="webcam-user">
    <xsd:complexType>
      <xsd:attributeGroup
        ref="ifmap:singleValueMetadataAttributes"/>
    </xsd:complexType>
  </xsd:element>

  <!-- webcam-ip is attached to the link between a device
    identifier and an ip-address identifier, and is used to
    associate an IP address with a computer that has a
    webcam -->
  <xsd:element name="webcam-ip">
    <xsd:complexType>
      <xsd:attributeGroup
        ref="ifmap:singleValueMetadataAttributes"/>
    </xsd:complexType>
  </xsd:element>
```

```
</xsd:schema>
```

When the user Joe starts the video conferencing application, the application publishes information about Joe, the webcam, and the IP address of Joe's computer into IF-MAP. In order to do this, the application generates a device name to be used in a device identifier to represent Joe's computer. Once the application has a device identifier for Joe's computer, it publishes the following metadata:

- webcam-capabilities attached to the device identifier
- webcam-user attached to the link between the device identifier and Joe's identity identifier
- webcam-ip attached to the link between the device identifier and the ip-address identifier that represents the IP address of Joe's computer.

The publish request looks like this:

```
<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:wc="urn:example.com:webcam">
  <env:Body>
    <ifmap:publish session-id="222">
      <update>
        <device><name>XYZ:1234</name></device>
        <metadata>
          <wc:webcam-capabilities
            ifmap-cardinality="singleValue">
            <enabled>1</enabled>
            <video-format>VGA</video-format>
          </wc:webcam-capabilities>
        </metadata>
      </update>
      <update>
        <device><name>222:1234</name></device>
        <identity name="Joe" type="username"/>
        <metadata>
          <wc:webcam-user ifmap-cardinality="singleValue"/>
        </metadata>
      </update>
      <update>
        <device><name>222:1234</name></device>
        <ip-address value="192.0.2.11" type="IPv4"/>
        <metadata>
          <wc:webcam-ip ifmap-cardinality="singleValue"/>
        </metadata>
      </update>
    </ifmap:publish>
  </env:Body>
</env:Envelope>
```

The IF-MAP server responds with a publishReceived message:

```

<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"

  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2">
  <env:Body>
    <ifmap:response>
      <publishReceived/>
    </ifmap:response>
  </env:Body>
</env:Envelope>

```

Sally is also running the video conferencing application, and decides to place a video call to Joe. Sally enters the user name "Joe" into the application. The application performs an IF-MAP search to determine the IP address and webcam capabilities of Joe's computer. The search starts with the identity identifier for Joe, and then follows the webcam-user link to find Joe's device. The search continues by following the webcam-ip link to find the IP address of Joe's computer. Along the way, the search picks up the webcam-capabilities metadata.

The search request looks like this:

```

<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:wc="urn:example.com:webcam">
  <env:Body>
    <ifmap:search session-id="223">
      match-links="wc:webcam-user or wc:webcam-ip"
      max-depth="2" result-filter="wc:webcam-capabilities"
      <identity name="Joe" type="username"/>
    </ifmap:search>
  </env:Body>
</env:Envelope>

```

The IF-MAP server responds with a searchResult message:

```

<?xml version="1.0"?>
<env:Envelope
  xmlns:env="http://www.w3.org/2003/05/soap-envelope"
  xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  xmlns:wc="urn:example.com:webcam">
  <env:Body>
    <ifmap:response>
      <searchResult>
        <resultItem>
          <device><name>222:1234</name></device>
          <metadata>
            <wc:webcam-capabilities
              ifmap-cardinality="singleValue" ifmap-timestamp=""
              ifmap-publisher-id="XYZ">
              <enabled>1</enabled>
              <video-format>VGA</video-format>
            </wc:webcam-capabilities>
          </metadata>
        </resultItem>
      </searchResult>
    </ifmap:response>
  </env:Body>
</env:Envelope>

```

```
</resultItem>
<resultItem>
  <device><name>222:1234</name></device>
  <identity name="Joe" type="username"/>
</resultItem>
<resultItem>
  <device><name>222:1234</name></device>
  <ip-address value="192.0.2.11" type="IPv4"/>
</resultItem>
</searchResult>
</ifmap:response>
</env:Body>
</env:Envelope>
```

The video conferencing application on Sally's computer determines from the webcam-capabilities that a video call is possible. The application uses the IP address returned in the search to contact Joe's computer and start the video call.

## 10 IF-MAP Schema

### 10.1 Identifier Types, Requests and Responses

```
<?xml version="1.0"?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
  targetNamespace="http://www.trustedcomputinggroup.org/2010/IFMAP/
  2">

  <!-- top-level elements represent all the possible
    requests and responses -->
  <xsd:element name="publish" type="PublishRequestType"/>
  <xsd:element name="search" type="SearchRequestType"/>
  <xsd:element name="subscribe" type="SubscribeRequestType"/>
  <xsd:element name="poll" type="PollRequestType"/>
  <xsd:element name="purgePublisher"
  type="PurgePublisherRequestType"/>
  <xsd:element name="newSession" type="NewSessionRequestType"/>
  <xsd:element name="renewSession" type="SessionRequestType"/>
  <xsd:element name="endSession" type="SessionRequestType"/>
  <xsd:element name="response" type="ResponseType"/>

  <!-- AccessRequestType Identifier represents an endpoint
    which is attempting to gain entry to the network-->
  <xsd:complexType name="AccessRequestType">
    <xsd:attribute name="administrative-domain"
  type="xsd:string"/>
    <xsd:attribute name="name" type="xsd:string" use="required"/>
  </xsd:complexType>

  <!-- DeviceType Identifier represents a physical asset
    which is attempting to gain entry to the network -->
  <xsd:complexType name="DeviceType">
    <xsd:choice>
      <xsd:element name="aik-name" type="xsd:string"/>
      <xsd:element name="name" type="xsd:string"/>
    </xsd:choice>
  </xsd:complexType>

  <!-- IdentityType Identifier represents an end-user -->
  <xsd:complexType name="IdentityType">
    <xsd:attribute name="administrative-domain"
  type="xsd:string"/>
    <xsd:attribute name="name" type="xsd:string" use="required"/>
    <xsd:attribute name="type" use="required">
      <xsd:simpleType>
        <xsd:restriction base="xsd:string">
          <xsd:enumeration value="aik-name"/>
          <xsd:enumeration value="distinguished-name"/>
          <xsd:enumeration value="dns-name"/>
          <xsd:enumeration value="email-address"/>
          <xsd:enumeration value="hip-hit"/>
          <xsd:enumeration value="kerberos-principal"/>
          <xsd:enumeration value="username"/>
        </xsd:restriction>
      </xsd:simpleType>
    </xsd:attribute>
  </xsd:complexType>

```

```
        <xsd:enumeration value="sip-uri"/>
        <xsd:enumeration value="tel-uri"/>
        <xsd:enumeration value="other"/>
    </xsd:restriction>
</xsd:simpleType>
</xsd:attribute>
<xsd:attribute name="other-type-definition"
type="xsd:string"/>
</xsd:complexType>

<!-- IPAddressType Identifier represents a single IP address -->
<xsd:complexType name="IPAddressType">
    <xsd:attribute name="administrative-domain"
type="xsd:string"/>
    <xsd:attribute name="value" type="xsd:string"
use="required"/>
    <xsd:attribute name="type">
        <xsd:simpleType>
            <xsd:restriction base="xsd:string">
                <xsd:enumeration value="IPv4"/>
                <xsd:enumeration value="IPv6"/>
            </xsd:restriction>
        </xsd:simpleType>
    </xsd:attribute>
</xsd:complexType>

<!-- MACAddressType Identifier represents an Ethernet MAC
address -->
<xsd:complexType name="MACAddressType">
    <xsd:attribute name="administrative-domain"
type="xsd:string"/>
    <xsd:attribute name="value" type="xsd:string"
use="required"/>
</xsd:complexType>

<!-- MetadataListType is a container for metadata within
other elements -->
<xsd:complexType name="MetadataListType">
    <xsd:sequence>
        <xsd:any minOccurs="0" maxOccurs="unbounded"/>
    </xsd:sequence>
</xsd:complexType>

<!-- FilterType is a subset of XPath -->
<xsd:simpleType name="FilterType">
    <xsd:restriction base="xsd:string"/>
</xsd:simpleType>

<xsd:complexType name="NewSessionRequestType">
    <xsd:attribute name="max-poll-result-size" type="xsd:integer"
use="optional"/>
</xsd:complexType>

<xsd:attributeGroup name="validationAttributes">
    <xsd:attribute name="validation" use="optional">
        <xsd:simpleType>
```

```
<xsd:restriction base="xsd:string">
  <xsd:enumeration value="None"/>
  <xsd:enumeration value="BaseOnly"/>
  <xsd:enumeration value="MetadataOnly"/>
  <xsd:enumeration value="All"/>
</xsd:restriction>
</xsd:simpleType>
</xsd:attribute>
</xsd:attributeGroup>

<xsd:attributeGroup name="sessionAttributes">
  <xsd:attribute name="session-id" type="xsd:string"
use="required"/>
</xsd:attributeGroup>

<!-- SessionRequestType is stateful session handling -->
<xsd:complexType name="SessionRequestType">
  <xsd:attributeGroup ref="sessionAttributes"/>
</xsd:complexType>

<!-- UpdateType is the type for requests that update metadata -
->
<xsd:complexType name="UpdateType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
    <xsd:element name="metadata" type="MetadataListType"
minOccurs="1" maxOccurs="1"/>
  </xsd:sequence>
  <xsd:attribute name="lifetime" default="session">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="session"/>
        <xsd:enumeration value="forever"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:complexType>

<!-- DeleteType is the type for the delete element of
a publish request, and specifies which metadata
to delete. -->
<xsd:complexType name="DeleteType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
  </xsd:sequence>
</xsd:complexType>
```



```

        </xsd:choice>
    </xsd:sequence>
    <xsd:attribute name="filter" type="FilterType"
use="optional"/>
</xsd:complexType>

<!-- PublishRequestType updates or deletes metadata -->
<xsd:complexType name="PublishRequestType">
    <xsd:sequence>
        <xsd:choice minOccurs="1" maxOccurs="unbounded">
            <xsd:element name="update" type="UpdateType"/>
            <xsd:element name="notify" type="UpdateType"/>
            <xsd:element name="delete" type="DeleteType"/>
        </xsd:choice>
    </xsd:sequence>
    <xsd:attributeGroup ref="sessionAttributes"/>
    <xsd:attributeGroup ref="validationAttributes"/>
</xsd:complexType>

<!-- SearchType specifies the parameters for a search, and is
used
for the search element as well as the update sub-element
of a
subscribe element -->
<xsd:complexType name="SearchType">
    <xsd:sequence>
        <xsd:choice minOccurs="1" maxOccurs="1">
            <xsd:element name="access-request"
type="AccessRequestType"/>
            <xsd:element name="identity" type="IdentityType"/>
            <xsd:element name="ip-address" type="IPAddressType"/>
            <xsd:element name="mac-address" type="MACAddressType"/>
            <xsd:element name="device" type="DeviceType"/>
        </xsd:choice>
    </xsd:sequence>
    <xsd:attribute name="match-links" type="FilterType"/>
    <xsd:attribute name="max-depth" type="xsd:unsignedInt"/>
    <xsd:attribute name="terminal-identifier-type"
type="xsd:string"/>
    <xsd:attribute name="max-size" type="xsd:unsignedInt"/>
    <xsd:attribute name="result-filter" type="FilterType"/>
</xsd:complexType>

<!-- SearchRequestType queries the server for matching
metadata -->
<xsd:complexType name="SearchRequestType">
    <xsd:complexContent>
        <xsd:extension base="SearchType">
            <xsd:attributeGroup ref="sessionAttributes"/>
            <xsd:attributeGroup ref="validationAttributes"/>
        </xsd:extension>
    </xsd:complexContent>
</xsd:complexType>

<!-- DeleteSearchRequestType is for removing subscriptions -->
<xsd:complexType name="DeleteSearchRequestType">
    <xsd:attribute name="name" type="xsd:string" use="required"/>

```

```

</xsd:complexType>

<!-- SubscribeRequestType is for managing subscriptions -->
<xsd:complexType name="SubscribeRequestType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="unbounded">
      <xsd:element name="update">
        <xsd:complexType>
          <xsd:complexContent>
            <xsd:extension base="SearchType">
              <xsd:attribute name="name" type="xsd:string"
use="required"/>
            </xsd:extension>
          </xsd:complexContent>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="delete"
type="DeleteSearchRequestType"/>
    </xsd:choice>
  </xsd:sequence>
  <xsd:attributeGroup ref="sessionAttributes"/>
  <xsd:attributeGroup ref="validationAttributes"/>
</xsd:complexType>

<!-- PollRequestType is for polling for notification of
  metadata changes that match subscriptions -->
<xsd:complexType name="PollRequestType">
  <xsd:attributeGroup ref="validationAttributes"/>
  <xsd:attributeGroup ref="sessionAttributes"/>
</xsd:complexType>

<!-- PurgePublisherRequestType is for removing all metadata
  published by a particular publisher -->
<xsd:complexType name="PurgePublisherRequestType">
  <xsd:attribute name="ifmap-publisher-id" type="xsd:string"/>
  <xsd:attributeGroup ref="sessionAttributes"/>
</xsd:complexType>

<!-- ResultItemType is for search or poll results showing
  metadata attached to identifiers and links-->
<xsd:complexType name="ResultItemType">
  <xsd:sequence>
    <xsd:choice minOccurs="1" maxOccurs="2">
      <xsd:element name="access-request"
type="AccessRequestType"/>
      <xsd:element name="identity" type="IdentityType"/>
      <xsd:element name="ip-address" type="IPAddressType"/>
      <xsd:element name="mac-address" type="MACAddressType"/>
      <xsd:element name="device" type="DeviceType"/>
    </xsd:choice>
    <xsd:element name="metadata" type="MetadataListType"
minOccurs="0" maxOccurs="1"/>
  </xsd:sequence>
</xsd:complexType>

<!-- SearchResultType contains the identifiers and links

```

```

    along with associated metadata -->
    <xsd:complexType name="SearchResultType">
      <xsd:sequence>
        <xsd:element name="resultItem" type="ResultItemType"
minOccurs="0" maxOccurs="unbounded"/>
      </xsd:sequence>
      <xsd:attribute name="name"/>
    </xsd:complexType>

    <!-- PollResultType contains a searchResult for each
    subscription that had changes since the last poll -->
    <xsd:complexType name="PollResultType">
      <xsd:sequence>
        <xsd:choice minOccurs="0" maxOccurs="unbounded">
          <xsd:element name="searchResult"
type="SearchResultType"/>
          <xsd:element name="updateResult"
type="SearchResultType"/>
          <xsd:element name="deleteResult"
type="SearchResultType"/>
          <xsd:element name="notifyResult"
type="SearchResultType"/>
          <xsd:element name="errorResult" type="ErrorResultType"/>
        </xsd:choice>
      </xsd:sequence>
    </xsd:complexType>

    <!-- ErrorResultType indicates the cause of an error -->
    <xsd:complexType name="ErrorResultType">
      <xsd:sequence>
        <xsd:element name="errorString" type="xsd:string"/>
      </xsd:sequence>
      <xsd:attribute name="errorCode" use="required">
        <xsd:simpleType>
          <xsd:restriction base="xsd:string">
            <xsd:enumeration value="AccessDenied"/>
            <xsd:enumeration value="Failure"/>
            <xsd:enumeration value="InvalidIdentifier"/>
            <xsd:enumeration value="InvalidIdentifierType"/>
            <xsd:enumeration value="IdentifierTooLong"/>
            <xsd:enumeration value="InvalidMetadata"/>
            <xsd:enumeration value="InvalidSchemaVersion"/>
            <xsd:enumeration value="InvalidSessionID"/>
            <xsd:enumeration value="MetadataTooLong"/>
            <xsd:enumeration value="SearchResultsTooBig"/>
            <xsd:enumeration value="PollResultsTooBig"/>
            <xsd:enumeration value="SystemError"/>
          </xsd:restriction>
        </xsd:simpleType>
      </xsd:attribute>
      <xsd:attribute name="name"/>
    </xsd:complexType>

    <!-- SessionResultType is for stateful session handling -->
    <xsd:complexType name="SessionResultType">
      <xsd:attribute name="session-id" type="xsd:string"
use="required"/>

```

```

    <xsd:attribute name="ifmap-publisher-id" type="xsd:string"
use="required"/>
  </xsd:complexType>

  <xsd:complexType name="NewSessionResultType">
    <xsd:attribute name="session-id" type="xsd:string"
use="required"/>
    <xsd:attribute name="ifmap-publisher-id" type="xsd:string"
use="required"/>
    <xsd:attribute name="max-poll-result-size" type="xsd:integer"
use="optional"/>
  </xsd:complexType>

  <!-- ResponseType encapsulates results from all the different
requests -->
  <xsd:complexType name="ResponseType">
    <xsd:choice>
      <xsd:element name="errorResult" type="ErrorResultType"/>
      <xsd:element name="pollResult" type="PollResultType"/>
      <xsd:element name="searchResult" type="SearchResultType"/>
      <xsd:element name="subscribeReceived"/>
      <xsd:element name="publishReceived"/>
      <xsd:element name="purgePublisherReceived"/>
      <xsd:element name="newSessionResult"
type="NewSessionResultType"/>
      <xsd:element name="renewSessionResult"/>
      <xsd:element name="endSessionResult"/>
    </xsd:choice>
    <xsd:attributeGroup ref="validationAttributes"/>
  </xsd:complexType>

  <!-- metadataAttributes specifies attributes on metadata which
are used by MAP servers.

  ifmap-publisher-id and ifmap-timestamp are added to all
metadata by the server before storage in the database.
MAP clients MUST NOT include ifmap-publisher-id or
ifmap-timestamp in published metadata.

  cardinality is used by the MAP client to indicate to
the server whether the metadata can have multiple values.

  anyAttribute enables metadata elements which include the
ifmap:metadataAttributes attributeGroup to add new
attributes for use with future versions of IF-MAP. -->

  <xsd:attributeGroup name="metadataAttributes">
    <xsd:attribute name="ifmap-publisher-id"/>
    <xsd:attribute name="ifmap-timestamp" type="xsd:dateTime"/>
    <xsd:anyAttribute/>
  </xsd:attributeGroup>

  <xsd:attributeGroup name="singleValueMetadataAttributes">
    <xsd:attributeGroup ref="metadataAttributes"/>
    <xsd:attribute name="ifmap-cardinality" use="required">
      <xsd:simpleType>
        <xsd:restriction base="xsd:string">

```

```
        <xsd:enumeration value="singleValue"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:attributeGroup>

<xsd:attributeGroup name="multiValueMetadataAttributes">
  <xsd:attributeGroup ref="metadataAttributes"/>
  <xsd:attribute name="ifmap-cardinality" use="required">
    <xsd:simpleType>
      <xsd:restriction base="xsd:string">
        <xsd:enumeration value="multiValue"/>
      </xsd:restriction>
    </xsd:simpleType>
  </xsd:attribute>
</xsd:attributeGroup>
</xsd:schema>
```