

Network Working Group
Request for Comments: 5209
Category: Informational

P. Sangster
Symantec
H. Khosravi
Intel
M. Mani
Avaya
K. Narayan
Cisco Systems
J. Tardo
Nevis Networks
June 2008

Network Endpoint Assessment (NEA): Overview and Requirements

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document defines the problem statement, scope, and protocol requirements between the components of the NEA (Network Endpoint Assessment) reference model. NEA provides owners of networks (e.g., an enterprise offering remote access) a mechanism to evaluate the posture of a system. This may take place during the request for network access and/or subsequently at any time while connected to the network. The learned posture information can then be applied to a variety of compliance-oriented decisions. The posture information is frequently useful for detecting systems that are lacking or have out-of-date security protection mechanisms such as: anti-virus and host-based firewall software. In order to provide context for the requirements, a reference model and terminology are introduced.

Table of Contents

| | |
|--|----|
| 1. Introduction | 3 |
| 1.1. Requirements Language | 4 |
| 2. Terminology | 5 |
| 3. Applicability | 7 |
| 3.1. Scope | 7 |
| 3.2. Applicability of Environments | 8 |
| 4. Problem Statement | 9 |
| 5. Reference Model | 10 |
| 5.1. NEA Client and Server | 12 |
| 5.1.1. NEA Client | 12 |
| 5.1.1.1. Posture Collector | 12 |
| 5.1.1.2. Posture Broker Client | 14 |
| 5.1.1.3. Posture Transport Client | 15 |
| 5.1.2. NEA Server | 15 |
| 5.1.2.1. Posture Validator | 15 |
| 5.1.2.2. Posture Broker Server | 17 |
| 5.1.2.3. Posture Transport Server | 18 |
| 5.2. Protocols | 18 |
| 5.2.1. Posture Attribute Protocol (PA) | 18 |
| 5.2.2. Posture Broker Protocol (PB) | 19 |
| 5.2.3. Posture Transport Protocol (PT) | 19 |
| 5.3. Attributes | 20 |
| 5.3.1. Attributes Normally Sent by NEA Client: | 21 |
| 5.3.2. Attributes Normally Sent by NEA Server: | 21 |
| 6. Use Cases | 22 |
| 6.1. Initial Assessment | 22 |
| 6.1.1. Triggered by Network Connection or Service Request | 22 |
| 6.1.1.1. Example | 23 |
| 6.1.1.2. Possible Flows and Protocol Usage | 23 |
| 6.1.1.3. Impact on Requirements | 25 |
| 6.1.2. Triggered by Endpoint | 25 |
| 6.1.2.1. Example | 25 |
| 6.1.2.2. Possible Flows and Protocol Usage | 26 |
| 6.1.2.3. Impact on Requirements | 28 |
| 6.2. Posture Reassessment | 28 |
| 6.2.1. Triggered by NEA Client | 28 |
| 6.2.1.1. Example | 28 |
| 6.2.1.2. Possible Flows & Protocol Usage | 29 |
| 6.2.1.3. Impact on Requirements | 30 |
| 6.2.2. Triggered by NEA Server | 30 |
| 6.2.2.1. Example | 30 |
| 6.2.2.2. Possible Flows and Protocol Usage | 31 |
| 6.2.2.3. Impact on Requirements | 33 |

| | |
|---|----|
| 7. Requirements | 34 |
| 7.1. Common Protocol Requirements | 34 |
| 7.2. Posture Attribute (PA) Protocol Requirements | 35 |
| 7.3. Posture Broker (PB) Protocol Requirements | 36 |
| 7.4. Posture Transport (PT) Protocol Requirements | 38 |
| 8. Security Considerations | 38 |
| 8.1. Trust | 39 |
| 8.1.1. Endpoint | 40 |
| 8.1.2. Network Communications | 41 |
| 8.1.3. NEA Server | 42 |
| 8.2. Protection Mechanisms at Multiple Layers | 43 |
| 8.3. Relevant Classes of Attack | 43 |
| 8.3.1. Man-in-the-Middle (MITM) | 44 |
| 8.3.2. Message Modification | 45 |
| 8.3.3. Message Replay or Attribute Theft | 45 |
| 8.3.4. Other Types of Attack | 46 |
| 9. Privacy Considerations | 46 |
| 9.1. Implementer Considerations | 47 |
| 9.2. Minimizing Attribute Disclosure | 49 |
| 10. References | 50 |
| 10.1. Normative References | 50 |
| 10.2. Informative References | 50 |
| 11. Acknowledgments | 51 |

1. Introduction

Endpoints connected to a network may be exposed to a wide variety of threats. Some protection against these threats can be provided by ensuring that endpoints conform to security policies. Therefore, the intent of NEA is to assess these endpoints to determine their compliance with security policies so that corrective measures can be provided before they are exposed to those threats. For example, if a system is determined to be out of compliance because it is lacking proper defensive mechanisms such as host-based firewalls, anti-virus software, or the absence of critical security patches, the NEA protocols provide a mechanism to detect this fact and indicate appropriate remediation actions to be taken. Note that an endpoint that is deemed compliant may still be vulnerable to threats that may exist on the network.

NEA typically involves the use of special client software running on the requesting endpoint that observes and reports on the configuration of the system to the network infrastructure. The infrastructure has corresponding validation software that is capable of comparing the endpoint's configuration information with network compliance policies and providing the result to appropriate authorization entities that make decisions about network and application access. Some endpoints may be incapable of running the

NEA Client software (e.g., printer) or be unwilling to share information about their configuration. This situation is outside the scope of NEA and is subject to local policies.

The result of an endpoint assessment may influence an access decision that is provisioned to the enforcement mechanisms on the network and/or endpoint requesting access. While the NEA Working Group recognizes there may be a link between an assessment and the enforcement of a resulting access decision, the mechanisms and protocols for enforcement are not in scope for this specification.

Architectures, similar to NEA, have existed in the industry for some time and are present in shipping products, but do not offer adequate interoperability. Some examples of such architectures include: Trusted Computing Group's Trusted Network Connect [TNC], Microsoft's Network Access Protection [NAP], and Cisco's Cisco Network Admission Control [CNAC]. These technologies assess the software and/or hardware configuration of endpoint devices for the purposes of monitoring or enforcing compliance to an organization's policy.

The NEA Working Group is developing standard protocols that can be used to communicate compliance information between a NEA Client and a NEA Server. This document provides the context for NEA including: terminology, applicability, problem statement, reference model, and use cases. It then identifies requirements for the protocols used to communicate between a NEA Client and NEA server. Finally, this document discusses some potential security and privacy considerations with the use of NEA. The majority of this specification provides informative text describing the context of NEA.

1.1. Requirements Language

Use of each capitalized word within a sentence or phrase carries the following meaning during the NEA WG's protocol selection process:

MUST - indicates an absolute requirement

MUST NOT - indicates something absolutely prohibited

SHOULD - indicates a strong recommendation of a desired result

SHOULD NOT - indicates a strong recommendation against a result

MAY - indicates a willingness to allow an optional outcome

Lower case use of "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", and "MAY" carry their normal meaning and are not subject to these definitions.

2. Terminology

This section defines a set of terms used throughout this document. In some cases these terms have been used in other contexts with different meanings so this section attempts to describe each term's meaning with respect to the NEA WG activities.

Assessment - The process of collecting posture for a set of capabilities on the endpoint (e.g., host-based firewall) such that the appropriate validators may evaluate the posture against compliance policy.

Assertion Attributes - Attributes that include reusable information about the success of a prior assessment of the endpoint. This could be used to optimize subsequent assessments by avoiding a full posture reassessment. For example, this classification of attribute might be issued specifically to a particular endpoint, dated and signed by the NEA Server allowing that endpoint to reuse it for a time period to assert compliance to a set of policies. The NEA Server might accept this in lieu of obtaining posture information.

Attribute - Data element including any requisite meta-data describing an observed, expected, or the operational status of an endpoint feature (e.g., anti-virus software is currently in use). Attributes are exchanged as part of the NEA protocols (see section 5.2). NEA recognizes a variety of usage scenarios where the use of an attribute in a particular type of message could indicate:

- o previously assessed status (Assertion Attributes),
- o observed configuration or property (Posture Attributes),
- o request for configuration or property information (Request Attributes),
- o assessment decision (Result Attributes), or
- o repair instructions (Remediation Attributes).

The NEA WG will standardize a subset of the attribute namespace known as standard attributes. Those attributes not standardized are referred to in this specification as vendor-specific.

Dialog - Sequence of request/response messages exchanged.

Endpoint - Any computing device that can be connected to a network. Such devices normally are associated with a particular link layer address before joining the network and potentially an IP address once on the network. This includes: laptops, desktops, servers, cell phones, or any device that may have an IP address.

Message - Self contained unit of communication between the NEA Client and Server. For example, a posture attribute message might carry a set of attributes describing the configuration of the anti-virus software on an endpoint.

Owner - the role of an entity who is the legal, rightful possessor of an asset (e.g., endpoint). The owner is entitled to maintain control over the policies enforced on the device even if the asset is not within the owner's possession. The owner may permit user override or augmentation of control policies or may choose to not assert any policies limiting use of asset.

Posture - Configuration and/or status of hardware or software on an endpoint as it pertains to an organization's security policy.

Posture Attributes - Attributes describing the configuration or status (posture) of a feature of the endpoint. For example, a Posture Attribute might describe the version of the operating system installed on the system.

Request Attributes - Attributes sent by a NEA Server identifying the posture information requested from the NEA Client. For example, a Request Attribute might be an attribute included in a request message from the NEA Server that is asking for the version information for the operating system on the endpoint.

Remediation Attributes - Attributes containing the remediation instructions for how to bring an endpoint into compliance with one or more policies. The NEA WG will not define standard remediation attributes, but this specification does describe where they are used within the reference model and protocols.

Result Attributes - Attributes describing whether the endpoint is in compliance with NEA policy. The Result Attribute is created by the NEA Server normally at the conclusion of the assessment to indicate whether or not the endpoint was considered compliant. More than one of these attributes may be used allowing for more granular feature level decisions to be communicated in addition to an overall, global assessment decision.

Session - Stateful connection capable of carrying multiple message exchanges associated with (an) assessment(s) of a particular endpoint. This document defines the term session at a conceptual level and does not describe the properties of the session or specify requirements for the NEA protocols to manage these sessions.

User - Role of a person that is making use of the services of an endpoint. The user may not own the endpoint so he or she might need to operate within the acceptable use constraints defined by the endpoint's owner. For example, an enterprise employee might be a user of a computer provided by the enterprise (owner) for business purposes.

3. Applicability

This section discusses the scope of the technologies being standardized and the network environments where it is envisioned that the NEA technologies might be applicable.

3.1. Scope

The priority of the NEA Working Group is to develop standard protocols at the higher layers in the reference model (see section 5): the Posture Attribute protocol (PA) and the Posture Broker protocol (PB). PA and PB will be designed to be carried over a variety of lower layer transport (PT) protocols. The NEA WG will identify standard PT protocol(s) that are mandatory to implement. PT protocols may be defined in other WGs because the requirements may not be specific to NEA. When used with a standard PT protocol (e.g., Extensible Authentication Protocol (EAP), Transport Layer Security (TLS) [TLS]), the PA and PB protocols will allow interoperability between a NEA Client from one vendor and a NEA Server from another. This specification will not focus on the other interfaces between the functional components of the NEA reference model nor requirements on their internals. Any discussion of these aspects is included to provide context for understanding the model and resulting requirements.

Some tangent areas not shown in the reference model that are also out of scope for the NEA working group, and thus this specification, include:

- o Standardizing the protocols and mechanisms for enforcing restricted network access,
- o Developing standard protocols for remediation of non-compliant endpoints,

- o Specifying protocols used to communicate with remote portions of the NEA Client or Server (e.g., remote collectors or validators of posture),
- o Supporting a NEA Client providing posture for other endpoints (e.g., a NEA Client on an Intrusion Detection System (IDS) providing posture for an endpoint without a NEA Client),
- o Defining the set of events or situations that might trigger a NEA Client or NEA Server to request an assessment,
- o Detecting or handling lying endpoints (see section 8.1.1 for more information).

3.2. Applicability of Environments

Because the NEA model is based on NEA-oriented software being present on the endpoint and in the network infrastructure, and due to the nature of the information being exposed, the use of NEA technologies may not apply in a variety of situations possible on the Internet. Therefore, this section discusses some of the scenarios where NEA is most likely to be applicable and some where it may not be. Ultimately, the use of NEA within a deployment is not restricted to just these scenarios. The decision of whether to use NEA technologies lies in the hands of the deployer (e.g., network provider) based upon the expected relationship they have with the owners and users of potential endpoints.

NEA technologies are largely focused on scenarios where the owner of the endpoint is the same as the owner of the network. This is a very common model for enterprises that provide equipment to employees to perform their duties. These employees are likely bound under an employment contract that outlines what level of visibility the employer expects to have into the employee's use of company assets and possibly activities during work hours. This contract may establish the expectation that the endpoint needs to conform to policies set forth by the enterprise.

Some other environments may be in a similar situation and thus find NEA technologies to be beneficial. For example, environments where the endpoint is owned by a party (possibly even the user) that has explicitly expressed a desire to conform to the policies established by a network or service provider in exchange for being able to access its resources. An example of this might be an independent contractor with a personal laptop, working for a company imposing NEA assessment policies on its employees, who may wish a similar level of access and is willing to conform to the company's policies. NEA technologies may be applicable to this situation.

Conversely, some environments where NEA is not expected to be applicable would be environments where the endpoint is owned by a user that has not agreed to conform to a network provider's policies. An example might include when the above contractor visits any public area like the local coffee shop that offers Internet access. This coffee shop would not be expected to be able to use NEA technologies to assess the posture of the contractor's laptop. Because of the potentially invasive nature of NEA technology, such an assessment could amount to an invasion of privacy of the contractor.

It is more difficult to determine whether NEA is applicable in other environments, so the NEA WG will consider them to be out of scope for consideration and specification. In order for an environment to be considered applicable for NEA, the owner or user of an endpoint must have established a clear expectation that it will comply with the policies of the owner and operator of the network. Such an expectation likely includes a willingness to disclose appropriate information necessary for the network to perform compliance checks.

4. Problem Statement

NEA technology may be used for a variety of purposes. This section highlights some of the major situations where NEA technologies may be beneficial.

One use is to facilitate endpoint compliance checking against an organization's security policy when an endpoint connects to the network. Organizations often require endpoints to run an IT-specified Operating System (OS) configuration and have certain security applications enabled, e.g., anti-virus software, host intrusion detection/prevention systems, personal firewalls, and patch management software. An endpoint that is not compliant with IT policy may be vulnerable to a number of known threats that might exist on the network.

Without NEA technology, ensuring compliance of endpoints to corporate policy is a time-consuming and difficult task. Not all endpoints are managed by a corporation's IT organization, e.g., lab assets and contractor machines. Even for assets that are managed, they may not receive updates in a timely fashion because they are not permanently attached to the corporate network, e.g., laptops. With NEA technology, the network is able to assess an endpoint as soon as it requests access to the network or at any time after joining the network. This provides the corporation an opportunity to check compliance of all NEA-capable endpoints in a timely fashion and facilitate endpoint remediation potentially while quarantined when needed.

NEA technology can be used to provide posture assessment for a range of ways of connecting to the network including (but not limited to) wired and wireless LAN access such as using 802.1X [802.1X], remote access via IPsec [IPSEC], or Secure Socket Layer (SSL) VPN, or gateway access.

Endpoints that are not NEA-capable or choose not to share sufficient posture to evaluate compliance may be subject to different access policies. The decision of how to handle non-compliant or non-participating endpoints can be made by the network administrator possibly based on information from other security mechanisms on the network (e.g., authentication). For example, remediation instructions or warnings may be sent to a non-compliant endpoint with a properly authorized user while allowing limited access to the network. Also, network access technologies can use the NEA results to restrict or deny access to an endpoint, while allowing vulnerabilities to be addressed before an endpoint is exposed to attack. The communication and representation of NEA assessment results to network access technologies on the network is out of scope for this document.

Reassessment is a second important use of NEA technology as it allows for additional assessments of previously considered compliant endpoints to be performed. This might become necessary because network compliance policies and/or endpoint posture can change over time. A system initially assessed as being compliant when it joined the network may no longer be in compliance after changes occur. For example, reassessment might be necessary if a user disables a security protection (e.g., host-based firewall) required by policy or when the firewall becomes non-compliant after a firewall patch is issued and network policy is changed to require the patch.

A third use of NEA technology may be to verify or supplement organization asset information stored in inventory databases.

NEA technology can also be used to check and report compliance for endpoints when they try to access certain mission critical applications within an enterprise, employing service (application) triggered assessment.

5. Reference Model

This section describes the reference model for Network Endpoint Assessment. This model is provided to establish a context for the discussion of requirements and may not directly map to any particular product or deployment architecture. The model identifies the major

functionality of the NEA Client and Server and their relationships, as well as the protocols they use to communicate at various levels (e.g., PA is carried by the PB protocol).

While the diagram shows 3 layered protocols, it is envisioned that PA is likely a thin message wrapper around a set of attributes and that it is batched and encapsulated in PB. PB is primarily a lightweight message batching protocol, so the protocol stack is mostly the transport (PT). The vertical lines in the model represent APIs and/or protocols between components within the NEA Client or Server. These interfaces are out of scope for standardization in the NEA WG.

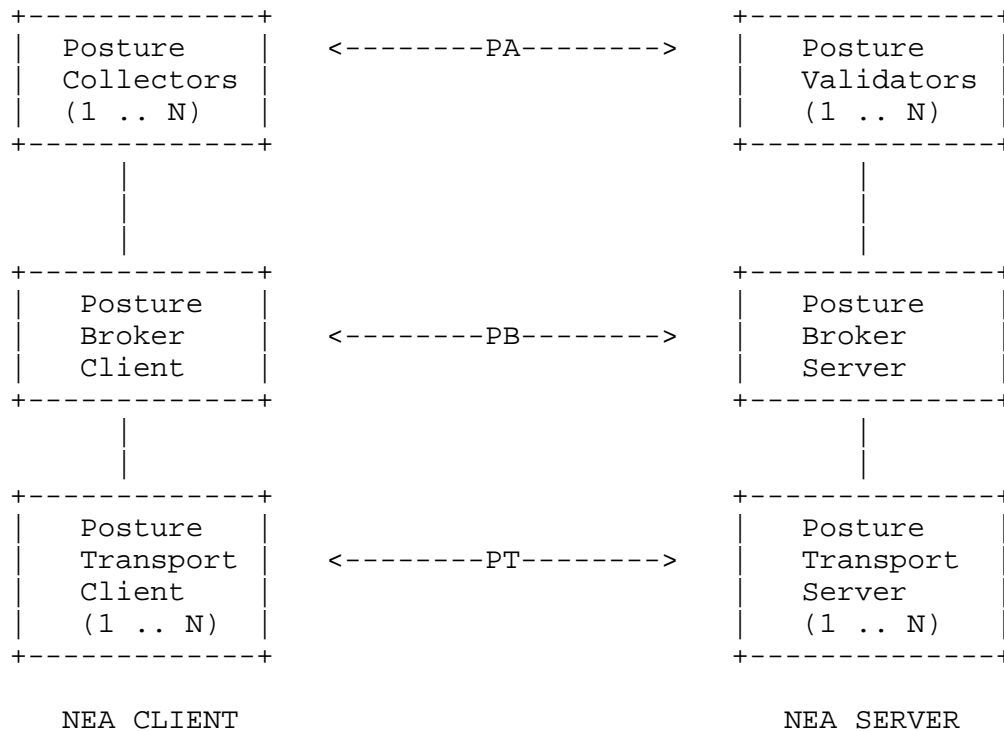


Figure 1: NEA Reference Model

The NEA reference model does not include mechanisms for discovery of NEA Clients and NEA Servers. It is expected that NEA Clients and NEA Servers are configured with information that allows them to reach each other. The specific methods of referencing the configuration and establishing the communication channel are out of scope for the NEA reference model and should be covered in the specifications of candidate protocols such as the Posture Transport (PT) protocol.

5.1. NEA Client and Server

5.1.1. NEA Client

The NEA Client is resident on an endpoint device and comprised of the following functionality:

- o Posture Collector(s)
- o Posture Broker Client
- o Posture Transport Client(s)

The NEA Client is responsible for responding to requests for attributes describing the configuration of the local operating domain of the client and handling the assessment results including potential remediation instructions for how to conform to policy. A NEA Client is not responsible for reporting on the posture of entities that might exist on the endpoint or over the network that are outside the domain of execution (e.g., in other virtual machine domains) of the NEA Client.

For example, a network address translation (NAT) device might route communications for many systems behind it, but when the NAT device joins the network, its NEA Client would only report its own (local) posture. Similarly, endpoints with virtualization capabilities might have multiple independent domains of execution (e.g., OS instances). Each NEA Client is only responsible for reporting posture for its domain of execution, but this information might be aggregated by other local mechanisms to represent the posture for multiple domains on the endpoint. Such posture aggregation mechanisms are outside the focus of this specification.

Endpoints lacking NEA Client software (which is out of NEA scope) or choosing not to provide the attributes required by the NEA Server could be considered non-compliant. The NEA model includes capabilities to enable the endpoint to update its contents in order to become compliant.

5.1.1.1. Posture Collector

The Posture Collector is responsible for responding to requests for posture information in Request Attributes from the NEA Server. The Posture Collector is also responsible for handling assessment decisions in Result Attributes and remediation instructions in Remediation Attributes. A single NEA Client can have several Posture Collectors capable of collecting standard and/or vendor-specific Posture Attributes for particular features of the endpoint. Typical

examples include Posture Collectors that provide information about Operating System (OS) version and patch levels, anti-virus software, and security mechanisms on the endpoint such as host-based Intrusion Detection System (IDS) or firewall.

Each Posture Collector will be associated with one or more identifiers that enable it to be specified as the destination in a PA message. The Posture Broker Client uses these identifiers to route messages to this Collector. An identifier might be dynamic (e.g., generated by the Posture Broker Client at run-time during registration) or more static (e.g., pre-assigned to the Posture Collector at install-time and passed to the Posture Broker Client during registration) or a function of the attribute messages the Collector desires to receive (e.g., message type for subscription).

The NEA model allocates the following responsibilities to the Posture Collector:

- o Consulting with local privacy and security policies that may restrict what information is allowed to be disclosed to a given NEA Server.
- o Receiving Request Attributes from a Posture Validator and performing the local processing required to respond appropriately. This may include:
 - Collecting associated posture information for particular features of the endpoint and returning this information in Posture Attributes.
 - Caching and recognizing the applicability of recently issued attributes containing reusable assertions that might serve to prove compliance and returning this attribute instead of posture information.
- o Receiving attributes containing remediation instructions on how to update functionality on the endpoint. This could require the Collector to interact with the user, owner, and/or a remediation server.
- o Monitoring the posture of (a) particular features(s) on the endpoint for posture changes that require notification to the Posture Broker Client.
- o Providing cryptographic verification of the attributes received from the Validator and offering cryptographic protection to the attributes returned.

The above list describes the model's view of the possible responsibilities of the Posture Collector. Note that this is not a set of requirements for what each Posture Collector implementation must support, nor is it an exhaustive list of all the things a Posture Collector may do.

5.1.1.2. Posture Broker Client

The Posture Broker Client is both a PA message multiplexer and a de-multiplexer. The Posture Broker Client is responsible for de-multiplexing the PB message received from the NEA Server and distributing each encapsulated PA message to the corresponding Posture Collector(s). The model also allows for the posture information request to be pre-provisioned on the NEA Client to improve performance by allowing the NEA Client to report posture without receiving a request for particular attributes from the NEA Server.

The Posture Broker Client also multiplexes the responses from the Posture Collector(s) and returns them to the NEA Server. The Posture Broker Client constructs one or more PB messages using the PA message(s) it obtains from the Posture Collector(s) involved in the assessment. The quantity and ordering of Posture Collector responses (PA message(s)) multiplexed into the PB response message(s) can be determined by the Posture Broker Client based on many factors including policy or characteristics of the underlying network transport (e.g., MTU). A particular NEA Client will have one Posture Broker Client.

The Posture Broker Client also handles the global assessment decision from the Posture Broker Server and may interact with the user to communicate the global assessment decision and aid in any necessary remediation steps.

The NEA model allocates the following responsibilities to the Posture Broker Client:

- o Maintaining a registry of known Posture Collectors and allowing for Posture Collectors to dynamically register and deregister.
- o Multiplexing and de-multiplexing attribute messages between the NEA Server and the relevant Posture Collectors.
- o Handling posture change notifications from Posture Collectors and triggering reassessment.
- o Providing user notification about the global assessment decision and other user messages sent by the NEA Server.

5.1.1.3. Posture Transport Client

The Posture Transport Client is responsible for establishing a reliable communication channel with the NEA Server for the message dialog between the NEA Client and NEA Server. There might be more than one Posture Transport Client on a particular NEA Client supporting different transport protocols (e.g., 802.1X, VPN). Certain Posture Transport Clients may be configured with the address of the appropriate Posture Transport Server to use for a particular network.

The NEA model allocates the following responsibilities to the Posture Transport Client:

- o Initiating and maintaining the communication channel to the NEA Server. The Posture Transport Client hides the details of the underlying carrier that could be a Layer 2 or Layer 3 protocol.
- o Providing cryptographic protection for the message dialog between the NEA Client and NEA Server.

5.1.2. NEA Server

The NEA Server is typically comprised of the following NEA functionality:

- o Posture Validator(s)
- o Posture Broker Server
- o Posture Transport Server(s)

The Posture Validators might be located on a separate server from the Posture Broker Server, requiring the Posture Broker Server to deal with both local and remote Posture Validators.

5.1.2.1. Posture Validator

A Posture Validator is responsible for handling Posture Attributes from corresponding Posture Collector(s). A Posture Validator can handle Posture Attributes from one or more Posture Collectors and vice-versa. The Posture Validator performs the posture assessment for one or more features of the endpoint (e.g., anti-virus software) and creates the result and, if necessary, the remediation instructions, or it may choose to request additional attributes from one or more Collectors.

Each Posture Validator will be associated with one or more identifiers that enable it to be specified as the destination in a PA message. The Posture Broker Server uses this identifier to route messages to this Validator. This identifier might be dynamic (e.g., generated by the Posture Broker Server at run-time during registration) or more static (e.g., pre-assigned to a Posture Validator at install-time and passed to the Posture Broker Server during registration) or a function of the attribute messages the Validator desires to receive (e.g., message type for subscription).

Posture Validators can be co-located on the NEA Server or can be hosted on separate servers. A particular NEA Server is likely to need to handle multiple Posture Validators.

The NEA model allocates the following responsibilities to the Posture Validator:

- o Requesting attributes from a Posture Collector. The request may include:
 - Request Attributes that indicate to the Posture Collector to fetch and provide Posture Attributes for particular functionality on the endpoint.
- o Receiving attributes from the Posture Collector. The response from the Posture Collector may include:
 - Posture Attributes collected for the requested functionality.
 - Assertion Attributes that indicate the compliance result from a prior assessment.
- o Assessing the posture of endpoint features based on the attributes received from the Collector.
- o Communicating the posture assessment result to the Posture Broker Server.
- o Communicating the posture assessment results to the Posture Collector; this attribute message may include:
 - Result Attributes that communicate the posture assessment result.
 - Remediation Attributes that communicate the remediation instructions to the Posture Collector.
- o Monitoring out-of-band updates that trigger reassessment and require notifications to be sent to the Posture Broker Server.

- o Providing cryptographic protection for attributes sent to the Posture Collector and offering cryptographic verification of the attributes received from the Posture Collector.

The above list describes the model's view of the possible responsibilities of the Posture Validator. Note that this is not a set of requirements for what each Posture Validator implementation must support, nor is it an exhaustive list of all the things a Posture Validator may do.

5.1.2.2. Posture Broker Server

The Posture Broker Server acts as a multiplexer and a de-multiplexer for attribute messages. The Posture Broker Server parses the PB messages received from the NEA Client and de-multiplexes them into PA messages that it passes to the associated Posture Validators. The Posture Broker Server multiplexes the PA messages (e.g., messages containing (a) Request Attribute(s) from the relevant Posture Validator(s)) into one or more PB messages and sends them to the NEA Client via the Posture Transport protocol. The quantity and ordering of Posture Validator responses (PA messages) and global assessment decision multiplexed into the PB response message(s) can be determined by the Posture Broker Server based on many factors including policy or characteristics of the underlying network transport (e.g., MTU).

The Posture Broker Server is also responsible for computing the global assessment decision based on individual posture assessment results from the various Posture Validators. This global assessment decision is sent back to the NEA Client in Result Attributes within a PB message. A particular NEA Server will have one Posture Broker Server, and this Posture Broker Server will handle all the local and remote Posture Validators.

The NEA model allocates the following responsibilities to the Posture Broker Server:

- o Maintaining a registry of Posture Validators and allowing for Posture Validators to register and deregister.
- o Multiplexing and de-multiplexing posture messages from and to the relevant Posture Validators.
- o Computing the global assessment decision based on posture assessment results from the various Posture Validators and compliance policy. This assessment decision is sent to the Posture Broker Client in a PB message.

5.1.2.3. Posture Transport Server

The Posture Transport Server is responsible for establishing a reliable communication channel with the NEA Client for the message dialog between the NEA Client and NEA Server. There might be more than one Posture Transport Server on a particular NEA Server to support different transport protocols. A particular Posture Transport Server will typically handle requests from several Posture Transport Clients and may require local configuration describing how to reach the NEA Clients.

The NEA model allocates the following responsibilities to the Posture Transport Server:

- o Initiating and maintaining a communication channel with, potentially, several NEA Clients.
- o Providing cryptographic protection for the message dialog between the NEA Client and NEA Server.

5.2. Protocols

The NEA reference model includes three layered protocols (PA, PB, and PT) that allow for the exchange of attributes across the network. While these protocols are intended to be used together to fulfill a particular role in the model, they may offer overlapping functionality. For example, each protocol should be capable of protecting its information from attack (see section 8.2 for more information).

5.2.1. Posture Attribute Protocol (PA)

PA is a protocol that carries one or more attributes between Posture Collectors and their associated Posture Validator. The PA protocol is a message-oriented lightweight wrapper around a set of attributes being exchanged. This wrapper may indicate the purpose of attributes within the message. Some of the types of messages expected include: requests for posture information (Request Attributes), posture information about the endpoint (Posture Attributes), results of an assessment (Result Attributes), reusable compliance assertions (Assertion Attributes), and instructions to remediate non-compliant portions of the endpoint (Remediation Attributes). The PA protocol also provides the requisite encoding and cryptographic protection for the Posture Attributes.

5.2.2. Posture Broker Protocol (PB)

PB is a protocol that carries aggregate attribute messages between the Posture Collectors on the NEA Client and the corresponding Posture Validators on the NEA Server involved in a particular assessment. The PB protocol provides a session allowing for message dialogs for every assessment. This PB session is then used to bind multiple Posture Attribute requests and responses from the different Posture Collectors and Posture Validators involved in a particular assessment. The PB protocol may also carry the global assessment decision in the Result Attribute from the Posture Broker Server to the Posture Broker Client. PB may be used to carry additional types of messages for use by the Posture Broker Client and Server (e.g., information about user preferred interface settings such as language).

5.2.3. Posture Transport Protocol (PT)

PT is a transport protocol between the NEA Client and the NEA Server responsible for carrying the messages generated by the PB protocol. The PT protocol(s) transport(s) PB messages during the network connection request or after network connectivity has been established.

In scenarios where an initial assessment needs to occur during the network connection, the PT protocol (e.g., EAP within 802.1X) may have constrained use of the network, so deployments may choose to limit the amount and/or size of the attributes exchanged. The NEA Client and NEA Server should be able to detect when a potentially constrained situation exists prior to the assessment based upon properties of the underlying network protocol. Using this information, NEA policy could dictate what aspects of the endpoint to include in the initial assessment and potentially limit the PA message attributes exchanged. This could be followed up by a full reassessment after the endpoint is placed on the network. Alternatively, deployments can choose not to limit their assessment by configuring their network access technology to temporarily grant restricted IP connectivity prior to the assessment and use an unconstrained, high bandwidth IP-based transport during the assessment. Some of the constraints that may exist for protocols involved in the network connection phase include:

- o Limited maximum transmission unit (MTU) size and ability to negotiate larger MTUs,
- o Inability to perform multiple roundtrips,
- o Lack of support for piggybacking attributes for other protocols,

- o Low bandwidth or high latency limitations precluding exchanges of large amounts of data,
- o Inability of servers to initiate messages except during the network connection phase.

The PT protocol selection process needs to consider the impact of selecting a particular PT and set of underlying protocols on the deployment needs of PA and PB. PA and PB will be selected prior to PT so the needs of PA and PB will be known. Certain underlying protocol stacks may be too constrained to support adequate NEA assessments during network connection.

The PT protocol provides reliable message delivery, mutual authentication, and cryptographic protection for the PB messages as specified by local policy.

5.3. Attributes

The PA protocol is responsible for the exchange of attributes between a Posture Collector and Posture Validator. The PB protocol may also carry the global assessment decision attributes from the Posture Broker Server. Attributes are effectively the reserved word 'nouns' of the posture assessment. The NEA Server is only able to ask for information that has a corresponding attribute, thus bounding what type of posture can be obtained. The NEA WG will define a common (standard) set of attributes that are expected to be widely applicable to Posture Collectors and thus used for maximum interoperability, but Posture Collectors may support additional vendor-specific attributes when necessary.

Depending on the deployment scenario, the purpose of the attributes exchanged may be different (e.g., posture information vs. asserted compliance). This section discusses the originator and expected situation resulting in the use of each classification of attributes in a PA message. These classifications are not intended to dictate how the NEA WG will specify the attributes when defining the attribute namespace or schema.

5.3.1. Attributes Normally Sent by NEA Client:

- o Posture Attributes - Attributes and values sent to report information about a particular aspect (based on semantic of the attribute) of the system. These attributes are typically sent in response to Request Attributes from the NEA Server. For example, a set of Posture Attributes might describe the status of the host-based firewall (e.g., if running, vendor, version). The NEA Server would base its decision on comparing this type of attribute against policy.
- o Assertion Attributes - Attributes stating recent prior compliance to policy in hopes of avoiding the need to recollect the posture and send it to the NEA Server. Examples of assertions include (a) NEA Server provided attributes (state) describing a prior evaluation (e.g., opaque to endpoint, signed, time stamped items stating specific results) or (b) NEA Client identity information used by the NEA Server to locate state about prior decisions (e.g., system-bound cookie). These might be returned in lieu of, or in addition to, Posture Attributes.

5.3.2. Attributes Normally Sent by NEA Server:

- o Request Attributes - Attributes that define the specific posture information desired by the NEA Server. These attributes might effectively form a template that the Posture Collector fills in (subject to local policy restrictions) with the specific value corresponding to each attribute. The resulting attributes are typically Posture or Assertion Attributes from the NEA Client.
- o Result Attributes - Attributes that contain the decisions of the Posture Validators and/or Posture Broker Server. The level of detail provided may vary from which individual attributes were compliant or not through just the global assessment decision.
- o Remediation Attributes - Attributes that explain to the NEA Client and its user how to update the endpoint to become compliant with the NEA Server policies. These attributes are sent when the global assessment decision was that the endpoint is not currently compliant. Remediation and Result Attributes may both exist within a NEA Server attribute message.
- o Assertion Attributes - Attributes containing NEA Server assertions of compliance to a policy for future use by the NEA Client. See section 5.3.1 for more information.

6. Use Cases

This section discusses several of the NEA use cases with intent to describe and collectively bound the NEA problem space under consideration. The use cases provide a context and general rationale for the defined requirements. In order to ease understanding of each use case and how it maps to the reference model, each use case will be accompanied by a simple example and a discussion of how this example relates to the NEA protocols. It should be emphasized that the provided examples are not intended to indicate the only approach to addressing the use case but rather are included to ease understanding of how the flows might occur and impact the NEA protocols.

We broadly classify the use cases into two categories, each with its own set of trigger events:

- o Initial assessment - evaluation of the posture of an endpoint that has not recently been assessed and thus is not in possession of any valid proof that it should be considered compliant. This evaluation might be triggered by a request to join a network, a request to use a service, or a desire to understand the posture of a system.
- o Reassessment - evaluation of the posture of an endpoint that has previously been assessed. This evaluation could occur for a variety of reasons including the NEA Client or Server recognizing an occurrence affecting the endpoint that might raise the endpoint's risk level. This could be as simple as it having been a long time since the endpoint's prior reassessment.

6.1. Initial Assessment

An initial assessment occurs when a NEA Client or Server event occurs that causes the evaluation of the posture of the endpoint for the first time. Endpoints do not qualify for this category of use case if they have been recently assessed and the NEA Client or Server has maintained state (or proof) that the endpoint is compliant and therefore does not need to have its posture evaluated again.

6.1.1. Triggered by Network Connection or Service Request

This use case focuses on assessments performed at the time an endpoint attempts to join a network or request use of a service that requires a posture evaluation. This use case is particularly interesting because it allows the NEA Server to evaluate the posture of an endpoint before allowing it access to the network or service.

This approach could be used to help detect endpoints with known vulnerabilities and facilitate their repair before they are admitted to the network and potentially exposed to threats on the network.

A variety of types of endpoint actions could result in this class of assessment. For example, an assessment could be triggered by the endpoint trying to access a highly protected network service (e.g., financial or HR application server) where heightened security checking is required. A better known example could include requesting entrance to a network that requires systems to meet compliance policy. This example is discussed in more detail in the following section.

6.1.1.1. Example

An IT employee returning from vacation boots his office desktop computer that generates a request to join the wired enterprise network. The network's security policy requires the system to provide posture information in order to determine whether the desktop's security features are enabled and up to date. The desktop sends its patch, firewall, and anti-virus posture information. The NEA Server determines that the system is lacking a recent security patch designed to fix a serious vulnerability and the system is placed on a restricted access network. The desktop follows the provided remediation instructions to download and install the necessary patch. Later, the desktop requests again to join the network and this time is provided full access to the enterprise network after a full assessment.

6.1.1.2. Possible Flows and Protocol Usage

The following describes typical message flows through the NEA reference model for this example use case:

1. The IT employee's desktop computer connects to the network through an access gateway in the wired enterprise network.
2. The Posture Broker Server on the NEA Server is instructed to assess the endpoint joining the wired network.
3. Based upon compliance policy, the Posture Broker Server contacts the operating system patch, host-based firewall, and anti-virus Posture Validators to request the necessary posture. Each Posture Validator creates a PA message containing the desired attributes to be requested for assessment from the desktop system.

4. The Posture Broker Server aggregates the PA messages from the Posture Validators into a PB message. The Posture Broker Server passes the PB message to the Posture Transport Server that uses the PT protocol to send the PB message to the NEA Client on the desktop computer.
5. The Posture Transport Client receives the message from the NEA Server and passes it to the Posture Broker Client for message delivery.
6. The Posture Broker Client de-multiplexes the PB message and delivers the PA messages with the requests for attributes to the firewall, operating system patch, and anti-virus Posture Collectors.
7. Each Posture Collector involved consults local privacy policy to determine what information is allowed to be disclosed and then returns the requested attributes that are authorized in a PA message to the Posture Broker Client.
8. The Posture Broker Client aggregates these PA messages into a single PB message and sends it to the Posture Broker Server using the Posture Transport Client to Server session.
9. The Posture Transport Server provides the PB message to the Posture Broker Server that de-multiplexes the message and sends the appropriate attributes to the corresponding Posture Validator.
10. Each Posture Validator compares the values of the attributes it receives with the expected values defined in its policy.
11. The anti-virus and firewall Posture Validators return attributes to the Posture Broker Server stating the desktop computer is compliant, but the operating system patch Posture Validator returns non-compliant. The operating system patch Posture Validator creates a PA message that contains attributes with remediation instructions in addition to the attribute indicating non-compliance result.
12. The Posture Broker Server aggregates the PA messages and sends them in a PB message to the Posture Broker Client via the Posture Transport Server and Posture Transport Client.

13. The Posture Broker Client delivers the PA messages with the results from the various Posture Validators to the Posture Collectors including the PA message containing attributes with remediation instructions to the operating system patch Posture Collector. This Posture Collector then interacts with the user to download and install the needed patches, potentially while the endpoint remains quarantined.
14. Upon completion of the remediation, the above steps 1-10 are repeated (triggered by the NEA Client repeating its request to join the network).
15. This time each involved Posture Validator (including the operating system patch Posture Validator) returns a compliant status and the Posture Broker Server returns a compliant result indicating a global success.
16. The Posture Broker Client receives the compliant result and the IT employee's desktop is now on the network.

6.1.1.3. Impact on Requirements

The following are several different aspects of the use case example that potentially need to be factored into the requirements.

- o Posture assessment before endpoint allowed on network
- o Endpoint sends attributes containing posture information
- o NEA Server sends remediation instructions
- o NEA Client causes a reassessment after remediation

6.1.2. Triggered by Endpoint

This use case highlights that an endpoint (possibly at the request of a user) may wish to trigger an assessment of its posture to determine whether its security protective mechanisms are running and up to date.

6.1.2.1. Example

A student goes to the terminal room to work on a project. The terminal room contains shared systems owned by the school that are on the network. These systems have been previously used by other students so their security posture is unknown. The student wishes to check whether a system is currently in compliance with the school's security policies prior to doing work, so she requests a posture

assessment. The NEA Server performs an initial assessment of the system and determines it is compliant but the anti-virus protection is not in use. The student receives an advisory response indicating the system's anti-virus software is turned off but that otherwise it complies with the school's policy. The student turns on the anti-virus software, initiates a scan, and upon completion decides to trust the system with her work.

6.1.2.2. Possible Flows and Protocol Usage

The following describes the message flows through the NEA reference model for the student using a terminal room shared system example:

1. Student triggers the Posture Broker Client on the computer system in the terminal room to initiate a posture assessment.
2. The Posture Broker Client establishes a session with the Posture Broker Server that causes an assessment to be triggered.
3. The Posture Broker Server detects the new session and consults policy to determine that Posture Validators to involve in the assessment. The Posture Broker Server decides to employ several Posture Validators including the anti-virus Posture Validator.
4. The Posture Validators involved create PA messages containing requests for particular attributes containing information about the desired terminal room computer based on the school's security policy.
5. The Posture Broker Server assembles a PB message including each of the PA messages from the Posture Validators.
6. The Posture Transport Server sends the PB message to the Posture Transport Client where it is passed on to the Posture Broker Client.
7. The Posture Broker Client on the student's computer de-multiplexes the PA messages and delivers them to the corresponding Posture Collectors.
8. The Posture Collectors consult privacy policy to decide what information to share with the Server. If allowable, the Collectors each return a PA message containing the requested posture to the Posture Broker Client.

9. The Posture Broker Client aggregates the returned PA messages into a PB message and hands it to the Posture Transport Client for transmission to the Posture Transport Server.
10. The Posture Broker Server separates and distributes the Posture Collector PA messages to the associated Posture Validators.
11. The Posture Validators determine whether the attributes containing the posture included in the PA message are compliant with their policies and returns a posture assessment decision to the Posture Broker Server. In this case, the anti-virus Posture Validator returns a PA message indicating a non-compliant result because the anti-virus software is not running and includes attributes describing how to activate the software.
12. The Posture Broker Server determines the overall compliance decision based on all of the Validators' assessment results and sends a PB message containing an attribute expressing the global assessment decision and the anti-virus Validator's PA message. In this case, the global assessment decision indicates the system is compliant (despite the anti-virus Validator's result) because the Posture Broker Server policy allowed for the anti-virus to not be running as long as the system was properly patched and running a firewall (which was the case according to the other Posture Validators).
13. The Posture Transport Server sends the PB message to the Posture Transport Client that provides the message to the Posture Broker Client.
14. The Posture Broker Client on the terminal room computer examines the PB message's global assessment decision attribute and reports to the student that the system was deemed to be compliant, but that an advisory was included.
15. The Posture Broker Client provides the PA message with the remediation attributes to the anti-virus Posture Collector that interacts with the user to explain how to turn on anti-virus to improve the local protections.
16. The student turns on the anti-virus software and on completion steps 1-10 are repeated.
17. This time the anti-virus Posture Validator returns a success status and the Posture Broker Server returns a successful global assessment decision in the PB message.

18. The Posture Broker Client receives the successful global assessment decision in the PB message and the student now uses the computer for her assignment.

6.1.2.3. Impact on Requirements

The following are several different aspects of the use case example that potentially need to be factored into the requirements.

- o Voluntary endpoint requested initial assessment,
- o Successful (compliant) global assessment decision included in PB message with a PA message containing an advisory set of attributes for remediation.

6.2. Posture Reassessment

Reassessment(s) of endpoints can happen anytime after being admitted to the network after a successful initial NEA assessment. These reassessments may be event-based, such as driven by posture changes detected by the NEA Client, or changes detected by network infrastructure such as detection of suspicious behavior or network policy updates on the NEA Server. They may also be periodic (timer-driven) to reassess the health of the endpoint.

6.2.1. Triggered by NEA Client

This use case allows for software on the endpoint or a user to determine that a reassessment of the system is required. There are a variety of reasons why such a reassessment might be beneficial including: changes in its previously reported posture, detection of potentially suspicious behavior, or even to enable the system to periodically poll the NEA Server to assess its condition relative to the latest policies.

6.2.1.1. Example

The desktops within a company's HR department have a history of poor security practices and eventual compromise. The HR department administrator decides to deploy software on each desktop to monitor the use of security protective mechanisms to assure their use. One day, an HR person accidentally turns off the desktop firewall. The monitoring process detects the lack of a firewall and contacts the NEA Server to request a reassessment of the firewall compliance. The NEA Server returns a decision that the firewall must be reactivated to stay on the network. The NEA Client explains the decision to the user and how to reactivate the firewall. The HR person restarts the firewall and initiates a request to rejoin the network.

6.2.1.2. Possible Flows & Protocol Usage

The following describes the message flows through the NEA reference model for the HR department example:

1. The desktop monitoring software that typically might act as a Posture Collector triggers the Posture Broker Client to initiate a posture reassessment. The Posture Broker Client creates a PB message that contains a PA message indicating the desktop firewall has been disabled.
2. The Posture Broker Client sends the PB message to the Posture Broker Server.
3. The Posture Transport Client sends the PB message to the Posture Transport Server over the PT protocol.
4. The Posture Broker Server receives the PB message and forwards the PA message to the firewall Posture Validator for evaluation.
5. The firewall Posture Validator determines that the endpoint is no longer compliant because its firewall has been disabled.
6. The Posture Validator generates a PA message that contains attributes indicating a non-compliant posture assessment result and remediation instructions for how to reactivate the firewall.
7. The Posture Validator communicates the PA message with the posture assessment result to the Posture Broker Server to respond back to the NEA Client.
8. The Posture Broker Server generates a PB message including a global assessment decision of non-compliant and the PA message from the firewall Posture Validator.
9. The Posture Transport Server transports the PB message to the Posture Transport Client where it is passed to the Posture Broker Client.
10. The Posture Broker Client processes the attribute containing the global assessment decision received from the NEA Server and displays the non-compliance messages to the user.

11. The Posture Broker Client forwards the PA message to the firewall Posture Collector; the Posture Collector displays the remediation instructions for how to enable the desktop firewall.
12. The user is prompted to initiate a reassessment after completing the remediation.
13. Upon completion of the remediation, the NEA Client reinitiates a request for reassessment and steps 1-4 are repeated. This time the firewall Posture Validator determines the endpoint is compliant and returns a successful posture assessment decision.
14. The Posture Broker Server generates a PB message with a global assessment decision of compliant and returns this to the NEA Client.

6.2.1.3. Impact on Requirements

The following are several different aspects of the use case example that potentially need to be factored into the requirements.

- o Voluntary, endpoint (software) initiated posture reassessment request
- o NEA Server requests specific firewall-oriented Posture Attributes
- o NEA Client (firewall Posture Collector) interacts with user to remediate problem

6.2.2. Triggered by NEA Server

In many cases, especially for reassessment, the NEA Server may initiate specific or complete reassessment of one or more endpoints triggered by:

- o Time (periodic)
- o Event occurrence
- o Policy updates

6.2.2.1. Example

An enterprise requires employees on the network to always stay up to date with security critical operating system patches. A marketing employee joins the network and performs an initial assessment. The assessment determines the employee's laptop is compliant. Several

hours later, a major operating system vendor releases a set of patches preventing a serious vulnerability that is being exploited on the Internet.

The enterprise administrators make available the patches and change the network policy to require them to be installed by 5 PM. This policy change causes the NEA Server to request a reassessment to determine which endpoints are impacted and lacking the patches. The marketing employee's laptop is reassessed and determined to need the patches. A remediation advisory is sent and presented to the employee explaining how to obtain the patches and that they must be installed by 5 PM. The marketing employee immediately downloads and installs the patches and obtains an assertion that all patches are now installed.

At 5 PM, the enterprise performs another reassessment of all impacted endpoints to determine if they are now in compliance. The marketing employee's laptop is reassessed and presents the assertion that it has the patches installed and thus is determined to be compliant.

6.2.2.2. Possible Flows and Protocol Usage

The following describes the message flows through the NEA reference model for the above example:

1. Marketing employee joins network and completes an initial assessment resulting in a compliant decision.
2. The Enterprise Administrator configures an operating system patch policy indicating that recent patches are required on all endpoints by 5 PM to prevent serious vulnerabilities.
3. The NEA Server's operating system patch Posture Validator becomes aware of this policy change and creates a PA message requesting attributes describing OS patches in use and triggers the Posture Broker Server to initiate a posture reassessment of all endpoints connected to the network.
4. The Posture Broker creates a PB message that includes the PA message from the operating system patch Posture Validator.
5. The Posture Broker Server gradually establishes a session with each available NEA Client.
6. The Posture Broker Server sends the PB message to the Posture Broker Client.

7. The Posture Transport Server carries the PB message to the Posture Transport Client over the PT protocol.
8. The Posture Broker Client receives the PB message and forwards the PA message to the operating system patch Posture Collector.
9. The operating system patch Posture Collector determines the OS patches present on the endpoint and if authorized by its disclosure policy creates a PA message containing the patch information attributes.
10. The Posture Broker Client sends a PB message that includes the operating system patch PA message.
11. The Posture Transport Client transports the PB message to the Posture Transport Server where it is passed to the Posture Broker Server.
12. The Posture Broker Server receives the PB message and delivers the PA message to the operating system patch Posture Validator.
13. The operating system patch Posture Validator extracts the attributes describing the current OS patches from the PA message and uses the values to determine whether the endpoint is compliant with the new policy. The Posture Validator determines that the endpoint is not compliant since it does not have the new OS patches installed.
14. The Posture Validator generates a PA message that includes attributes stating the posture assessment decision is non-compliant and attributes containing the remediation instructions to enable the endpoint to download the required OS patches.
15. The Posture Validator communicates the posture assessment result to the Posture Broker Server along with its PA message.
16. The Posture Broker Server generates a global assessment decision and sends a PB message with the decision and the operating system patch Posture Validator's PA message.
17. The Posture Transport Server transports the PB message to the Posture Transport Client where it is passed to the Posture Broker Client.
18. The Posture Broker Client processes the Result Attribute received from the NEA Server and displays the non-compliance decision to the user.

19. The Posture Broker Client forwards the PA message containing the remediation instructions to the operating system patch Posture Collector; the Posture Collector guides the user with instructions on how to become compliant that include downloading the appropriate OS patches to prevent the vulnerability.
20. The marketing employee installs the required patches and now is in compliance.
21. The NEA Client triggers a reassessment of the operating system patches that causes a repeat of many of the steps above. This time, in step 13 the operating system patch Posture Validator determines the marketing employee's laptop is compliant. It returns a reusable (e.g., signed and dated) set of attributes that assert OS patch compliance to the latest policy. These OS patch compliance assertions can be used in a future PA message from the operating system patch Collector instead of determining and providing the specific patch set posture as before.
22. This time when the operating system patch Posture Collector receives the PA message that contains reusable attributes asserting compliance, it caches those attributes for future use.
23. Later at 5 PM, the NEA Server triggers a gradual reassessment to determine compliance to the patch advisory. When the operating system patch Posture Collector receives the request for posture information (like in step 9 above) it returns the cached set of assertions (instead of specific OS patch information) to indicate that the patches have been installed instead of determining all the patches that have been installed on the system.
24. When the operating system patch Posture Validator receives the PA message containing the assertions, it is able to determine that they are authentic and acceptable assertions instead of specific posture. It returns a posture assessment decision of compliant thus allowing the laptop to remain on the network.

6.2.2.3. Impact on Requirements

The following are several different aspects of the use case example that potentially need to be factored into the requirements.

- o Server-initiated reassessment required due to urgent patch availability

- o NEA Client submits reusable assertion attributes instead of posture that patch is installed
- o NEA Server capable of recognizing previously issued assertion attributes are sufficient instead of posture

7. Requirements

This section describes the requirements that will be used by the NEA WG to assess and compare candidate protocols for PA, PB, and PT. These requirements frequently express features that a candidate protocol must be capable of offering so that a deployer can decide whether to make use of that feature. This section does not state requirements about what features of each protocol must be used during a deployment.

For example, a requirement (MUST, SHOULD, or MAY) might exist for cryptographic security protections to be available from each protocol but this does not require that a deployer make use of all or even any of them should they deem their environment to offer other protections that are sufficient.

7.1. Common Protocol Requirements

The following are the common requirements that apply to the PA, PB, and PT protocols in the NEA reference model:

- C-1 NEA protocols MUST support multiple round trips between the NEA Client and NEA Server in a single assessment.
- C-2 NEA protocols SHOULD provide a way for both the NEA Client and the NEA Server to initiate a posture assessment or reassessment as needed.
- C-3 NEA protocols including security capabilities MUST be capable of protecting against active and passive attacks by intermediaries and endpoints including prevention from replay based attacks.
- C-4 The PA and PB protocols MUST be capable of operating over any PT protocol. For example, the PB protocol must provide a transport independent interface allowing the PA protocol to operate without change across a variety of network protocol environments (e.g., EAP/802.1X, TLS, and Internet Key Exchange Protocol version 2 (IKEv2)).

- C-5 The selection process for NEA protocols MUST evaluate and prefer the reuse of existing open standards that meet the requirements before defining new ones. The goal of NEA is not to create additional alternative protocols where acceptable solutions already exist.
- C-6 NEA protocols MUST be highly scalable; the protocols MUST support many Posture Collectors on a large number of NEA Clients to be assessed by numerous Posture Validators residing on multiple NEA Servers.
- C-7 The protocols MUST support efficient transport of a large number of attribute messages between the NEA Client and the NEA Server.
- C-8 NEA protocols MUST operate efficiently over low bandwidth or high latency links.
- C-9 For any strings intended for display to a user, the protocols MUST support adapting these strings to the user's language preferences.
- C-10 NEA protocols MUST support encoding of strings in UTF-8 format [UTF8].
- C-11 Due to the potentially different transport characteristics provided by the underlying candidate PT protocols, the NEA Client and NEA Server MUST be capable of becoming aware of and adapting to the limitations of the available PT protocol. For example, some PT protocol characteristics that might impact the operation of PA and PB include restrictions on: which end can initiate a NEA connection, maximum data size in a message or full assessment, upper bound on number of roundtrips, and ordering (duplex) of messages exchanged. The selection process for the PT protocols MUST consider the limitations the candidate PT protocol would impose upon the PA and PB protocols.

7.2. Posture Attribute (PA) Protocol Requirements

The Posture Attribute (PA) protocol defines the transport and data model to carry posture and validation information between a particular Posture Collector associated with the NEA Client and a Posture Validator associated with a NEA Server. The PA protocol carries collections of standard attributes and vendor-specific attributes. The PA protocol itself is carried inside the PB protocol.

The following requirements define the desired properties that form the basis for comparison and evaluation of candidate PA protocols. These requirements do not mandate the use of these properties, but merely that the candidate protocols are capable of offering the property if it should be needed.

PA-1 The PA protocol MUST support communication of an extensible set of NEA standards defined attributes. These attributes will be distinguishable from non-standard attributes.

PA-2 The PA protocol MUST support communication of an extensible set of vendor-specific attributes. These attributes will be segmented into uniquely identified vendor-specific namespaces.

PA-3 The PA protocol MUST enable a Posture Validator to make one or more requests for attributes from a Posture Collector within a single assessment. This enables the Posture Validator to reassess the posture of a particular endpoint feature or to request additional posture including from other parts of the endpoint.

PA-4 The PA protocol MUST be capable of returning attributes from a Posture Validator to a Posture Collector. For example, this might enable the Posture Collector to learn the specific reason for a failed assessment and to aid in remediation and notification of the system owner.

PA-5 The PA protocol SHOULD provide authentication, integrity, and confidentiality protection for attributes communicated between a Posture Collector and Posture Validator. This enables end-to-end security across a NEA deployment that might involve traversal of several systems or trust boundaries.

PA-6 The PA protocol MUST be capable of carrying attributes that contain non-binary and binary data including encrypted content.

7.3. Posture Broker (PB) Protocol Requirements

The PB protocol supports multiplexing of Posture Attribute messages (based on PA protocol) between the Posture Collectors on the NEA Client to and from the Posture Validators on the NEA Server (in either direction).

The PB protocol carries the global assessment decision made by the Posture Broker Server, taking into account the results of the Posture Validators involved in the assessment, to the Posture Broker Client.

The PB protocol also aggregates and transports advisories and notifications such as remediation instructions (e.g., patch references) from one or more Posture Validators.

The requirements for the PB protocol are:

- PB-1 The PB protocol MUST be capable of carrying attributes from the Posture Broker Server to the Posture Broker Client. This enables the Posture Broker Client to learn the posture assessment decision and if appropriate to aid in remediation and notification of the endpoint owner.
- PB-2 The PB protocol MUST NOT interpret the contents of PA messages being carried, i.e., the data it is carrying must be opaque to it.
- PB-3 The PB protocol MUST carry unique identifiers that are used by the Posture Brokers to route (deliver) PA messages between Posture Collectors and Posture Validators. Such message routing should facilitate dynamic registration or deregistration of Posture Collectors and Validators. For example, a dynamically registered anti-virus Posture Validator should be able to subscribe to receive messages from its respective anti-virus Posture Collector on NEA Clients.
- PB-4 The PB protocol MUST be capable of supporting a half-duplex PT protocol. However this does not preclude PB from operating full-duplex when running over a full-duplex PT.
- PB-5 The PB protocol MAY support authentication, integrity and confidentiality protection for the attribute messages it carries between a Posture Broker Client and Posture Broker Server. This provides security protection for a message dialog of the groupings of attribute messages exchanged between the Posture Broker Client and Posture Broker Server. Such protection is orthogonal to PA protections (which are end to end) and allows for simpler Posture Collector and Validators to be implemented, and for consolidation of cryptographic operations possibly improving scalability and manageability.
- PB-6 The PB protocol MUST support grouping of attribute messages optimize transport of messages and minimize round trips.

7.4. Posture Transport (PT) Protocol Requirements

The Posture Transport (PT) protocol carries PB protocol messages between the Posture Transport Client and the Posture Transport Server. PT is responsible for providing a protected transport for the PB protocol. The PT protocol may itself be transported by one or more concatenated sessions using lower layer protocols, such as 802.1X, RADIUS [RADIUS], TLS, or IKE.

This section defines the requirements that candidate PT protocols must be capable of supporting.

PT-1 The PT protocol MUST NOT interpret the contents of PB messages being transported, i.e., the data it is carrying must be opaque to it.

PT-2 The PT protocol MUST be capable of supporting mutual authentication, integrity, confidentiality, and replay protection of the PB messages between the Posture Transport Client and the Posture Transport Server.

PT-3 The PT protocol MUST provide reliable delivery for the PB protocol. This includes the ability to perform fragmentation and reassembly, detect duplicates, and reorder to provide in-sequence delivery, as required.

PT-4 The PT protocol SHOULD be able to run over existing network access protocols such as 802.1X and IKEv2.

PT-5 The PT protocol SHOULD be able to run between a NEA Client and NEA Server over TCP or UDP (similar to Lightweight Directory Access Protocol (LDAP)).

8. Security Considerations

This document defines the functional requirements for the PA, PB, and PT protocols used for Network Endpoint Assessment. As such, it does not define a specific protocol stack or set of technologies, so this section will highlight security issues that may apply to NEA in general or to particular aspects of the NEA reference model.

Note that while a number of topics are outside the scope of the NEA WG and thus this specification (see section 3.1), it is important that those mechanisms are protected from attack. For example, the methods of triggering an assessment or reassessment are out of scope but should be appropriately protected from attack (e.g., an attacker hiding the event indicating a NEA Server policy change has occurred).

NEA intends to facilitate detection and corrective actions for cooperating endpoints to become compliant with network compliance policies. For example, it is envisioned that these policies will allow deployers to detect out-of-date, inactive, or absent security mechanisms on the endpoint that might leave it more vulnerable to known attacks. If an endpoint is more vulnerable to compromise, then it is riskier to have this endpoint present on the network with other valuable assets. By proactively assessing cooperating endpoints before their entrance to the network, deployers can improve their resilience to attack prior to network access. Similarly, reassessments of cooperating endpoints on the network may be helpful in assuring that security mechanisms remain in use and are up to date with the latest policies.

NEA fully recognizes that not all endpoints will be cooperating by providing their valid posture (or any posture at all). This might occur if malware is influencing the NEA Client or policies, and thus a trustworthy assessment isn't possible. Such a situation could result in the admission of an endpoint that introduces threats to the network and other endpoints despite passing the NEA compliance assessment.

8.1. Trust

Network Endpoint Assessment involves assessing the posture of endpoints entering or already on the network against compliance policies to assure they are adequately protected. Therefore, there must be an implied distrusting of endpoints until there is reason to believe (based on posture information) that they are protected from threats addressed by compliance policy and can be trusted to not propagate those threats to other endpoints. On the network provider side, the NEA Client normally is expected to trust the network infrastructure systems to not misuse any disclosed posture information (see section 9) and any remediation instructions provided to the endpoint. The NEA Client normally also needs to trust that the NEA Server will only request information required to determine whether the endpoint is safe to access the network assets.

Between the NEA Client and Server there exists a network that is not assumed to be trustworthy. Therefore, little about the network is implicitly trusted beyond its willingness and ability to transport the exchanged messages in a timely manner. The amount of trust given to each component of the NEA reference model is deployment specific. The NEA WG intends to provide security mechanisms to reduce the amount of trust that must be assumed by a deployer. The following sections will discuss each area in more detail.

8.1.1. Endpoint

For NEA to properly operate, the endpoint needs to be trusted to accurately represent the requested security posture of the endpoint to the NEA Server. By NEA WG charter, the NEA reference model does not explicitly specify how to detect or prevent lying endpoints that intentionally misrepresent their posture. Similarly, the detection of malware (e.g., root kits) that are able to trick the Posture Collectors into returning incorrect information is the subject for research and standardization outside the IETF (e.g., Trusted Computing Group [TCG]) and is not specifically addressed by the model. However, if such mechanisms are used in a deployment, the NEA reference model should be able to accommodate these technologies by allowing them to communicate over PA to Posture Validators or work orthogonally to protect the NEA Client from attack and assure the ability of Posture Collectors to view the actual posture.

Besides having to trust the integrity of the NEA Client and its ability to accurately collect and report Posture Attributes about the endpoint, we try to limit other assumed trust. Most of the usage models for NEA expect the posture information to be sent to the NEA Server for evaluation and decision making. When PA and/or PT level security protections are used, the endpoint needs to trust the integrity and potentially confidentiality of the trust anchor information (e.g., public key certificates) used by the Posture Collector and/or Posture Transport Client. However, NEA implementations may choose to send or pre-provision some policies to the endpoint for evaluation that would assume more trust in the endpoint. In this case, the NEA Server must trust the endpoint's policy storage, evaluation, and reporting mechanisms to not falsify the results of the posture evaluation.

Generally the endpoint should not trust network communications (e.g., inbound connection requests) unless this trust has been specifically authorized by the user or owner defined policy or action. The NEA reference model assumes the entire NEA Client is local to the endpoint. Unsolicited communications originating from the network should be inspected by normal host-based security protective mechanisms (e.g., firewalls, security protocols, Intrusion Detection/Prevention System (IDS/IPS), etc.). Communications associated with a NEA assessment or reassessment requires some level of trust particularly when initiated by the NEA Server (reassessment). The degree of trust can be limited by use of strong security protections on the messages as dictated by the network deployer and the endpoint user/owner policy.

8.1.2. Network Communications

Between the NEA Client and Server, there may exist a variety of types of devices to facilitate the communication path. Some of the devices may serve as intermediaries (e.g., simple L2 switches) so they may have the opportunity to observe and change the message dialogs.

The intermediary devices may fall into a few major categories that impact our degree of trust in their operation. First, some intermediary devices may act as message forwarders or carriers for PT (e.g., L2 switches, L3 routers). For these devices we trust them not to drop the messages or actively attempt to disrupt (e.g., denial of service (DoS)) the NEA deployment.

Second, some intermediary devices may be part of the access control layer of the network and as such, we trust them to enforce policies including remediation, isolation, and access controls given to them as a result on a NEA assessment. These devices may also fill other types of roles described in this section.

Third, some devices may act as a termination point or proxy for the PT carrier protocol. Frequently, it is expected that the carrier protocol for PT will terminate on the NEA Client and Server so will be co-resident with the PT endpoints. If this expectation is not present in a deployment, we must trust the termination device to accurately proxy the PT messages without alteration into the next carrier protocol (e.g., if inner EAP method messages are transitioned from an EAP [EAP] tunnel to a RADIUS session).

Fourth, many networks include infrastructure such as IDS/IPS devices that monitor and take corrective action when suspicious behavior is observed on the network. These devices may have a relationship with the NEA Server that is not within scope for this specification. Devices trusted by the NEA Server to provide security information that might affect the NEA Server's decisions are trusted to operate properly and not cause the NEA Server to make incorrect decisions.

Finally, other types of intermediary devices may exist on the network between the NEA Client and Server that are present to service other network functions beside NEA. These devices might be capable of passively eavesdropping on the network, archiving information for future purposes (e.g., replay or privacy invasion), or more actively attacking the NEA protocols. Because these devices do not play a role in facilitating NEA, it is essential that NEA deployers not be forced to trust them for NEA to reliably operate. Therefore, it is required that NEA protocols offer security protections to assure these devices can't steal, alter, spoof or otherwise damage the reliability of the message dialogs.

8.1.3. NEA Server

The NEA Server (including potentially remote systems providing posture validation services) is generally trusted to apply the specified assessment policies and must be protected from compromise. It is essential that NEA Server deployments properly safeguard these systems from a variety of attacks from the network and endpoints to assure their proper operation.

While there is a need to trust the NEA Server operation to some degree, rigorous security architecture, analysis, monitoring, and review should assure its network footprint and internal workings are protected from attack. The network footprint would include communications over the network that might be subject to attack such as policy provisioning from the policy authoring systems and general security and system management protocols. Some examples of internal workings include protections from malware attacking the intra-NEA Server communications, NEA Server internal logic, or policy stores (particularly those that would change the resulting decisions or enforcements). The NEA Server needs to trust the underlying NEA and lower layer network protocols to properly behave and safeguard the exchanged messages with the endpoint. The NEA reference model does not attempt to address integrity protection of the operating system or other software supporting the NEA Server.

One interesting example is where some components of the NEA Server physically reside in different systems. This might occur when a Posture Validator (or a remote backend server used by a local Posture Validator) exists on another system from the Posture Broker Server. Similarly, the Posture Broker Server might exist on a separate system from the Posture Transport Server. When there is a physical separation, the communications between the remote components of the NEA Server must ensure that the PB session and PA message dialogs are resistant to active and passive attacks, in particular, guarded against eavesdropping, forgery and replay. Similarly, the Posture Validators may also wish to minimize their trust in the Posture Broker Server beyond its ability to properly send and deliver PA messages. The Posture Validators could employ end-to-end PA security to verify the authenticity and protect the integrity and/or confidentiality of the PA messages exchanged.

When PA security is used, each Posture Validator must be able to trust the integrity and potentially confidentiality of its trust anchor policies.

8.2. Protection Mechanisms at Multiple Layers

Inherent in the requirements is a desire for NEA candidate protocols throughout the reference model to be capable of providing strong security mechanisms as dictated by the particular deployment. In some cases, these mechanisms may appear to provide overlapping or redundant protections. These apparent overlaps may be used in combination to offer a defense in depth approach to security. However, because of the layering of the protocols, each set of protections offers slightly different benefits and levels of granularity.

For example, a deployer may wish to encrypt traffic at the PT layer to protect against some forms of traffic analysis or interception by an eavesdropper. Additionally, the deployer may also selectively encrypt messages containing the posture of an endpoint to achieve end-to-end confidentiality to its corresponding Posture Validator. In particular, this might be desired when the Posture Validator is not co-located with the NEA Server so the information will traverse additional network segments after the PT protections have been enforced or so that the Posture Validator can authenticate the corresponding Posture Collector (or vice versa).

Different use cases and environments for the NEA technologies will likely influence the selection of the strength and security mechanisms employed during an assessment. The goal of the NEA requirements is to encourage the selection of technologies and protocols that are capable of providing the necessary protections for a wide variety of types of assessment.

8.3. Relevant Classes of Attack

A variety of attacks are possible against the NEA protocols and assessment technologies. This section does not include a full security analysis, but wishes to highlight a few attacks that influenced the requirement definition and should be considered by deployers selecting use of protective mechanisms within the NEA reference model.

As discussed, there are a variety of protective mechanisms included in the requirements for candidate NEA protocols. Different use cases and environments may cause deployers to decide not to use some of these mechanisms; however, this should be done with an understanding that the deployment may become vulnerable to some classes of attack. As always, a balance of risk vs. performance, usability, manageability, and other factors should be taken into account.

The following types of attacks are applicable to network protocols defined in the reference model and thus should be considered by deployers.

8.3.1. Man-in-the-Middle (MITM)

MITM attacks against a network protocol exist when a third party can insert itself between two communicating entities without detection and gain benefit from involvement in their message dialog. For example, a malware infested system might wish to join the network replaying posture observed from a clean endpoint entering the network. This might occur by the system inserting itself into and actively proxying an assessment message dialog. The impact of the damage caused by the MITM can be limited or prevented by selection of appropriate protocol protective mechanisms.

For example, the requirement for PT to be capable of supporting mutual authentication prior to any endpoint assessment message dialogs prevents the attacker from inserting itself as an active participant (proxy) within the communications without detection (assuming the attacker lacks credentials convincing either party it is legitimate). Reusable credentials should not be exposed on the network to assure the MITM doesn't have a way to impersonate either party. The PT requirement for confidentiality-protected (encrypted) communications linked to the above authentication prevents a passive MITM from eavesdropping by observing the message dialog and keeping a record of the conformant posture values for future use. The PT requirement for replay prevention stops a passive MITM from later establishing a new session (or hijacking an existing session) and replaying previously observed message dialogs.

If a non-compliant, active MITM is able to trick a clean endpoint to give up its posture information, and the MITM has legitimate credentials, it might be able to appear to a NEA Server as having compliant posture when it does not. For example, a non-compliant MITM could connect and authenticate to a NEA Server and as the NEA Server requests posture information, the MITM could request the same posture from the clean endpoint. If the clean endpoint trusts the MITM to perform a reassessment and is willing to share the requested posture, the MITM could obtain the needed posture from the clean endpoint and send it to the NEA Server. In order to address this form of MITM attack, the NEA protocols would need to offer a strong (cryptographic) binding between the posture information and the authenticated session to the NEA Server so the NEA Server knows the posture originated from the endpoint that authenticated. Such a strong binding between the posture's origin and the authenticating endpoint may be feasible so should be preferred by the NEA WG.

8.3.2. Message Modification

Without message integrity protection, an attacker capable of intercepting a message might be capable of modifying its contents and causing an incorrect decision to be made. For example, the attacker might change the Posture Attributes to always reflect incorrect values and thus prevent a compliant system from joining the network. Unless the NEA Server could detect this change, the attacker could prevent admission to large numbers of clean systems. Conversely, the attacker could allow a malware infested machine to be admitted by changing the sent Posture Attributes to reflect compliant values, thus hiding the malware from the Posture Validator. The attacker could also infect compliant endpoints by sending malicious remediation instructions that, when performed, would introduce malware on the endpoint or deactivate security mechanisms.

In order to protect against such attacks, the PT includes a requirement for strong integrity protection (e.g., including a protected hash like a Hashed Message Authentication Code (HMAC) [HMAC] of the message) so any change to a message would be detected. PA includes a similar requirement to enable end-to-end integrity protection of the attributes, extending the protection all the way to the Posture Validator even if it is located on another system behind the NEA Server.

It is important that integrity protection schemes leverage fresh secret information (not known by the attacker) that is bound to the authenticated session such as an HMAC using a derived fresh secret associated with the session. Inclusion of freshness information allows the parties to protect against some forms of message replay attacks using secret information from prior sessions.

8.3.3. Message Replay or Attribute Theft

An attacker might listen to the network, recording message dialogs or attributes from a compliant endpoint for later reuse to the same NEA Server or just to build an inventory of software running on other systems watching for known vulnerabilities. The NEA Server needs to be capable of detecting the replay of posture and/or the model must assure that the eavesdropper cannot obtain the information in the first place. For this reason, the PT protocol is required to provide confidentiality and replay prevention.

The cryptographic protection from disclosure of the PT, PB, or PA messages prevents the passive listener from observing the exchanged messages and thus prevents theft of the information for future use. However, an active attacker might be able to replay the encrypted message if there is no strong link to the originating party or

session. By linking the encrypted message dialog to the authentication event and leveraging per-transaction freshness and keying exchanges, this prevents a replay of the encrypted transaction.

8.3.4. Other Types of Attack

This section doesn't claim to present an exhaustive list of attacks against the NEA reference model. Several types of attack will become easier to understand and analyze once the NEA WG has created specifications describing the specific selected technologies and protocols to be used within NEA. One such area is Denial of Service (DoS). At this point in time, it is not practical to try to define all of the potential exposures present within the NEA protocols, so such an analysis should be included in the Security Considerations sections of the selected NEA protocols.

However, it is important that the NEA Server be resilient to DoS attacks as an outage might affect large numbers of endpoints wishing to join or remain on the network. The NEA reference model expects that the PT protocol would have some amount of DoS resilience and that the PA and PB protocols would need to build upon that base with their own protections. To help narrow the window of attack by unauthenticated parties, it is envisioned that NEA Servers would employ PT protocols that enable an early mutual authentication of the requesting endpoint as one technique for filtering out attacks.

Attacks occurring after the authentication would at least come from sources possessing valid credentials and could potentially be held accountable. Similarly, NEA protocols should offer strong replay protection to prevent DoS-based attacks based on replayed sessions and messages. Posture assessment should be strongly linked with the Posture Transport authentications that occurred to assure the posture came from the authenticated party. Cryptographic mechanisms and other potentially resource intensive operations should be used sparingly until the validity of the request can be established. This and other resource/protocol based attacks can be evaluated once the NEA technologies and their cryptographic use have been selected.

9. Privacy Considerations

While there are a number of beneficial uses of the NEA technology for organizations that own and operate networks offering services to similarly owned endpoints, these same technologies might enhance the potential for abuse and invasion of personal privacy if misused. This section will discuss a few of the potential privacy concerns raised by the deployment of this technology and offer some guidance to implementers.

The NEA technology enables greater visibility into the configuration of an endpoint from the network. Such transparency enables the network to take into consideration the strength of the endpoint's security mechanisms when making access control decisions to network resources. However, this transparency could also be used to enforce restrictive policies to the detriment of the user by limiting their choice of software or prying into past or present uses of the endpoint.

The scope of the NEA WG was limited to specifying protocols targeting the use cases where the endpoints and network are owned by the same party or the endpoint owner has established a clear expectation of disclosure/compliance with the network owner. This is a familiar model for governments, institutions, and a wide variety of enterprises that provide endpoints to their employees to perform their jobs. In many of these situations, the endpoint is purchased and owned by the enterprise and they often reserve the right to audit and possibly dictate the allowable uses of the device. The NEA technologies allow them to automate the inspection of the contents of an endpoint and this information may be linked to the access control mechanisms on the network to limit endpoint use should the endpoint not meet minimal compliance levels.

In these environments, the level of personal privacy the employee enjoys may be significantly reduced subject to local laws and customs. However, in situations where the endpoint is owned by the user or where local laws protect the rights of the user even when using endpoints owned by another party, it is critical that the NEA implementation enable the user to control what endpoint information is shared with the network. Such controls imposed by the user might prevent or limit their ability to access certain networks or protected resources, but this must be a user choice.

9.1. Implementer Considerations

The NEA WG is not defining NEA Client policy content standards nor defining requirements on aspects of an implementation outside of the network protocols; however, the following guidance is provided to encourage privacy friendly implementations for broader use than just the enterprise-oriented setting described above.

NEA Client implementations are encouraged to offer an opt-in policy to users prior to sharing their endpoint's posture information. The opt-in mechanism should be on a per-user, per-NEA Server basis so each user can control which networks can access any posture information on their system. For those networks that are allowed to assess the endpoint, the user should be able to specify granular restrictions on what particular types and specific attributes Posture

Collectors are allowed to disclose. Posture Validator implementations are discouraged from having the default behavior of using wild carded requests for posture potentially leading to overexposure of information (see section 9.2). Instead Posture Validators, by default, should only request the specific attributes that are required to perform their assessment.

Requests for attributes that are not explicitly allowed (or specifically disallowed) to be shared should result in a user notification and/or log record so the user can assess whether the service is doing something undesirable or whether the user is willing to share this additional information in order to gain access. Some products might consider policy-driven support for prompting the user for authorization with a specific description of the posture information being requested prior to sending it to the NEA Server.

It is envisioned that the owner of the endpoint is able to specify disclosure policies that may override or influence the user's policies on the attributes visible to the network. If the owner disclosure policy allows for broader posture availability than the user policy, the implementation should provide a feedback mechanism to the user so they understand the situation and can choose whether to use the endpoint in those circumstances.

In such a system, it is important that the user's policy authoring interface is easy to understand and clearly articulates the current disclosure policy of the system including any influences from the owner policy. Users should be able to understand what posture is available to the network and the general impact of this information being known. In order to minimize the list of restrictions enumerated, use of a conservative default disclosure policy such as "that which is not explicitly authorized for disclosure is not allowed" might make sense to avoid unintentional leakage of information.

NEA Server implementations should provide newly subscribing endpoints with a disclosure statement that clearly states:

- o What information is required
- o How this information will be used and protected
- o What local privacy policies are applicable

This information will empower subscribing users to decide whether the disclosure of this information is acceptable considering local laws and customs.

9.2. Minimizing Attribute Disclosure

One important issue in the design of the NEA reference model and protocols is enabling endpoints to disclose minimal information required to establish compliance with network policies. There are several models that could be considered as to how the disclosed attribute set is established. Each model has privacy related benefits and issues that should be considered by product developers. This section summarizes three potential models for how attribute disclosure might be provided within NEA products and some privacy implications potentially associated with each model.

The first model is easy to implement and deploy but has privacy and potentially latency and scalability implications. This approach effectively defaults the local policy to send all known NEA Posture Attributes when an assessment occurs. While this might simplify deployment, it exposes a lot of information that is potentially not relevant to the security assessment of the system and may introduce privacy issues. For example, is it really important that the enterprise know whether Firefox is being used on a system instead of other browsers during the security posture assessment?

The second model involves an out-of-band provisioning of the disclosure policy to all endpoints. This model may involve the enterprise establishing policy that a particular list of attributes must be provided when a NEA exchange occurs. Endpoint privacy policy may filter this attribute list, but such changes could cause the endpoint not to be given network or resource access. This model simplifies the network exchange as the endpoint always sends the filtered list of attributes when challenged by a particular network. However, this approach requires an out-of-band management protocol to establish and manage the NEA disclosure policies of all systems.

The third model avoids the need for pre-provisioning of a disclosure policy by allowing the NEA Server to specifically request what attributes are required. This is somewhat analogous to the policy being provisioned during the NEA exchanges so is much easier to manage. This model allows for the NEA Server to iteratively ask for attributes based on the values of prior attributes. Note, even in this model the NEA protocols are not expected to be a general purpose query language, but rather allow the NEA Server to request specific attributes as only the defined attributes are possible to request. For example, an enterprise might ask about the OS version in the initial message dialog and after learning the system is running Linux ask for a different set of attributes specific to Linux than it would if the endpoint was a Windows system. It is envisioned that this

approach might minimize the set of attributes sent over the network if the assessment is of a complex system (such as trying to understand what patches are missing from an OS).

In each model, the user could create a set of per-network privacy filter policies enforced by the NEA Client to prevent the disclosure of attributes felt to be personal in nature or not relevant to a particular network. Such filters would protect the privacy of the user but might result in the user not being allowed access to the desired asset (or network) or being provided limited access.

10. References

10.1. Normative References

- [UTF8] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

10.2. Informative References

- [802.1X] IEEE Standards for Local and Metropolitan Area Networks: Port based Network Access Control, IEEE Std 802.1X-2001, June 2001.
- [CNAC] Cisco, Cisco's Network Admission Control Main Web Site, <http://www.cisco.com/go/nac>
- [EAP] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [HMAC] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [IPSEC] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [NAP] Microsoft, Network Access Protection Main Web Site, <http://www.microsoft.com/nap>
- [RADIUS] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

- [TLS] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", RFC 4346, April 2006.
- [TCG] Trusted Computing Group, Main TCG Web Site, <http://www.trustedcomputinggroup.org/>
- [TNC] Trusted Computing Group, Trusted Network Connect Main Web Site, <https://www.trustedcomputinggroup.org/groups/network/>

11. Acknowledgments

The authors of this document would like to acknowledge the NEA Working Group members who have contributed to previous requirements and problem statement documents that influenced the direction of this specification: Kevin Amarin, Parvez Anandam, Diana Arroyo, Uri Blumenthal, Alan DeKok, Lauren Giroux, Steve Hanna, Thomas Hardjono, Tim Polk, Ravi Sahita, Joe Salowey, Chris Salter, Mauricio Sanchez, Yaron Sheffer, Jeff Six, Susan Thompson, Gary Tomlinson, John Vollbrecht, Nancy Winget, Han Yin, and Hao Zhou.

Authors' Addresses

Paul Sangster
Symantec Corporation
6825 Citrine Dr
Carlsbad, CA 92009 USA
Phone: +1 760 438-5656
EMail: Paul_Sangster@symantec.com

Hormuzd Khosravi
Intel
2111 NE 25th Avenue
Hillsboro, OR 97124 USA
Phone: +1 503 264 0334
EMail: hormuzd.m.khosravi@intel.com

Mahalingam Mani
Avaya Inc.
1033 McCarthy Blvd.
Milpitas, CA 95035 USA
Phone: +1 408 321-4840
EMail: mmani@avaya.com

Kaushik Narayan
Cisco Systems Inc.
10 West Tasman Drive
San Jose, CA 95134
Phone: +1 408 526-8168
EMail: kaushik@cisco.com

Joseph Tardo
Nevis Networks
295 N. Bernardo Ave., Suite 100
Mountain View, CA 94043 USA
EMail: joseph.tardo@nevisnetworks.com

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.

