

Network Working Group  
Internet-Draft  
Updates: [3411](#), 3412, 3414, 3417  
(if approved)  
Intended status: Standards Track  
Expires: August 28, 2008

D. Harrington  
Huawei Technologies (USA)  
J. Schoenwaelder  
Jacobs University Bremen  
February 25, 2008

## **Transport Subsystem for the Simple Network Management Protocol (SNMP) draft-ietf-isms-tmsm-12**

### Status of This Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 28, 2008.

### Copyright Notice

Copyright (C) The IETF Trust (2008).

### Abstract

This document defines a Transport Subsystem, extending the Simple Network Management Protocol (SNMP) architecture defined in [RFC 3411](#). This document defines a subsystem to contain Transport Models, comparable to other subsystems in the [RFC3411](#) architecture. As work is being done to expand the transport to include secure transport such as SSH and TLS, using a subsystem will enable consistent design

and modularity of such Transport Models. This document identifies and describes some key aspects that need to be considered for any Transport Model for SNMP.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">The Internet-Standard Management Framework . . . . .</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Where this Extension Fits . . . . .</a>	<a href="#">3</a>
<a href="#">1.3.</a>	<a href="#">Conventions . . . . .</a>	<a href="#">5</a>
<a href="#">2.</a>	<a href="#">Motivation . . . . .</a>	<a href="#">5</a>
<a href="#">3.</a>	<a href="#">Requirements of a Transport Model . . . . .</a>	<a href="#">7</a>
<a href="#">3.1.</a>	<a href="#">Message Security Requirements . . . . .</a>	<a href="#">7</a>
<a href="#">3.1.1.</a>	<a href="#">Security Protocol Requirements . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.</a>	<a href="#">SNMP Requirements . . . . .</a>	<a href="#">8</a>
<a href="#">3.2.1.</a>	<a href="#">Architectural Modularity Requirements . . . . .</a>	<a href="#">8</a>
<a href="#">3.2.2.</a>	<a href="#">Access Control Requirements . . . . .</a>	<a href="#">11</a>
<a href="#">3.2.3.</a>	<a href="#">Security Parameter Passing Requirements . . . . .</a>	<a href="#">12</a>
<a href="#">3.2.4.</a>	<a href="#">Separation of Authentication and Authorization . . . . .</a>	<a href="#">13</a>
<a href="#">3.3.</a>	<a href="#">Session Requirements . . . . .</a>	<a href="#">14</a>
<a href="#">3.3.1.</a>	<a href="#">Session Establishment Requirements . . . . .</a>	<a href="#">14</a>
<a href="#">3.3.2.</a>	<a href="#">Session Maintenance Requirements . . . . .</a>	<a href="#">15</a>
<a href="#">3.3.3.</a>	<a href="#">Message security versus session security . . . . .</a>	<a href="#">16</a>
<a href="#">4.</a>	<a href="#">Scenario Diagrams and the Transport Subsystem . . . . .</a>	<a href="#">17</a>
<a href="#">5.</a>	<a href="#">Cached Information and References . . . . .</a>	<a href="#">17</a>
<a href="#">5.1.</a>	<a href="#">securityStateReference . . . . .</a>	<a href="#">18</a>
<a href="#">5.2.</a>	<a href="#">tmStateReference . . . . .</a>	<a href="#">18</a>
<a href="#">6.</a>	<a href="#">Abstract Service Interfaces . . . . .</a>	<a href="#">19</a>
<a href="#">6.1.</a>	<a href="#">sendMessage ASI . . . . .</a>	<a href="#">19</a>
<a href="#">6.2.</a>	<a href="#">Other Outgoing ASIs . . . . .</a>	<a href="#">20</a>
<a href="#">6.3.</a>	<a href="#">The receiveMessage ASI . . . . .</a>	<a href="#">22</a>
<a href="#">6.4.</a>	<a href="#">Other Incoming ASIs . . . . .</a>	<a href="#">22</a>
<a href="#">7.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">24</a>
<a href="#">7.1.</a>	<a href="#">Coexistence, Security Parameters, and Access Control . . . . .</a>	<a href="#">25</a>
<a href="#">8.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">26</a>
<a href="#">9.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">26</a>
<a href="#">10.</a>	<a href="#">References . . . . .</a>	<a href="#">26</a>
<a href="#">10.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">26</a>
<a href="#">10.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">27</a>
<a href="#">Appendix A.</a>	<a href="#">Why tmStateReference? . . . . .</a>	<a href="#">28</a>
<a href="#">A.1.</a>	<a href="#">Define an Abstract Service Interface . . . . .</a>	<a href="#">28</a>
<a href="#">A.2.</a>	<a href="#">Using an Encapsulating Header . . . . .</a>	<a href="#">29</a>
<a href="#">A.3.</a>	<a href="#">Modifying Existing Fields in an SNMP Message . . . . .</a>	<a href="#">29</a>
<a href="#">A.4.</a>	<a href="#">Using a Cache . . . . .</a>	<a href="#">29</a>
<a href="#">Appendix B.</a>	<a href="#">Open Issues . . . . .</a>	<a href="#">30</a>
<a href="#">Appendix C.</a>	<a href="#">Change Log . . . . .</a>	<a href="#">30</a>



## **1. Introduction**

This document defines a Transport Subsystem, extending the Simple Network Management Protocol (SNMP) architecture defined in [[RFC3411](#)]. This document identifies and describes some key aspects that need to be considered for any Transport Model for SNMP.

### **1.1. The Internet-Standard Management Framework**

For a detailed overview of the documents that describe the current Internet-Standard Management Framework, please refer to [section 7 of RFC 3410](#) [[RFC3410](#)].

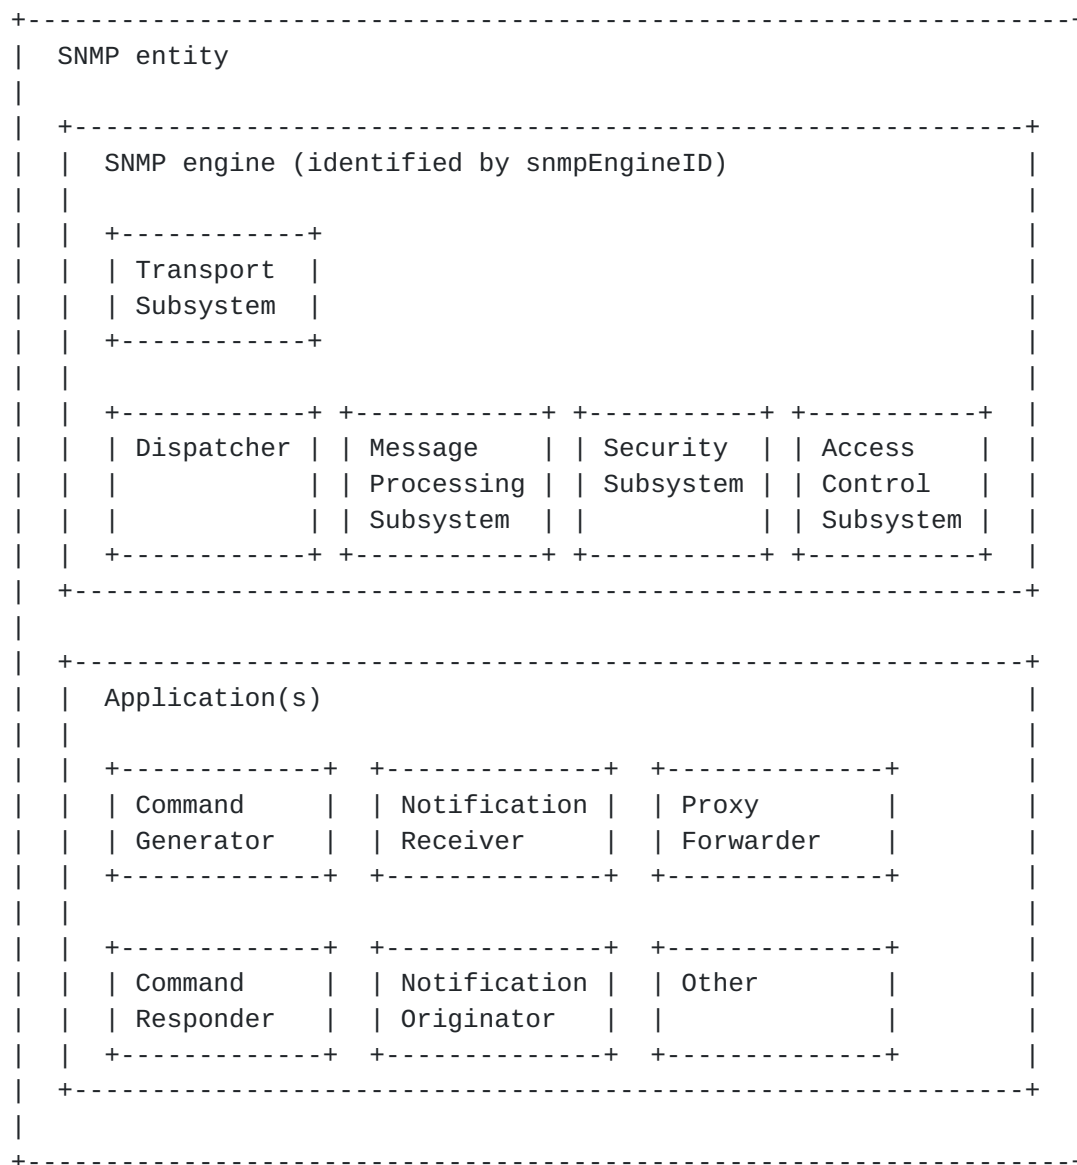
### **1.2. Where this Extension Fits**

It is expected that readers of this document will have read [RFC3410](#) and [RFC3411](#), and have a general understanding of the functionality defined in RFCs 3412-3418.

The "Transport Subsystem" is an additional component for the SNMP Engine depicted in [RFC3411, section 3.1](#).



The following diagram depicts its place in the [RFC3411](#) architecture.:



The transport mappings defined in [RFC3417](#) do not provide lower-layer security functionality, and thus do not provide transport-specific security parameters. This document updates [RFC3411](#) and [RFC3417](#) by defining an architectural extension and ASIs that transport mappings (models) can use to pass transport-specific security parameters to other subsystems, including transport-specific security parameters translated into transport-independent securityName and securityLevel parameters

The Transport Security Model [[I-D.ietf-isms-transport-security-model](#)] and the Secure Shell Transport Model [[I-D.ietf-isms-secshell](#)] utilize



the Transport Subsystem. The Transport Security Model is an alternative to the existing SNMPv1 Security Model [[RFC3584](#)], the SNMPv2c Security Model [[RFC3584](#)], and the User-based Security Model [[RFC3414](#)]. The Secure Shell Transport Model is an alternative to existing transport mappings (or models) as described in [[RFC3417](#)].

### **1.3. Conventions**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Non uppercased versions of the keywords should be read as in normal English. They will usually, but not always, be used in a context relating to compatibility with the [RFC3411](#) architecture or the subsystem defined here, but which might have no impact on on-the-wire compatibility. These terms are used as guidance for designers of proposed IETF models to make the designs compatible with [RFC3411](#) subsystems and Abstract Service Interfaces (see [section 3.2](#)). Implementers are free to implement differently. Some usages of these lowercase terms are simply normal English usage.

For consistency with SNMP-related specifications, this document favors terminology as defined in STD62 rather than favoring terminology that is consistent with non-SNMP specifications that use different variations of the same terminology. This is consistent with the IESG decision to not require the SNMPv3 terminology be modified to match the usage of other non-SNMP specifications when SNMPv3 was advanced to Full Standard.

## **2. Motivation**

Just as there are multiple ways to secure one's home or business, in a continuum of alternatives, there are multiple ways to secure a network management protocol. Let's consider three general approaches.

In the first approach, an individual could sit on his front porch waiting for intruders. In the second approach, he could hire an employee, schedule the employee, position the employee to guard what he wants protected, hire a second guard to cover if the first gets sick, and so on. In the third approach, he could hire a security company, tell them what he wants protected, and they could hire employees, train them, position the guards, schedule the guards, send a replacement when a guard cannot make it, etc., thus providing the desired security, with no significant effort on his part other than identifying requirements and verifying the quality of the service being provided.





The User-based Security Model (USM) as defined in [\[RFC3414\]](#) largely uses the first approach - it provides its own security. It utilizes existing mechanisms (e.g., SHA), but provides all the coordination. USM provides for the authentication of a principal, message encryption, data integrity checking, timeliness checking, etc.

USM was designed to be independent of other existing security infrastructures. USM therefore requires a separate principal and key management infrastructure. Operators have reported that deploying another principal and key management infrastructure in order to use SNMPv3 is a deterrent to deploying SNMPv3. It is possible to use external mechanisms to handle the distribution of keys for use by USM. The more important issue is that operators wanted to leverage a single user base that wasn't specific to SNMP.

A solution based on the second approach might use a USM-compliant architecture, but combine the authentication mechanism with an external mechanism, such as RADIUS [\[RFC2865\]](#), to provide the authentication service. It might be possible to utilize an external protocol to encrypt a message, to check timeliness, to check data integrity, etc. It is difficult to cobble together a number of subcontracted services and coordinate them however, because it is difficult to build solid security bindings between the various services, and potential for gaps in the security is significant.

A solution based on the third approach might utilize one or more lower-layer security mechanisms to provide the message-oriented security services required. These would include authentication of the sender, encryption, timeliness checking, and data integrity checking. There are a number of IETF standards available or in development to address these problems through security layers at the transport layer or application layer, among them TLS [\[RFC4346\]](#), SASL [\[RFC4422\]](#), and SSH [\[RFC4251\]](#).

From an operational perspective, it is highly desirable to use security mechanisms that can unify the administrative security management for SNMPv3, command line interfaces (CLIs) and other management interfaces. The use of security services provided by lower layers is the approach commonly used for the CLI, and is also the approach being proposed for NETCONF [\[RFC4741\]](#).

This document defines a Transport Subsystem extension to the [RFC3411](#) architecture based on the third approach. This extension specifies how other lower layer protocols with common security infrastructures can be used underneath the SNMP protocol and the desired goal of unified administrative security can be met.

This extension allows security to be provided by an external protocol



connected to the SNMP engine through an SNMP Transport Model [RFC3417]. Such a Transport Model would then enable the use of existing security mechanisms such as (TLS) [RFC4346] or SSH [RFC4251] within the RFC3411 architecture.

There are a number of Internet security protocols and mechanisms that are in wide spread use. Many of them try to provide a generic infrastructure to be used by many different application layer protocols. The motivation behind the Transport Subsystem is to leverage these protocols where it seems useful.

There are a number of challenges to be addressed to map the security provided by a secure transport into the SNMP architecture so that SNMP continues to provide interoperability with existing implementations. These challenges are described in detail in this document. For some key issues, design choices are described that might be made to provide a workable solution that meets operational requirements and fits into the SNMP architecture defined in [RFC3411].

### **3. Requirements of a Transport Model**

#### **3.1. Message Security Requirements**

Transport security protocols SHOULD provide protection against the following message-oriented threats [RFC3411]:

1. modification of information
2. masquerade
3. message stream modification
4. disclosure

These threats are described in [section 1.4 of \[RFC3411\]](#). It is not required to protect against denial of service or traffic analysis, but it should not make those threats significantly worse.

##### **3.1.1. Security Protocol Requirements**

There are a number of standard protocols that could be proposed as possible solutions within the Transport Subsystem. Some factors SHOULD be considered when selecting a protocol.

Using a protocol in a manner for which it was not designed has numerous problems. The advertised security characteristics of a protocol might depend on it being used as designed; when used in other ways, it might not deliver the expected security characteristics. It is recommended that any proposed model include a description of the applicability of the Transport Model.



A Transport Model SHOULD require no modifications to the underlying protocol. Modifying the protocol might change its security characteristics in ways that would impact other existing usages. If a change is necessary, the change SHOULD be an extension that has no impact on the existing usages. Any Transport Model SHOULD include a description of potential impact on other usages of the protocol.

Transport Models MUST be able to coexist with each other.

## **3.2. SNMP Requirements**

### **3.2.1. Architectural Modularity Requirements**

SNMP version 3 (SNMPv3) is based on a modular architecture (defined in [\[RFC3411\] section 3](#)) to allow the evolution of the SNMP protocol standards over time, and to minimize side effects between subsystems when changes are made.

The [RFC3411](#) architecture includes a Security Subsystem for enabling different methods of providing security services, a Message Processing Subsystem permitting different message versions to be handled by a single engine, Applications(s) to support different types of application processors, and an Access Control Subsystem for allowing multiple approaches to access control. The [RFC3411](#) architecture does not include a subsystem for Transport Models, despite the fact there are multiple transport mappings already defined for SNMP. This document addresses the need for a Transport Subsystem compatible with the [RFC3411](#) architecture. As work is being done to expand the transport to include secure transport such as SSH and TLS, using a subsystem will enable consistent design and modularity of such Transport Models.

The design of this Transport Subsystem accepts the goals of the [RFC3411](#) architecture defined in [section 1.5 of \[RFC3411\]](#). This Transport Subsystem uses a modular design that will permit Transport Models to be advanced through the standards process independently of other Transport Models, and independent of other modular SNMP components as much as possible.

Parameters have been added to the ASIs to pass model-independent transport address information.

IETF standards typically require one mandatory to implement solution, with the capability of adding new mechanisms in the future. Part of the motivation of developing Transport Models is to develop support for secure transport protocols, such as a Transport Model that utilizes the Secure Shell protocol. Any Transport Model SHOULD define one minimum-compliance security mechanism, such as



certificates, to ensure a basic level of interoperability, but should also be able to support additional existing and new mechanisms.

The Transport Subsystem permits multiple transport protocols to be "plugged into" the [RFC3411](#) architecture, supported by corresponding Transport Models, including models that are security-aware.

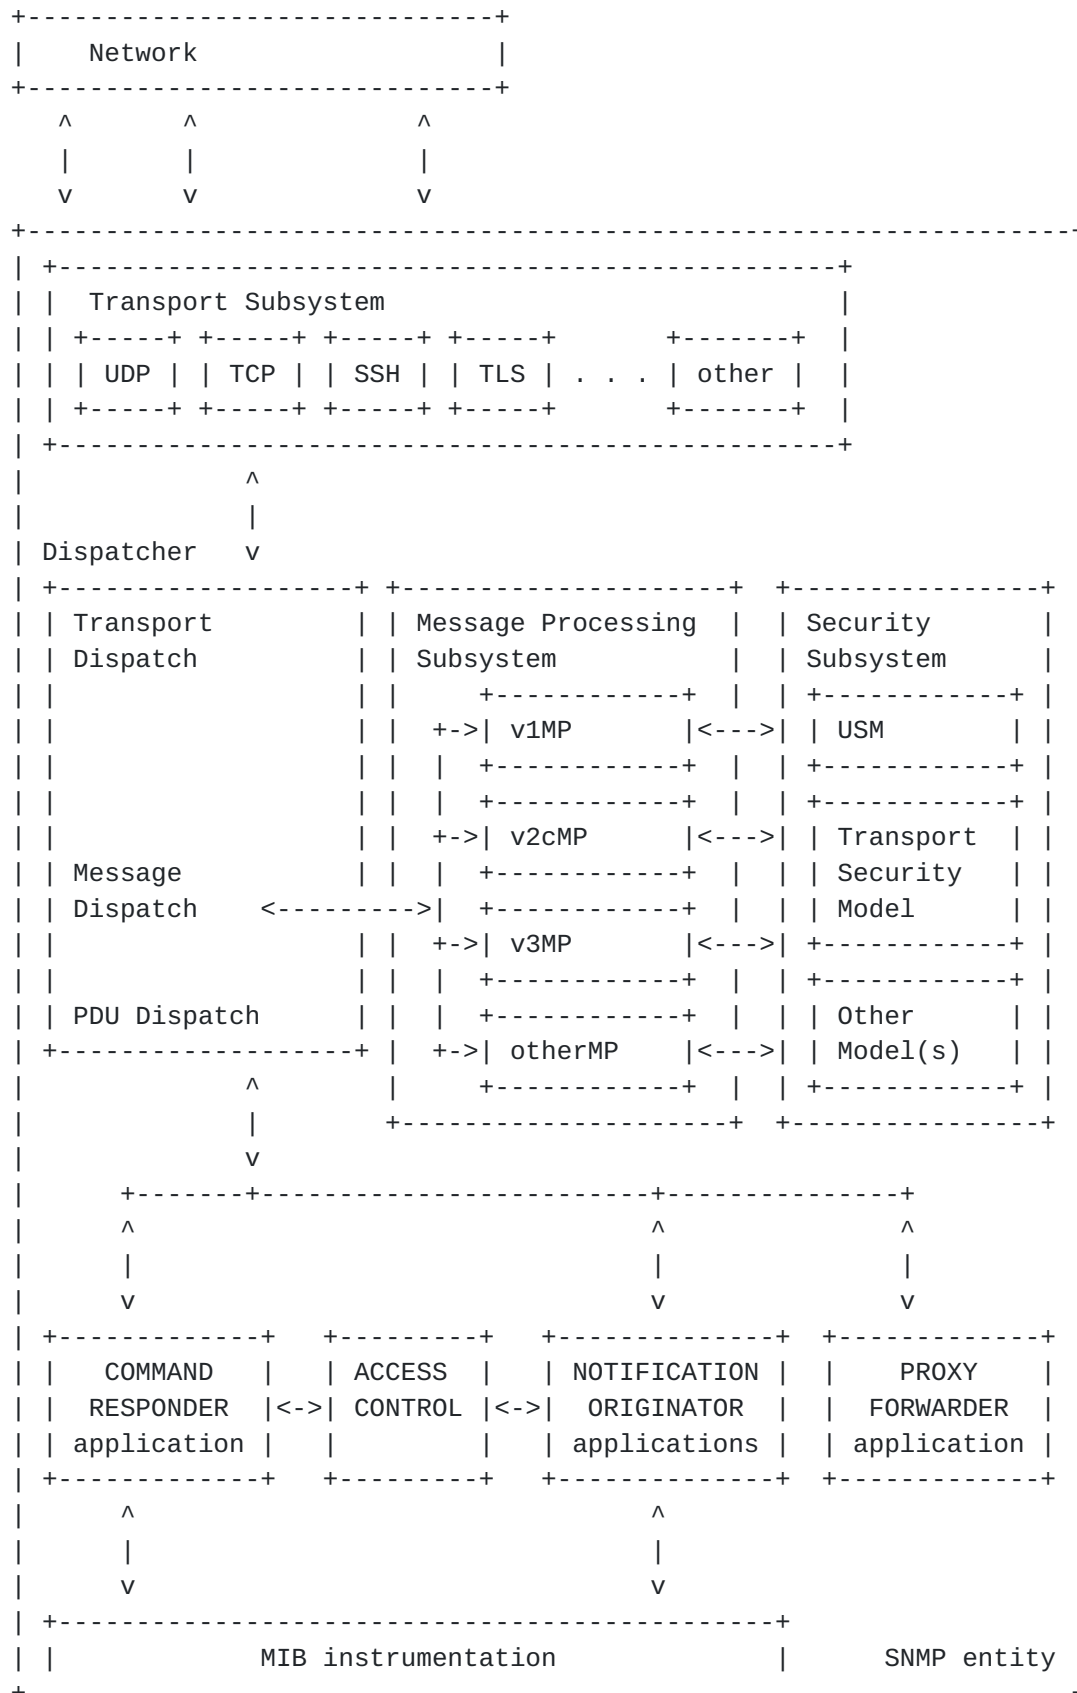
The [RFC3411](#) architecture and the Security Subsystem assume that a Security Model is called by a Message Processing Model and will perform multiple security functions within the Security Subsystem. A Transport Model that supports a secure transport protocol might perform similar security functions within the Transport Subsystem. A Transport Model might perform the translation of transport security parameters to/from security-model-independent parameters.

To accommodate this, an implementation-specific cache of transport-specific information will be described (not shown), and the data flows between the Transport Subsystem and the Transport Dispatch, between the Message Dispatch and the Message Processing Subsystem, and between the Message Processing Subsystem and the Security Subsystem will be extended to pass security-model-independent values. New Security Models may also be defined that understand how to work with the modified ASIs and the cache. One such Security Model, the Transport Security Model, is defined in [\[I-D.ietf-isms-transport-security-model\]](#)

The following diagram depicts the SNMPv3 architecture including the new Transport Subsystem defined in this document, and a new Transport Security Model defined in [\[I-D.ietf-isms-transport-security-model\]](#).









#### **3.2.1.1. Processing Differences between USM and Secure Transport**

USM and secure transports differ in the processing order and responsibilities within the [RFC3411](#) architecture. While the steps are the same, they occur in a different order, and may be done by different subsystems. With USM and some other Security Models, security processing starts when the Message Processing Model decodes portions of the encoded message to extract security parameters and header parameters that identify which Security Model should process the message to perform authentication, decryption, timeliness checking, integrity checking, and translation of parameters to model-independent parameters. A secure transport performs those security functions on the message, before the message is decoded.

#### **3.2.1.2. Passing Information between Engines**

A secure Transport Model will establish an authenticated and/or encrypted tunnel between the Transport Models of two SNMP engines. After a transport layer tunnel is established, then SNMP messages can be sent through the tunnel from one SNMP engine to the other SNMP engine. Transport Models MAY support sending multiple SNMP messages through the same tunnel.

#### **3.2.2. Access Control Requirements**

[RFC3411](#) made some design decisions related to the support of an Access Control Subsystem. These include establishing and passing in a model-independent manner the securityModel, securityName and securityLevel parameters, and separating message authentication from data access authorization.

##### **3.2.2.1. securityName and securityLevel Mapping**

SNMP data access controls are expected to work on the basis of who can perform what operations on which subsets of data, and based on the security services that will be provided to secure the data in transit. The securityModel and securityLevel parameters establish the protections for transit - whether authentication and privacy services will be or have been applied to the message. The securityName is a model-independent identifier of the security "principal",

The Message Processing Subsystem relies on a Security Model, such as USM, to play a role in security that goes beyond protecting the message - it provides a mapping between the security-model-specific principal for an incoming message to a security-model independent securityName which can be used for subsequent processing, such as for access control. The securityName is mapped from a mechanism-specific



identity, and this mapping must be done for incoming messages by the Security Model before it passes `securityName` to the Message Processing Model via the `processIncoming` ASI.

A Security Model is also responsible to specify, via the `securityLevel` parameter, whether incoming messages have been authenticated and/or encrypted, and to ensure that outgoing messages are authenticated and/or encrypted based on the value of `securityLevel`.

A translation from a mechanism-specific identity to a `securityName` might be done by a Transport Model, and the proposed `securityName` and a proposed `securityLevel` might then be made available to a Security Model via the `tmStateReference`. A Security Model may have multiple sources for determining the principal and desired security services, and a particular Security Model may or may not utilize the `securityName` mapping and `securityLevel` made available by the Transport Model when deciding the value of the `securityName` and `securityLevel` to be passed to the Message Processing Model.

### **3.2.3. Security Parameter Passing Requirements**

[RFC3411 section 4](#) describes abstract data flows between the subsystems, models and applications within the architecture. Abstract Service Interfaces describe the flow of data, passing model-independent information between subsystems within an engine. The [RFC3411](#) architecture has no ASI parameters for passing security information between the Transport Subsystem and the dispatcher, or between the dispatcher and the Message Processing Model. This document defines or modifies ASIs for this purpose.

A Message Processing Model might unpack SNMP-specific security parameters from an incoming message before calling a specific Security Model to authenticate and decrypt an incoming message, perform integrity checking, and translate security-model-specific parameters into model-independent parameters. When using a secure Transport Model, some security parameters might be provided through means other than carrying them in the SNMP message; some of the parameters for incoming messages might be extracted from the transport layer by the Transport Model before the message is passed to the Message Processing Subsystem.

This document describes a cache mechanism (see [Section 5](#)), into which the Transport Model puts information about the transport and security parameters applied to a transport connection or an incoming message, and a Security Model may extract that information from the cache. A `tmStateReference` is passed as an extra parameter in the ASIs of the Transport Subsystem and the Message Processing and Security



Subsystems, to identify the relevant cache. This approach of passing a model-independent reference is consistent with the securityStateReference cache already being passed around in the [RFC3411](#) ASIs.

For outgoing messages, even when a secure Transport Model will provide the security services, a Message Processing Model might have a Security Model actually create the message from its component parts. Whether there are any security services provided by the Security Model for an outgoing message is security-model-dependent. For incoming messages, even when a secure Transport Model provides security services, a Security Model might provide some security functionality that can only be provided after the message version or other parameters are extracted from the message.

#### **3.2.4. Separation of Authentication and Authorization**

The [RFC3411](#) architecture defines a separation of authentication and the authorization to access and/or modify MIB data. A set of model-independent parameters (securityModel, securityName, and securityLevel) are passed between the Security Subsystem, the applications, and the Access Control Subsystem.

This separation was a deliberate decision of the SNMPv3 WG, to allow support for authentication protocols which did not provide data access authorization capabilities, and to support data access authorization schemes, such as VACM, that do not perform their own authentication. This decision also permits different types of data access policies, such as one built on UNIX groups or Windows domains. The VACM approach is based on administrator-defined groups of users.

A Message Processing Model determines which Security Model is used, either based on the message version, e.g., SNMPv1 and SNMPv2c, and possibly by a value specified in the message, e.g., SNMPv3.

The Security Model makes the decision which securityName and securityLevel values are passed as model-independent parameters to an application, which then passes them via the isAccessAllowed ASI to the Access Control Subsystem.

An Access Control Model performs the mapping from the model-independent security parameters to a policy within the Access Control Model that is access-control-model-dependent.

A Transport Model does not know which securityModel will be used for an incoming message, so a Transport Model cannot know how the securityName and securityLevel parameters are determined. A Transport Model can provide a mapping from a transport-specific





identity and provide candidate values for the `securityName` and `securityLevel`, but there is no guarantee the transport-provided values will be used by the Security Model.

For example, the SNMPv1 Message Processing Model described in [RFC3584](#) always selects the SNMPv1 Security Model. This is true even if the SNMPv1 message was protected in transit using a secure Transport Model, such as one based on SSH or TLS. The SNMPv1 Security Model does not know the `tmStateReference` exists.

### **[3.3.](#) Session Requirements**

Some secure transports might have a notion of sessions, while other secure transports might provide channels or other session-like mechanism. Throughout this document, the term session is used in a broad sense to cover sessions, channels, and session-like mechanisms. Session refers to an association between two SNMP engines that permits the transmission of one or more SNMP messages within the lifetime of the session. How the session is actually established, opened, closed, or maintained is specific to a particular Transport Model.

Sessions are not part of the SNMP architecture defined in [\[RFC3411\]](#), but are considered desirable because the cost of authentication can be amortized over potentially many transactions.

The architecture defined in [\[RFC3411\]](#) does not include a session selector in the Abstract Service Interfaces, and neither is that done for the Transport Subsystem, so an SNMP application has no mechanism to select a session using the ASIs except by passing a unique combination of `transportDomain`, `transportAddress`, `securityName`, and `securityLevel`. Implementers, of course, might provide non-standard mechanisms to select sessions. The `transportDomain` and `transportAddress` identify the transport connection to a remote network node; the `securityName` identifies which security principal to communicate with at that address (e.g., different NMS applications), and the `securityLevel` might permit selection of different sets of security properties for different purposes (e.g., encrypted SETs vs. non-encrypted GETs).

To reduce redundancy, this document describes aspects that are expected to be common to all Transport Model sessions.

#### **[3.3.1.](#) Session Establishment Requirements**

SNMP has no mechanism to specify a transport session using the ASIs except by passing a unique combination `transportDomain`, `transportAddress`, `securityName`, and `securityLevel` to be used to



identify a session in a transport-independent manner. SNMP applications provide the transportDomain, transportAddress, securityName, and securityLevel to be used to create a session.

For an outgoing message, securityLevel is the requested security for the message, passed in the ASIs. If the Transport Model cannot provide at least the requested level of security, the Transport Model SHOULD discard the message and notify the dispatcher that establishing a session and sending the message failed.

A Transport Model determines whether an appropriate session exists (transportDomain, transportAddress, securityName, and securityLevel) for an outgoing message. If an appropriate session does not yet exist, the Transport Model attempts to establish a session for delivery. If a session cannot be established then the message is discarded and the dispatcher should be notified that sending the message failed.

Transport session establishment might require provisioning authentication credentials at an engine, either statically or dynamically. How this is done is dependent on the transport model and the implementation.

The Transport Subsystem has no knowledge of pduType, so cannot distinguish between a session created to carry different pduTypes. To differentiate a session established for different purposes, such as a notification session versus a request-response session, an application can use different securityNames or transport addresses. For example, in SNMPv1, UDP ports 161 and 162 were used to differentiate types of traffic. New transport models may define a single well-known port for all traffic types. Administrators might choose to define one port for SNMP request-response traffic, but configure notifications to be sent to a different port.

#### **3.3.2. Session Maintenance Requirements**

A Transport Model can tear down sessions as needed. It might be necessary for some implementations to tear down sessions as the result of resource constraints, for example.

The decision to tear down a session is implementation-dependent. While it is possible for an implementation to automatically tear down each session once an operation has completed, this is not recommended for anticipated performance reasons. How an implementation determines that an operation has completed, including all potential error paths, is implementation-dependent.

The elements of procedure describe when cached information can be



discarded, in some circumstances, and the timing of cache cleanup might have security implications, but cache memory management is an implementation issue.

If a Transport Model defines MIB module objects to maintain session state information, then the Transport Model **MUST** define what **SHOULD** happen to the objects when a related session is torn down, since this will impact interoperability of the MIB module.

### **3.3.3. Message security versus session security**

A Transport Model session is associated with state information that is maintained for its lifetime. This state information allows for the application of various security services to multiple messages. Cryptographic keys associated with the transport session **SHOULD** be used to provide authentication, integrity checking, and encryption services, as needed, for data that is communicated during the session. The cryptographic protocols used to establish keys for a Transport Model session **SHOULD** ensure that fresh new session keys are generated for each session. In addition sequence information might be maintained in the session which can be used to prevent the replay and reordering of messages within a session. If each session uses new keys, then a cross-session replay attack will be unsuccessful; that is, an attacker cannot successfully replay on one session a message he observed from another session. A good security protocol will also protect against replay attacks within a session; that is, an attacker cannot successfully replay a message observed earlier in the same session.

A Transport Model session will have a single transportDomain, transportAddress, securityName and securityLevel associated with it. If an exchange between communicating engines requires a different securityLevel or is on behalf of a different securityName, then another session would be needed. An immediate consequence of this is that implementations **SHOULD** be able to maintain some reasonable number of concurrent sessions.

For Transport Models, securityName should be specified during session setup, and associated with the session identifier.

SNMPv3 was designed to support multiple levels of security, selectable on a per-message basis by an SNMP application, because, for example, there is not much value in using encryption for a Commander Generator to poll for potentially non-sensitive performance data on thousands of interfaces every ten minutes; the encryption might add significant overhead to processing of the messages.

Some Transport Models might support only specific authentication and



encryption services, such as requiring all messages to be carried using both authentication and encryption, regardless of the security level requested by an SNMP application. A Transport Model may upgrade the requested security level, i.e. noAuthNoPriv and authNoPriv MAY be sent over an authenticated and encrypted session.

#### **4. Scenario Diagrams and the Transport Subsystem**

[RFC3411 section 4.6.1](#) and 4.6.2 provide scenario diagrams to illustrate how an outgoing message is created, and how an incoming message is processed. [RFC3411](#) does not define ASIs for "Send SNMP Request Message to Network" or "Receive SNMP Response Message from Network", and does not define ASIs for "Receive SNMP Message from Network" or "Send SNMP message to Network".

This document defines a sendMessage ASI to send SNMP messages to the network, regardless of pduType, and a receiveMessage ASI to receive SNMP messages from the network, regardless of pduType.

#### **5. Cached Information and References**

The [RFC3411](#) architecture uses caches to store dynamic model-specific information, and uses references in the ASIs to indicate in a model-independent manner which cached information flows between subsystems.

There are two levels of state that might need to be maintained: the security state in a request-response pair, and potentially long-term state relating to transport and security.

This state is maintained in caches. To simplify the elements of procedure, the release of state information is not always explicitly specified. As a general rule, if state information is available when a message being processed gets discarded, the state related to that message should also be discarded, and if state information is available when a relationship between engines is severed, such as the closing of a transport session, the state information for that relationship might also be discarded.

This document differentiates the tmStateReference from the securityStateReference. This document does not specify an implementation strategy, only an abstract description of the data that flows between subsystems. An implementation might use one cache and one reference to serve both functions, but an implementer must be aware of the cache-release issues to prevent the cache from being released before a security or Transport Model has had an opportunity to extract the information it needs.





### **5.1. securityStateReference**

The securityStateReference parameter is defined in [RFC3411](#). securityStateReference is not accessible to models of the Transport Subsystem.

### **5.2. tmStateReference**

For each transport session, information about the message security is stored in a cache to pass model- and mechanism-specific parameters. The state referenced by tmStateReference may be saved across multiple messages, in a Local Configuration Datastore (LCD), as compared to securityStateReference which is usually only saved for the life of a request-response pair of messages.

For security reasons, if a secure transport session is closed between the time a request message is received and the corresponding response message is sent, then the response message SHOULD be discarded, even if a new session has been established. The SNMPv3 WG decided that this should be a SHOULD architecturally, and it is a security-model-specific decision whether to REQUIRE this.

Since a transport model does not know whether a message contains a response, and transport session information is transport-model-specific, the tmStateReference contains two pieces of information for performing the request-response transport session pairing.

Each Security Model that supports the tmStateReference cache SHOULD pass a tmSameSecurity parameter in the tmStateReference cache for outgoing messages to indicate whether the same security parameters MUST be used for the outgoing message as was used for the corresponding incoming message.

Each transport model that supports sessions and supports the tmStateReference cache SHOULD include a transport-specific session identifier in the cache for an incoming message, so that if a security model requests tmSameSecurity, the transport model can determine whether the current existing transport session is the same as the transport session used for the incoming request.

When processing an outgoing message, if the tmSameSecurity requirement is indicated by the security model, but the session identified in the tmStateReference does not match the current established transport session, i.e., it is not the same transport session, then the message MUST be discarded, and the dispatcher should be notified the sending of the message failed.

Since the contents of a cache are meaningful only within an



implementation, and not on-the-wire, the format of the cache and the LCD are implementation-specific.

## 6. Abstract Service Interfaces

Abstract service interfaces have been defined by [RFC 3411](#) to describe the conceptual data flows between the various subsystems within an SNMP entity, and to help keep the subsystems independent of each other except for the common parameters.

This document follows the example of [RFC3411](#) regarding the release of state information, and regarding error indications.

1) The release of state information is not always explicitly specified in a transport model. As a general rule, if state information is available when a message gets discarded, the message-state information should also be released, and if state information is available when a session is closed, the session state information should also be released. Note that keeping sensitive security information longer than necessary might introduce potential vulnerabilities to an implementation.

2) An error indication in statusInformation may include an OID and value for an incremented counter and a value for securityLevel, and values for contextEngineID and contextName for the counter, and the securityStateReference if the information is available at the point where the error is detected.

### 6.1. sendMessage ASI

The sendMessage ASI is used to pass a message from the Dispatcher to the appropriate Transport Model for sending.

In the diagram in [section 4.6.1 of RFC 3411](#), the sendMessage ASI replaces the text "Send SNMP Request Message to Network". In [section 4.6.2](#), the sendMessage ASI replaces the text "Send SNMP Message to Network"

If present and valid, the tmStateReference refers to a cache containing transport-model-specific parameters for the transport and transport security. How the information in the cache is used is transport-model-dependent and implementation-dependent. How a tmStateReference is determined to be present and valid is implementation-dependent.

This may sound underspecified, but a transport model might be something like SNMP over UDP over IPv6, where no security is provided, so it might have no mechanisms for utilizing a securityName



and securityLevel.

```
statusInformation =
sendMessage(
IN    destTransportDomain      -- transport domain to be used
IN    destTransportAddress     -- transport address to be used
IN    outgoingMessage         -- the message to send
IN    outgoingMessageLength    -- its length
IN    tmStateReference         -- reference to transport state
)
```

## **6.2. Other Outgoing ASIs**

A tmStateReference parameter has been added to the prepareOutgoingMessage, prepareResponseMessage, generateRequestMsg, and generateResponseMsg ASIs as an OUT parameter. The transportDomain and transportAddress parameters have been added to the generateRequestMsg, and generateResponseMsg ASIs as IN parameters (not shown).

```
statusInformation =      -- success or errorIndication
prepareOutgoingMessage(
IN  transportDomain      -- transport domain to be used
IN  transportAddress     -- transport address to be used
IN  messageProcessingModel -- typically, SNMP version
IN  securityModel        -- Security Model to use
IN  securityName         -- on behalf of this principal
IN  securityLevel        -- Level of Security requested
IN  contextEngineID      -- data from/at this entity
IN  contextName          -- data from/in this context
IN  pduVersion           -- the version of the PDU
IN  PDU                  -- SNMP Protocol Data Unit
IN  expectResponse       -- TRUE or FALSE
IN  sendPduHandle        -- the handle for matching
                          incoming responses
OUT destTransportDomain  -- destination transport domain
OUT destTransportAddress -- destination transport address
OUT outgoingMessage     -- the message to send
OUT outgoingMessageLength -- its length
OUT tmStateReference     -- (NEW) reference to transport state
)
```



```
statusInformation =          -- success or errorIndication
prepareResponseMessage(
IN  messageProcessingModel    -- typically, SNMP version
IN  securityModel             -- Security Model to use
IN  securityName              -- on behalf of this principal
IN  securityLevel             -- Level of Security requested
IN  contextEngineID           -- data from/at this entity
IN  contextName               -- data from/in this context
IN  pduVersion                -- the version of the PDU
IN  PDU                       -- SNMP Protocol Data Unit
IN  maxSizeResponseScopedPDU  -- maximum size able to accept
IN  stateReference            -- reference to state information
                              -- as presented with the request
IN  statusInformation         -- success or errorIndication
                              -- error counter OID/value if error

OUT destTransportDomain       -- destination transport domain
OUT destTransportAddress      -- destination transport address
OUT outgoingMessage           -- the message to send
OUT outgoingMessageLength     -- its length
OUT tmStateReference          -- (NEW) reference to transport state
)
```

The `tmStateReference` parameter of `generateRequestMsg` or `generateResponseMsg` is passed in the OUT parameters of the Security Subsystem to the Message Processing Subsystem. If a cache exists for a session identifiable from `transportDomain`, `transportAddress`, `securityModel`, `securityName`, and `securityLevel`, then an appropriate Security Model might create a `tmStateReference` to the cache and pass that as an OUT parameter.

If one does not exist, the Security Model might create a cache referenced by `tmStateReference`. This information might include `transportDomain`, `transportAddress`, the `securityLevel`, and the `securityName`, plus any model or mechanism-specific details. The contents of the cache may be incomplete until the Transport Model has established a session. What information is passed, and how this information is determined, is implementation and security-model-specific.

The `prepareOutgoingMessage` ASI passes `tmStateReference` from the Message Processing Subsystem to the dispatcher. How or if the Message Processing Subsystem modifies or utilizes the contents of the cache is message-processing-model-specific.

This may sound underspecified, but a message processing model might have access to all the information from the cache and from the message, and an application might specify a Security Model such as USM to authenticate and secure the SNMP message, but also specify a





secure transport such as that provided by the SSH Transport Model to send the message to its destination.

### **6.3. The receiveMessage ASI**

If one does not exist, the Transport Model might create a cache referenced by `tmStateReference`. If present, this information might include `transportDomain`, `transportAddress`, `securityLevel`, and `securityName`, plus model or mechanism-specific details. How this information is determined is implementation and transport-model-specific.

In the diagram in [section 4.6.1 of RFC 3411](#), the `receiveMessage` ASI replaces the text "Receive SNMP Response Message from Network". In [section 4.6.2](#), the `receiveMessage` ASI replaces the text "Receive SNMP Message from Network"

This may sound underspecified, but a transport model might be something like SNMP over UDP over IPv6, where no security is provided, so it might have no mechanisms for determining a `securityName` and `securityLevel`.

The Transport Model does not know the `securityModel` for an incoming message; this will be determined by the Message Processing Model in a message-processing-model-dependent manner.

The `receiveMessage` ASI is used to pass a message from the Transport Subsystem to the Dispatcher.

```
statusInformation =
receiveMessage(
IN    transportDomain          -- origin transport domain
IN    transportAddress         -- origin transport address
IN    incomingMessage          -- the message received
IN    incomingMessageLength    -- its length
IN    tmStateReference         -- reference to transport state
)
```

### **6.4. Other Incoming ASIs**

To support the Transport Subsystem, the `tmStateReference` is added to the `prepareDataElements` ASI (from the Dispatcher to the Message Processing Subsystem), and to the `processIncomingMsg` ASI (from the Message Processing Subsystem to the Security Model Subsystem). How or if a Message Processing Model or Security Model uses `tmStateReference` is message-processing-model-dependent and security-model-dependent.



```

result =                -- SUCCESS or errorIndication
prepareDataElements(
IN   transportDomain    -- origin transport domain
IN   transportAddress    -- origin transport address
IN   wholeMsg           -- as received from the network
IN   wholeMsgLength     -- as received from the network
IN   tmStateReference   -- (NEW) from the Transport Model
OUT  messageProcessingModel -- typically, SNMP version
OUT  securityModel      -- Security Model to use
OUT  securityName       -- on behalf of this principal
OUT  securityLevel      -- Level of Security requested
OUT  contextEngineID    -- data from/at this entity
OUT  contextName        -- data from/in this context
OUT  pduVersion         -- the version of the PDU
OUT  PDU                -- SNMP Protocol Data Unit
OUT  pduType            -- SNMP PDU type
OUT  sendPduHandle      -- handle for matched request
OUT  maxSizeResponseScopedPDU -- maximum size sender can accept
OUT  statusInformation  -- success or errorIndication
                        -- error counter OID/value if error
OUT  stateReference     -- reference to state information
                        -- to be used for possible Response
)

```

```

statusInformation = -- errorIndication or success
                  -- error counter OID/value if error

processIncomingMsg(
IN   messageProcessingModel -- typically, SNMP version
IN   maxMessageSize         -- of the sending SNMP entity
IN   securityParameters     -- for the received message
IN   securityModel          -- for the received message
IN   securityLevel          -- Level of Security
IN   wholeMsg               -- as received on the wire
IN   wholeMsgLength         -- length as received on the wire
IN   tmStateReference       -- (NEW) from the Transport Model
OUT  securityEngineID       -- authoritative SNMP entity
OUT  securityName           -- identification of the principal
OUT  scopedPDU,             -- message (plaintext) payload
OUT  maxSizeResponseScopedPDU -- maximum size sender can handle
OUT  securityStateReference -- reference to security state
)                            -- information, needed for response

```

The tmStateReference parameter of prepareDataElements is passed from the dispatcher to the Message Processing Subsystem. How or if the Message Processing Subsystem modifies or utilizes the contents of the cache is message-processing-model-specific.



The processIncomingMessage ASI passes tmStateReference from the Message Processing Subsystem to the Security Subsystem.

If tmStateReference is present and valid, an appropriate Security Model might utilize the information in the cache. How or if the Security Subsystem utilizes the information in the cache is security-model-specific.

This may sound underspecified, but a message processing model might have access to all the information from the cache and from the message. The Message Processing Model might determine that the USM Security Model is specified in an SNMPv3 message header; the USM Security Model has no need of values in the tmStateReference cache to authenticate and secure the SNMP message, but an application might have specified to use a secure transport such as that provided by the SSH Transport Model to send the message to its destination.

## **7. Security Considerations**

This document defines an architectural approach that permits SNMP to utilize transport layer security services. Each proposed Transport Model should discuss the security considerations of the Transport Model.

It is considered desirable by some industry segments that SNMP Transport Models should utilize transport layer security that addresses perfect forward secrecy at least for encryption keys. Perfect forward secrecy guarantees that compromise of long term secret keys does not result in disclosure of past session keys. Each proposed Transport Model should include a discussion in its security considerations of whether perfect forward security is appropriate for the Transport Model.

Since the cache and LCD will contain security-related parameters, implementers should store this information (in memory or in persistent storage) in a manner to protect it from unauthorized disclosure and/or modification.

Care must be taken to ensure that a SNMP engine is sending packets out over a transport using credentials that are legal for that engine to use on behalf of that user. Otherwise an engine that has multiple transports open might be "tricked" into sending a message through the wrong transport.

A Security Model may have multiple sources from which to define the securityName and securityLevel. The use of a secure Transport Model does not imply that the securityName and securityLevel chosen by the Security Model represent the transport-authenticated identity or the



transport-provided security services. The `securityModel`, `securityName`, and `securityLevel` parameters are a related set, and an administrator should understand how the specified `securityModel` selects the corresponding `securityName` and `securityLevel`.

### **7.1. Coexistence, Security Parameters, and Access Control**

In the [RFC3411](#) architecture, the Message Processing Model makes the decision about which Security Model to use. The architectural change described by this document does not alter that.

The architecture change described by this document does however, allow SNMP to support two different approaches to security - message-driven security and transport-driven security. With message-driven security, SNMP provides its own security, and passes security parameters within the SNMP message; with transport-driven security, SNMP depends on an external entity to provide security during transport by "wrapping" the SNMP message.

Security models defined before the Transport Security Model (i.e., SNMPv1, SNMPv2c, and USM) do not support transport-based security, and only have access to the security parameters contained within the SNMP message. They do not know about the security parameters associated with a secure transport. As a result, the Access Control Subsystem bases its decisions on the security parameters extracted from the SNMP message, not on transport-based security parameters.

Implications of coexistence of older security models with secure transport models are known. The `securityName` used for access control decisions represents an SNMP-authenticated identity, not the transport-authenticated identity. (I can transport-authenticate as guest and then simply use a community name for root, or a USM non-authenticated identity.)

- o An SNMPv1 message will always be paired with an SNMPv1 Security Model (per [RFC3584](#)), regardless of the transport mapping or transport model used, and access controls will be based on the community name.
- o An SNMPv2c message will always be paired with an SNMPv2c Security Model (per [RFC3584](#)), regardless of the transport mapping or transport model used, and access controls will be based on the community name.
- o An SNMPv3 message will always be paired with the `securityModel` specified in the `msgSecurityParameters` field of the message (per [RFC3412](#)), regardless of the transport mapping or transport model used. If the SNMPv3 message specifies the User-based Security Model (USM), access controls will be based on the USM user. If the SNMPv3 message specifies the Transport Security Model (TSM), access controls will be based on the principal authenticated by





the transport.

## **8. IANA Considerations**

This document requires no action by IANA.

## **9. Acknowledgments**

The Integrated Security for SNMP WG would like to thank the following people for their contributions to the process:

The authors of submitted Security Model proposals: Chris Elliot, Wes Hardaker, David Harrington, Keith McCloghrie, Kaushik Narayan, David Perkins, Joseph Salowey, and Juergen Schoenwaelder.

The members of the Protocol Evaluation Team: Uri Blumenthal, Lakshminath Dondeti, Randy Presuhn, and Eric Rescorla.

WG members who performed detailed reviews: Jeffrey Hutzelman, Bert Wijnen, Tom Petch.

## **10. References**

### **10.1. Normative References**

- |           |   |
|-----------|---|
| [RFC2119] | Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <a href="#">BCP 14</a> , <a href="#">RFC 2119</a> , March 1997.  |
| [RFC3411] | Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, <a href="#">RFC 3411</a> , December 2002. |
| [RFC3412] | Case, J., Harrington, D., Presuhn, R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", STD 62, <a href="#">RFC 3412</a> , December 2002. |
| [RFC3414] | Blumenthal, U. and B.   |



Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", STD 62, [RFC 3414](#), December 2002.

[RFC3417]

Presuhn, R., "Transport Mappings for the Simple Network Management Protocol (SNMP)", STD 62, [RFC 3417](#), December 2002.

## **[10.2.](#) Informative References**

[RFC2865]

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

[RFC3410]

Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", [RFC 3410](#), December 2002.

[RFC3584]

Frye, R., Levi, D., Routhier, S., and B. Wijnen, "Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework", [BCP 74](#), [RFC 3584](#), August 2003.

[RFC4346]

Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", [RFC 4346](#), April 2006.

[RFC4422]

Melnikov, A. and K. Zeilenga, "Simple Authentication and Security



- Layer (SASL)", [RFC 4422](#), June 2006.
- [RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture", [RFC 4251](#), January 2006.
- [RFC4741] Enns, R., "NETCONF Configuration Protocol", [RFC 4741](#), December 2006.
- [I-D.ietf-isms-transport-security-model] Harrington, D., "Transport Security Model for SNMP", draft-ietf-isms-transport-security-model-07 (work in progress), November 2007.
- [I-D.ietf-isms-secshell] Harrington, D. and J. Salowey, "Secure Shell Transport Model for SNMP", [draft-ietf-isms-secshell-09](#) (work in progress), November 2007.

## **[Appendix A.](#) Why tmStateReference?**

This appendix considers why a cache-based approach was selected for passing parameters.

There are four approaches that could be used for passing information between the Transport Model and a Security Model.

1. one could define an ASI to supplement the existing ASIs, or
2. one could add a header to encapsulate the SNMP message,
3. one could utilize fields already defined in the existing SNMPv3 message, or
4. one could pass the information in an implementation-specific cache or via a MIB module.

### **[A.1.](#) Define an Abstract Service Interface**

Abstract Service Interfaces (ASIs) are defined by a set of primitives that specify the services provided and the abstract data elements that are to be passed when the services are invoked. Defining additional ASIs to pass the security and transport information from the Transport Subsystem to Security Subsystem has the advantage of being consistent with existing [RFC3411](#)/3412 practice, and helps to



ensure that any Transport Model proposals pass the necessary data, and do not cause side effects by creating model-specific dependencies between itself and other models or other subsystems other than those that are clearly defined by an ASI.

#### **A.2. Using an Encapsulating Header**

A header could encapsulate the SNMP message to pass necessary information from the Transport Model to the dispatcher and then to a Message Processing Model. The message header would be included in the wholeMessage ASI parameter, and would be removed by a corresponding Message Processing Model. This would imply the (one and only) messaging dispatcher would need to be modified to determine which SNMP message version was involved, and a new Message Processing Model would need to be developed that knew how to extract the header from the message and pass it to the Security Model.

#### **A.3. Modifying Existing Fields in an SNMP Message**

[RFC3412] defines the SNMPv3 message, which contains fields to pass security related parameters. The Transport Subsystem could use these fields in an SNMPv3 message, or comparable fields in other message formats to pass information between Transport Models in different SNMP engines, and to pass information between a Transport Model and a corresponding Message Processing Model.

If the fields in an incoming SNMPv3 message are changed by the Transport Model before passing it to the Security Model, then the Transport Model will need to decode the ASN.1 message, modify the fields, and re-encode the message in ASN.1 before passing the message on to the message dispatcher or to the transport layer. This would require an intimate knowledge of the message format and message versions so the Transport Model knew which fields could be modified. This would seriously violate the modularity of the architecture.

#### **A.4. Using a Cache**

This document describes a cache, into which the Transport Model puts information about the security applied to an incoming message, and a Security Model can extract that information from the cache. Given that there might be multiple TM-security caches, a tmStateReference is passed as an extra parameter in the ASIs between the Transport Subsystem and the Security Subsystem, so the Security Model knows which cache of information to consult.

This approach does create dependencies between a specific Transport Model and a corresponding specific Security Model. However, the approach of passing a model-independent reference to a model-





dependent cache is consistent with the securityStateReference already being passed around in the [RFC3411](#) ASIs.

## [Appendix B](#). Open Issues

NOTE to RFC editor: If this section is empty, then please remove this open issues section before publishing this document as an RFC. (If it is not empty, please send it back to the editor to resolve.

o

## [Appendix C](#). Change Log

NOTE to RFC editor: Please remove this change log before publishing this document as an RFC.

Changes from -09- to -10-

- o Pointed to companion documents
- o Wordsmithed extensively
- o Modified the note about SNMPv3-consistent terminology
- o Modified the note about [RFC2119](#) terminology.
- o Modified discussion of cryptographic key generation.
- o Added security considerations about coexistence with older security models
- o Expanded discussion of same session functionality
- o Described how sendMessage and receiveMessage fit into [RFC3411](#) diagrams
- o Modified prepareResponseMessage ASI
- o

Changes from -08- to -09-

- o A question was raised that notifications would not work properly, but we could never find the circumstances where this was true.
- o removed appendix with parameter matrix
- o Added a note about terminology, for consistency with SNMPv3 rather than with [RFC2828](#).

Changes from -07- to -08-

- o Identified new parameters in ASIs.
- o Added discussion about well-known ports.

Changes from -06- to -07-

- o Removed discussion of double authentication



- o Removed all direct and indirect references to pduType by Transport Subsystem
- o Added warning regarding keeping sensitive security information available longer than needed.
- o Removed knowledge of securityStateReference from Transport Subsystem.
- o Changed transport session identifier to not include securityModel, since this is not known for incoming messages until the message processing model.

Changes from revision -05- to -06-

mostly editorial changes  
removed some paragraphs considered unnecessary  
added Updates to header  
modified some text to get the security details right  
modified text re: ASIs so they are not API-like  
cleaned up some diagrams  
cleaned up [RFC2119](#) language  
added section numbers to citations to [RFC3411](#)  
removed gun for political correctness

Changes from revision -04- to -05-

removed all objects from the MIB module.  
changed document status to "Standard" rather than the xml2rfc default of informational.

changed mention of MD5 to SHA  
moved addressing style to TDomain and TAddress  
modified the diagrams as requested  
removed the "layered stack" diagrams that compared USM and a Transport Model processing  
removed discussion of speculative features that might exist in future Transport Models  
removed openSession and closeSession ASIs, since those are model-dependent  
removed the MIB module  
removed the MIB boilerplate intro (this memo defines a SMIPv2 MIB ...)  
removed IANA considerations related to the now-gone MIB module  
removed security considerations related to the MIB module  
removed references needed for the MIB module  
changed receiveMessage ASI to use origin transport domain/address  
updated Parameter CSV [appendix](#)

[C](#)hanges from revision -03- to -04-



changed title from Transport Mapping Security Model Architectural Extension to Transport Subsystem  
modified the abstract and introduction  
changed TMSM to TMS  
changed MPSP to simply Security Model  
changed SMSP to simply Security Model  
changed TMSP to Transport Model  
removed MPSP and TMSP and SMSP from Acronyms section  
modified diagrams  
removed most references to dispatcher functionality  
worked to remove dependencies between transport and security models.  
defined snmpTransportModel enumeration similar to snmpSecurityModel, etc.  
eliminated all reference to SNMPv3 msgXXXX fields  
changed tmSessionReference back to tmStateReference

#### Changes from revision -02- to -03-

- o removed session table from MIB module
- o removed sessionID from ASIs
- o reorganized to put ASI discussions in EOP section, as was done in SSHSM
- o changed user auth to client auth
- o changed tmStateReference to tmSessionReference
- o modified document to meet consensus positions published by JS
  - \* authoritative is model-specific
  - \* msgSecurityParameters usage is model-specific
  - \* msgFlags vs. securityLevel is model/implementation-specific
  - \* notifications must be able to cause creation of a session
  - \* security considerations must be model-specific
  - \* TDomain and TAddress are model-specific
  - \* MPSP changed to SMSP (Security Model security processing)

#### Changes from revision -01- to -02-

- o wrote text for session establishment requirements section.
- o wrote text for session maintenance requirements section.
- o removed section on relation to SNMPv2-MIB
- o updated MIB module to pass smilint
- o Added Structure of the MIB module, and other expected MIB-related sections.
- o updated author address
- o corrected spelling
- o removed msgFlags appendix
- o Removed section on implementation considerations.



- o started modifying the security boilerplate to address TMS and MIB security issues
- o reorganized slightly to better separate requirements from proposed solution. This probably needs additional work.
- o removed section with sample protocols and sample tmSessionReference.
- o Added section for acronyms
- o moved section comparing parameter passing techniques to appendix.
- o Removed section on notification requirements.

Changes from revision -00-

- o changed SSH references from I-Ds to RFCs
- o removed parameters from tmSessionReference for DTLS that revealed lower layer info.
- o Added TMS-MIB module
- o Added Internet-Standard Management Framework boilerplate
- o Added Structure of the MIB Module
- o Added MIB security considerations boilerplate (to be completed)
- o Added IANA Considerations
- o Added ASI Parameter table
- o Added discussion of Sessions
- o Added Open issues and Change Log
- o Rearranged sections

Authors' Addresses

David Harrington  
Huawei Technologies (USA)  
1700 Alma Dr. Suite 100  
Plano, TX 75075  
USA

Phone: +1 603 436 8634  
EMail: dharrington@huawei.com

Juergen Schoenwaelder  
Jacobs University Bremen  
Campus Ring 1  
28725 Bremen  
Germany

Phone: +49 421 200-3587  
EMail: j.schoenwaelder@iu-bremen.de





## Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

