Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 14, 2007

D. Harrington
Huawei Technologies (USA)
J. Schoenwaelder
International University Bremen
October 11, 2006

Transport Subsystem for the Simple Network Management Protocol (SNMP)
draft-ietf-isms-tmsm-04.txt

Status of This Memo

Copyright Notice

Abstract

This document describes a Transport Subsystem, extending the Simple Network Management Protocol (SNMP) architecture defined in RFC 3411. This document describes 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 discusses some key aspects that need to be considered for any transport model for SNMP.

This memo also defines a portion of the Management Information Base (MIB) for managing models in the Transport Subsystem.

Table of Contents

1. Introduction

This document describes a Transport Subsystem, extending the Simple Network Management Protocol (SNMP) architecture defined in [RFC3411]. This document identifies and discusses 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].

Managed objects are accessed via a virtual information store, termed the Management Information Base or MIB. MIB objects are generally accessed through the Simple Network Management Protocol (SNMP). Objects in the MIB are defined using the mechanisms defined in the Structure of Management Information (SMI). This memo specifies a MIB module that is compliant to the SMIv2, which is described in STD 58, RFC 2578 [RFC2578], STD 58, RFC 2579 [RFC2579] and STD 58, RFC 2580 [RFC2580].

1.2. 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].

Some points requiring further WG research and discussion are identified by [discuss] markers in the text. Some points requiring further editing by the editors are marked [todo] in the text.

1.3. Acronyms

This section contains a list of acronyms used within the document and references to where in the document the acronym is defined, for easy lookup.

- [todo]

1.4. Motivation

There are multiple ways to secure one's home or business, in a continuum of alternatives. Let's consider three general approaches. In the first approach, an individual could buy a gun, learn to use it, and sit on your front porch waiting for intruders. In the second approach, one could hire an employee with a gun, schedule the employee, position the employee to guard what you want protected, hire a second guard to cover if the first gets sick, and so on. In

the third approach, you could hire a security company, tell them what you want protected, and they could hire employees, train them, buy the guns, position the guards, schedule the guards, send a replacement when a guard cannot make it, etc., thus providing the security you want, with no significant effort on your 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 (MD5=the gun), 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 but difficult to define external mechanisms that handle the distribution of keys for use by the USM approach.

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 [RFC4366], 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 [I-D.ietf-netconf-ssh].

This document describes a Transport Subsystem extension to the RFC3411 architecture, that 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) [RFC4366] 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 work without any surprises. These challenges are discussed in detail in this document. For some key issues, design choices are discussed that may be made to provide a workable solution that meets operational requirements and fits into the SNMP architecture defined in [RFC3411].

2. Requirements of a Transport Model

2.1. Message Security Requirements

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

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

According to [RFC3411], it is not required to protect against denial of service or traffic analysis.

2.1.1. 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 may depend on its being used as designed; when used in other ways, it may not deliver the expected security characteristics. It is recommended that any proposed model include a discussion of the applicability statement of the protocols to be used.

A transport model should require no modifications to the underlying protocol. Modifying the protocol may 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. It is recommended that any transport model include a discussion of potential impact on other usages of the protocol.

It has been a long-standing requirement that SNMP be able to work when the network is unstable, to enable network troubleshooting and repair. The UDP approach has been considered to meet that need well, with an assumption that getting small messages through, even if out of order, is better than getting no messages through. There has been a long debate about whether UDP actually offers better support than TCP when the underlying IP or lower layers are unstable. There has been recent discussion of whether operators actually use SNMP to troubleshoot and repair unstable networks.

There has been discussion of ways SNMP could be extended to better support management/monitoring needs when a network is running just fine. Use of a TCP transport, for example, could enable larger message sizes and more efficient table retrievals.

Transport models MUST be able to coexist with other transport models, and may be designed to utilize either TCP or UDP or SCTP.

2.2. SNMP Requirements

2.2.1. Architectural Modularity Requirements

SNMP version 3 (SNMPv3) is based on a modular architecture (described 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 messaging subsystem permitting different message versions to be handled by a single engine, an application subsystem 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.

In SNMPv2, there were many problems of side effects between subsystems caused by the manipulation of MIB objects, especially

those related to authentication and authorization, because many of the parameters were stored in shared MIB objects, and different models and protocols could assign different values to the objects. Contributors assumed slightly different shades of meaning depending on the models and protocols being used. As the shared MIB module design was modified to accommodate a specific model, other models which used the same MIB objects would be broken.

Abstract Service Interfaces (ASIs) were developed to pass model-independent parameters. The models were required to translate from their model-dependent formats into a model-independent format, defined using model-independent semantics, which would not impact other models.

Parameters have been provided in the ASIs to pass model-independent information about the authentication that has been provided. These parameters include a model-independent identifier of the security "principal", the security model used to perform the authentication, and which SNMP-specific security features were applied to the message (authentication and/or privacy).

Parameters have been provided in the ASIs to pass model-independent transport address information. These parameters utilize the TransportType and TransportAddress

The design of a transport subsystem must abide the goals of the RFC3411 architecture defined in [RFC3411]. To that end, this transport subsystem proposal 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.

IETF standards typically require one mandatory to implement solution, with the capability of adding new mechanisms in the future. Part of the motivstion 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, preferably one which is already widely used to secure the transport layer protocol.

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 USM 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 may perform similar

security functions within the transport subsystem. A transport model may perform the translation of transport security parameters to/from security-model-independent parameters. To accommodate this, the ASIs for the transport subsystem, the messaging subsystem, and the security subsystem will be extended to pass security-model-independent values, and a cache of transport-specific information.

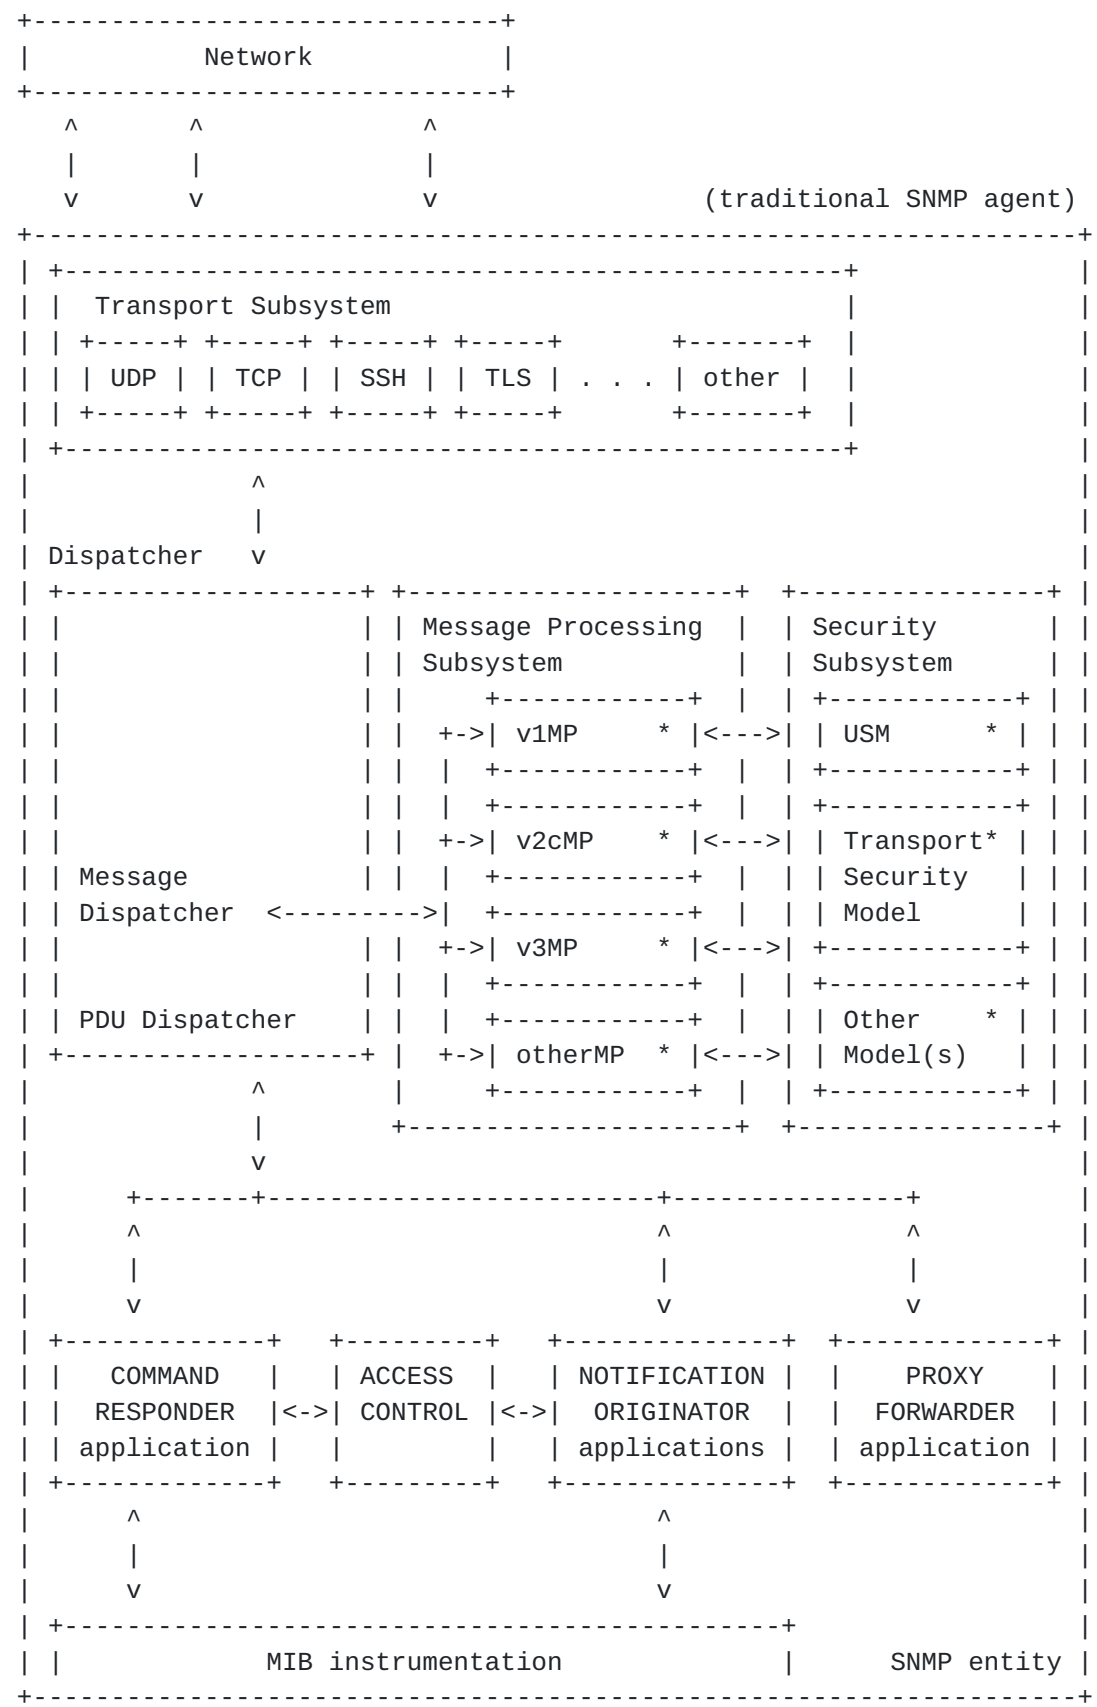
Network
(traditional SNMP agent)
Transport Subsystem
UDP
TCP
SSH
TLS
. . .
other
Dispatcher
Message Dispatcher
PDU Dispatcher
Message Processing Subsystem
v1MP
v2cMP
v3MP
otherMP
Security Subsystem
USM
Transport Security Model
Other Model(s)
COMMAND RESPONDER application
ACCESS CONTROL
NOTIFICATION ORIGINATOR applications
PROXY FORWARDER application
MIB instrumentation
SNMP entity

2.2.1.1. USM and the RFC3411 Architecture

The following diagrams illustrate the difference in the security processing done by the USM model and the security processing potentially done by a transport model.

The USM security model is encapsulated by the messaging model, because the messaging model needs to perform the following steps (for incoming messages)

1) decode the ASN.1 (messaging model)
2) determine the SNMP security model and parameters (messaging model)
3) decrypt the encrypted portions of the message (security model)
4) translate parameters to model-independent parameters (security model)
5) determine which application should get the decrypted portions (messaging model), and
6) pass on the decrypted portions with model-independent parameters.

The USM approach uses SNMP-specific message security and parameters.

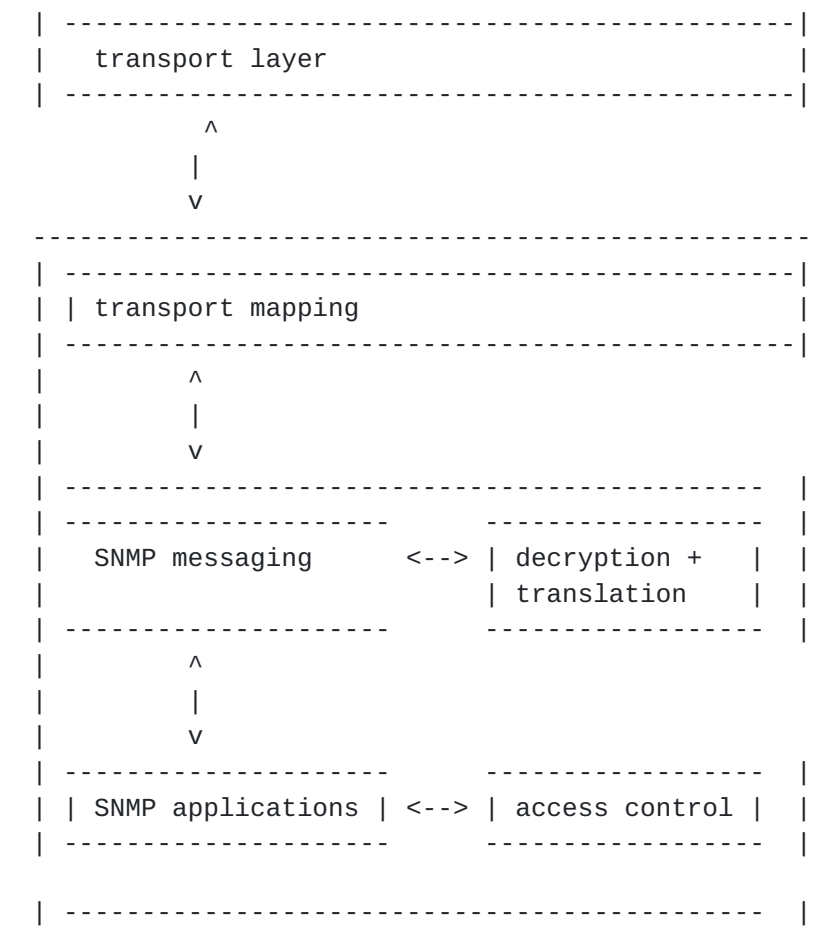

2.2.1.2. Transport Subsystem and the RFC3411 Architecture

With the Transport Subsystem, the order of the steps may differ and may be handled by different subsystems:

1) decrypt the encrypted portions of the message (transport layer)
2*) translate parameters to model-independent parameters (transport model)
3) determine the SNMP security model and parameters (messaging model)
4) decode the ASN.1 (messaging model)
5) determine which application should get the decrypted portions (messaging model)
7) pass on the decrypted portions with model-independent security parameters

If a message is secured using non-SNMP-specific message security and parameters, then the transport model should provide the translation from e.g., an SSH user name to the securityName in step 3,

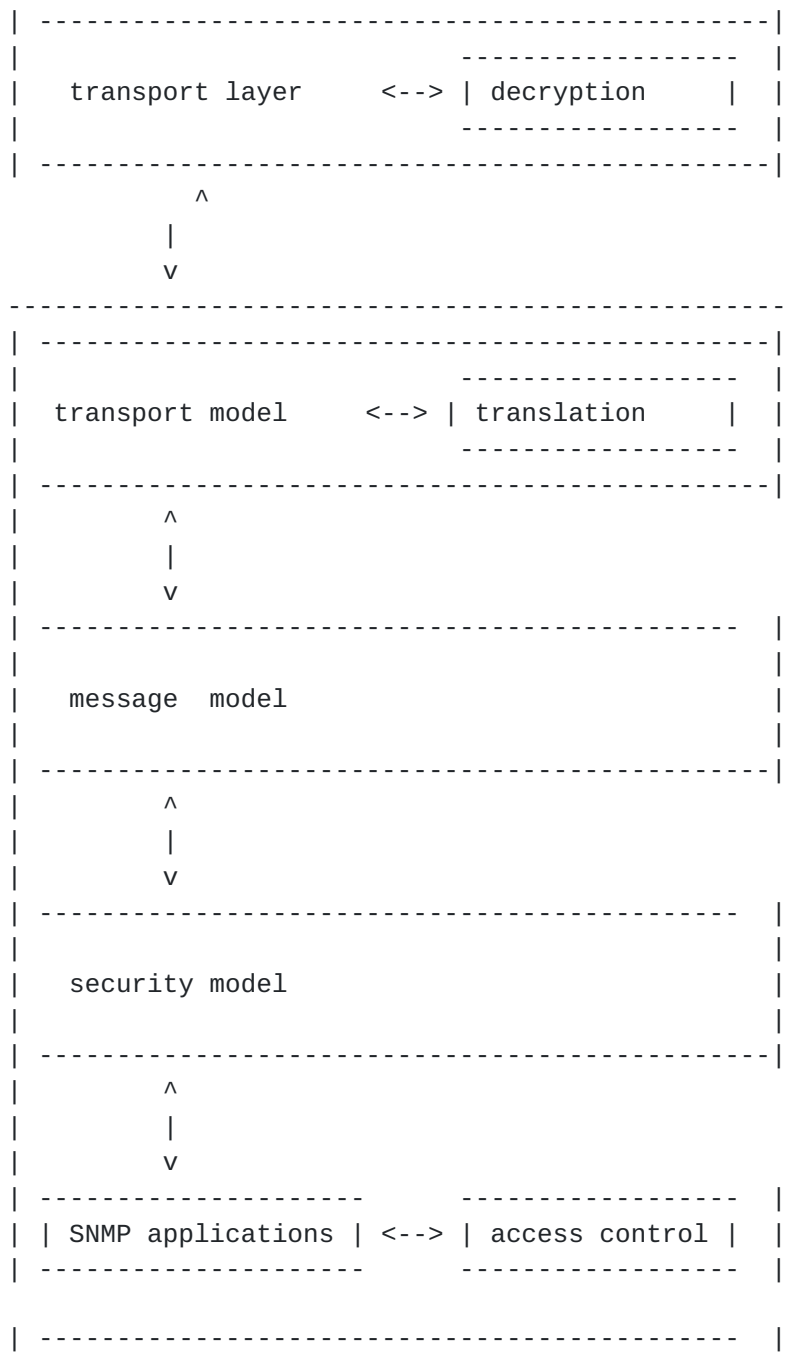

2.2.1.3. Passing Information between Engines

A secure transport model will establish an encrypted tunnel between the transport models of two SNMP engines. One transport model instance encrypts all messages, and the other transport model instance decrypts the messages.

After a transport layer tunnel is established, then SNMP messages can conceptually be sent through the tunnel from one SNMP engine to another SNMP engine. Once the tunnel is established, multiple SNMP messages may be able to be passed through the same tunnel.

2.2.2. Access Control Requirements

2.2.2.1. securityName Binding

For SNMP access control to function properly, security processing must establish a securityModel identifier, a securityLevel, and a securityName, which is the security model independent identifier for a 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 USM-specific principal to a security-model independent securityName which can be used for subsequent processing, such as for access control.

The securityName MUST be bound to the mechanism-specific authenticated identity, and this mapping MUST be done for incoming messages before the security model passes securityName to the message processing model via the processIncoming() ASI. This translation from a mechanism-specific authenticated identity to a securityName MAY be done by the transport model, and the securityname is then provided to the security model to be passed to the message processing model..

If the type of authentication provided by the transport layer (e.g. TLS) is considered adequate to secure and/or encrypt the message, but inadequate to provide the desired granularity of access control (e.g. user-based), then a second authentication (e.g., one provided via a RADIUS server) MAY be used to provide the authentication identity which is bound to the securityName. This approach would require a good analysis of the potential for man-in-the-middle attacks or masquerade possibilities.

2.2.2.2. Separation of Authentication and Authorization

A transport model that provides security services should take care to not violate the separation of authentication and authorization in the RFC3411 architecture. The isAccessAllowed() primitive is used for passing security-model independent parameters between the subsystems of the architecture.

Mapping of (securityModel, securityName) to an access control policy should be handled within the access control subsystem, not the transport or security subsystems, to be consistent with the modularity of the RFC3411 architecture. This separation was a

deliberate decision of the SNMPv3 WG, to allow support for authentication protocols which did not provide authorization capabilities, and to support authorization schemes, such as VACM, that do not perform their own authentication.

An authorization model (in the access control subsystem) MAY require authentication by certain securityModels and a minimum securityLevel to allow access to the data.

Transport models that provide secure transport are an enhancement for the SNMPv3 privacy and authentication, but they are not a significant improvement for the authorization (access control) needs of SNMPv3. Only the model-independent parameters for the isAccessAllowed() primitive [RFC3411] are provided by the transport and security subsystems.

A transport model must not specify how the securityModel and securityName could be dynamically mapped to an access control mechanism, such as a VACM-style groupName. The mapping of (securityModel, securityName) to a groupName is a VACM-specific mechanism for naming an access control policy, and for tying the named policy to the addressing capabilities of the data modeling language (e.g. SMIv2 [RFC2578]), the operations supported, and other factors. Providing a binding outside the Access Control subsystem might create dependencies that could make it harder to develop alternate models of access control, such as one built on UNIX groups or Windows domains. The preferred approach is to pass the model-independent security parameters via the isAccessAllowed() ASI, and perform the mapping from the model-independent security parameters to an authorization-model-dependent access policy within the access control model.

To provide support for protocols which simultaneously send information for authentication and authorization, such as RADIUS [RFC2865], model-specific authorization information MAY be cached or otherwise made available to the access control subsystem, e.g., via a MIB table similar to the vacmSecurityToGroupTable, so the access control subsystem can create an appropriate binding between the model-independent securityModel and securityName and a model-specific access control policy. This may be highly undesirable, however, if it creates a dependency between a transport model or a security model and an access control model, just as it is undesirable for a transport model to create a dependency between an SNMP message version and the security provided by a transport model.

2.2.3. Security Parameter Passing Requirements

RFC3411 section 4 describes primitives to describe the abstract data flows between the various subsystems, models and applications within the architecture. Abstract Service Interfaces describe the flow of data between subsystems within an engine. The ASIs generally pass model-independent information.

Within an engine using a transport model, outgoing SNMP messages are passed unencrypted from the message dispatcher to the transport model, and incoming messages are passed unencrypted from the transport model to the message dispatcher.

The security parameters include a model-independent identifier of the security "principal", the security model used to perform the authentication, and which SNMP-specific security services were (should be) applied to the message (authentication and/or privacy).

In the RFC3411 architecture, which reflects the USM security model design, the messaging model must 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 model-specific security parameters into model-independent parameters.

When using a secure transport model, security parameters MAY be provided through means other than carrying them in the SNMP message. The parameters MAY be provided by SNMP applications for outgoing messages, and the parameters for incoming messages MAY be extracted from the transport layer by the transport model before the message is passed to the message processing subsystem.

For outgoing messages, even when a secure transport model will provide the security services, it is necessary to have an security model because it is the security model that actually creates the message from its component parts. Whether there are any security services provided by the security model for an outgoing message is model-dependent.

For incoming messages, even when a secure transport model provides security services, a security model is necessary because there might be some security functionality that can only be provided after the message version is known. The message version is determined by the Message Processing model and passed to the security model via the processIncoming() ASI.

The RFC3411 architecture has no ASI parameters for passing security information between a transport mapping (a transport model) and the

dispatcher, and between the dispatcher and the message processing model.

This document describes a cache mechanism, 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 messaging 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. [todo: can we avoid dependencies?]

2.3. Session Requirements

Throughout this document, the term session is used. Some underlying secure transports will have a notion of session. Some underlying secure transports might enable the use of channels or other session-like thing. In this document the term session refers to an association between two SNMP engines that permits the secure 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 described in [RFC3411], but are considered desirable because the cost of authentication can be amortized over potentially many transactions.

It is important to note that the architecture described 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 cannot select the session except by passing a unique combination of transport type, transport address, securityName, securityModel, and securityLevel.

All transport models should discuss the impact of sessions on SNMP usage, including how to establish/open a transport session (i.e., how it maps to the concepts of session-like things of the underlying protocol), how to behave when a session cannot be established, how to close a session properly, how to behave when a session is closed improperly, the session security properties, session establishment overhead, and session maintenance overhead.

To reduce redundancy, this document will discuss aspects that are expected to be common to all transport model sessions.

2.3.1. Session Establishment Requirements

SNMP applications must provide the transport type, transport address, securityName, securityModel, and securityLevel to be used for a session.

SNMP Applications typically have no knowledge of whether the session that will be used to carry commands was initially established as a notification session, or a request-response session, and SHOULD NOT make any assumptions based on knowing the direction of the session. If an administrator or transport model designer wants to differentiate a session established for different purposes, such as a notification session versus a request-response session, the application can use different securityNames or transport addresses (e.g., port 161 vs. port 162) for different purposes.

An SNMP engine containing an application that initiates communication, e.g., a Command Generator or Notification Originator, MUST be able to attempt to establish a session for delivery if a session does not yet exist. If a session cannot be established then the message is discarded.

Sessions are usually established by the transport model when no appropriate session is found for an outgoing message, but sessions may be established in advance to support features such as notifications. How sessions are established in advance is beyond the scope of this document.

Sessions are initiated by notification originators when there is no currently established connection that can be used to send the notification. For a client-server security protocol, this may require provisioning authentication credentials on the agent, either statically or dynamically, so the client/agent can successfully authenticate to a notification receiver.

A transport model must be able to determine whether a session does or does not exist, and must be able to determine which session has the appropriate security characteristics (transport type, transport address, securityName, securityModel, and securityLevel) for an outgoing message. [discuss: does the transport model have insight into the securityModel?]

A transport model implementation MAY reuse an already established session with the appropriate transport type, transport address, securityName, securityModel, and securityLevel characteristics for delivery of a message originated by a different type of application than originally caused the session to be created. For example, an implementation that has an existing session originally established to

receive a request may use that session to send an outgoing notification, and may use a session that was originally established to send a notification to send a request. Responses are expected to be returned using the same session that carried the corresponding request message. Reuse of sessions is not required for conformance.

If a session can be reused for a different type of message, but a receiver is not prepared to accept different message types over the same session, then the message MAY be dropped by the receiver. This may strongly affect the usefulness of session reuse.

2.3.2. Session Maintenance Requirements

A transport model can tear down sessions as needed. It may 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.

Implementations should be careful to not tear down a session between the time a request is received and the time the response is sent. The elements of procedure for transport models should be sure to describe the expected behavior when no session exists for a response. [todo: do we already say that the message should be discarded, or is that just in the ssh transport model?]

The elements of procedure may discuss when cached information can be discarded, and the timing of cache cleanup may 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 describe what happens to the objects when a related session is torn down, since this will impact interoperability of the MIB module.

2.3.3. 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 established at the beginning of the session SHOULD be used to provide authentication, integrity checking, and encryption services 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. If each session uses new session keys, then messages cannot be replayed from one session to another. In addition sequence information MAY be maintained in the session which can be used to prevent the replay and reordering of messages within a session.

A transport model session will typically have a single transport type, ransport address, securityModel, securityName and securityLevel associated with it. If an exchange between communicating engines would require a different securityLevel or would be on behalf of a different securityName, or to use a different securityModel, 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 is typically 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 there is not much value in using encryption for a Commander Generator to poll for non-sensitive performance data on thousands of interfaces every ten minutes; the encryption adds significant overhead to processing of the messages.

Some transport models MAY 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.

Some transport models may use an underlying transport that provides a per-message requested level of authentication and encryption services. For example, if a session is created as 'authPriv', then keys for encryption could still be negotiated once at the beginning of the session. But if a message is presented to the session with a security level of authNoPriv, then that message could simply be authenticated and not encrypted within the same transport session. Whether this is possible depends on the transport model and the secure transport used.

If the underlying transport layer security is configurable on a per-message basis, a transport model could have a transport-model MIB module with configurable maxSecurityLevel and a minSecurityLevel objects to identify the range of possible levels. A session's maxSecurityLevel would identify the maximum security it could provide, and a session created with a minSecurityLevel of authPriv would reject an attempt to send an authNoPriv message. The elements

of procedure of the transport model would need to describe the procedures to enable this determination. [discuss: this is a feature I find questionable. It can be developed as a feature of a specific transport model. Do we need this discussion here?]

For transport models that do not support variable security services in one session, multiple sessions could be established with different security levels, and for every packet the SNMP engine could select the appropriate session based on the requested securityLevel. Some SNMP entities are resource-constrained. Adding sessions increases the need for resources, but so does encrypting unnecessarily. Designers of transport models should consider the trade offs for resource-constrained devices.

3. Scenario Diagrams for the Transport Subsystem

RFC3411 section 4.6 provides scenario diagrams to illustrate how an outgoing message is created, and how an incoming message is processed. Both diagrams are incomplete, however. In section 4.6.1, the diagram doesn't show the ASI for sending an SNMP request to the network or receiving an SNMP response message from the network. In section 4.6.2, the diagram doesn't illustrate the interfaces required to receive an SNMP message from the network, or to send an SNMP message to the network.

3.1. Command Generator or Notification Originator

This diagram from RFC3411 4.6.1 shows how a Command Generator or Notification Originator application [RFC3413] requests that a PDU be sent, and how the response is returned (asynchronously) to that application.

3.2. Command Responder

This diagram shows how a Command Responder or Notification Receiver application registers for handling a pduType, how a PDU is dispatched to the application after an SNMP message is received, and how the Response is (asynchronously) send back to the network.

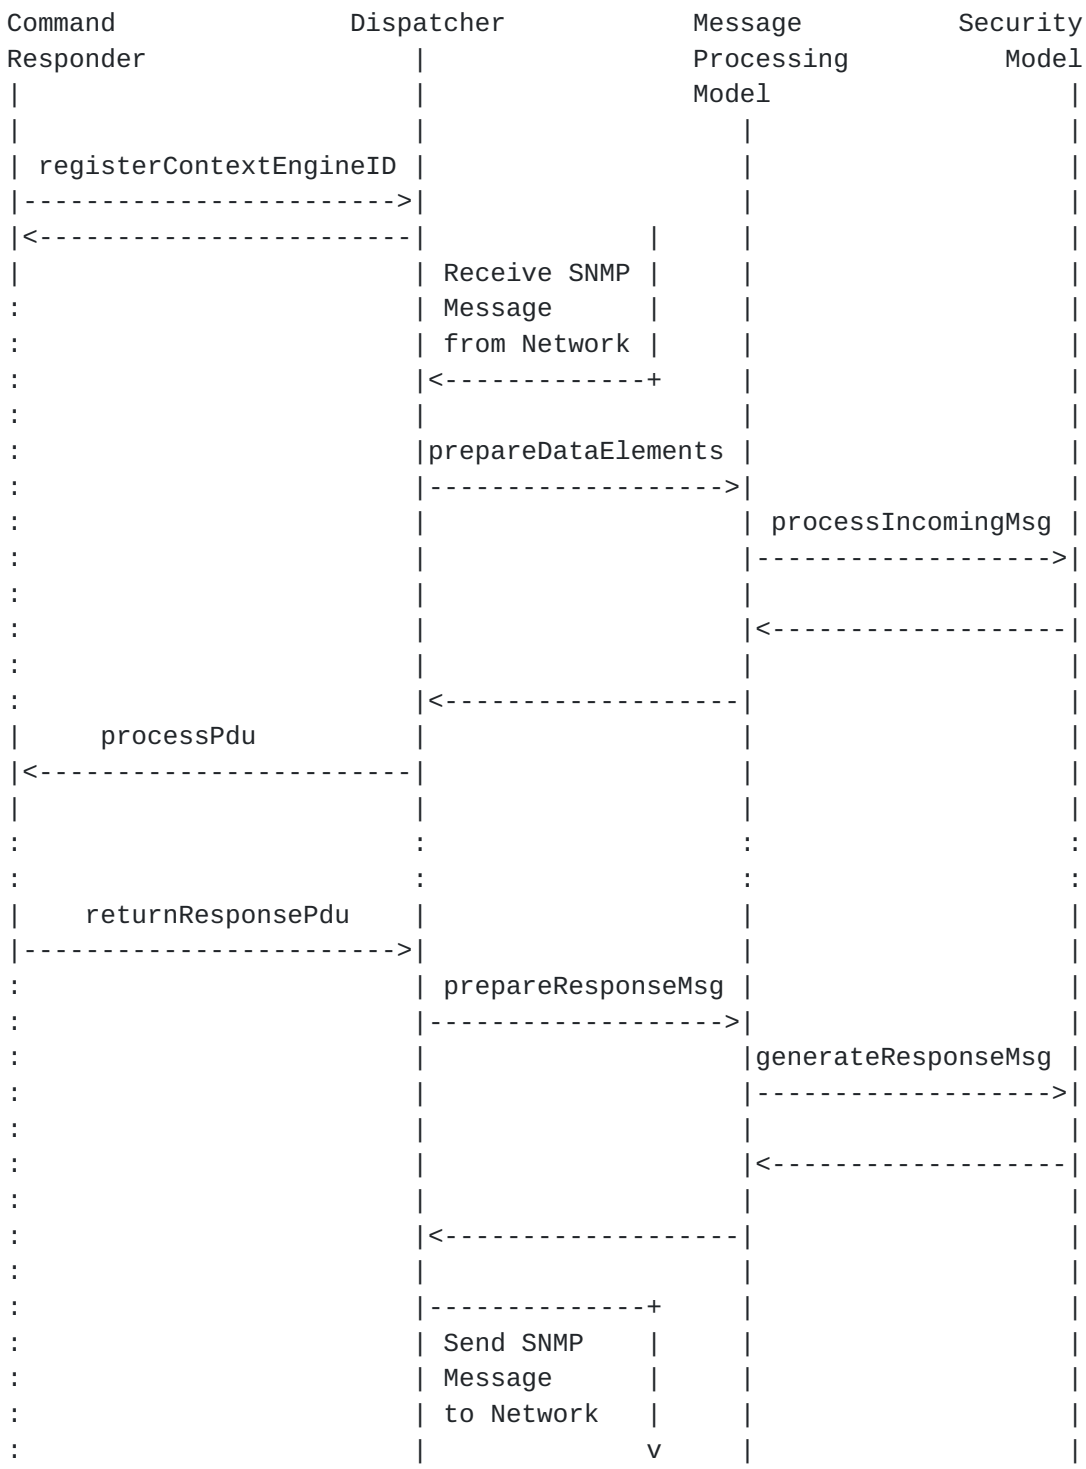

4. Cached Information and References

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 must flow between

subsystems.

There are two levels of state that may 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 and a Local Configuration Datastore (LCD). 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 discussion of the data that must flow between subsystems. An implementation MAY 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.

4.1. securityStateReference

From RFC3411: "For each message received, the Security Model caches the state information such that a Response message can be generated using the same security information, even if the Local Configuration Datastore is altered between the time of the incoming request and the outgoing response.

A Message Processing Model has the responsibility for explicitly releasing the cached data if such data is no longer needed. To enable this, an abstract securityStateReference data element is passed from the Security Model to the Message Processing Model. The cached security data may be implicitly released via the generation of a response, or explicitly released by using the stateRelease primitive, as described in RFC3411 section 4.5.1."

The information saved should include the model-independent parameters (transportType, transportAddress, securityName, securityModel, and securityLevel), related security parameters, and other information needed to imatch the response with the request. The Message Processing Model has the responsibility for explicitly releasing the securityStateReference when such data is no longer needed. The securityStateReference cached data may be implicitly released via the

generation of a response, or explicitly released by using the stateRelease primitive, as described in RFC 3411 section 4.5.1."

If the transport model connection is closed between the time a Request is received and a Response message is being prepared, then the Response message MAY be discarded.

4.2. tmStateReference

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

The format of the cache and the LCD are implementation-specific. For ease of explanation, this document defines a MIB module to conceptually represent the LCD, but this is not meant to contrain implementations from doing it differently.

It is expected that the LCD will allow lookup based on the combination of transportType, transportAddress, securityName, securityModel, and securityLevel. It is expected that the cache contain these values or contain pointers/references to entries in the LCD.

It is expected that a transport model may store transport-specific parameters in the LCD for subsequent usage.

5. Abstract Service Interfaces

[todo: the discussion of ASIs that are not directly related to the transport or security models was added to the document because it was difficult to understand what information was available at what points, and who provided the information. The presence of this expository text can make it hard to find the relevant ASIs for the transport subsystem, and can be confusing because it talks about things that the transport subsystem should not know about. This text should be reduced.

Abstract service interfaces have been defined by RFC 3411 to describe the conceptual data flows between the various subsystems within an SNMP entity.

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 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.

An error indication may return 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.

5.1. Generating an Outgoing SNMP Message

This section describes the procedure followed by an RFC3411-compatible system whenever it generates a message containing a management operation (such as a request, a response, a notification, or a report) on behalf of a user.

```
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
            )
```

Note that tmStateReference has been added to this ASI.

The IN parameters of the prepareOutgoingMessage() ASI are used to pass information from the dispatcher (for the application subsystem) to the message processing subsystem.

The abstract service primitive from a Message Processing Model to a Security Model to generate the components of a Request message is generateRequestMsg().

The abstract service primitive from a Message Processing Model to a Security Model to generate the components of a Response message is generateResponseMsg().

Upon completion of processing, the Security Model returns statusInformation. If the process was successful, the completed message is returned. If the process was not successful, then an errorIndication is returned.

The OUT parameters of the prepareOutgoingMessage() ASI are used to pass information from the message processing model to the dispatcher and on to the transport model:

5.2. Processing for an Outgoing Message

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

```
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
 )
```

The Transport Subsystem provides the following primitives to pass data back and forth between the dispatcher and specific transport models, which provide the interface to the underlying secure transport service. Each transport model should define the elements of procedure for the openSession() and closeSession() interfaces.

```
 statusInformation =
openSession(
IN   transportDomain             -- transport domain to be used
IN   transportAddress            -- transport address to be used
IN   tmStateReference
 )


statusInformation =
closeSession(
IN   tmStateReference
 )
```

5.3. Processing an Incoming SNMP Message

5.3.1. Processing an Incoming Message

If one does not exist, the Transport Model will need to create an entry in a Local Configuration Datastore referenced by tmStateReference. This information will include transportDomain, transportAddress, the securityModel, the securityLevel, and the securityName, plus any model or mechanism-specific details. How this information is determined is model-specific.

The recvMessage ASI is used to pass a message from the transport subsystem to the Dispatcher.

```
statusInformation =
recvMessage(
IN   destTransportDomain          -- transport domain to be used
IN   destTransportAddress         -- transport address to be used
IN   incomingMessage              -- the message received
IN   incomingMessageLength        -- its length
IN   tmStateReference
 )
```

5.3.2. Prepare Data Elements from Incoming Messages

The abstract service primitive from the Dispatcher to a Message Processing Model for a received message is:

```
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         -- 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
)
```

Note that tmStateReference has been added to this ASI.

5.3.3. Processing an Incoming Message

This section describes the procedure followed by the Security Model whenever it receives an incoming message containing a management operation on behalf of a user from a Message Processing model.

The Message Processing Model extracts some information from the wholeMsg. The abstract service primitive from a Message Processing Model to the Security Subsystem for a received message is::

```
   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        -- 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
```

1) The securityEngineID is set to a value in a model-specific manner. If the securityEngineID is not utilized by the specific model, then it should be set to the local snmpEngineID, to satisfy the SNMPv3 message processing model in RFC 3412 section 7.2 13a).

2) Extract the value of securityName from the Local Configuration Datastore entry referenced by tmStateReference.

3) The scopedPDU component is extracted from the wholeMsg.

4) The maxSizeResponseScopedPDU is calculated. This is the maximum size allowed for a scopedPDU for a possible Response message.

5)The security data is cached as cachedSecurityData, so that a possible response to this message can and will use the same security parameters. Then securityStateReference is set for subsequent reference to this cached data.

4) The statusInformation is set to success and a return is made to the calling module passing back the OUT parameters as specified in the processIncomingMsg primitive.

6. The Transport-Subsystem-MIB Module

This memo defines a portion of the Management Information Base (MIB) for statistics in the Transport Subsystem.

6.1. Structure of the MIB Module

Objects in this MIB module are arranged into subtrees. Each subtree is organized as a set of related objects. The overall structure and assignment of objects to their subtrees, and the intended purpose of each subtree, is shown below.

6.1.1. The tmsmStats Subtree

This subtree contains security-model-independent counters which are applicable to all security models based on the .Transport Subsystem. This subtree provides information for identifying fault conditions and performance degradation.

6.2. Relationship to Other MIB Modules

Some management objects defined in other MIB modules are applicable to an entity implementing this MIB. In particular, it is assumed that an entity implementing the Transport-Subsystem-MIB module will also implement the SNMPv2-MIB [RFC3418].

This MIB module is expected to be used with the MIB modules defined for managing specific transport models within the transport subsystem. This MIB module is designed to be transport-model independent and security-model independent, and contains objects useful for managing common aspects of any transport model. Specific transport models may define a MIB module to contain transport-model dependent information.

6.2.1. Textual Conventions

Generic and Common Textual Conventions used in this document can be found summarized at http://www.ops.ietf.org/mib-common-tcs.html

6.2.2. MIB Modules Required for IMPORTS

The. following MIB module imports items from [RFC2578], [RFC2579], [RFC2580], [RFC3411], and [RFC3419]

6.3. Definitions

```
Transport-Subsystem-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE,
    mib-2, Integer32, Unsigned32, Gauge32
      FROM SNMPv2-SMI
    TestAndIncr, StorageType, RowStatus
```

```
        FROM SNMPv2-TC
      MODULE-COMPLIANCE, OBJECT-GROUP
        FROM SNMPv2-CONF
      SnmpSecurityModel,
      SnmpAdminString,  SnmpSecurityLevel, SnmpEngineID
         FROM SNMP-FRAMEWORK-MIB
      TransportAddress, TransportAddressType
        FROM TRANSPORT-ADDRESS-MIB
      ;

  tmsMIB MODULE-IDENTITY
      LAST-UPDATED "200610060000Z"
      ORGANIZATION "ISMS Working Group"
      CONTACT-INFO "WG-EMail:   isms@lists.ietf.org
                    Subscribe:  isms-request@lists.ietf.org

                 Chairs:
                   Juergen Quittek
                   NEC Europe Ltd.
                   Network Laboratories
                   Kurfuersten-Anlage 36
                   69115 Heidelberg
                   Germany
                   +49 6221 90511-15
                    quittek@netlab.nec.de

                    Juergen Schoenwaelder
                    International University Bremen
                    Campus Ring 1
                    28725 Bremen
                    Germany
                    +49 421 200-3587
                    j.schoenwaelder@iu-bremen.de

                 Editor:
                    David Harrington
                    FutureWei Technologies
                    1700 Alma Drive, Suite 100
                    Plano, Texas 75075
                    USA
                    +1 603-436-8634
                    dharrington@huawei.com
                      "
         DESCRIPTION  "The Transport Subsystem MIB Module

                       Copyright (C) The Internet Society (2006). This
                       version of this MIB module is part of RFC XXXX;
                       see the RFC itself for full legal notices.
```

```
  -- NOTE to RFC editor: replace XXXX with actual RFC number
  --                     for this document and remove this note
                     "

         REVISION     "200610060000Z"         -- 20 April 2006
         DESCRIPTION  "The initial version, published in RFC XXXX.
  -- NOTE to RFC editor: replace XXXX with actual RFC number
  --                     for this document and remove this note
                     "

      ::= { mib-2 xxxx }
  -- RFC Ed.: replace xxxx with IANA-assigned number and
  --          remove this note

  -- ------------------------------------------------------------ --
  -- subtrees in the Transport-Subsystem-MIB
  -- ------------------------------------------------------------ --

  tmsNotifications OBJECT IDENTIFIER ::= { tmsMIB 0 }
  tmsObjects       OBJECT IDENTIFIER ::= { tmsMIB 1 }
  tmsConformance   OBJECT IDENTIFIER ::= { tmsMIB 2 }

  -- ------------------------------------------------------------
  -- Objects
  -- ------------------------------------------------------------

  -- Textual Conventions

  SnmpTransportModel ::= TEXTUAL-CONVENTION
      STATUS       current
      DESCRIPTION "An identifier that uniquely identifies a
                   Transport Model of the Transport Subsystem within
                   the SNMP Management Architecture.

                   The values for transportModel are allocated as
                   follows:

                   - The zero value does not identify any particular
                     transport model.

                   - Values between 1 and 255, inclusive, are reserved
                     for standards-track Transport Models and are
                     managed by the Internet Assigned Numbers Authority
                     (IANA).
                   - Values greater than 255 are allocated to
                     enterprise-specific Transport Models.  An
                     enterprise-specific transportModel value is defined
                     to be:
```

```
      enterpriseID * 256 + transport model within
      enterprise

      For example, the fourth Transport Model defined by
      the enterprise whose enterpriseID is 1 would be
      260.

    This scheme for allocation of transportModel
    values allows for a maximum of 255 standards-
    based Transport Models, and for a maximum of
    256 Transport Models per enterprise.

    It is believed that the assignment of new
    transportModel values will be rare in practice
    because the larger the number of simultaneously
    utilized Transport Models, the larger the
    chance that interoperability will suffer.
    Consequently, it is believed that such a range
    will be sufficient.  In the unlikely event that
    the standards committee finds this number to be
    insufficient over time, an enterprise number
    can be allocated to obtain an additional 256
    possible values.

    Note that the most significant bit must be zero;
    hence, there are 23 bits allocated for various
    organizations to design and define non-standard
    transportModels.  This limits the ability to
    define new proprietary implementations of Transport
    Models to the first 8,388,608 enterprises.

    It is worthwhile to note that, in its encoded
    form, the transportModel value will normally
    require only a single byte since, in practice,
    the leftmost bits will be zero for most messages
    and sign extension is suppressed by the encoding
    rules.

    As of this writing, there are several values
    of transportModel defined for use with SNMP or
    reserved for use with supporting MIB objects.
    They are as follows:

        0  reserved for 'any'
        1  reserved for UDP
        2  reserved for TCP
        3  SSH Transport Model
"
```

```
      SYNTAX       INTEGER(0 .. 2147483647)

  -- Notifications for the Transport Subsystem

  -- Statistics for the Transport Subsystem


  tmsStats         OBJECT IDENTIFIER ::= { tmsObjects 1 }


  -- ---------------------------------------------------------
  -- Conformance Information
  -- ---------------------------------------------------------

  tmsGroups OBJECT IDENTIFIER ::= { tmsConformance 1 }

  tmsCompliances OBJECT IDENTIFIER ::= { tmsConformance 2 }

  -- ---------------------------------------------------------
  -- Units of conformance
  -- ---------------------------------------------------------
  tmsGroup OBJECT-GROUP
      OBJECTS {

      }
      STATUS      current
      DESCRIPTION "A collection of objects for maintaining session
                   information of an SNMP engine which implements the
                   Transport subsystem.
                  "

      ::= { tmsGroups 2 }

  -- ---------------------------------------------------------
  -- Compliance statements
  -- ---------------------------------------------------------

  tmsCompliance MODULE-COMPLIANCE
      STATUS      current
      DESCRIPTION
          "The compliance statement for SNMP engines that support the
          Transport-Subsystem-MIB"
      MODULE
          MANDATORY-GROUPS { tmsGroup }
      ::= { tmsCompliances 1 }

  END
```

7. Security Considerations

This document describes an architectural approach and multiple proposed configurations that would permit SNMP to utilize transport layer security services. Each section containing a proposal should discuss the security considerations.

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.

There are no management objects defined in this MIB module that have a MAX-ACCESS clause of read-write and/or read-create. So, if this MIB module is implemented correctly, then there is no risk that an intruder can alter or create any management objects of this MIB module via direct SNMP SET operations.

Some of the readable objects in this MIB module (i.e., objects with a MAX-ACCESS other than not-accessible) may be considered sensitive or vulnerable in some network environments. It is thus important to control even GET and/or NOTIFY access to these objects and possibly to even encrypt the values of these objects when sending them over the network via SNMP. These are the tables and objects and their sensitivity/vulnerability:

- [todo] list the tables and objects and state why they are sensitive.

SNMP versions prior to SNMPv3 did not include adequate security. Even if the network itself is secure (for example by using IPSec), even then, there is no control as to who on the secure network is allowed to access and GET/SET (read/change/create/delete) the objects in this MIB module.

It is RECOMMENDED that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Further, deployment of SNMP versions prior to SNMPv3 is NOT RECOMMENDED. Instead, it is RECOMMENDED to deploy SNMPv3 and to enable cryptographic security. It is then a customer/operator responsibility to ensure that the SNMP entity giving access to an instance of this MIB module is properly configured to give access to the objects only to those principals (users) that have legitimate rights to indeed GET or SET (change/create/delete) them.

8. IANA Considerations

IANA is requested to create a new registry in the Simple Network Management Protocol (SNMP) Number Spaces for SnmpTransportModels, as described in the Transport-Subsystem-MIB defined in this document. Values 0 through 255 are IANA-assigned by Standards Action, as defined in RFC2434. Values above 255 are assigned by Hierarchical allocation, using the algorithm defined in the definition of the SnmpTransportModels TEXTUAL-CONVENTION in the Transport-Subsystem-MIB in this document.

The MIB module in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
Transport-Subsystem-MIB	{ mib-2 XXXX }

Editor's Note (to be removed prior to publication): the IANA is requested to assign a value for "XXXX" under the 'mib-2' subtree and to record the assignment in the SMI Numbers registry. When the assignment has been made, the RFC Editor is asked to replace "XXXX" (here and in the MIB module) with the assigned value and to remove this note.

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, Dave Harrington, Keith McCloghrie, Kaushik Narayan, Dave 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 committed to and performed detailed reviews: Jeffrey Hutzelman

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4366] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", RFC 4366, April 2006.

[RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.

[RFC2579] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Textual Conventions for SMIv2", STD 58, RFC 2579, April 1999.

[RFC2580] McCloghrie, K., Perkins, D., and J. Schoenwaelder, "Conformance Statements for SMIv2", STD 58, RFC 2580, April 1999.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, 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, RFC 3412, 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.

[RFC3416] Presuhn, R., "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3416, December 2002.

[RFC3417] Presuhn, R., "Transport Mappings for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3417, December 2002.

[RFC3418] Presuhn, R., "Management Information Base (MIB) for the

Simple Network Management Protocol (SNMP)", STD 62, RFC 3418, December 2002.

[RFC3419] Daniele, M. and J. Schoenwaelder, "Textual Conventions for Transport Addresses", RFC 3419, December 2002.

[RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture", RFC 4251, January 2006.

10.2. Informative References

[RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.

[RFC3413] Levi, D., Meyer, P., and B. Stewart, "Simple Network Management Protocol (SNMP) Applications", STD 62, RFC 3413, December 2002.

[RFC4422] Melnikov, A. and K. Zeilenga, "Simple Authentication and Security Layer (SASL)", RFC 4422, June 2006.

[I-D.ietf-netconf-ssh] Wasserman, M. and T. Goddard, "Using the NETCONF Configuration Protocol over Secure Shell (SSH)", draft-ietf-netconf-ssh-06 (work in progress), March 2006.

Appendix A. Parameter Table

Following is a CSV formatted matrix useful for tracking data flows into and out of the dispatcher, message, and security subsystems. Import this into your favorite spreadsheet or other CSV compatible application. You will need to remove lines feeds from the second and third lines, which needed to be wrapped to fit into RFC limits.

A.1. ParameterList.csv

```
,Dispatcher,,,,Messaging,,,Security,,

,sendPdu,returnResponse,processPdu,processResponse
,prepareOutgoingMessage,prepareResponseMessage,prepareDataElements
,generateRequest,processIncoming,generateResponse

transportDomain,In,,,,In,,In,,,
```

```
   transportAddress,In,,,,In,,In,,,

   destTransportDomain,,,,,Out,Out,,,,

   destTransportAddress,,,,,Out,Out,,,,

   messageProcessingModel,In,In,In,In,In,In,Out,In,In,In

   securityModel,In,In,In,In,In,In,Out,In,In,In

   securityName,In,In,In,In,In,In,Out,In,Out,In

   securityLevel,In,In,In,In,In,In,Out,In,In,In

   contextEngineID,In,In,In,In,In,In,Out,,,

   contextName,In,In,In,In,In,In,Out,,,

   expectResponse,In,,,,In,,,,,

   PDU,In,In,In,In,In,In,Out,,,

   pduVersion,In,In,In,In,In,In,Out,,,

   statusInfo,Out,In,,In,,In,Out,Out,Out,Out

   errorIndication,Out,Out,,,,,Out,,,

   sendPduHandle,Out,,,In,In,,Out,,,

   maxSizeResponsePDU,,In,In,,,In,Out,,Out,

   stateReference,,In,In,,,In,Out,,,

   wholeMessage,,,,,Out,Out,,Out,In,Out

   messageLength,,,,,Out,Out,,Out,In,Out

   maxMessageSize,,,,,,,,In,In,In

   globalData,,,,,,,,In,,In

   securityEngineID,,,,,,,,In,Out,In

   scopedPDU,,,,,,,,In,Out,In

   securityParameters,,,,,,,,Out,,Out
```

securityStateReference,,,,,,,,,Out,In

pduType,,,,,,,Out,,,

tmStateReference,,,,,,Out,In,,In,

Appendix B. Why tmStateReference?

This appendix considers why a cache-based approach was selected for passing parameters. This section may be removed from subsequent revisions of the document.

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

B.1. Define an Abstract Service Interface

Abstract Service Interfaces (ASIs) [RFC3411] 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.

B.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 messaging security model. The message header would be included in the wholeMessage ASI parameter, and would be removed by a corresponding messaging 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.

B.3. Modifying Existing Fields in an SNMP Message

[RFC3412] describes 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 messaging security 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.

B.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 an Security Model extracts that information from the cache. Given that there may 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. This approach of passing a model-independent reference is consistent with the securityStateReference cache already being passed around in the RFC3411 ASIs.

Appendix C. Open Issues

Appendix D. Change Log

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

Changes 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-

- removed session table from MIB module
- removed sessionID from ASIs
- reorganized to put ASI discussions in EOP section, as was done in SSHSM
- changed user auth to client auth
- changed tmStateReference to tmSessionReference
- 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-

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

- reorganized slightly to better separate requirements from proposed solution. This probably needs additional work.
- removed section with sample protocols and sample tmSessionReference.
- Added section for acronyms
- moved section comparing parameter passing techniques to appendix.
- Removed section on notification requirements.

Changes from revision -00-

- changed SSH references from I-Ds to RFCs
- removed parameters from tmSessionReference for DTLS that revealed lower layer info.
- Added TMS-MIB module
- Added Internet-Standard Management Framework boilerplate
- Added Structure of the MIB Module
- Added MIB security considerations boilerplate (to be completed)
- Added IANA Considerations
- Added ASI Parameter table
- Added discussion of Sessions
- Added Open issues and Change Log
- 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
International University Bremen
Campus Ring 1
28725 Bremen
Germany

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

Full Copyright Statement

Intellectual Property

Acknowledgement

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