TCG Trusted Network Communications

IF-MAP Metadata for ICS Security

Specification Version 1.0

Revision 46

15 September 2014

Published

Contact:

admin@trustedcomputinggroup.org

TCG PUBLISHED

TCG

Disclaimers, Notices, and License Terms

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.

Contact the Trusted Computing Group at 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

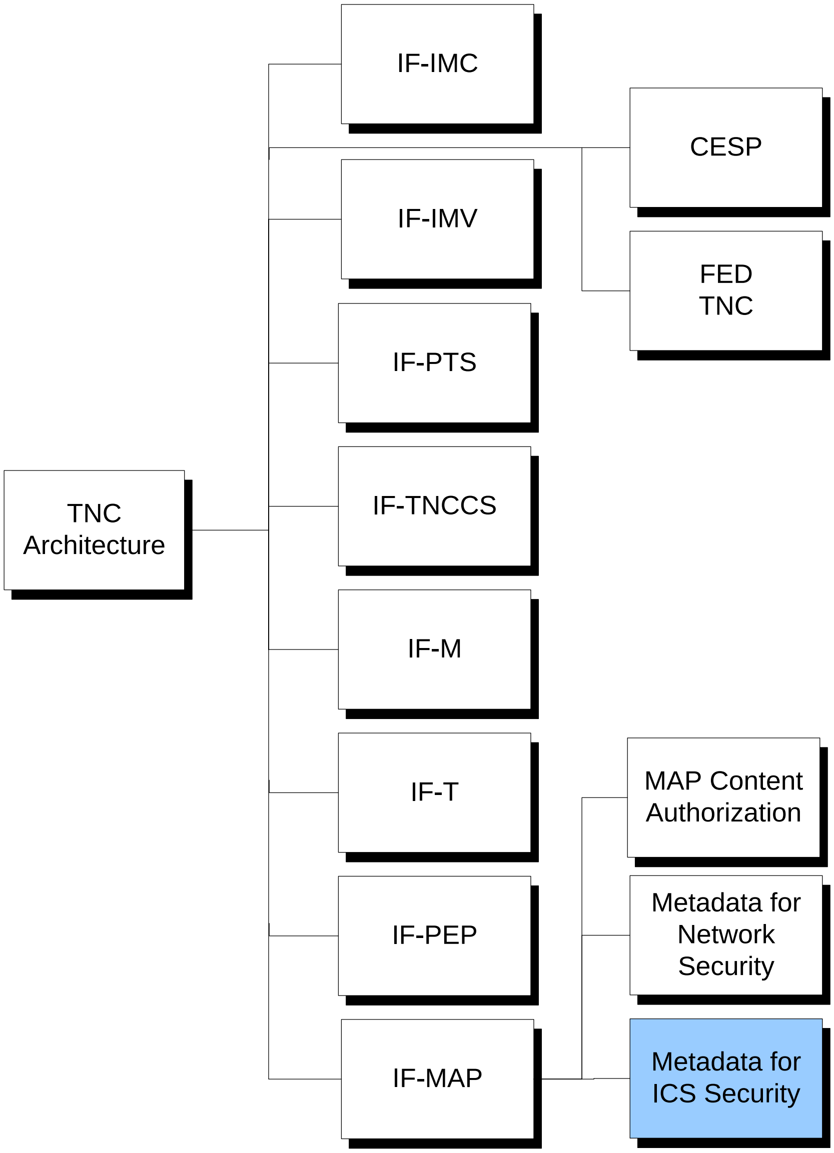
IF-IMC
CESP
IF-IMV
FED TNC
IF-PTS
TNC Architecture
IF-TNCCS
IF-M
MAP Content Authorization
IF-T
Metadata for Network Security
IF-PEP
IF-MAP
Metadata for ICS Security

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.

Frédéric Guihéry	Amossys
David Louin	Amossys
Maxime Olivier	Amossys
David Mattes (Co-Editor)	Asguard Networks
Padma Krishnaswamy	Batelle Memorial Institute
Eric Fleischman	Boeing
Richard Hill	Boeing
Steve Venema (Co-Editor)	Boeing
John Tolbert	Boeing
Steven Venema	Boeing
Eric Byres (Invited Expert)	Byres Security
Nancy Cam-Winget	Cisco Systems
Josef von Helden	Fachhochschule Hannover
Arne Welzel	Fachhochschule Hannover
Ronald Marx	Fraunhofer SIT
Gerald Maunier	Gemalto N.V.
Chris Daly	General Dynamics C4 Systems
Graeme Proudler	Hewlett-Packard
Andreas Steffen	HSR University of Applied Sciences Rapperswil
James Tan	Infoblox
David Vigier	Infoblox
Cliff Kahn	Juniper Networks
Lisa Lorenzin (Co-Editor)	Juniper Networks
Steve Hanna	Juniper Networks
Mark Labbancz	Lumeta Corporation
Atul Shah	Microsoft
Trevor Freeman	Microsoft
Charles Schmidt	MITRE
Jim Banoczi	US Government

Chris Bean	US Government
Julie Haney	US Government
Jessica Fitzgerald-McKay	US Government
Gloria Serrao	US Government
Theresa Thomas	US Government
Josef Allen	Oak Ridge National Lab
Ira McDonald	Samsung
Carolin Latze	SwissCom
Paul Sangster	Symantec
Anne-Rose Gratadour	Thales

Table of Contents

1 Introduction

1.1 Scope and Audience

The Trusted Network Communications Working Group (TNC-WG) has defined an open solution architecture that enables security coordination between components of network-connected industrial control systems. 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 Metadata for Industrial Control Systems (**ICS**) Security.

Architects, designers, developers and technologists who wish to implement, use, or understand IF-MAP in an Industrial Control Systems Security context should read this document carefully. Before reading this document any further, the reader should review and understand the TNC architecture as described in the TNC Architecture specification[1] and the TNC IF-MAP Binding for SOAP[2] and TNC IF-MAP Metadata for Network Security[3].

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 IETF RFC 2119[4]. This specification does not distinguish blocks of informative comments and normative requirements. Therefore, for the sake of clarity, note that lower case instances of must, should, etc. do not indicate normative requirements.

2 Background

2.1 Purpose of IF-MAP Metadata for ICS Security

The purpose of the IF-MAP Metadata for ICS Security specification is to facilitate secure deployment and management of large-scale industrial control systems by creating virtual OSI layer 2 and/or layer 3 overlay networks on top of standard shared IP network infrastructure - particularly (though not necessarily) TNC-compliant IP network infrastructure. This specification addresses use cases for identity-based access policy, device identity, and certificate lifecycle management; it anticipates use in retrofitting this new security functionality into existing industrial control systems as well as incorporation into new ICS products. The use cases and functionality standardized in this specification can operate independently of, but are intended to interoperate with, the existing IF-MAP Metadata for Network Security specification and broader TNC and TCG enabled capabilities for increased security and flexibility.

2.2 Context

2.2.1 Industrial Control Systems and Overlay Networks

Industrial control systems are used to operate and manage complex physical processes. They generally contain a combination of:

- *Sensors* for measuring parameters related to physical processes (e.g., temperature, flow rate, pressure, volume, etc.).
- *Actuators* (e.g., valves, motors, etc.) which affect physical processes.
- *Controllers* (e.g. programmable logic controllers (PLCs)) which contain program logic that adjusts actuator controls based on sensor data and supervisory input.
- *Supervisory systems* (e.g. Supervisory Control And Data Acquisition (SCADA) systems) that monitor and log control system data for trend analysis and as-needed diagnostics and intervention.

While some control systems are completely localized in one physical location (e.g., an automated work cell in a manufacturing plant), many of these systems have physically disjoint components that require interconnection over some kind of network infrastructure. For example, the geographically distributed physical processes in oil refineries, oil pipelines, energy transmission, etc. will typically have sensors (pressure, flow, voltage, temperature, etc.) distributed over a wide area that are not co-located with actuators (valves, switches, etc.). Moreover, the supervisory data for even single-locale control systems are typically carried to distant control centers over network infrastructure.

Communications between control system components use a variety of proprietary and standard protocols, which are often implemented on top of standard layer 2 link protocols (IEEE 802.3, IEEE 802.11, IEEE RS485, etc.). Unfortunately, many of these communications protocols were never designed for security; the trend towards operation in shared, untrusted communications media environments interconnecting physically disjoint control system components is a significant security issue for the control systems industry. Furthermore, some of the control system communications protocols do not use standard Internet Protocols (IPv4 or IPv6), making interconnection of control system components over standard IP network infrastructure difficult. The interconnection of networks not natively designed for an IP stack poses particular challenges, especially around authorization and privacy/confidentiality, when migration to an IP-based infrastructure is introduced.

Overlay networks provide an important layer of protection for control systems by isolating ICS components into one or more protected virtual enclaves while allowing those components to leverage the connectivity offered by increasingly ubiquitous shared, untrusted commodity IP infrastructure. This specification describes how to utilize TNC architectural components - IF-MAP in particular - to define and coordinate the operation of virtual overlay networks in order to create effectively-isolated layer 2 and layer 3/4 network enclaves that communicate securely over shared wide-area untrusted IP networks. In the control systems industry, these commodity network substrates are often called “Backhaul Networks” to denote their function of providing connectivity between ICS components over some distance. The backhaul network may be owned by the end users (e.g., an Ethernet or Wi-Fi LAN), or by a third party network provider (e.g., a satellite communications service provider). Some backhauls may be combinations of both. A backhaul network is generally a single IP routing domain, so all devices connected to a given backhaul network can talk with each other directly at layer 3.

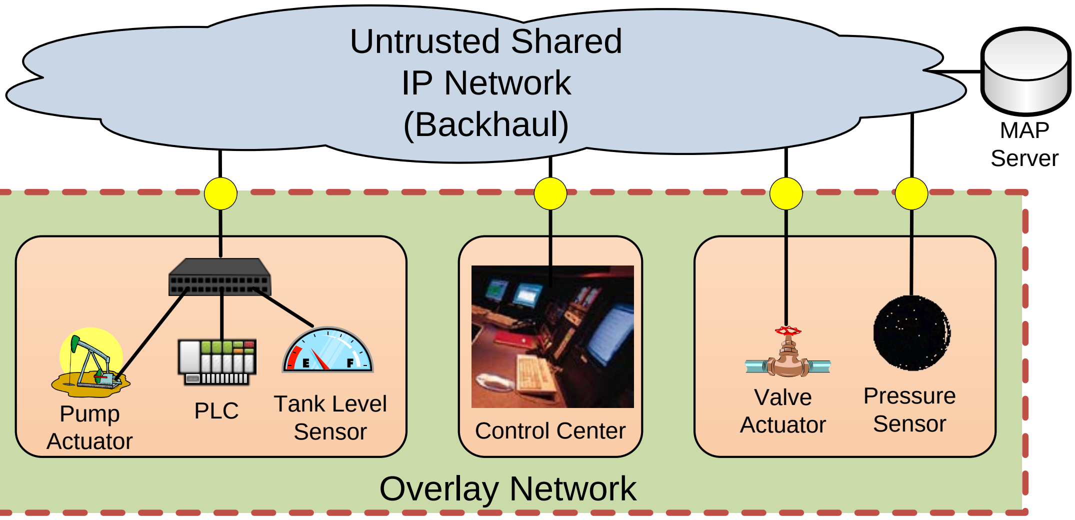

Figure 1: Example of an ICS Overlay Network

Figure 1 shows a notional example of an ICS overlay network. Here we see three clusters of ICS devices (inner boxes), in separate physical locations, that are all part of a single overlay network (outer box). The yellow circles in the diagram denote “Backhaul Interface” (BHI) functionality that implements the necessary authentication, encryption, translation, and authorization policy enforcement capabilities to create the overlay network. This overlay network is a virtual construct whereby ICS components in different ICS device clusters communicate with one another (subject to policy controls expressed in the MAP graph) as though they were on a shared, private layer 2 or layer 3 network. Any ICS device communications that flow between BHIs over the backhaul network are protected by encryption and integrity checks so that other users of the backhaul network are unable to view or modify the ICS data. In effect, the four BHIs shown in Figure 1 form a meshed, private virtual network with MAP-coordinated authorization policies.

Simply put, BHIs communicate with each other and with the MAP Server (operated by the end user) over the backhaul network; ICS devices communicate with each other over a given overlay network, as coordinated and controlled by the BHIs. BHIs utilize two types of encrypted communication: IF-MAP communication with the MAP Server over HTTPS/TLS, and VPLS tunnels from BHI to BHI over HIP. The BHI is intended to restrict which traffic is allowed onto the overlay network, and which traffic originating from its protected overlay network segment is allowed outbound to devices on the overlay that are protected by other BHIs. This specification supports the use of multiple overlay networks all using the same backhaul network while requiring no configuration changes to the backhaul network. Instead, changes to the MAP graph can immediately (re)define the structure and connectivity policies of individual overlay networks and their ICS components.

In the above example, the MAP service provides the necessary operational coordination functionality for these BHIs to provide the overlay network functionality. This includes:

- Coordination between the BHIs themselves (current IP addresses, identity, certificates, etc.)
- Administrative policy that defines with which BHIs a given BHI can communicate
- BHI overlay policies controlling which ICS devices the BHI allows to communicate across the overlay
- Administration policies related to the user identities that are allowed to access and alter the configuration of the overlay network and the BHIs

In the case of control systems, device lifecycles are often measured in decades, and system downtime may not occur for as much as a year or more. During that time, configuration of the underlying IP network may well change; indeed, in the case of wireless and/or mobile control system components, much of this coordination information regarding the backhaul network interfaces may vary from one minute to the next due to roaming-triggered address changes. One advantage that the MAP service has to offer over other directory services such as LDAP is lower-latency (a few seconds), low-bandwidth, asynchronous updates of subscribed state changes in coordination data. In addition, if the metadata associated with this specification shares a MAP graph with specified metadata from other use cases (e.g., TNC IF-MAP Metadata for Network Security[3]), then new opportunities for combined use cases become possible. However, such extended uses cases are outside of the scope of this specification.

This specification anticipates that different implementations of the BHI functionality may take different forms. In general, these implementations will fall into the following three categories:

- BHI functionality embedded within ICS devices.
- BHI functionality embedded within the points-of-presence supplied by the backhaul network provider.
- BHI functionality embedded in a security gateway device between the backhaul network and the ICS devices on the overlay network. In this situation, the gateway device is often called an “Endbox”.

These categories of implementations will generally be supplied by ICS device vendors, backhaul network providers, and third-party security appliance providers respectively. However, in order to ensure interoperability, all BHI implementations MUST comply with this specification. The intent of this specification is to support implementations that employ any combination of these categories (and perhaps some we haven’t yet considered); interoperability between different implementations is the focus of this specification.

2.2.2 Related Standards Work

From a TNC perspective, control systems components that lack a TNC Client (such as sensors, actuators, SCADA systems, controllers, etc.) are all clientless endpoints. The traditional TNC use cases have focused primarily on protecting the network against intrusion or other inappropriate client behavior or configuration. For control systems, allowed behaviors must often be defined very rigorously in order to meet operational reliability and security requirements. For example, there is generally no reason why a view-only operator station should be sending programming data to a controller. Additionally, more granular control over which kinds of commands and other data can be sent between particular components is often desirable. These needs, along with the necessity to ensure integrity, authenticity, and availability in SCADA environments, drive a general requirement for a strong authentication, authorization, and accounting (AAA) environment in order to have correct policy enforcement.

While the specific implementation details of AAA applications are outside the scope of this document (except perhaps as examples), a clear definition of metadata and its use in coordinating between control system components and their network access mediators (the BHIs) is a valuable contribution to this problem domain.

It is also important to note other related public standards activities that may overlap with the scope of this specification. The following three standards activities are mentioned for reference only. It is the intent of this specification to align with - but remain independent of - these other standards activities.

2.2.2.1 International Society for Automation Work Group 15: “Wireless Backhaul”

The nomenclature used in the industrial control systems community may be somewhat unfamiliar to the reader who is more familiar with other elements of the TNC standards family. Figure 2 shows the high level architecture used by the ISA100.15 WG[1]. The definitions for the labels used in Figure 2 are found in Table 1.

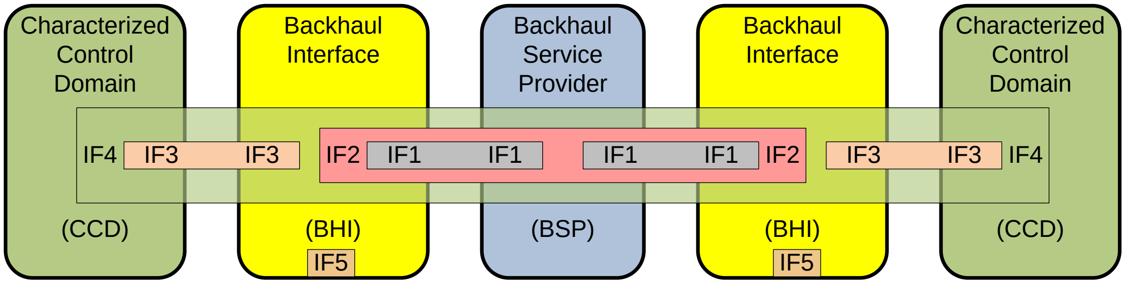

Figure 2: ISA100.15 Functional Architecture

Table 1: Definitions for labels used in Figure 2

Acronym	Description
BSP	Backhaul Service Provider
BHI	Backhaul Interface
CCD	Characterized Control Domain
IF1	Interface for the specific requirements of a particular BSP technology/provider
IF2	BHI interoperability and coordination interface
IF3	Multiple layer 2 (e.g., 802.3) interfaces for control systems data links
IF4	Transparent connectivity interface between CCD's; no interpretation of application protocols
IF5	Configuration, security and operation management interface

[1] https://www.isa.org/isa100/

From a TNC perspective, the BHI functionality serves as a Flow Controller (see the TNC Architecture[1]) providing access control functionality that protects the ICS equipment. This functionality may be implemented within individual networked ICS devices, provided as a part of the customer site interface for a BSP, or implemented as a gateway device in front of one or more ICS devices (e.g., “Endbox”). The BHI may also include TNC client functionality in the case where the BSP uses TNC services. As a result, the BHI is required to be a MAP Client, and may also be - but is not required to be - a TNC Client. The implementation described in this specification is consistent with the above abstract architecture, but uses somewhat different terminology. It is important to relate both sets of terminology, since it would otherwise be difficult to understand the highly relevant relationships between this specification and those of the ISA (as well as the IETF and The Open Group, as described below). Table 2 describes the relationships between these terminologies.

Table 2: Equivalence of terms in ISA and TCG

ISA Term	TCG Term	Description
Backhaul Network	**Backhaul Network**	A shared, untrusted commodity IP network providing communication between ICS devices, possibly in separate geographic locations (e.g. over Ethernet, WLAN, WAN links / MPLS, satellite links, etc.)
IF4 Connections	**Overlay Network**	A logically private network overlaid on top of the Backhaul Network. Multiple Overlay Networks can co-exist on the Backhaul Network.
Characterized Control Domain (CCD)	**Overlay Network Segment**	A group of one or more control system devices that are part of the same control system application in a (generally) single geographic locale. Multiple Overlay Network Segments are interconnected into a single Overlay Network over a unified Backhaul Network.
Backhaul Interface (BHI)	**Backhaul Interface (BHI)**	A device which is compliant with this specification and can act as a gateway between a Backhaul Network and an Overlay Network, participating in the creation and control of an Overlay Network. This may be a compliant ICS device, or a proxy for one or more non-compliant ICS devices, or both.
IF5 Connections	**Overlay Manager**	One or more entities that manage the configuration of the relevant aspects of the Backhaul and Overlay networks, including communication policies for ICS devices.

2.2.2.2 IETF Host Identity Protocol (HIP), Virtual Private LAN Service (VPLS)

The Host Identity Protocol is an emerging IETF standard that provides encrypted IPsec-like tunnels between endpoints that are bound to cryptographic host identities rather than IP addresses. This work is being done in the IETF HIP workgroup; the reader is invited to review various RFCs that have been and are being developed within this workgroup.[2]

A Virtual Private LAN Service (VPLS) Internet Draft[6] has been defined in the HIP workgroup. The VPLS creates private overlay networks using the HIP protocol suite such that communications patterns and roles map well to the ISA functional architecture shown in Figure 2 above. An open-source implementation of this functionality has been implemented[3] and is already in extensive use in at least one large company. For more details on HIP-based VPLS, see [7] and [8].

[2] http://datatracker.ietf.org/wg/hip/

[3] http://www.openhip.org

2.2.2.3 The Open Group's Security Forum: Secure Mobile Architecture (SMA)

The patterns defined in the ISA architecture above are not necessarily unique to the ICS use cases in this specification. Other use cases involving the transportation of trusted data between dynamic nodes over untrusted dynamic networks are common. The Open Group[4] is developing a Secure Mobile Architecture document[5] currently published as a reference architecture by that organization. The document includes ICS use cases similar to those in this document, but also focuses on a wider variety of use cases around meshed communications between moving transportation systems (something that is quickly appearing in automotive products today), medical systems (networks of medical devices exchanging information in, say, an operating theater), and cloud computing (networks of networks containing dynamically provisioned computing services).

The reader is invited to contact The Open Group for more information on this work and participate in its development if there is sufficient interest.

2.3 Lexicon

The following terms have specific meanings throughout the remainder of this document.

Word	Definition
Entity	A user, computing system, asset, application, or process
Actor	An entity available to perform a role
Agent	An entity with delegated authority to act on behalf of another entity
Principal	An entity whose identity can be authenticated
Group	A named collection of entities which share a particular set of attributes or roles
Role	An action or functional behavior that may be performed by one or more actors

2.4 Supported Use Cases

Use cases that this version of IF-MAP Metadata for ICS Security supports:

- A BHI acts as an enforcement point, enforcing both backhaul and overlay connectivity policies which can change dynamically.
- BHIs use metadata published by IF-MAP Sensors (e.g., physical security sensors, IDS/IPS's, PDP's, etc.) to determine what ICS Authorization Policies (overlay policies and backhaul policies - see section 3.1.4) to apply in reaction to changes in the environment.
- A set of BHIs whose IP addresses may change dynamically can locate each other and achieve secure communications despite address changes.

[4] http://www.theopengroup.org

[5] https://www2.opengroup.org/ogsys/catalog/C142

- An overlay management function can dynamically alter the authorization policies for BHIs and the ICS devices on their overlay networks, and BHIs will quickly respond to the changes.

The contents of this document are intended to span these use cases but are not intended to be limited to these use cases.

2.5 Requirements

The following are the requirements which IF-MAP Metadata for ICS Security 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.

2.5.1 General Requirements

1. **Extends TNC IF-MAP Metadata for Network Security**

   As much as makes sense, the IF-MAP Metadata for ICS Security specification will coordinate with the metadata types specified in [3] and [14]. Additional metadata types will be defined in a new namespace to preclude confusion with existing applications which are unaware of the ICS use cases and this specification.

2. **Easy to use and implement**

   IF-MAP Metadata for ICS Security should be easy for vendors and customers to use. It should allow them to enhance existing products to support the TNC architecture and integrate with legacy functionality and products without requiring substantial changes. It should allow customers to integrate the functionality described in this specification into existing infrastructure and allow interoperation with other compliant products.

3. **Unambiguous**

   There should be clarity and lack of ambiguity for identification of specific entities for which metadata exists and which are interacting with the MAP Server. For example, users, endpoints, and all other instances of TNC elements should be uniquely identifiable.

2.5.2 ICS Domain-Specific Network Security Requirements

4. **Usable with existing ICS devices and infrastructure**

   ICS infrastructure has a very long lifecycle (decades), so a mix of legacy and new ICS technologies is common. To the extent possible, IF-MAP Metadata for ICS Security must support both new and legacy ICS technologies on the same network. In some cases this requirement may result in the use of proxy BHIs in front of the legacy devices.

5. **Enables isolation between ICS security domains and/or individual ICS components**

   IF-MAP Metadata for ICS Security must enable multiple control systems to share a (potentially untrusted) backhaul network while providing isolation between control systems or subsystems as defined by the ICS administrators. The specification must include any necessary normative language to enable coordination between multiple BHIs as well as administration applications to provide for and manage this isolation capability.

6. **Enables the use of cryptographically bound identities for ICS devices and PEP's**

   IF-MAP Metadata for ICS Security must facilitate the interoperable use of cryptographic identities in the form of X.509 certificates. Since BHIs are often used in hazardous environments or difficult-to-reach locations, a standard model for MAP-based automated certificate management would be useful. Such a model is currently unspecified, but may be specified in a future TCG specification.

7. **Enables the use of overlay networks**

   The lack of support for dynamic IP address allocation (e.g. via dynamic host configuration protocol (DHCP)) in many ICS devices, coupled with their legacy of being interconnected on isolated networks, drives the need for groups of ICS devices to be isolated in IP address space. For example, it is not uncommon for multiple sets of identical manufacturing systems in a factory to all use the same set of IP addresses. Layer 2 and layer 3 isolation techniques are commonly used to accommodate this limitation. This requirement is supported by the use of metadata for rendezvous among compliant BHIs.

8. **Accommodates the creation of extended metadata definitions for deep packet inspection**

   The IF-MAP Metadata for ICS Security specification must not preclude the future creation of extended metadata types for particular ICS protocols (e.g., Profinet™, Fieldbus™, etc.) that allow ICS protocol-specific deep packet inspection capabilities to be implemented at BHIs or elsewhere.

2.6 Assumptions

Here are the assumptions that IF-MAP Metadata for ICS Security makes about other components in the TNC architecture and ICS environment.

1. **Operational access control**

   In order to increase operational reliability and security, a MAP Server MUST support access controls to constrain the IF-MAP operations a given MAP Client may use to manipulate certain types or instances of metadata. In other words, this is not a trusted environment where all clients which have MAP Server access credentials are trusted to do any IF-MAP operation.

2. **Self-Provisioning Backhaul Interfaces**

   A BHI is expected to be able to configure itself upon initial installation into a backhaul network: acquire a network address, locate an appropriate MAP Server, and obtain information about itself that may be pre-provisioned into the MAP Server. Methods of autoconfiguration are currently unspecified, but may be specified in a future TCG specification. This implies the availability of DHCP and DNS services provided by the backhaul network to enable the BHI to discover its IP network parameters and the IP address of DNS servers.

3 IF-MAP Metadata for ICS Security

This specification defines new IF-MAP extended identifiers and metadata supporting authentication, authorization, and coordination of components in an ICS environment. See section 3.4 of the TNC IF-MAP Binding for SOAP[2] for details on construction and usage of these elements.

Subsequent versions of this specification or other TNC specifications may define new XML attributes for IF-MAP elements. To anticipate this, the schema defined by this specification includes the XML element “<anyAttribute/>” within an element’s declaration. 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 other TNC specifications. Any unrecognized attributes beginning with the ifmap- prefix MUST be ignored by MAP Clients and MAP Servers.

3.1 Authentication and Authorization

3.1.1 MAP Client Roles

There are two MAP Client roles involved in the use cases specified in this document. Authorization constraints associated with these roles are discussed separately in 3.1.4.

3.1.1.1 BHI Role

A BHI utilizes IF-MAP client functionality in order to publish, search, and subscribe to information in the MAP graph. BHIs publish information about their configuration and state, and coordinate with other BHIs in this role as well as other roles (described below). Each BHI serves as its own agent and is expected to have its own unique identity and associated certificate.

In TNC MAP Content Authorization[14], a BHI is assigned *BHI Tasks*, which are associated, one-to-one, with overlay networks.

3.1.1.2 ICS Administrative MAP Client Role

An ICS Administrative MAP Client is an agent that acts on behalf of the principal roles.

3.1.2 Principal Roles

Two other roles are involved in the use cases specified in this document. The principals that hold these roles are humans, applications, etc. These principals act on MAP content indirectly, by giving orders to an ICS Administrative MAP Client. Authorization constraints associated with these roles are discussed separately in 3.1.4.

3.1.2.1 Overlay Manager Role

An Overlay Manager role manages the following relationships expressed in the MAP graph for one or more overlay networks:

- **Overlay Membership:** which BHIs are members of a given overlay network
- **BHI-to-BHI policy (backhaul policy):** which BHI pairs may communicate for a given overlay network; see section 3.5.1
- **Device-to-device policy (overlay policy):** which ICS devices (specified by IP address and/or MAC address) may send communications out through a given BHI for a given overlay network; see section 3.5.12

Unlike the BHIs, which serve as their own agents, IF-MAP operations related to the Overlay Manager role will be performed by an ICS Administrative MAP Client - a software agent (e.g., provisioning application) - on behalf of one or more principals (users). Multiple Overlay Manager roles may be defined so that different groups of one or more overlays may be managed by different user groups. See section 4.2.2 for constraints on Overlay Managers.

3.1.2.2 Administrator Role

The Administrator role manages the following relationships expressed in the MAP graph for all overlay networks:

- Overlay network coordination; creation of Overlay Manager tasks for overlay networks.
- Assignment of particular principals as managers of particular overlay networks.
- MAP graph monitoring and repair to detect and correct unusual conditions (e.g., orphaned overlay networks)

This role is defined separately from the Overlay Manager role because overlay networks are group objects in a given MAP graph with a global namespace; in other words, the names of all of the overlay networks live in a shared namespace in the graph. This namespace is best managed separately (e.g., at an enterprise level) by an Administrator role that oversees the Overlay Manager assignments and overlay network coordination, in order to prevent Overlay Manager roles from potentially interfering with each other through namespace collisions, renaming of overlay networks, etc. For example, if an organization has a thousand overlays represented in the MAP graph, all sharing the same namespace, coordination is required to ensure that Overlay Manager A doesn’t delete an Overlay Network 1 that was created by Overlay Manager B, or vice versa.

Also, a large organization may have different users / groups of users (perhaps even using different tools) managing different overlay networks / groups of overlay networks. The Administrator role performs the function of delegation of responsibility for different overlay networks to different Overlay Managers. Coordination between different overlay manager groups is the duty of the Administrator; members of an overlay manager group are expected to coordinate among themselves.

Similar to the Overlay Manager role, IF-MAP operations related to the Administrator role will be performed by an ICS Administrative MAP Client on behalf of one or more principals (users).

3.1.2.3 Indirect Action of Overlay Managers and Administrators

Throughout this specification, any statement that an Overlay Manager or Administrator acts upon (attaches, publishes, deletes, etc.) metadata should be understood to refer to an ICS Administrative MAP Client performing the operation on behalf of the Overlay Manager or Administrator.

3.1.3 Identifiers and Authentication

Three types of authenticated communication occur in this specification’s use cases:

1. MAP Client actors and agents (BHIs and ICS Administrative MAP Clients) authenticating to the MAP Server when establishing an IF-MAP session.
2. BHI actors authenticating to each other while establishing the pairwise HIP tunnels used for carrying secured communications between BHIs across untrusted BSP(s).
3. Principal actors (Overlay Managers, Administrators) communicating with an ICS Administrative MAP Client. The protocol for this communication is out of scope of this specification, except for security considerations.

The BHI authentication MAY also facilitate additional BHI functionality that is out of scope for this specification (e.g., remote logging).

Each of the above-described roles will be associated with one or more distinct actors or agents (BHIs, Overlay Managers, and Administrators); each instance of these MUST have an identifier, and BHIs MUST have an associated X.509 certificate, to facilitate authentication by and secure communications with peer actors or agents.

Implementations of IF-MAP Metadata for ICS Security MUST support the use of cryptographic identities in the form of X.509 certificates for authentication, and BHIs MUST be authenticated in this manner; for Overlay Managers and Administrators, other forms of authentication (e.g., username/password) MAY be supported in addition to certificates, but are not recommended. When using certificate-based authentication, provisioning of an explicit chain of trust is required prior to establishing secure communications; actors and agents MUST validate peer X.509 certificates and SHOULD perform out-of-band certificate validation (e.g., OCSP and/or certificate revocation lists) when validating peers. (In addition, MAP Servers are required to perform validation of MAP Client identities as specified in TNC IF-MAP Binding for SOAP[2], section 6.3.1). The MAP service may be used to distribute and share certificate validation and Customer PKI trust chains, but these scenarios are out of scope of this document.

The specific methods used by principals in Overlay Manager and Administrator roles to authenticate to an ICS Administrative MAP client are out of scope of this specification. However, it is recommended that a robust authentication method be used, since these roles are privileged and can significantly alter the functionality of overlay networks. See section 5.2.4 for additional information.

Each BHI has two or three related identifiers:

1) A unique backhaul interface identifier associated with the physical BHI device
2) Either one or two certificate identifiers used by the BHI for authentication purposes.

Details on these identifiers are discussed further in each of the subsections below.

3.1.3.1 Backhaul Interface Identifier

The `backhaul-interface` identifier is generated based on attributes assigned by the manufacturer. It MUST be permanently associated with the BHI and MUST be globally unique. It is represented in the MAP graph by the `backhaul-interface` identifier type. In order to facilitate identifier uniqueness, the `backhaul-interface` device identifier MUST be formatted according to one of the two following patterns:

`BHI@`*`PEN`*`#`*`SN`*

`BHI@`*`PEN`*`#`*`MODEL-SN`*

where

- *"PEN"* is replaced with the registered Private Enterprise Number of the manufacturer [6]
- *"SN"* is replaced with a string representing a manufacturer-specific unique serial number.
- *"MODEL"* is replaced with a string representing a manufacturer-specific model number

When constructing the backhaul-interface identifier, any non-alphanumeric characters in the serial number and model number MUST be omitted, in order to avoid potential issues with whitespace and eliminate the possibility of collisions with the # and - delimiters.

[6] Currently assigned Private Enterprise Numbers (PENs) may be found at http://www.iana.org/assignments/enterprise-numbers . New PENs may be obtained from IANA at zero cost using the application form found at http://www.iana.org/protocols/apply/

Note that standalone “SN” numbers may not be unique across the entire manufacture’s product line; the manufacturer may need to prepend a product model number if internal serial numbers are only unique to a given product or model. Examples are shown below in Table 3.

Table 3: Examples of Backhaul Interface Identifiers (BHI IDs)

Example	Description
BHI@123#23A53B	A BHI identifier for the company associated with PEN “123” with manufacturer-unique designation (e.g., serial number) “23A53B”.
BHI@123#325-23A53B	Similar to the above example, except that a model number “325” is prepended to the SN field.

3.1.3.2 Certificate Identifiers

This specification uses the terms *Initial Secure Device Identifier* (IDevID) and *Locally Significant Secure Device Identifier* (LDevID) to describe the two possible types of certificates associated with the BHI (manufacturer-provisioned certificate and customer-provisioned certificate). This terminology was selected for its alignment with the terminology used in IEEE 802.1AR Secure Device Identity[15]. Note that IEEE 802.1AR is not incorporated as a normative reference in this version of the IF-MAP Metadata for ICS Security specification, but it may be included in a future revision.

IDevID and LDevID are each represented by a `distinguished-name` identifier in the MAP graph. MAP Clients MUST use the administrative domain "ifmap:client" with the IDevID and LDevID `distinguished-name` identifiers; this ensures compatibility with TNC MAP Content Authorization[14]. Domain-specific information can be expressed in non-CN fields of the distinguished name.

3.1.3.2.1 BHI Factory Certificate (IDevID) Identifier

The manufacturer MUST provision the BHI with an IDevID certificate. This certificate’s “subject” field is a Distinguished Name (DN) which MUST contain a common name (CN) field that matches the associated `backhaul-interface` identifier name described in section 3.1.3.1. (For match criteria, see section 3.3.2.3.1 of the TNC IF-MAP Binding for SOAP[2].) The certificate subject field is expected to be unique for a given BHI due to the uniqueness constraints of the backhaul-interface identifier (see section 3.1.3.1). Other fields of the certificate subject’s DN are manufacturer-specific and are undefined by this specification.

The validity of the IDevID is recommended to be effectively infinite, which SHOULD be achieved by assigning the notAfter date to the GeneralizedTime value of 99991231235959Z, as described in RFC 5280[16]; if not, the expiration date MAY be set to any suitably large value (bearing in mind that the lifetime of ICS devices can extend to multiple decades).

The IDevID MAY be signed by a manufacturer’s intermediate CA, in order to simplify MAP Content Authorization policies. For example, different BHI models could have certificates signed by different CAs in order to simplify the specification of a MAP Content Authorization policy that allows limited rights for a particular CA chain, rather than having to set a policy for each individual BHI.

The manufacturer MAY provide a method whereby IDevIDs can be refreshed as part of a BHI certificate lifecycle management process.

3.1.3.2.2 BHI Customer Certificate (LDevID) Identifier

The manager / operator of the BHI MAY provision the BHI with a unique customer certificate (LDevID). The purpose of this certificate is to facilitate use of a PKI that is specific to the BHI's deployed operating environment rather than the manufacturer's PKI. If a valid LDevID exists, the BHI MUST use this certificate instead of the IDevID. However, the creation of an LDevID on a given BHI MUST NOT cause the IDevID to be deleted.

The specifics of the CN field in the LDevID's distinguished name are out of scope of this specification. However, it is recommended that the CN contents also be reflected on a physically visible label on the outside of the BHI package, where practical, to help minimize potential confusion during provisioning, deployment, and maintenance operations. For example, the CN might contain a company asset number that matches an asset tag applied to the outside of the BHI package.

3.1.3.2.3 Active BHI Certificate

If a valid LDevID exists on a given BHI, that BHI MUST use the LDevID instead of the IDevID. If the LDevID is later deleted, expires, or becomes invalid for some other reason (e.g. revocation), then the BHI MUST revert to using its current IDevID. As a result, a BHI will only be associated with one DN identifier at a given time - either the LDevID if present and valid, or the IDevID if no valid LDevID is present.

The active certificate is used by the BHI to authenticate both to the MAP server and to peer BHIs; if the current certificate is changed, the BHI MUST establish new sessions with the MAP server and peer BHIs. Instances in which this is necessary might include, but are not limited to:

- LDevID is provisioned, thus superseding the IDevID
- LDevID expires, thus switching back to using the original IDevID
- LDevID is replaced or updated by a new LDevID certificate

3.1.4 Authorization Policy

There is an important distinction to be made between two types of authorization policy specified by TNC. The first type, called MAP Content Authorization Policy, defines authorization policy enforced by the MAP Server. This policy controls which MAP operations a given MAP Client may perform on particular types of instances of metadata, and is specified in TNC MAP Content Authorization[14]. In many cases, the enforcement policies will depend upon the identity of the MAP Client and/or the roles and tasks assigned to that MAP Client identity.

The MAP Server MUST implement some form of MAP content authorization policy to limit MAP Client interactions with the MAP graph as described in this document. The MAP Server SHOULD implement TNC MAP Content Authorization as specified in [14]; we expect that this requirement will become mandatory in a future version of this specification. Suggested TNC MAP Content Authorization policies for ICS Security Metadata can be found on the TCG website.[7]

ICS Authorization Policies are the second type of policies; they define policies which are enforced by various actors and agents defined in this specification - according to the contents and structure of specific metadata in the MAP graph. These policies as they apply to this specification's use cases are defined below:

- **Backhaul Policy:** Defines the set of BHIs with which a given BHI may communicate.

[7] http://www.trustedcomputinggroup.org/resources/tnc_map_content_authorization

- **Overlay Policy:** For the set of ICS devices being protected by a given BHI, defines which devices that BHI will allow to communicate across the overlay network.
- **BHI Management Policy:** Specifies which BHIs each Overlay Management principal may manage.
- **Overlay Management Policy:** Specifies which overlays each Overlay Manager principal may manage.
- **Additional Policies:** Other policies out of scope of this specification. These might include deep packet inspection policies that enforce ICS protocol-specific authorization rules (e.g., a particular identity or role may only send "read" commands, not "write" commands, to a particular control system device).

In some cases, ICS Authorization Policies will be static; in others the policies will change over time. All MAP Clients described in this specification MUST create IF-MAP subscriptions for relevant policy metadata and maintain an active poll on these subscriptions, so that timely notification of policy changes occurs and locally enforced policies immediately reflect those changes.

For each of the roles defined in section 3.1.1, there are agents (applications and/or devices) that serve in those roles for a given implementation; those agents MUST include enforcement functionality that enforces specified policies based on the identity (and perhaps other attributes) of the agents with whom it is interacting. For example, a MAP Server must enforce MAP Content Authorization Policies that limit what a particular BHI can do to the MAP graph; the BHI must enforce ICS Authorization Policies related to what data it passes or blocks for ICS device data flows; the ICS Administrative MAP Client must enforce ICS Authorization Policies governing Overlay Managers and ICS Administrators. These policies may change over time, so the agents MUST monitor for changes in policy (using subscriptions, in the case of ICS Authorization Policies) and enforce the new policies as they become available.

3.2 Overview of Target Functionality

The metadata defined in this specification will facilitate secure coordination for the provisioning and operation of - and secure communication between - compliant BHI devices, as well as for associated configuration management functions. This will enable secure coordination in general operation.

From an architectural standpoint, four classes of client coordination activities (as described below) are present in the ICS security domain, two of which are supported by the metadata standardized in this specification. These classes of activities are described here for convenience in understanding usage, but are not actually part of the metadata and do not constrain its use.

3.2.1 Configuration and Operational Management

This specification addresses functionality for the configuration and operational management of BHIs. It is likely that additional functionality is needed for BHI diagnostics, logging, firmware upgrades, etc.; such additional functionality is outside of the scope of this specification.

3.2.2 Operational Communications Facilitation

During run-time operations, BHIs need to communicate with each other. These communications MUST support protection from man-in-the-middle attacks (message privacy and integrity). IF-MAP coordination facilitates establishment of communication, as well as policy enforcement capabilities that restrict which BHIs may send messages to other BHIs.

3.2.3 Credential Management

Historically, the ICS community has often had no authentication capability, or perhaps a shared credential (i.e., shared password) for authentication. However, backhaul-based interconnection drives a requirement for X.509 certificate-based credentials in order to provide traceable, cryptographically bound identities for all MAP Clients (as well as the MAP Server) described in this specification, as well as for BHI-to-BHI secure communications. This allows policies to be specified and enforced down to per-device granularity, enables non-repudiation of device-specific logging, allows device credentials to be revoked if necessary, and prevents the compromise of one device's credentials from compromising the rest. In devices where a TPM is available, TPM-based certificate storage increases security by tying a certificate to a physical device. The provisioning and management of certificates is out of scope of this document but may be specified in a future TCG specification.

3.2.4 MAP Server Auto-Discovery

MAP Clients SHOULD attempt to use pre-configured information to locate and connect to a MAP Server. However, if that server is unavailable or if the client is unconfigured, the MAP Client MAY use an alternate method such as DNS-based Service discovery using SRV resource records as specified in IETF RFC 6335[9]. Such a method is currently unspecified, but may be specified in a future TCG specification.

3.3 IF-MAP Identifiers Used in This Specification

This specification uses a variety of identifiers defined in the TNC IF-MAP Binding for SOAP[2], including extended identifier types. Table 4 in section 3.4 gives a summary of these identifier types, whether they are extended identifiers, and how they are used in the ICS Security use cases.

3.3.1 Use of administrative-domain Identifier Attribute

The TNC IF-MAP Binding for SOAP[2] recommends that the `administrative-domain` attribute of an `ip-address` identifier represent the routing domain in which the IP address occurs, and the `administrative-domain` attribute of an `mac-address` identifier represent the local network (LAN/VLAN segment) in which the MAC address occurs (see section 3.3.1.2 of [2]).

Within an ICS environment, it is beneficial to utilize the `administrative-domain` attribute to identify the overlay network in which the `ip-address` identifier or `mac-address` identifier occurs. In order to accomplish that goal, the following naming convention for IP routing domains and local networks in overlay networks facilitates communication of the overlay network name in the `administrative-domain` attribute:

- The `administrative-domain` attribute of an `ip-address` identifier referring to an ICS device on the overlay network MUST be the string concatenation of the name value of the associated `overlay-network-group` followed by a string that uniquely denotes the routing domain, separated by a colon.
- The `administrative-domain` attribute of a `mac-address` identifier referring to an ICS device on the overlay network MUST be the string concatenation of the name value of the associated `overlay-network-group` followed by a string that uniquely denotes the local LAN/VLAN segment, separated by a colon.

For example, if MAC address xx:xx:xx:xx:xx:xx occurs in VLAN 100 on overlay network Overlay1, and MAC address yy:yy:yy:yy:yy:yy occurs in VLAN 200 on overlay network Overlay1, then the `mac-address` identifier with name "xx:xx:xx:xx:xx:xx" could have `administrative-domain` attribute "Overlay1:100" and the `mac-address` identifier with name "yy:yy:yy:yy:yy:yy" could have `administrative-domain` attribute "Overlay1:200".

For overlay networks that have a single routing domain or single LAN/VLAN segment, MAP Clients SHOULD use the same string for both the overlay network name and the IP routing domain or LAN/VLAN segment, (e.g. Overlay1:Overlay1).

As specified in section 3.1.3.2, MAP Clients are required to use the administrative domain "ifmap:client" with the IDevID and LDevID `distinguished-name` identifiers; this ensures compatibility with TNC MAP Content Authorization[14].

A special `overlay-manager-group` identifier instance with administrative-domain attribute "tcg-reserved" and name attribute "unassigned" is used by BHIs to link to a known identifier when they are otherwise unassigned to a particular `overlay-manager-group` (see section 4.1.7). MAP Clients MUST NOT use the administrative domain "tcg-reserved" for any purpose other than those specified in this document.

The `administrative-domain` identifier attribute MUST NOT be used for any identifiers detailed in this specification except as specified in this section.

3.4 Extended Identifiers

Details for each of the three extended identifiers listed in Table 4 are given in the subsections below. As defined in section 3.2.3 of the TNC IF-MAP Binding for SOAP[2], extended identifiers are defined as extensions of the `identity` type `other`. In future IF-MAP specifications, extended identifier types may be represented in a different manner for IF-MAP operations, but the definitions of the identifiers themselves should otherwise remain unchanged.

Table 4: Identifier Types Used in this Specification

Identifier Type	Ext. ID	Usage
backhaul-interface	Y	A globally unique identifier associated with a given BHI (see section 3.1.3.1).
distinguished-name	N	The full DN of an identity certificate's subject for an actor or agent (see section 3.1.3.2 for DNs associated with BHIs).
hip-hit	N	A HIP Host Identity Tag associated with the public key in a given certificate.
ifmap-client-catalog	Y	A single-instance identifier providing a location for MAP Clients to publish self-describing `device-characteristic` metadata, as specified in TNC MAP Content Authorization[14]
ip-address	N	An IPv4 or IPv6 address associated with a BHI or with a device on the overlay network. Used for BHIs to define one of possibly multiple IP addresses (in the case of multi-homed BHIs). Used in defining overlay policies.
mac-address	N	A MAC identifier associated with a device on the overlay network. Used in defining overlay policies.

overlay-manager-group	Y	The name of a group that can be assigned to manage one or more overlay networks and/or BHIs.
overlay-network-group	Y	The name of an overlay network to which various BHIs are bound.

3.4.1 backhaul-interface

The `backhaul-interface` extended identifier type contains a single `name` attribute which specifies the unique name of a given BHI; this name is based upon attributes assigned by the manufacturer, as described in section 3.1.3.1. Further context is provided by metadata attached to this identifier (see sections 3.5.1, 3.5.2, 3.5.4, 3.5.9.2, 3.5.10.1, 3.5.11, 3.5.12, and 3.5.13).

```
<xsd:element name="backhaul-interface">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string"
          use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.4.2 overlay-manager-group

The `overlay-manager-group` extended identifier type contains a single `name` attribute which specifies the name of a particular Overlay Manager group. Further context is provided by metadata attached to this identifier (see sections 3.5.7, 3.5.9, and 3.5.10.2).

```
<xsd:element name="overlay-manager-group">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string"
          use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.4.3 overlay-network-group

The `overlay-network-group` extended identifier type contains a single `name` attribute which specifies the name of the overlay network. Further context is provided by metadata attached to this identifier (see sections 3.5.1, 3.5.8, 3.5.9.1, and 3.5.10.1).

```
<xsd:element name="overlay-network-group">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string"
          use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5 IF-MAP Metadata for ICS Security Types

The IF-MAP Metadata for ICS Security schema contains the following metadata types, each of which is defined in detail in the subsections below. Figure 3 shows a graphical summary of these metadata types used in context with associated identifiers in the MAP graph.

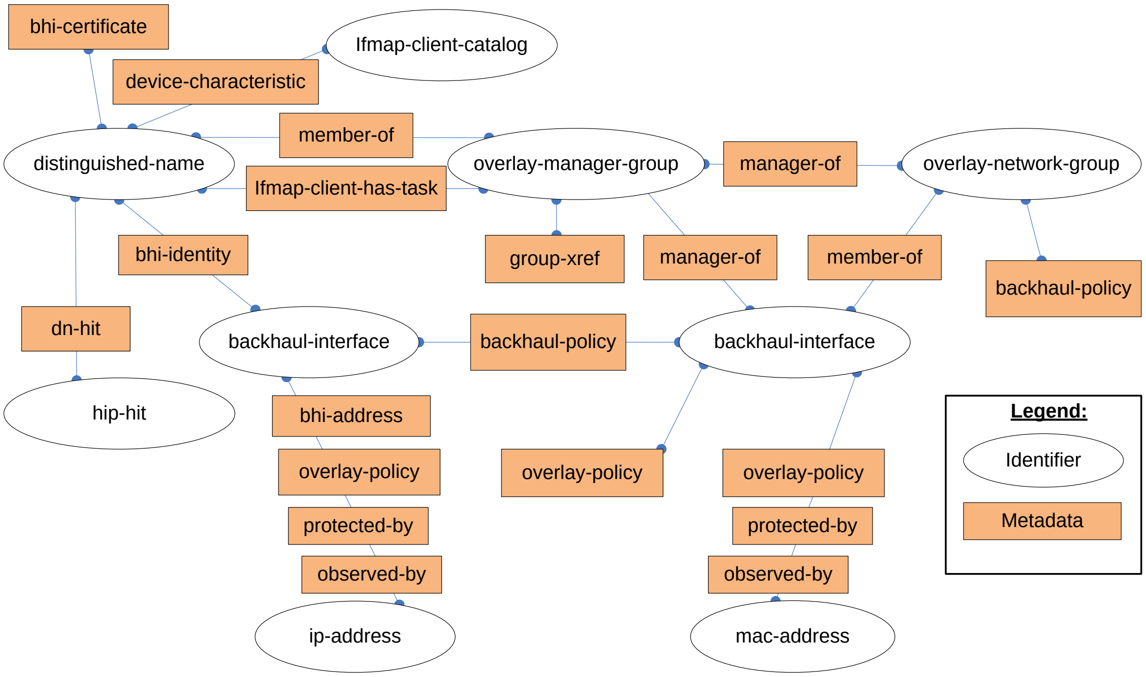

Figure 3: Graphical summary of ICS security metadata types

This graphical summary of all possible ICS security metadata types is not intended to represent an actual MAP graph. Rather, it shows to what types of identifiers and links each metadata type is generally attached. For instance, a `backhaul-interface` identifier would never be linked to a single `ip-address` identifier via both `bhi-address` and `protected-by`; the `bhi-address` link would be to an `ip-address` identifier representing the IP address of the BHI itself, and the `protected-by` link would be to an `ip-address` identifier representing the IP address of another system that the BHI has detected on the network.

Table 5 lists in alphabetical order all metadata types used in this specification; details on the usage of each of the new metadata types are given in the subsections below. The metadata types described in this section, which are defined in their own namespace that is specific to this specification, MUST NOT be used for purposes other than what is specified in this document. These metadata types MUST be attached only to prescribed kinds of identifiers or to links adjacent to prescribed kinds of identifiers, as specified in Table 5, so that match-links filters will work as intended.

Table 5: Metadata Types Used in this Specification

Metadata Type	Cardinality	Where Attached
backhaul-policy	Multi	Identifier: overlay-network-group Link: backhaul-interface ⬄ backhaul-interface

bhi-address	Single	Link: backhaul-interface ⬄ ip-address
bhi-certificate	Multi	Identifier: distinguished-name
bhi-identity	Single	Link: backhaul-interface ⬄ distinguished-name
device-characteristic	Multi	Link: distinguished-name ⬄ map-client-catalog; see [14]
dn-hit	Single	Link: distinguished-name ⬄ hip-hit
group-xref	Multi	Identifier: overlay-manager-group
ifmap-client-has-task	Single	Link: distinguished-name ⬄ overlay-network-group; see [14]
manager-of	Single	Link: overlay-manager-group ⬄ backhaul-interface Link: overlay-manager-group ⬄ overlay-network-group
member-of	Single	Link: overlay-network-group ⬄ backhaul-interface Link: overlay-manager-group ⬄ distinguished-name
observed-by	Multi	Link: backhaul-interface ⬄ ip-address Link: backhaul-interface ⬄ mac-address
overlay-policy	Single	Link: backhaul-interface ⬄ ip-address Link: backhaul-interface ⬄ mac-address Identifier: backhaul-interface
protected-by	Single	Link: backhaul-interface ⬄ ip-address Link: backhaul-interface ⬄ mac-address

MAP Clients MUST NOT publish, modify, or delete any metadata type on a target identifier or link for which their role is not listed in the Publisher column of Table 6.

Table 6: Permitted Publishers of Metadata Types

Metadata Type	Where Attached	Publisher
backhaul-policy	Identifier: overlay-network-group Link: backhaul-interface ⇔ backhaul-interface	Overlay Manager
bhi-address	Link: backhaul-interface ⇔ ip-address	BHI*
bhi-certificate	Identifier: distinguished-name	BHI*
bhi-identity	Link: backhaul-interface ⇔ distinguished-name	BHI*
device-characteristic	Link: distinguished-name ⇔ ifmap-client-catalog; see [14]	BHI*
dn-hit	Link: distinguished-name ⇔ hip-hit	BHI*
group-xref	Identifier: overlay-manager-group	Administrator
ifmap-client-has-task	Link: distinguished-name ⇔ overlay-network-group; see [14]	Overlay Manager
manager-of	Link: overlay-manager-group ⇔ backhaul-interface	Administrator, BHI**
manager-of	Link: overlay-manager-group ⇔ overlay-network-group	Administrator
member-of	Link: overlay-network-group ⇔ backhaul-interface	Overlay Manager
member-of	Link: overlay-manager-group ⇔ distinguished-name	Administrator
observed-by	Link: device backhaul-interface ⇔ ip-address Link: device backhaul-interface ⇔ mac-address	BHI
overlay-policy	Link: backhaul-interface ⇔ ip-address Link: backhaul-interface ⇔ mac-address Identifier: backhaul-interface	Overlay Manager
protected-by	Link: device backhaul-interface ⇔ ip-address Link: device backhaul-interface ⇔ mac-address	BHI

* A BHI MUST publish these metadata types only on identifiers representing itself - not other BHIs.

** A BHI MUST NOT publish manager-of metadata on any identifier other than the TCG reserved `overlay-manager-group` identifier instance with administrative-domain attribute “tcg-reserved” and name attribute “unassigned”. (See section 4.1.7 for details on locator metadata for unassigned BHIs.)

3.5.1 backhaul-policy

Clients MUST publish this only on overlay-network-group, or between backhaul-interface and backhaul-interface

Overlay Managers use the multi-valued `backhaul-policy` metadata type to control whether some, all, or none of the BHIs on a given backhaul network communicate with one another over the backhaul network. For this purpose, each BHI acts as a TNC Flow Controller (see the TNC Architecture[1]). This metadata type has a `name` string attribute that MUST contain the name of an overlay network, as well as a `policy` string attribute that MUST have a value of either “allow” or “deny”.

The `backhaul-policy` metadata type SHOULD only be published by the Overlay Manager role associated with affected BHIs.

3.5.1.1 Default Backhaul Policy

An Overlay Manager MAY attach the `backhaul-policy` metadata type to an `overlay-network-group` identifier, as shown on the left side of Figure 4, to indicate a default policy for BHIs that are members of that overlay network on a given backhaul network.

3.5.1.2 Pairwise Backhaul Policy

An Overlay Manager MAY attach the `backhaul-policy` metadata type to a link between the `backhaul-interface` identifiers of a given pair of BHIs, as shown at the bottom of Figure 4, to indicate policy for communications between that pair of BHIs on a per-overlay network basis.

3.5.1.3 Considerations for use of Backhaul Policy

Each BHI MUST enforce policies dictated by `backhaul-policy` metadata in a bidirectional manner. In other words, if a backhaul policy permits traffic between two BHIs, packets may flow in either direction between those BHIs.

The use of default `backhaul-policy` metadata will, in some configurations, result in a more or less efficient MAP graph versus pairwise use of `backhaul-policy` metadata. For example, a default `backhaul-policy` of “allow” will result in a full-mesh configuration; once pairwise `backhaul-policy` of “deny” is published on half of the possible links, the default policy becomes inefficient. In general, implementations should consider the effects of using `backhaul-policy` metadata in these two different ways.

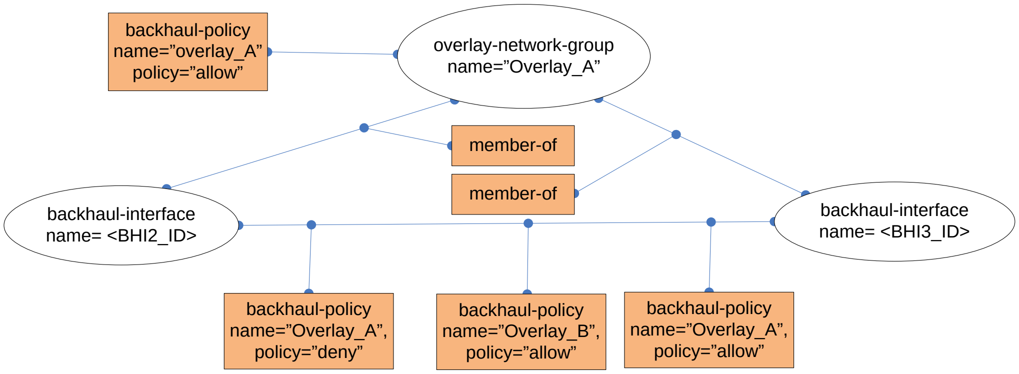

Figure 4: Example MAP graph for use of the backhaul-policy metadata type

Figure 4 shows several possible uses of `backhaul-policy` metadata for different scenarios as described above, but does not depict a MAP graph in an operational environment. As described in Section 4.1.9.1, pairwise backhaul-policy metadata takes precedence over default backhaul-policy metadata, and a policy attribute of “deny” takes precedence over a policy attribute of “allow”, so the outcome of the above MAP graph would be that BHI2 and BHI3 would deny communication of Overlay_A traffic across the backhaul network.

```
<xsd:element name="backhaul-policy">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="name" type="xsd:string"
        minOccurs="1" maxOccurs="1"/>
      <xsd:element name="policy" minOccurs="1" maxOccurs="1">
        <xsd:simpleType>
          <xsd:restriction base="xsd:string">
            <xsd:enumeration value="allow"/>
            <xsd:enumeration value="deny"/>
          </xsd:restriction>
        </xsd:simpleType>
      </xsd:element>
    </xsd:sequence>
    <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.2 bhi-address

Clients MUST publish this only between: backhaul-interface and ip-address

A BHI attaches the bhi-address metadata type to a link between a backhaul-interface identifier and an ip-address identifier representing the BHI’s IP address on the backhaul network, as shown in the example in Figure 5. This metadata is used by a BHI to advertise its available IP address (or addresses, in the case of multi-homed BHIs).

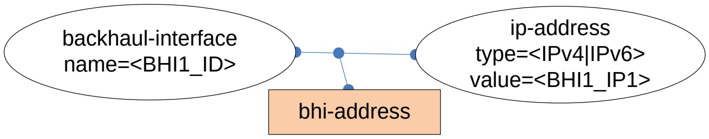

Figure 5: Example MAP graph showing use of bhi-address metadata type

```
<xsd:element name="bhi-address">
  <xsd:complexType>
    <xsd:attributeGroup
      ref="ifmap:singleValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.3 bhi-certificate

Clients MUST publish this only on: distinguished-name

A BHI attaches the multi-valued bhi-certificate metadata type, which contains a base64-encoded (see RFC 2045[10]) X.509 certificate, to a distinguished-name identifier as shown in Figure 6. bhi-certificate metadata MUST only be published on distinguished-name identifiers that have `bhi-identity` metadata present on a link between that `distinguished-name` identifier and a `backhaul-interface` identifier.

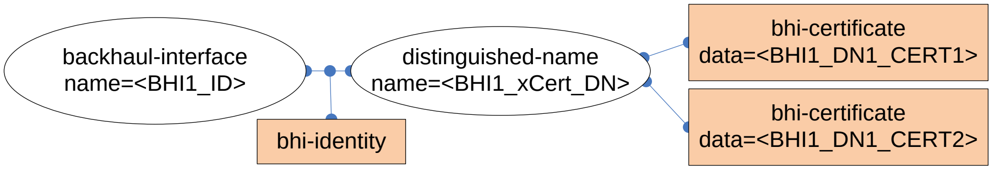

Figure 6: Example MAP graph showing use of bhi-certificate metadata type

Multiple certificates are supported, using this multi-valued metadata type, to allow for certificate lifecycle management activities (e.g., issuing a new certificate before an existing certificate has expired). All authentications based upon `bhi-certificate` metadata SHOULD be validated using out-of-band certificate validation (e.g., OCSP) and/or certificate revocation lists; the specifics of such validation are out of scope of this specification.

This metadata type provides MAP Clients with the capability to look up certificates associated with a distinguished name and, by further reference via a link between that `distinguished-name` and a `backhaul-interface` identifier, the device using one of those certificates. The “Subject” field of each certificate stored in a given `bhi-certificate` MUST match the associated `distinguished-name` value (see section 3.3.2.3.1 of [2] for equivalence rules).

```
<xsd:element name="bhi-certificate">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="data" type="xsd:base64Binary"
        minOccurs="1" maxOccurs="1"/>
    </xsd:sequence>
    <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.4 bhi-identity

Clients MUST publish this only between: backhaul-interface and distinguished-name

A BHI attaches the `bhi-identity` metadata type to a link between a `backhaul-interface` identifier and the `distinguished-name` identifier associated with a certificate installed on that BHI, as shown in Figure 7. This metadata type and its associated link SHOULD only be published by a MAP Client associated with the referenced BHI.

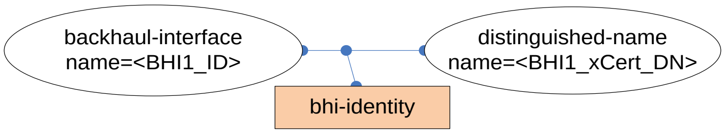

Figure 7: Example MAP graph showing use of the bhi-identity metadata type

Due to certificate lifecycle management activities, BHIs MAY have more than one certificate identity associated with a given Distinguished Name identity (as in the case of certificate renewal, when it is desirable to provision a certificate with a new expiration date before the expiration date of the old certificate). An example of this situation is shown in Figure 9.

```
<xsd:element name="bhi-identity">
  <xsd:complexType>
    <xsd:attributeGroup
      ref="ifmap:singleValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.5 device-characteristic

Clients MUST publish this only between: distinguished-name and ifmap-client-catalog

When a new BHI is plugged into the network, an ICS Administrator or Overlay Manager needs to assign it a task so that it can participate in an overlay network. `device-characteristic` metadata helps this workflow, by providing device characterization and connection to a rendezvous point that ICS Administrative MAP Clients can use to identify BHIs and determine whether they are already assigned to an overlay network. See [14] for more details.

This metadata type is defined in TNC IF-MAP Metadata for Network Security[3]. This specification uses it as prescribed by TNC MAP Content Authorization[14]. Accordingly, every BHI MUST, upon establishing a session with a MAP Server, ensure that the MAP graph contains accurate `device-characteristic` metadata describing itself, attached to a link between its self identifier (which will be a `distinguished-name` identifier) and the `ifmap-client-catalog` identifier as illustrated in Figure 8.

An ICS Administrative MAP Client SHOULD determine whether TNC MAP Content Authorization is in use, using the mechanism provided in section 3.4.5.2 of the TNC IF-MAP Binding for SOAP[2]. If so, the ICS Administrative MAP Client SHOULD subscribe to learn about `device-characteristic` metadata that has a `device-type` attribute containing the value "tcg-ics-backhaul-interface" and is attached to a link adjacent to the `ifmap-client-catalog` identifier. On finding such metadata, the ICS Administrative MAP Client should check to see if the identifier at the other end of that link (representing a BHI) has `ifmap-client-has-task` metadata attached to a link to an `overlay-network-group` identifier, as described in section 3.5.8. If not, the ICS Administrative MAP Client SHOULD log an error.

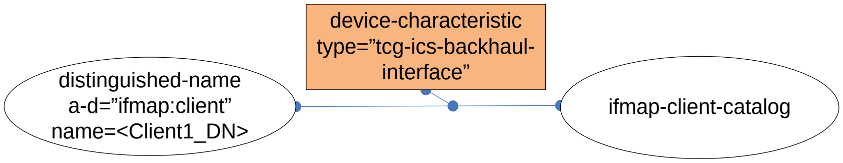

Figure 8: Example MAP graph showing use of device-characteristic metadata type for MAP Content Authorization

3.5.6 dn-hit

Clients MUST publish this only between: distinguished-name and hip-hit

A BHI attaches the single-valued `dn-hit` metadata type to a link between a `distinguished-name` identifier representing that BHI and a `hip-hit` identifier. A given HIT (Host Identity Tag) is computed from the public key of a certificate as described in IETF RFC 4423[11]. See section 3.1.3.2.3 for guidance on when to use LDevID vs. IDevID. Figure 9 shows an example of these relationships where there is a separate `hip-hit` identifier (and therefore a separate `dn-hit` metadata link) for each of the certificates associated with a particular `distinguished-name`.

The purpose of this metadata type and its associated links is to allow MAP Clients to determine distinguished names and associated certificates for a given `hip-hit`. This is valuable for HIP-enabled MAP Clients when all that is known about an incoming HIP protocol flow is the `hip-hit` name.

It is also important to note that variations on Figure 9 are possible, depending on the particulars of a given PKI implementation. For example, when a new certificate is issued before an existing one expires (a PKI best practice), the same key material may be used in both the existing and new certificates. In this case, there would be only one `hip-hit` identifier and associated `dn-hit` metadata link that refers to both certificates via the `distinguished-name` identifier. In another scenario a certificate authority may issue two certificates with different distinguished names yet the same key material. This could happen in situations where there is an organizational change that affects the structure of the Distinguished Name while it is undesirable to issue new key material (e.g., distribute new private keys or service new CSR's). In this situation, a given `hip-hit` identifier will have two `dn-hit` metadata links to different `distinguished-name` identifiers.

In all cases, the MAP Client MUST resolve ambiguities associated with multiple matches to searches starting from a given `hip-hit` identifier. For example, when two `distinguished-name` identifiers are linked to the same `hip-hit` identifier, each `distinguished-name` and its associated `bhi-certificate` metadata could be tested for suitability to the current operating environment. How these ambiguities are resolved is implementation dependent.

BHIs publish this metadata and its associated link for use by peer BHIs as the starting point for finding the BHI associated with a particular HIT when connectivity is requested (i.e., the "responder" role in HIP, as described in [12]).

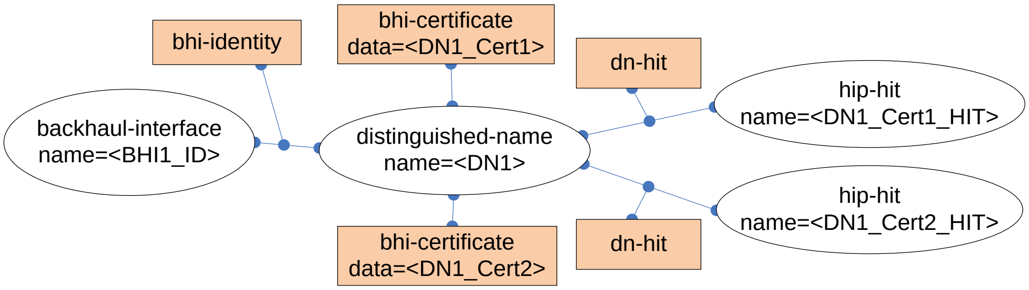
bhi-identity
bhi-certificate data=<DN1_Cert1>
dn-hit
hip-hit name=<DN1_Cert1_HIT>
backhaul-interface name=<BHI1_ID>
distinguished-name name=<DN1>
hip-hit name=<DN1_Cert2_HIT>
bhi-certificate data=<DN1_Cert2>
dn-hit

Figure 9: Example MAP graph showing use of the dn-hit metadata type

```
<xsd:element name="dn-hit">
  <xsd:complexType>
    <xsd:attributeGroup
      ref="ifmap:singleValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.7 group-xref

Clients MUST publish this only on: overlay-manager-group

The multi-valued `group-xref` metadata type is used as an external reference to group members provided by an LDAP URI (including Active Directory group memberships, since Active Directory supports LDAP queries for group membership). An Administrator attaches `group-xref` metadata to an `overlay-manager-group` identifier to include a list of principals who are members of a particular overlay manager group. When this metadata is present, an Overlay Manager MUST use this information to access the appropriate group list in LDAP for purposes of policy decisions and enforcement when managing overlay manager groups. Any credentials needed to access the LDAP service (password, certificate) are outside of the scope of this specification.

The `uri` string attribute of this metadata type contains an LDAP directory search specification (see examples in IETF RFC 4516[13]) that will result in a list of principals. The structure of the search specification (distinguished name, attributes, filter, etc.) will depend upon the schema used by the selected LDAP server. For implementations where the LDAP server is on separate system from the MAP server, clients MUST use a secure protocol for communications between the two services. Examples include (but are not limited to) StartTLS (IETF RFC 2830)[17] and LDAPS (LDAP over SSL) protocol[8].

In order to minimize search polling of the LDAP server for changes to the associated LDAP group(s) as well as to propagate timely awareness of LDAP group changes by Overlay Managers, a MAP Client application closely associated with the MAP Server MAY publish a notify on a particular `group-xref` metadata instance upon changes to the LDAP group(s). Overlay Managers MAY maintain an active subscription for these notify publications. Overlay Managers MAY also recheck the LDAP server for changes to group membership at a configurable interval.

Figure 10 shows an example of this metadata type used in a MAP graph.

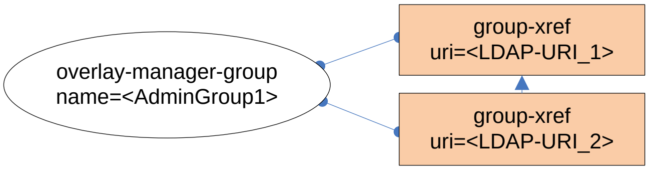

Figure 10: Example MAP graph showing the use of the group-xref metadata type

```
<xsd:element name="group-xref">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="uri" type="xsd:string"
        minOccurs="1" maxOccurs="1"/>
    </xsd:sequence>
    <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.8 ifmap-client-has-task

This metadata type is specified by, and used according to, TNC MAP Content Authorization[14].

If TNC MAP Content Authorization is enabled, an Administrator or Overlay Manager places this metadata type on a link between a BHI's self identifier (a `distinguished-name` identifier) and an `overlay-network-group` identifier as illustrated in Figure 11. Doing so grants the BHI the access to MAP content that allows it to perform its required function.

[8] Not standardized, but commonly supported on port 636

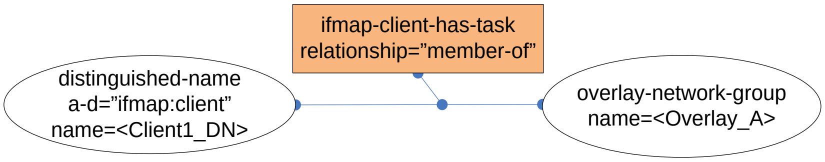

Figure 11: Example MAP graph showing use of ifmap-client-has-task metadata type for MAP Content Authorization

3.5.9 manager-of

Clients MUST publish this only between: overlay-manager-group and backhaul-interface, or overlay-manager-group and overlay-network

An Administrator or Overlay Manager attaches the single-valued `manager-of` metadata type to links to indicate a management relationship. This metadata type is used on two ways in this specification:

3.5.9.1 Manager of an Overlay Network Group

An Administrator attaches the `manager-of` metadata type to a link between an `overlay-manager-group` identifier and an `overlay-network-group` identifier to indicate that a given group will be managing a particular overlay network, as shown in Figure 12.

In this example, the group named OMGroup_1 is the manager of the overlay network named Overlay_A.

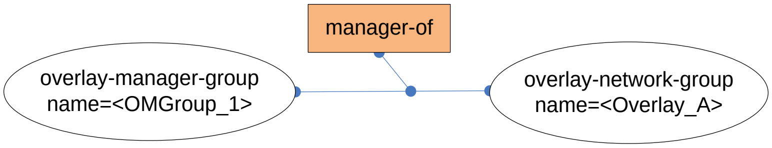

Figure 12: Example MAP graph showing use of manager-of metadata type for overlay management

3.5.9.2 Manager of a BHI

An Overlay Manager attaches the `manager-of` metadata type to a link between an `overlay-manager-group` identifier and a `backhaul-interface` identifier for a particular BHI to indicate that a given group will be managing a particular BHI, as shown in Figure 13.

In this example, the group named OMGroup_1 is the manager of the BHI represented by the `backhaul-interface` identifier BHI1_ ID.

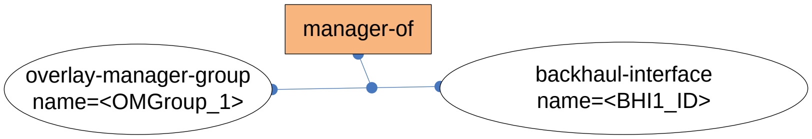

Figure 13: Example MAP graph showing use of manager-of metadata type for device management

```
<xsd:element name="manager-of">
  <xsd:complexType>
     <xsd:attributeGroup
       ref="ifmap:singleValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.10 member-of

Clients MUST publish this only between: overlay-network-group and backhaul-interface, or overlay-manager-group and distinguished-name

An Administrator or Overlay Manager attaches the single-valued `member-of` metadata type to a link to indicate a subordinate relationship. This metadata type is used in two different ways in this specification:

3.5.10.1 Member of an Overlay Network Group

An Overlay Manager attaches the `member-of` metadata type to a link between an `overlay-network-group` identifier and a `backhaul-interface` identifier to indicate that the BHI associated with that `backhaul-interface` name is a member of that overlay network. An example of this use of the `member-of` metadata type is shown in Figure 4 (section 3.5.1.3).

3.5.10.2 Member of an Overlay Manager Group

An Administrator attaches the `member-of` metadata type to a link between an `overlay-manager-group` identifier and a `distinguished-name` identifier representing a user or application authorized to act as an Overlay Manager to indicate that the principal represented by that Distinguished Name is a member of that Overlay Manager group. Figure 14 shows an example MAP graph where the Principal2 is a member of both OMGroup_1 and OMGroup_2 overlay manager groups while Principal1 and Principal3 are members only of OMGroup_1 or OMGroup_2 respectively.

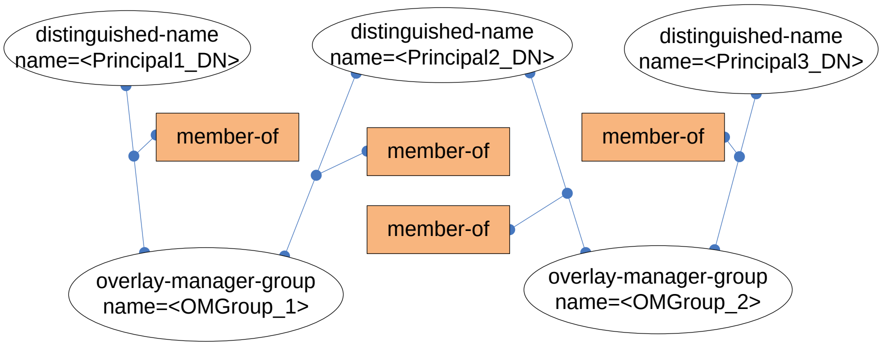

Figure 14: Example using member-of metadata type for principal membership in overlay manager groups

```
<xsd:element name="member-of">
  <xsd:complexType>
     <xsd:attributeGroup
       ref="ifmap:singleValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.11 observed-by

Clients MUST publish this only between: backhaul-interface and ip-address, or backhaul-interface and mac-address

A BHI attaches the `observed-by` metadata type to a link between the BHI's `backhaul-interface` identifier and either an `ip-address` or `mac-address` identifier representing an ICS device or other device within an overlay (e.g. DHCP server, etc.). `observed-by` metadata is a multi-valued metadata type, since more than one BHI may observe an ICS device on the same overlay network.

A BHI MUST publish this metadata type on a link between its own `backhaul-interface` identifier and the associated `ip-address` or `mac-address` identifier when it observes any inbound network traffic, originating on its backhaul interface, from ICS devices on overlay networks of which it is a member. The BHI MUST NOT publish `observed-by` metadata on devices it protects, i.e. devices on its protected overlay network segment. The BHI MUST publish `observed-by` metadata with a `lifetime` attribute of "session".

In the case where the BHI is a member of only one overlay network, the BHI MUST include an `administrative-domain` attribute (of the linked `ip-address` and `mac-address` identifiers, derived according to the rules in section 3.3.1) that reflects the name of that overlay network when publishing `observed-by` metadata on identifiers, in order to prevent confusion when the same IP addresses and MAC addresses are used in different overlay networks. In the case where the BHI is a member of more than one overlay network and is able to determine the source overlay network (i.e., the overlay network where the packets originated) of the ICS Device, the BHI MUST include an `administrative-domain` attribute that reflects the source overlay network when publishing `observed-by` metadata. In the case where the BHI is a member of more than one overlay network and is not able to determine the source overlay network of the ICS Device, the BHI MUST omit the administrative domain attribute when publishing `observed-by` metadata. In the case where the BHI is not assigned to an overlay network, the BHI MUST NOT publish `observed-by` metadata.

An Overlay Manager may use this metadata to better characterize a set of ICS devices within an overlay where the IP addresses and MAC addresses of these devices may not otherwise be known, and to identify which devices are protected by which BHIs on the overlay network. This metadata may also be used to trigger alarm events if unexpected IP addresses or MAC addresses are observed on an otherwise steady-state ICS system configuration. This metadata may become stale, as systems move, outages occur, etc.; ICS administrators are recommended to implement a plan to clear/refresh this information on a periodic basis. The BHI SHOULD be configurable with a deletion interval for automated deletion of `observed-by` metadata once it reaches a certain time since it last saw evidence to support that metadata, which SHOULD default to 14 days.

Figure 15 shows an example of this metadata type attached to links between BHI1's `backhaul-interface` identifier and an `ip-address` and a `mac-address` identifier, respectively, for two ICS devices (Dev1 and Dev2) which are present in the overlay network and not resident in the overlay network segment protected by that BHI. Since the `ip-address` and `mac-address` identifiers to which the `observed-by` links are attached in this example refer to ICS devices in an identifiable source overlay network, the BHI is required to include an `administrative-domain` attribute of those identifiers derived according to section 3.3.1.

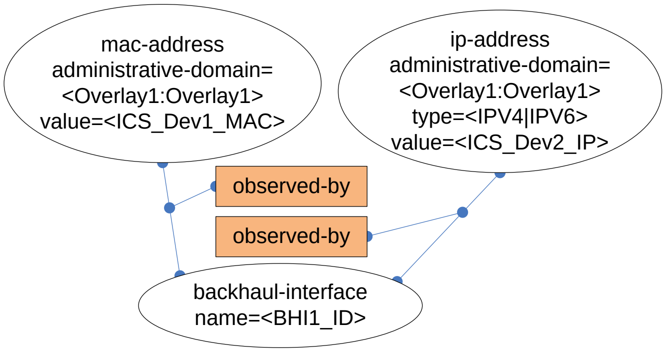

Figure 15: Example MAP graph showing use of observed-by metadata type

```
<xsd:element name="observed-by">
  <xsd:complexType>
   <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.12 overlay-policy

Clients MUST publish this only on backhaul-interface, or between: backhaul-interface and ip-address, or backhaul-interface and mac-address

Overlay Managers use the multi-valued `overlay-policy` metadata type to control whether some, all, or none of the ICS devices in an overlay network are allowed to send communications outbound on the overlay network. For this purpose, each BHI acts as a TNC Flow Controller. This metadata type has a `name` string attribute that MUST contain the name of an overlay network, as well as a `policy` string attribute that MUST have a value of either “allow” or “deny”.

3.5.12.1 Default Overlay Policy

An Overlay Manager MAY attach the `overlay-policy` metadata type to a `backhaul-interface` identifier, as shown in the bottom right corner of **Error! Reference source not found.**, to indicate a default policy for traffic related to a particular overlay network that is received by that particular BHI.

3.5.12.2 Device-specific Overlay Policy

An Overlay Manager MAY attach the `overlay-policy` metadata type to a link between a `backhaul-interface` identifier and an `ip-address` or a `mac-address` identifier of an ICS device within a given overlay network, as shown in the bottom left corner of Figure 19, to indicate that the referenced IP address or MAC address within the overlay will have its non-local (BHI-to-BHI) overlay communications governed by the given policy when sent out via the specified BHI.

```
<xsd:element name="overlay-policy">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="name" type="xsd:string"
        minOccurs="1" maxOccurs="1"/>
      <xsd:element name="policy" minOccurs="1" maxOccurs="1">
        <xsd:simpleType>
          <xsd:restriction base="xsd:string">
            <xsd:enumeration value="allow"/>
            <xsd:enumeration value="deny"/>
          </xsd:restriction>
        </xsd:simpleType>
      </xsd:element>
    </xsd:sequence>
    <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

3.5.13 protected-by

Clients MUST publish this only between: backhaul-interface and ip-address, or backhaul-interface and mac-address

A BHI attaches the `protected-by` metadata type to a link between the BHI's `backhaul-interface` identifier and either an `ip-address` or `mac-address` identifier representing an ICS device or other device within the overlay network (e.g. IPS sensor, etc.). The `protected-by` metadata type is multi-valued, since more than one BHI may protect an ICS device on the same overlay network (for example, if a single overlay network segment has two or more BHIs connecting it to redundant backhaul networks).

A BHI MUST publish this metadata type on a link between its own `backhaul-interface` identifier and the associated `ip-address` or `mac-address` identifier when it observes any network traffic (including both traffic internal to the overlay network segment it protects, and traffic outbound through the BHI onto the overlay network) originating on its protected overlay network segment interface. The BHI MUST NOT publish `protected-by` metadata on devices it does not protect, i.e. devices not on its protected overlay network segment. The BHI MUST publish `protected-by` metadata with a `lifetime` attribute of "session".

In the case where the BHI is a member of only one overlay network, the BHI MUST include an `administrative-domain` attribute (of the linked `ip-address` and `mac-address` identifiers, derived according to the rules in section 3.3.1) that reflects the name of the associated overlay network when publishing `protected-by` metadata on identifiers, in order to prevent confusion when the same IP addresses and MAC addresses are used in different overlay networks. In the case where the BHI is a member of more than one overlay network and is able to determine the source overlay network (i.e., the overlay network where the packets originated) of the ICS Device, the BHI MUST include an `administrative-domain` attribute that reflects the source overlay network when publishing `protected-by` metadata. In the case where the BHI is a member of more than one overlay network and is not able to determine the source overlay network of the ICS Device, the BHI MUST omit the administrative domain attribute when publishing `protected-by` metadata. In the case where the BHI is not assigned to an overlay network, the BHI MUST NOT publish `protected-by` metadata.

An Overlay Manager may use this metadata to better characterize a set of ICS devices within an overlay where the IP addresses and MAC addresses of these devices may not otherwise be known, and to identify which devices are protected by which BHIs on the overlay network for routing and other purposes. This metadata may also be used to trigger alarm events if unexpected IP addresses or MAC addresses are observed on an otherwise steady-state ICS system configuration. This metadata may become stale, as systems move, outages occur, etc.; ICS administrators are recommended to implement a plan to clear/refresh this information on a periodic basis. The BHI SHOULD be configurable with a deletion interval for automated deletion of `protected-by` metadata once it reaches a certain time since it last saw evidence to support that metadata, which SHOULD default to 14 days.

Figure 16 shows an example of this metadata type attached to links between BHI1's `backhaul-interface` identifier and an `ip-address` and a `mac-address` identifier, respectively, for two ICS devices (Dev1 and Dev2) which are present in the overlay network segment protected by that BHI. Since the `ip-address` and `mac-address` identifiers to which the `protected-by` links are attached in this example refer to ICS devices in an identifiable source overlay network, the BHI is required to include an `administrative-domain` attribute of those identifiers derived according to section 3.3.1.

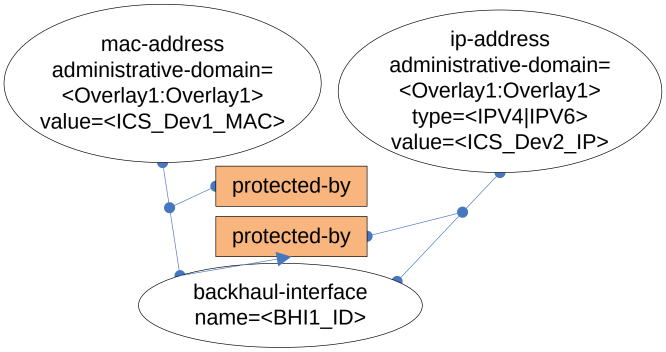

Figure 16: Example MAP graph showing use of protected-by metadata type

```
<xsd:element name="protected-by">
  <xsd:complexType>
    <xsd:attributeGroup
      ref="ifmap:multiValueMetadataAttributes"/>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

4 MAP Client Operations

The definition and operation of overlay networks requires coordination between actors and agents serving in the various roles described in section 3.1. For each of these roles, a subsection below will describe actions to be taken in various scenarios in order to implement one or more working overlay network environments that are coordinated and managed via a shared MAP graph.

4.1 BHI Role

In order to coordinate with other BHIs and other roles, a given BHI is responsible for advertising certain aspects of its own state by publishing metadata into the MAP graph. Additionally, the BHI searches and subscribes to portions of the MAP graph, in order to remain informed of relevant state of peer BHIs and current policy information, and enforce relevant policies found in the MAP graph. Specifics are described in the subsections below.

4.1.1 MAP Server Selection

The BHI has up to three possible methods by which to locate a MAP Server. This specification assumes that the environment contains one MAP Server maintaining a single MAP graph. A situation where multiple MAP Servers are coordinating maintenance of a single MAP graph can be envisioned, but is out of the scope of this specification.

1. **Pre-configuration:** A MAP Server definition, specified by DNS name and/or IP address, which has been provisioned into the BHI by an out of band method.
2. **Previous experience:** A MAP Server that the BHI had previously selected (using method 3 below) and successfully used.
3. **MAP Server discovery:** Selecting a MAP Server using the discovery method mentioned in section 3.2.4.

In the case where a MAP Server is defined using pre-configuration (method 1), the BHI MUST connect to the defined MAP Server. If no successful MAP session can be established to the preconfigured MAP Server, the BHI MUST NOT use methods 2 or 3 to select a MAP Server.

In the case where no MAP Server is pre-configured, the BHI MUST attempt to select the MAP Server using previous experience (method 2). If no information is available for previous experience, or no successful MAP session can be established to any MAP Server located by previous experience, then the BHI MUST attempt to select the MAP Server using MAP Server discovery (method 3).

For methods 1 and 2 listed above, the BHI MUST store its own copy of the expected root certificate and MUST validate the MAP Server at the time of session establishment; if CA chain validation of the server fails, the BHI MUST close and not use the TLS connection.

The Administrator role is responsible for maintaining a set of `distinguished-name` identifiers and associated certificate metadata in the MAP graph for a given MAP Server's CA chain. Note that in the case of the third information source listed above, the BHI may be connecting to a MAP Server for which it has no way to independently verify the server's CA chain. This poses a risk of the BHI connecting to a rogue MAP Server; details of impacts associated with this threat are discussed in section 5 ("Security Considerations").

4.1.2 Maintaining an IF-MAP Session

A BHI SHOULD maintain an IF-MAP session with a MAP Server at all times. When an IF-MAP session fails or is unavailable, the BHI MUST periodically attempt to renew an existing session or establish a new session with a MAP Server using the selection method described in 4.1.1. When the new session is established, the BHI will need to re-establish all of its required published metadata and subscriptions (including the metadata publication addressed in section 4.1.6).

4.1.3 Publishing BHI Internal States

Each BHI MUST publish and maintain the following information as session-lifetime metadata about itself in the MAP graph:

- Current LDevID (or IDevID if no valid LDevID is available) certificate(s) as `bhi-certificate` metadata on its own `distinguished-name` identifier (see section 3.1.3). This is the same certificate that the IF-MAP client uses to establish the IF-MAP session. As certificates are provisioned, updated, and/or revoked, the BHI MUST publish and delete this metadata accordingly.
- `bhi-identity` metadata on a link between its `distinguished-name` identifier and its `backhaul-interface` identifier.
- For each current LDevID (or IDevID if no LDevID is available), `dn-hit` metadata on a link between the `distinguished-name` identifier and the `hip-hit` identifier associated with that certificate.
- For each currently active physical network interface on the backhaul network, `bhi-address` metadata on a link between its `bhi-identity` identifier and each `ip-address` identifier that matches the interface's IP address(es). As interfaces go up and down, or as IP addresses change, the BHI MUST publish and delete metadata on these links accordingly. If a BHI maintains multiple active network interfaces and publishes multiple `bhi-address` metadata links, the BHI MUST respond to protected overlay traffic on any published network interface.
- Self-describing device-characteristic metadata, according to 3.5.5.

Figure 17 shows an example MAP graph with this published metadata for a given BHI called BHI1. Note that the BHI's ability to maintain an IF-MAP session may be affected by interface availability and/or network health. This metadata is published with a lifetime of "session", so that if and when the BHI becomes unavailable, the MAP graph is automatically updated (by the MAP Server purging metadata due to MAP Client disconnection) and subscribers to this information are notified.

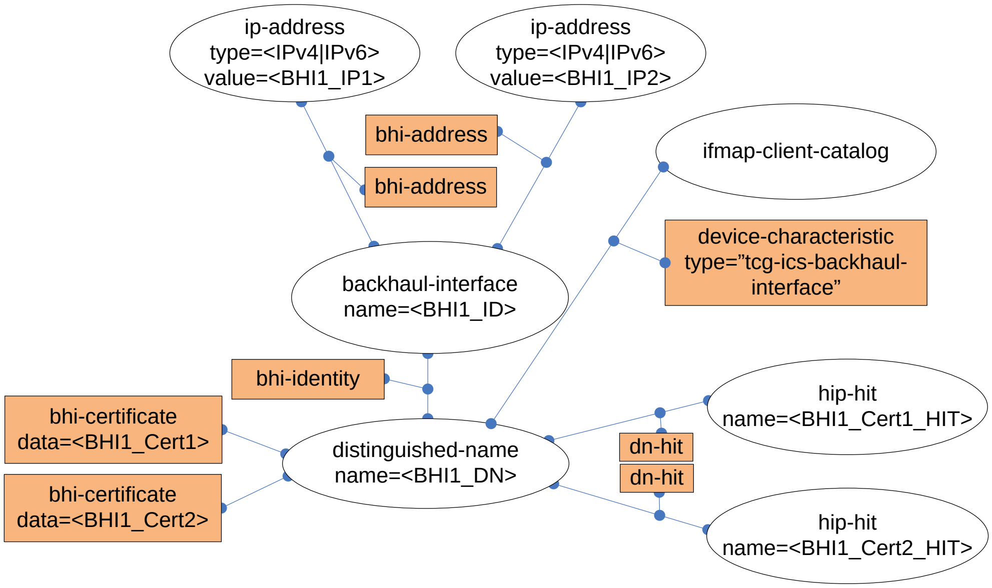

Figure 17: Example MAP Graph Showing BHI-Specific Metadata

4.1.4 Subscribing to Overlay Membership and Backhaul Policy

The BHI MUST create a subscription that allows it to maintain an accurate internal representation of overlay membership and backhaul policy. An example of the search criteria is shown in Table 7 for a BHI whose `backhaul-interface` identifier is BHI1_ID.

Table 7: BHI Search Criteria Example for Overlay Membership and Backhaul Policy

Search Parameter	Value
identifier:	BHI1_ID
match-links:	`member-of` or `backhaul-policy`
max-depth:	2
max-size:	N/A
result-filter:	`member-of` or `backhaul-policy`
terminal-identifier-type:	N/A

The BHI is required to use its MAP-maintained internal representation to enforce policies about which BHI peers it allows itself to communicate with. The details of these enforcement behaviors are described in section 4.1.9.

4.1.5 Subscribing to Overlay Policy

The BHI MUST create a subscription that allows it to maintain an accurate internal representation of the overlay policies for ICS devices on each overlay network segment it protects. An example of the search criteria is shown in Table 8 for a BHI whose `backhaul-interface` identifier is BHI1_ID.

Table 8: BHI Search Criteria Example for Overlay Policy

Search Parameter	Value
identifier:	BHI1_ID
match-links:	`overlay-policy`
max-depth:	1
max-size:	N/A
result-filter:	`overlay-policy`
terminal-identifier-type:	N/A

The BHI is required to use its MAP-maintained internal representation to enforce policies about which ICS devices may communicate through this BHI out onto the overlay network.

4.1.6 Publishing Metadata Describing ICS Devices

As specified in section 3.5.13, the BHI is required to publish `protected-by` metadata on links between its own `backhaul-interface` identifier and `ip-address` identifiers or `mac-address` identifiers of observed ICS devices on the overlay network segment(s) it protects. As specified in section 3.5.11, the BHI is required to publish `observed-by` metadata on links between its own `backhaul-interface` identifier and `ip-address` identifiers or `mac-address` identifiers of observed ICS devices on the rest of the overlay network. In both cases, the `administrative-domain` attribute of the `ip-address` identifier and/or `mac-address` identifier (if included) is required to reflect the overlay network in which the ICS device represented by that `ip-address` identifier or mac-address identifier resides, as well as the IP routing domain or local network segment in which the IP address or MAC address occurs. This information is intended to be of use in monitoring the state of the overlay network by principals operating in the Overlay Manager role.

Upon establishing a new session to a MAP Server, either an initial connection or after a disconnection, the BHI SHOULD re-publish all observed-by and protected-by metadata that is valid (i.e. hasn't been aged out by the deletion interval).

4.1.7 Publishing Locator Metadata for Unassigned BHIs

Each backhaul-interface identifier is expected to be linked by manager-of metadata to a single overlay-manager-group identifier to indicate which Overlay Manager group is managing that BHI. The BHI MUST subscribe to manager-of metadata on its own backhaul-interface identifier to maintain awareness of whether it is assigned to any overlay-manager-group at any given time. As described in section 3.1, MAP Content Authorization will need to allow the BHI to receive subscription results for manager-of metadata. When no such metadata exists, the BHI MUST publish manager-of metadata on a link to a reserved overlay-manager-group identifier with the administrative-domain attribute set to "tcg-reserved" and the name attribute set to "unassigned". This metadata will allow Overlay Managers to identify otherwise-unassigned BHIs that are available for assignment to a particular Overlay Manager group. Note that this unassigned metadata will be deleted by an Overlay Manager when assignment occurs; the BHI need only publish the unassigned metadata if and when it finds itself in that state.

An example of the search criteria with which the BHI can find existing manager-of metadata, to determine whether it is assigned to an Overlay Manager group or unassigned, is shown in Table 9. A BHI is expected to have only one manager at any given time, per section 4.2.2.2.

Table 9: BHI Search Criteria Example for Monitoring Reserved overlay-manager-group Metadata

Search Parameter	Value
identifier:	BHI1_ID
match-links:	manager-of
max-depth:	1
max-size:	N/A
result-filter:	N/A
terminal-identifier-type:	N/A

Any BHI in the unassigned state MUST NOT perform any overlay network functions even if it is a member of one or more overlay network groups. This situation is not expected to occur under normal operations, but this requirement provides consistent behavior if and when it occurs, such as when a BHI enters a factory-reset condition.

4.1.8 Protection of Overlay Network Traffic over the Backhaul Network

Because the backhaul network is untrusted, BHIs MUST protect the confidentiality and integrity of overlay network traffic when it transits the backhaul network. Further, BHIs MUST provide and verify source authentication and replay prevention for this traffic.

4.1.9 Policy Enforcement Responsibilities

In order for the BHI to send a packet across the overlay network, both backhaul policy and overlay policy are required to allow the packet to be sent. Each of these two policy controls is necessary, but neither is individually sufficient.

When evaluating policy, there are two general principles for overlay policy prioritization:

- More specific policy overrides more general policy
- Default policy applies in the absence of specific policy.

4.1.9.1 Backhaul Policy Prioritization

BHI backhaul policy enforcement for a given backhaul network (as defined by the associated `overlay-network-group` identifier) is controlled by `backhaul-policy` metadata whose names match the associated `overlay-network-group` identifier. BHI backhaul policy decisions are prioritized as follows (highest priority first):

1. When `backhaul-policy` metadata whose name value matches the given `backhaul-network-group` identifier name is attached to a link between `backhaul-interface` identifiers representing two BHIs, as defined in section 3.5.1.2 (Pairwise Backhaul Policy), each BHI MUST enforce the specified policy for communications with its peer BHI on that backhaul network.
2. When `backhaul-policy` metadata is attached to an `overlay-network-group` identifier, as defined in section 3.5.1.1 (Default Backhaul Policy), all BHIs that are members of that overlay network MUST enforce the specified policy for all communications on the backhaul network, unless overridden by specific `backhaul-policy` metadata attached to a link as described in priority 1 above.
3. In the absence of `backhaul-policy` metadata attached to a given `overlay-network-group` identifier (priority 2 above), all BHIs that are members of that overlay network MUST treat that overlay network as having an implicit Default Backhaul Policy of “deny”.

If multiple conflicting backhaul-policy metadata are present for priority 1 and/or priority 2 above, the BHI MUST enforce the more restrictive policy (“deny”) at that priority level.

In the case where two BHIs are both members of two overlay networks (e.g. Overlay_A and Overlay_B), and the two overlay networks have different default backhaul-policy metadata, the BHI is expected to enforce the appropriate policy for a given overlay as long as the BHI can determine whether specific traffic between the two BHIs is intended to be part of Overlay_A or Overlay_B, such as by IP address or VLAN tag of the originating ICS device. If the traffic cannot be differentiated, the BHI MUST enforce the more restrictive policy.

An example of the use of the backhaul-policy metadata type is shown in Figure 18. A default “allow” `backhaul-policy` is attached to the Overlay_A overlay network, which means that all BHIs who are members of this overlay may communicate with each other over the backhaul unless overridden for specific BHI pairs. For example, the `backhaul-policy` attached to the link between BHI1 and BHI2 overrides this default “allow” policy by denying communication between BHI1 and BHI2 for Overlay_A.

In contrast, no default `backhaul-policy` metadata is attached to the Overlay_B `overlay-network-group` identifier in Figure 18, which means that the default policy is “deny” unless overridden for particular BHI pairs. The presence of a pairwise backhaul policy - the link with the `backhaul-policy` metadata type between BHI1 and BHI3 in this figure - overrides the default “deny” policy for this overlay network and allows these two BHIs to communicate over the backhaul when carrying that overlay’s network traffic. `backhaul-policy` metadata containing a specification for a particular overlay network is meaningless for a given BHI if that BHI is not a member of that overlay (see section 3.5.10 for a description of how the `member-of` metadata type is used).

Also note the multiple `backhaul-policy` metadata on the link between BHI2 and BHI3 where both an “allow” and a “deny” policy are specified for Overlay_A. As defined earlier, this example will result in a “deny” policy (“deny” always wins in the case of ambiguity).

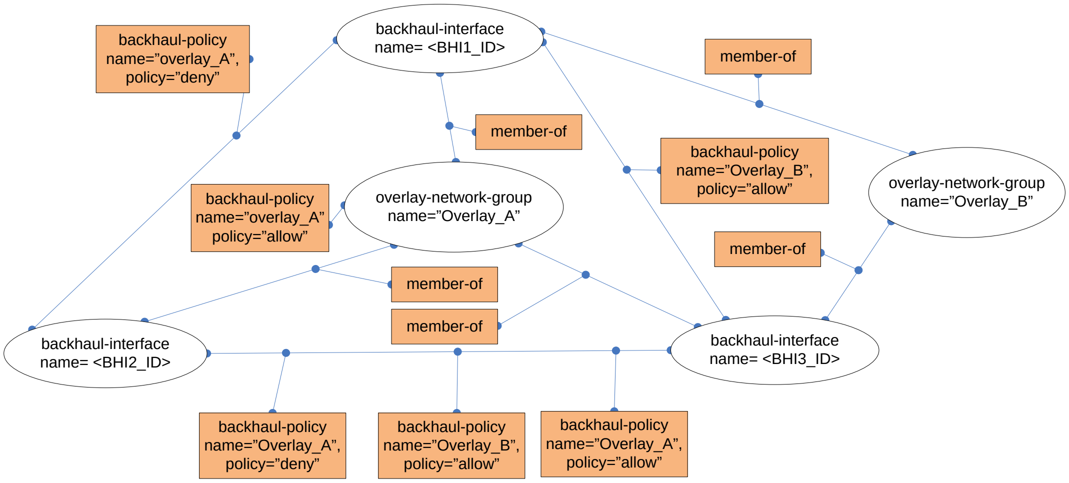
backhaul-policy name=”overlay_A”, policy=”deny”
backhaul-interface name= <BHI1_ID>
member-of
member-of
backhaul-policy name=”Overlay_B”, policy=”allow”
backhaul-policy name=”overlay_A” policy=”allow”
overlay-network-group name=”Overlay_A”
overlay-network-group name=”Overlay_B”
member-of
member-of
member-of
backhaul-interface name= <BHI2_ID>
backhaul-interface name= <BHI3_ID>
backhaul-policy name=”Overlay_A”, policy=”deny”
backhaul-policy name=”Overlay_B”, policy=”allow”
backhaul-policy name=”Overlay_A”, policy=”allow”

Figure 18: Example MAP graph for application of backhaul policy

4.1.9.2 Overlay Policy Prioritization

Based on the general principles of prioritization, overlay policies are prioritized as follows:

1. When `overlay-policy` metadata is attached to a link between an identifier representing a given BHI and an `ip-address` or `mac-address` identifier, that BHI MUST enforce the specified policy for communications from that IP address or MAC address.
2. When `overlay-policy` metadata is attached to a backhaul-interface identifier representing a given BHI, that BHI MUST enforce the specified policy for communications from all ICS devices on the identified overlay network, unless overridden by specific `overlay-policy` metadata attached to a link as described in 3.5.12.2.
3. In the absence of `overlay-policy` metadata attached to a given `backhaul-interface` identifier, that BHI MUST enforce a global default “deny” policy for communications from all ICS devices, unless overridden by specific `overlay-policy` metadata attached to a link as described in 3.5.12.2.

Note that as described in 2.2.1, the overlay policy is intended to restrict only outbound traffic (traffic originating from its protected overlay network segment and destined for devices on the overlay that are protected by other BHIs), not inbound traffic.

Since IP packets are a specific type of layer-2 packet associated with a given `mac-address` identifier, `overlay-policy` metadata attached to links to `ip-address` identifiers are considered more specific than `overlay-policy` metadata attached to links to `mac-address` identifiers. As a result, in the case where “allow” `overlay-policy` metadata is present on a link to an `ip-address` identifier and “deny” `overlay-policy` metadata is present on a link to a `mac-address` identifier associated with the same ICS device, then the BHI MUST only pass the allowed IP packets (which will generally contain the denied MAC address within their layer-2 header source address parameter), and the BHI MUST NOT forward all other layer-2 packets with that source MAC address value, because policy applied at L3 is more specific than – and thus overrides – policy applied at L2. Conversely, in the case where “deny” `overlay-policy` `metadata` is present on a link to an `ip-address` identifier and an “allow” `overlay-policy` is present on a link to a `mac-address` identifier associated with the same ICS device, then the BHI MUST pass all types of layer-2 traffic from that MAC address except IP traffic from that IP address, for the same reason..

In the case where multiple instances of this metadata type are present with different combinations of their name & policy values, the BHI MUST consider only instances whose name value matches the overlay network associated with affected ICS devices. In the case where the name value in the metadata does not match the name of the relevant overlay network, then the BHI MUST apply the default policy for that particular overlay network. In the case where multiple metadata instances match a given overlay name but have different policy values (i.e., some are “allow” and some are “deny”), then the BHI MUST interpret the resulting policy as “deny”.

Figure 19 shows an example of this metadata type in use. In this example, the presence of the `overlay-policy` metadata with `state` set to “allow” attached to the link between the BHI1_ID `backhaul-interface` identifier and the Overlay1:ICS_Dev1_IP `ip-address` identifier indicates that BHI1 will provide layer 3 (IP) overlay services for Dev1’s overlay IP address. Similarly, the presence of `overlay-policy` metadata with `state` set to “allow” attached to the link between the BHI_1ID `backhaul-interface` identifier and the Overlay1:ICS_Dev2_MAC `mac-address` identifier indicates that BHI1 will provide layer 2 overlay services for ICS_Dev2.

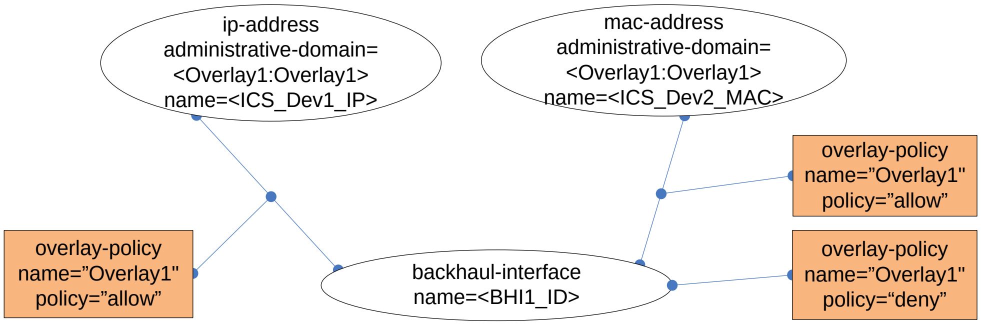

Figure 19: Example MAP graph for application of overlay policy

4.2 Overlay Manager Role

The Overlay Manager role enables the secure management of multiple overlay networks. Principals (users, applications, etc.) operating in this role will have authorization to modify the configuration of one or more independent overlay networks according to membership relations expressed in the MAP graph as well as authorization constraints described in this section.

4.2.1 Management Model

Authorization constraints are divided into four sets of constraints upon roles (discussed in section 4.2.2) based upon the concept of management by Overlay Manager groups. By allowing different Overlay Manager groups to be assigned to different overlay networks, distinct groups of principals are each able to manage their own overlay network(s) without risking one group negatively impacting another’s overlay. This is particularly important for large enterprise environments where thousands of BHIs may be deployed in large numbers of separate overlay networks.

In addition to the concept of management for overlay networks, a separate management type is defined for the BHIs. This is necessary because a single BHI may be used as a member of multiple overlay networks (e.g., a single BHI with multiple overlay network segment interface ports); if BHI management were simply part of the overlay network management, then principals from one

Overlay Manager could create, modify, or delete BHI membership in overlay networks of which they were not managers.

4.2.1.1 Manager of a BHI

BHI management is conferred by a `manager-of` metadata published by an Overlay Manager on a link between an `overlay-manager-group` identifier and a `backhaul-interface` identifier as described in section 3.5.9.2.

4.2.1.2 Manager of an Overlay Network

Overlay network management is conferred by a `manager-of` metadata published by an Administrator on a link between an `overlay-manager-group` identifier and an `overlay-network-group` identifier as described in section 3.5.9.1.

4.2.2 Authorization Requirements

An ICS Administrative MAP Client MUST NOT allow Overlay Managers to create or delete overlays or manage assignment of Overlay Manager groups to a given overlay; these functions are reserved for the Administrator role (see section 4.3).

Based upon the two types of management defined above, a set of four constraints upon roles are described in this section, each with restrictions on which metadata and associated links it may create, modify or delete.

4.2.2.1 Constraints upon Management of Overlay Network Management

Principals operating in an Overlay Manager role are not permitted to alter the management privileges assigned to managers of an overlay network. Specifically these principals MUST NOT create, modify, or delete `manager-of` metadata on links to an `overlay-network-group` identifier. These operations are instead reserved for the Administrator role (discussed in section 4.3).

4.2.2.2 Constraints upon Management of BHI Management

Releasing an assigned BHI: If a given BHI is currently assigned to a given Overlay Manager Group, a principal acting in the corresponding Overlay Manager role may release that assignment if and only if the BHI is no longer a member of any overlay network (i.e., it has no `member-of` metadata on links between its `backhaul-interface` identifier and any `overlay-network-group` identifiers). To release a BHI, the Overlay Manager MUST delete the existing `manager-of` metadata on the link to that particular `overlay-manager-group` identifier. The BHI will be responsible for publishing new metadata on a link to the reserved “unassigned” identifier (see section 4.1.7) when this deletion occurs.

Claiming an unassigned BHI: If a given BHI is currently unassigned (see section 4.1.7), then a principal who is a member of a particular Overlay Manager group may claim management of that BHI for that group. To claim a BHI, the Overlay Manager MUST perform the following MAP changes, which MUST be done as a single atomic IF-MAP request:

- Delete the existing `manager-of` metadata on the link between the BHI’s `backhaul-interface` identifier and the reserved “unassigned” `overlay-manager-group` identifier described in section 4.1.7.
- Publish new `manager-of` metadata on the link between the BHI’s `backhaul-interface` identifier and the `overlay-manager-group` identifier to which the BHI is being assigned.

The atomic operation is required to prevent the BHI from re-establishing its “unassigned” `manager-of` metadata after the Overlay Manager deletes it, but before the Overlay Manager publishes the new manager-of metadata. Essentially, this is a first-come, first-served approach whereby the first attempt to claim a given unassigned BHI will win and all others will fail.

Note that an Overlay Manager MUST NOT claim a BHI that has manager-of metadata on a link to the “unassigned” overlay-manager-group identifier if the BHI is also a member of one or more overlay networks. This situation is not expected to occur under normal operations, but this requirement will provide consistent behavior if and when it does occur, such as when a BHI enters a factory-reset condition. Only the Administrator role (defined in section 4.3) has the ability to correct this erroneous condition.

4.2.2.3 Constraints upon Management of Overlay Network Membership

The member-of metadata on links between backhaul-interface identifiers and overlay-network-group identifiers determine a given BHI’s membership in a given overlay network. An ICS Administrative MAP Client MUST allow a principal acting in the Overlay Manager role to add or delete the member-of metadata on these links only when the principal is a member of the overlay manager group(s) that own(s) both the overlay network and the BHI.

4.2.2.4 Constraints upon Management of Overlay Network Configuration

A principal acting in the Overlay Manager role who has management of a given overlay network may manage overlay-specific policy rules. To manage these rules, the Overlay Manager creates, modifies, and deletes backhaul-policy and overlay-policy metadata as needed. Under normal operations, these changes will immediately affect the behaviors of the associated BHIs according to the policies which these metadata define.

Note that as a result, an Overlay Manager who has manager rights to an overlay network but does not have manager rights to a particular BHI on that network can still publish overlay-policy metadata on links directly to that BHI. This is not problematic since a given policy would only be applicable to traffic over the overlay network specified in the overlay-policy metadata; thus, two managers of different networks are not going to create conflicting policies.

4.2.3 Example

Figure 20 shows an example MAP graph of an overlay network named Overlay_1 and three BHIs named BHI1, BHI2, and BHI3. Both BHI1 and BHI3 are members of Overlay_1, but BHI2 is not a member of any overlay network. The overlay network is owned by the Overlay Manager group named OMGroup_1, as are BHI1 and BHI2. BHI3 is owned by the Overlay Manager group named OMGroup_2.

As specified in 4.2.2.1, a principal who is a member of OMGroup_1 can add or remove BHI1 and/or BHI2 from Overlay_1, because the overlay network and the BHIs are all owned by the same Overlay Manager group. However, only Principal_2 in this example can add or delete BHI3 to/from Overlay_1, since only this principal is a member of both OMGroup_1 (which owns the overlay) and OMGroup_2 (which owns BHI3).

As specified in 4.2.2.1, no principals shown in this example can add or delete manager-of metadata on links between an Overlay Manager group and an overlay network (e.g., the manager-of metadata on the link between OMGroup_1 and Overlay_1).

As specified in 4.2.2.2**,** Principal_2 and Principal_3 can release management of BHI3, but only if it is first removed from the Overlay_1 overlay network. Only Principal_2 can do this removal from Overlay_1 since it is the only principal who is a member of the two overlay manager groups (OMGroup_1 and OMGroup_2) that own Overlay_1 and BHI3. An Administrator can also perform these operations.

As specified in 4.2.2.4, Principal_1, Principal_2, and any Principal in the LDAP group referenced by the Group1_URI can alter backhaul-policy and overlay-policy link metadata for any

BHI that is a member of Overlay_1. This is true even for BHI3, despite the fact that these principals are not managers of BHI3, because BHI management is only required for adding/deleting overlay membership and for releasing/claiming BHI management.

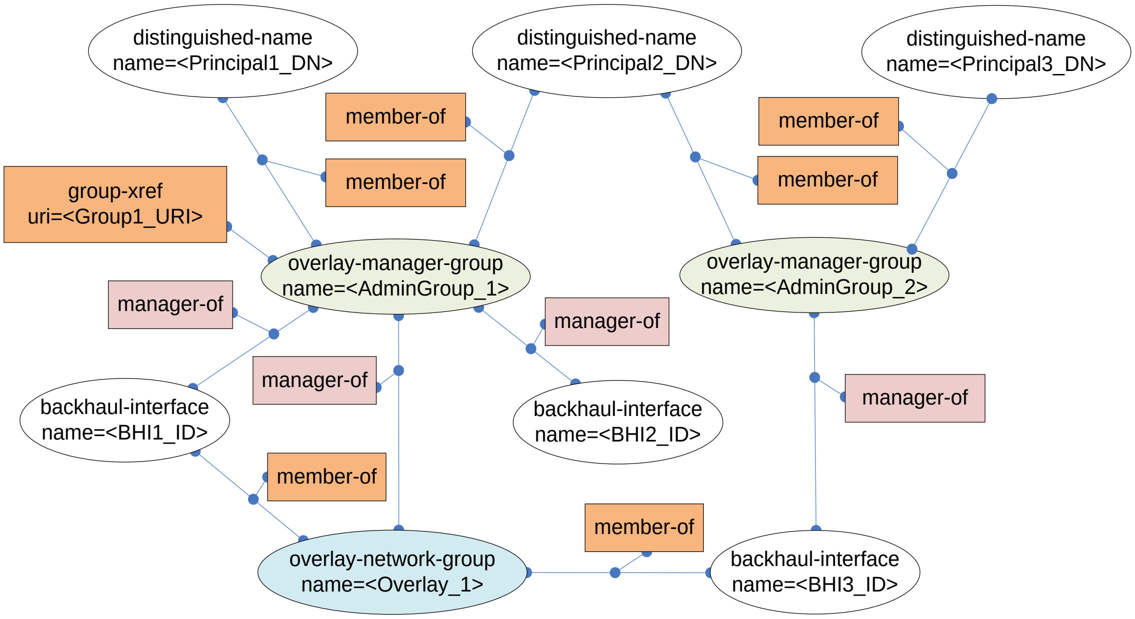

Figure 20: Example Graph for Overlay Manager Sub-roles

4.3 Administrator Role

The Administrator role defines overlay networks. To define an overlay network, the Administrator MUST publish manager-of metadata on links between overlay-manager-group identifiers and overlay-network-group identifiers. The Administrator MUST NOT delete manager-of metadata if the overlay-network-group identifier has any member-of metadata on links to backhaul-interface identifiers.

The Administrator role is also used to manage membership of principals in overlay manager groups. To assign individual principals as managers of an overlay network group, the Administrator MUST publish member-of metadata attached to links between overlay-manager-group identifiers and distinguished-name identifiers. To assign an externally-defined group of principals as managers of an overlay network group, the Administrator MUST publish group-xref metadata attached to an overlay-manager-group identifier.

In Figure 20, the manager-of metadata on the link between OMGroup_1 and Overlay_1 could only have been created by a principal operating in the Administrator role. Similarly, only principals operating with the Administrator role could have created the member-of metadata on links to Principal_1, Principal_2 and Principal_3, as well as the group-xref metadata in the figure.

The Administrator role is also used to correct any unusual conditions that may arise in the MAP graph. For example, if a BHI has no manager, but is a member of one or more overlay networks, the Administrator will need to either delete the BHI from those overlay networks, or manually assign management of the BHI to an appropriate Overlay Manager group. Another example (as described in section 4.2.2.2) is the case where a backhaul-interface identifier has manager-of metadata on a link to the “unassigned” overlay-manager-group identifier, but the BHI in question is a member of one or more overlay networks. In that case, the Administrator will need to either delete the BHI from those overlay networks, or delete the `manager-of` metadata on the link to the “unassigned” overlay manager group and assign the BHI to an appropriate overlay manager group.

These conditions should not occur, but the Administrator role is available to correct such conditions should they arise. This is a particularly important function to have available when facing the management of a large number of BHIs and overlay networks, as might be found in a typical large Enterprise deployment.

4.4 Implementation of Policy Enforcement for Roles

4.4.1 Policy Enforcement for the BHI Role

The prescribed MAP operations for the BHI role (section 4.1) will be performed by the BHI itself. The BHI MUST NOT attempt MAP operations described in this specification which are outside of the scope of MAP operations specified in section 4.1.

In addition, MAP Content Authorization policies SHOULD be defined in the MAP Server such that a given BHI cannot perform MAP operations described in this specification which are outside of its prescribed scope.

4.4.2 Policy Enforcement for the Overlay Manager Role

The prescribed MAP operations for the Overlay Manager role will be performed by ICS Administrative MAP Clients. These clients MUST NOT attempt MAP operations which are outside of the scope of the MAP operations described in section 4.2 and 4.3.

An ICS Administrative MAP Client MUST authenticate principals and perform MAP operations on their behalf according to the constraints defined for their roles (as specified in section 4.2.2). Such a MAP Client MUST NOT attempt MAP operations on behalf of an Overlay Manager principal if the operations are outside of the scope of the Overlay Manager’s role constraints for that particular principal.

5 Security Considerations

By securing communications sent by ICS devices and protecting them from unauthorized access without requiring changes to the devices themselves, IF-MAP Metadata for ICS Security should substantially increase the security of industrial control systems, thereby ensuring the safety of these systems while reducing configuration and operational management costs. However, no security system is perfect so it's important to understand the strengths, weaknesses, and assumptions in each system.

This section aims to provide a thorough security analysis of IF-MAP Metadata for ICS Security and of systems that employ this specification. 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). Some of these countermeasures are required. Others are optional, since different customers may have different threat environments and different countermeasures already in place.

5.1 Trust Model

The first step in analyzing the security of IF-MAP Metadata for ICS Security is to describe the trust model, listing what each architectural element is trusted to do. The elements here are assumed to comply with the described behaviors, but provisions are made in the Threat Model and Countermeasures sections for handling elements that fail to perform as they were trusted to do.

5.1.1 Backhaul Network

The backhaul network is trusted to:

- Perform best effort delivery of network traffic

The backhaul network is not expected (not trusted) to:

- Ensure the confidentiality or integrity of messages sent over it
- Provide timely or reliable service

5.1.2 Overlay Network

The overlay network is trusted to:

- Perform best effort delivery of network traffic
- Protect the confidentiality and integrity of messages sent over it from external attack

The overlay network is not expected (not trusted) to:

- Provide timely or reliable service

5.1.3 MAP Clients

As described in the TNC IF-MAP Binding for SOAP[2] Security Considerations, all authorized MAP Clients are trusted to:

- Verify the identity (via CA chain validation) of the MAP Server to which they connect
- Preserve the confidentiality of sensitive data retrieved from the MAP Server
- Ensure the accuracy of data in the MAP Server database, by avoiding database contamination 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

5.1.4 Backhaul Interfaces (BHIs)

BHIs, being MAP Clients, are trusted to meet all the expectations of MAP Clients and also to:

- Deliver network traffic sent by ICS devices within an overlay network in accordance with policies stored in the MAP
- Protect the confidentiality, integrity, and availability of that network traffic
- Update the MAP based on observations of ICS connections as well as its own overlay membership

5.1.5 Factory-reset BHI

When a BHI is reset to its factory default state, the BHI is trusted to meet all the expectations of MAP Clients and also to:

- Clear any cached list of preferred, trusted, or discovered MAP Servers
- Clear any cached metadata
- Remove all LDevIDs and associated private keys
- Disconnect from the MAP Server

5.1.6 ICS Administrative MAP Client

ICS Administrative MAP Clients are trusted to meet all of the expectations of MAP Clients and also to:

- Change the MAP graph on behalf of Overlay Manager and Administrator principals
- Do so only when the principal is authorized to perform the requested change

5.1.7 Overlay Manager

Overlay Managers are trusted to:

- Establish policies for overlay networks managed by the overlay manager and store those policies in the MAP
- Join BHIs to overlay networks
- Claim and release BHIs from their management purview in accordance with the rules set forth in this specification
- Manage the management lifecycle of BHIs

5.1.8 Administrators

Administrators are trusted to:

- Create, delete, and rename overlay networks
- Assign overlay managers to overlay networks
- Identify unusual conditions within the MAP graph described in this specification and correct them

5.1.9 ICS Devices

ICS devices are trusted to:

- Send traffic to any ICS device in its overlay network (except as constrained by BHI policy)
- Not impersonate other ICS devices protected by the same BHI
- Not bridge or forward traffic from a given overlay network to other networks
- Not allow traffic from other networks into the overlay network

5.1.10 Certification Authority (CA) Infrastructure

Certification Authorities (CAs) and their supporting infrastructure are responsible for issuing and revoking the certificates used to authenticate BHIs, Overlay Managers, Administrators, other MAP Clients, and the MAP Server. The CA infrastructure is trusted to:

- Certify BHIs and other MAP Clients only in accordance with CA policies

5.1.11 MAP Server

As described in the IF-MAP Security Considerations, 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.

The MAP Server SHOULD enforce TNC MAP Content Authorization[14].

5.2 Threat Model

To secure the IF-MAP Metadata for ICS Security protocol and the architectural elements that implement it, this section identifies the attacks that can be mounted against the protocol and elements. This section should not be considered a complete list of such attacks. Rather, it is an attempt to highlight the most significant attacks in each area.

5.2.1 Backhaul Network and Overlay Network Attacks

A variety of attacks can be mounted using a 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
- Network traffic may be relayed to or from another network in violation of expected norms
- 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
- Network services upon which the MAP Clients may depend may be compromised; for example, abuse of NTP in an attempt to subvert certificate validation of an expired certificates

All these attacks can be mounted against the backhaul network or the overlay network. However, the impact of attacks on the backhaul network is substantially reduced by the countermeasures recommended in section 5.3, especially those countermeasures provided by the BHIs.

5.2.2 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 manager is malicious, careless, or incompetent or because the MAP Client has been compromised and is now controlled by a party other than its manager); 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 corrupts data sent to the MAP); 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.

For a complete list of attacks that may be mounted by an authorized MAP Client, see the IF-MAP specification. Here is a more specific list of attacks relevant to the ICS security context:

- Act as an Overlay Manager or Administrator, mounting the attacks listed in sections 5.2.5 and 5.2.6
- Modify metadata that is to be published by a BHI, mounting some of the attacks listed in section 5.2.3
- Read some or all of the metadata described in this specification and use the information contained therein to plan and execute an attack on an industrial control system

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 such as compromising the authentication system or CA infrastructure).

5.2.3 Backhaul Interfaces (BHIs)

A compromised or malicious BHI can cause serious problems unless countermeasures are in place. Because an authorized BHI is also an authorized MAP client, it can mount the attacks described in section 5.2.2. In addition, a BHI may mount the following attacks relevant to its role in the ICS security context:

- Generate or modify network traffic on the overlay network or backhaul network
- Release network traffic from the overlay network to unauthorized parties
- Inject network traffic from unauthorized parties into the overlay network
- Permit network traffic to flow on the overlay network without the proper security constraints as described in section 4.1.8.
- Report false information about which ICS devices are connected to the BHI

Dependencies of or vulnerabilities of authorized BHIs may be exploited to effect these attacks. Another way to effect these attacks is to gain the ability to impersonate a BHI (through theft of the BHI's identity credentials or through other means such as compromising the authentication system or CA infrastructure).

5.2.4 ICS Administrative MAP Client

An Administrative MAP Client has elevated privileges for managing the lifecycle of overlay networks, policy settings, etc., as well as correcting any configuration problems that cannot be addressed by lower-privileged clients (BHIs, overlay managers). Therefore, a compromised or malicious Administrative MAP Client can cause serious problems such as the unauthorized creation, alteration or deletion of overlay networks, as well as unauthorized changes to backhaul and overlay policies that allow unauthorized communications to ICS devices and/or denial of service.

5.2.5 Overlay Managers

An Overlay Manager decides which BHIs should belong to which overlay networks and what policies should be used on those overlay networks. An Overlay Manager generally includes several software components and one or more human beings who operate these software components to effect the desired actions. The responsibilities and privileges of an Overlay Manager are generally limited to a small number of overlay networks.

A compromised or malicious Overlay Manager or an error or deliberate action by the human or humans in charge of the overlay network can cause serious problems unless countermeasures are in place. An Overlay Manager may mount the following attacks relevant to its role in the ICS security context:

- Add inappropriate BHIs to overlay networks. If these BHIs are compromised or malicious, those BHIs may mount the attacks listed in section 5.2.3.
- Add valid BHIs to improper overlay networks, causing those BHIs and perhaps the ICS devices behind them to be exposed to a variety of network-based attacks
- Improperly claim an unassigned BHI, depleting the stock of BHIs and exposing the ICS devices behind that BHI to a variety of network-based attacks
- Remove BHIs from the overlay networks to which they have been joined
- Install improper overlay policies or backhaul policies, resulting in incorrect communications on the overlay network. This may simply disrupt communications on the overlay network or it may expose BHIs and ICS devices to malicious attacks or disclose traffic on the overlay network to unauthorized parties.

Dependencies of or vulnerabilities of authorized Overlay Managers may be exploited to effect these attacks. If an authorized human is involved with operating or managing an Overlay Manager, this human may be a source of security vulnerabilities. The human may be malicious, or they may be pressured or tricked into making changes that result in the vulnerabilities listed above. Another way to effect attacks using an Overlay Manager is to gain the ability to impersonate an Overlay Manager (through theft of the Overlay Manager's identity credentials or through other means such as compromising the authentication system or CA infrastructure).

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

5.2.6 Administrators

An Administrator decides which overlay networks should be created, which Overlay Managers should be able to manage these networks, which conditions of the MAP graph are unusual, and how those unusual conditions should be handled. An Administrator generally includes several software components and one or more human beings who operate these software components to effect the desired actions. The privileges of an Administrator generally cover all Overlay Networks administered through a given MAP.

Compromised or malicious Administrator software or an error or deliberate action by the human in charge can cause serious problems unless countermeasures are in place. An Administrator may mount the following attacks relevant to its role in the ICS security context:

- Create unauthorized overlay networks
- Delete or rename overlay networks in an unauthorized manner
- Improperly assign groups of Overlay Managers to overlay networks
- Improperly configure the members of an Overlay Manager group

Because an Administrator has the ability to assign groups of Overlay Managers to overlay networks and configure the members of an Overlay Manager group, an Administrator can assign an unauthorized Overlay Manager to any overlay network. However, that Overlay Manager will still need to be accepted by the MAP Server as an authorized MAP Client.

Dependencies of or vulnerabilities of authorized Administrators may be exploited to effect these attacks. If an authorized human is involved with operating or managing an Administrator, this human may be a source of security vulnerabilities. The human may be malicious, or they may be pressured or tricked into making changes that result in the vulnerabilities listed above. Another way to effect attacks using an Administrator is to gain the ability to impersonate an Administrator (through theft of the Administrator's identity credentials or through other means such as compromising the authentication system or CA infrastructure).

An unauthorized Administrator (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.

5.2.7 ICS Devices

ICS devices without BHI functionality do not access the MAP Server. They simply send and receive network traffic on the overlay network. Still, each ICS device has access to the overlay network, so a compromised or malicious ICS device can cause trouble. Because an ICS device is a device on the overlay network, any of the attacks listed in section 5.2.1 may be relevant. However, here is a list of especially significant attacks that may be mounted by an ICS device:

- Attack other devices on the overlay network
- Impersonate other devices on the overlay network (by using their IP or MAC address)

- Bridge or forward traffic from the overlay network to or from other networks

Dependencies of or vulnerabilities of ICS devices may be exploited to effect these attacks. Because ICS devices are generally not authenticated, any device with access to the overlay network should be considered an ICS device. For this reason, suitable access controls should be applied to physical and network access to the overlay network, and the physical and logical integrity of ICS devices should be protected to prevent them from becoming compromised.

5.2.8 Certification Authority (CA) Infrastructure

If the CA infrastructure is compromised, several very serious attacks are possible:

- Issue certificates to the wrong parties, enabling them to impersonate BHIs, Overlay Managers, Administrators, or other MAP Clients to the MAP Server
- Issue certificates to an unauthorized MAP Server
- Revoke certificates issued to valid parties, resulting in denial of access

Effectively, compromise of the CA infrastructure enables all the attacks listed anywhere in this Security Considerations section to be mounted.

5.2.9 Registration Authority (RA)

Because the CA infrastructure takes its instructions from the RAs, compromise of an RA can result in the CA issuing a certificate to an unauthorized party, leading to all the consequences described in section 5.2.8.

5.2.10 MAP Server

As noted in the Security Considerations in the IF-MAP specification, a compromised or malicious MAP Server can cause all the problems that any authorized MAP Client could cause. It can even mount several other attacks described in that specification to cover its tracks or cause general confusion. Therefore, protecting the MAP Server from compromise is essential.

5.3 Countermeasures

Adopting certain fundamental best practices and countermeasures can provide substantial protection against the attacks listed in the previous section. All countermeasures outlined in [2] section 6.3 are included here by reference. Also, specific additional countermeasures apply in an Industrial Control Systems environment.

5.3.1 Securing the MAP Server

The IF-MAP specification provides best practices for securing the MAP Server. All of these countermeasures and best practices should be followed to the letter. As the foundation of the security provided by this specification, the MAP Server must be properly secured.

Because of the substantial security impact of using a compromised or malicious MAP Server (as described in section 5.2.10), the BHIs and all other MAP clients must be careful to connect only to authorized MAP Servers. Therefore, the requirements in section 4.1.1 must be followed carefully. MAP Server selection method 3 (as described in section 4.1.1) should be avoided since it provides little protection against choosing an unauthorized MAP Server.

5.3.2 Securing MAP Clients

As noted in the IF-MAP specification, MAP clients are often widely distributed with little physical security. This is especially true with IF-MAP Metadata for ICS Security. BHIs may be deployed in dangerous and unguarded physical environments. Fortunately, this danger can be reduced tremendously by restricting the privileges of MAP Clients with MAP Content Authorization policies. BHIs should have very limited access. As the level of access increases, the security measures imposed on these components should increase. Increased security controls are particularly important to overlay manager MAP Clients and administrative MAP Clients since they have additional privileges to alter overlay and backhaul network policies and even create and delete entire overlay networks. These additional security controls may include physical security requirements, monitoring and logging, rate limiting, and requirements for multiple authentications and authorizations. In addition, MAP Clients should use secure versions of common network infrastructure protocols, such as DNS and NTP, wherever possible.

5.3.3 Securing Backhaul Network

Due to the security protections provided by the BHIs, the impact of attacks from the backhaul network is greatly reduced. Tampering with this network may cause interference with ICS communications but this interference can be detected with logging. And even with malicious access to the backhaul network, the integrity and confidentiality of communications among ICS devices will be preserved, avoiding the possibility of compromise of these devices.

Still, interference with communications among ICS devices could cause considerable damage, so it's advisable to evaluate the impact of such interference. For example, if communications between an operator workstation and a PLC controlling a pump are disrupted, the tank regulated by that pump could overflow. If this impact is unacceptable, a backup for the backhaul network may be needed, or some mechanism (e.g. keepalives) for detecting disruptions to backhaul network communications and moving into a failsafe mode if that occurs.

5.3.4 Securing Overlay Network, BHIs, and ICS Devices

The BHIs protect the overlay network from external attacks. However, they do not protect against attacks from within the overlay network. Attacks from inside the overlay network can be avoided by securing the BHIs and ICS devices from compromise. Such attacks can be detected with the use of common network security techniques such as network intrusion detection systems, and their impact can be mitigated by restricting connectivity within the overlay network.

5.3.5 Securing Overlay Managers and Administrators

To reduce the likelihood of Overlay Manager or Administrator compromise, keep the number of Overlay Managers and Administrators to a minimum, ensure that the hardware and software employed are securely configured and managed, and conduct background checks on human operators. To detect malicious actions by human operators, monitor and log their actions employing anomaly detection to detect unusual patterns of behavior. To mitigate the impact of Overlay Manager or Administrator compromise, require multiple parties to independently authorize actions with a large impact such as creating or deleting an overlay network.

5.3.6 Securing the CA Infrastructure and RAs

Best practices for securing CAs are well established and documented across the industry. Among other things, CAs should be offline when possible. All CA actions should be logged with an immutable log. And different CAs should be employed for different types of certificates with different impact and assurance levels.

6 Privacy Considerations

Privacy considerations are not a substantial factor in an Industrial Control Systems environment.

7 References

[1] Trusted Computing Group, *TNC Architecture for Interoperability*, Revision 1.5, May 2012.

[2] Trusted Computing Group, *TNC IF-MAP Binding for SOAP*, Revision 2.2, March 2014

[3] Trusted Computing Group, *TNC IF-MAP Metadata for Network Security*, Revision 1.1, May 2012.

[4] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*, RFC 2119, Best Practices, March 1997, IETF.

[5] International Telecommunication Union, Telecommunication Sector (ITU-T), Recommendation X.690, “Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)”, July 2002.

[6] T. Henderson, S. Venema, D. Mattes, *HIP-based Virtual Private LAN Service (HIPLS)*, draft-henderson-hip-vpls-06, June 2013, IETF.

[7] T. Henderson, *HIP as a VPLS Solution*, http://www.ietf.org/proceedings/77/slides/HIPRG-6.pdf

[8] D. Mattes, *HIP VPLS at Boeing*, http://www.ietf.org/proceedings/81/slides/HIPRG-2.pdf

[9] M. Cotton, L. Eggert, J. Touch, M. Westerlund, S. Cheshire, Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry, RFC 6335, Best Current Practice, August 2011, IETF.

[10] N. Freed, N. Borenstein, Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies, RFC 2045, Standards Track, November 1996 IETF.

[11] R. Moskowitz, P. Nikander, *Host Identity Protocol (HIP) Architecture*, RFC 4423, Informational, May 2006, IETF.

[12] R. Moskowitz, P. Nikander, P. Jokela, T. Henderson, *Host Identity Protocol*, RFC 5201, Experimental, April 2008, IETF.

[13] M. Smith, T. Howes, *Lightweight Directory Access Protocol (LDAP): Uniform Resource Locator*, RFC 4516, June 2006, IETF.

[14] Trusted Computing Group, *TNC MAP Content Authorization*, Revision 1.0, June 2014.

[15] IEEE Std 802.1AR-2009, Ed. by M. Pritikin, *IEEE Standard for Local and Metropolitan Area Networks: Secure Device Identity*, December 2009, IEEE.

[16] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, W. Polk, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, RFC 5280, May 2008, IETF.

[17] J. Hodges, R. Morgan, M. Wahl, *Lightweight Directory Access Protocol (v3): Extension for Transport Layer Security*, RFC 2830, May 2000, IETF.

8 IF-MAP Metadata for ICS Security Schema

Metadata that a MAP Client publishes under this schema's name-space URI MUST comply with this schema. For example, XML requires that elements in a sequence in a schema be in the correct order; otherwise, the IF-MAP request would not pass XML validation.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
 xmlns:ifmap="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
xmlns:base-id="http://www.trustedcomputinggroup.org/2011/IFMAP-
IDENTIFIER/1"
 xmlns="http://www.trustedcomputinggroup.org/2010/IFMAP-ICS-METADATA/1"
 targetNamespace="http://www.trustedcomputinggroup.org/2010/IFMAP-ICS-
METADATA/1"
 elementFormDefault="qualified"
 attributeFormDefault="unqualified">

<!-- Importing from IFMAP-2.1r14 namespace for schema validation -->
  <xsd:import
   namespace="http://www.trustedcomputinggroup.org/2010/IFMAP/2"
   schemaLocation="IF-MAP-2.1r14-schema.xsd"/>

<!-- Importing from IFMAP-2.1r14-ExtendedIdentifier namespace for schema
validation -->
  <xsd:import
   namespace="http://www.trustedcomputinggroup.org/2011/IFMAP-
IDENTIFIER/1"
   schemaLocation="IF-MAP-2.1r14-ExtIdentifier-schema.xsd"/>


<!-- Schema for IF-MAP Metadata for ICS Security -->

<xsd:element name="backhaul-interface">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string" use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>

<xsd:element name="overlay-manager-group">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string" use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>

<xsd:element name="overlay-network-group">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:extension base="base-id:IdentifierType">
        <xsd:attribute name="name" type="xsd:string" use="required"/>
      </xsd:extension>
    </xsd:complexContent>
    <xsd:anyAttribute/>
  </xsd:complexType>
</xsd:element>
```

```
<!-- backhaul-policy metadata type -->
<xsd:element name="backhaul-policy">
   <xsd:complexType>
     <xsd:sequence>
       <xsd:element name="name" type="xsd:string" minOccurs="1"
          maxOccurs="1"/>
       <xsd:element name="policy" minOccurs="1" maxOccurs="1">
         <xsd:simpleType>
           <xsd:restriction base="xsd:string">
             <xsd:enumeration value="allow"/>
             <xsd:enumeration value="deny"/>
            </xsd:restriction>
          </xsd:simpleType>
        </xsd:element>
     </xsd:sequence>
     <xsd:attributeGroup ref="ifmap:multiValueMetadataAttributes"/>
     <xsd:anyAttribute/>
   </xsd:complexType>
 </xsd:element>

 <!-- bhi-identity metadata type -->
 <xsd:element name="bhi-identity">
   <xsd:complexType>
     <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
     <xsd:anyAttribute/>
   </xsd:complexType>
 </xsd:element>

 <!-- bhi-address metadata type -->
 <xsd:element name="bhi-address">
   <xsd:complexType>
     <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
     <xsd:anyAttribute/>
   </xsd:complexType>
 </xsd:element>

 <!-- bhi-certificate metadata type -->
 <xsd:element name="bhi-certificate">
   <xsd:complexType>
     <xsd:sequence>
       <xsd:element name="data" type="xsd:base64Binary" minOccurs="1"
         maxOccurs="1"/>
     </xsd:sequence>
     <xsd:attributeGroup ref="ifmap:multiValueMetadataAttributes"/>
     <xsd:anyAttribute/>
   </xsd:complexType>
 </xsd:element>

 <!-- dn-hit metadata type -->
 <xsd:element name="dn-hit">
   <xsd:complexType>
     <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
     <xsd:anyAttribute/>
   </xsd:complexType>
 </xsd:element>

 <!-- group-xref metadata type -->
 <xsd:element name="group-xref">
   <xsd:complexType>
```

```
      <xsd:sequence>
         <xsd:element name="uri" type="xsd:string" minOccurs="1"
           maxOccurs="1"/>
      </xsd:sequence>
      <xsd:attributeGroup ref="ifmap:multiValueMetadataAttributes"/>
      <xsd:anyAttribute/>
    </xsd:complexType>
  </xsd:element>

  <!-- manager-of metadata type -->
  <xsd:element name="manager-of">
    <xsd:complexType>
      <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
      <xsd:anyAttribute/>
    </xsd:complexType>
  </xsd:element>

  <!-- member-of metadata type -->
  <xsd:element name="member-of">
    <xsd:complexType>
      <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
      <xsd:anyAttribute/>
    </xsd:complexType>
  </xsd:element>

  <!-- observed-by metadata type -->
  <xsd:element name="observed-by">
    <xsd:complexType>
      <xsd:attributeGroup ref="ifmap:multiValueMetadataAttributes"/>
      <xsd:anyAttribute/>
    </xsd:complexType>
  </xsd:element>

 <!-- overlay-policy metadata type -->
 <xsd:element name="overlay-policy">
    <xsd:complexType>
      <xsd:sequence>
        <xsd:element name="name" type="xsd:string" minOccurs="1"
           maxOccurs="1"/>
        <xsd:element name="policy" minOccurs="1" maxOccurs="1">
          <xsd:simpleType>
            <xsd:restriction base="xsd:string">
              <xsd:enumeration value="allow"/>
              <xsd:enumeration value="deny"/>
             </xsd:restriction>
           </xsd:simpleType>
         </xsd:element>
      </xsd:sequence>
      <xsd:anyAttribute/>
      <xsd:attributeGroup ref="ifmap:multiValueMetadataAttributes"/>
    </xsd:complexType>
  </xsd:element>

  <!-- protected-by metadata type -->
  <xsd:element name="protected-by">
    <xsd:complexType>
      <xsd:attributeGroup ref="ifmap:singleValueMetadataAttributes"/>
      <xsd:anyAttribute/>
    </xsd:complexType>
  </xsd:element>
```

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