

Internet Engineering Task Force (IETF)
Request for Comments: 6024
Category: Informational
ISSN: 2070-1721

R. Reddy
National Security Agency
C. Wallace
Cygnacom Solutions
October 2010

Trust Anchor Management Requirements

Abstract

A trust anchor represents an authoritative entity via a public key and associated data. The public key is used to verify digital signatures, and the associated data is used to constrain the types of information for which the trust anchor is authoritative. A relying party uses trust anchors to determine if a digitally signed object is valid by verifying a digital signature using the trust anchor's public key, and by enforcing the constraints expressed in the associated data for the trust anchor. This document describes some of the problems associated with the lack of a standard trust anchor management mechanism and defines requirements for data formats and push-based protocols designed to address these problems.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6024>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
1.1. Terminology	4
1.2. Requirements Notation	4
2. Problem Statement	5
3. Requirements	6
3.1. Transport Independence	6
3.2. Basic Management Operations	7
3.3. Management Targets	7
3.4. Delegation of TA Manager Authority	8
3.5. RFC 5280 Support	9
3.6. Support Purposes other than Certification Path Validation	9
3.7. Trust Anchor Format	10
3.8. Source Authentication	10
3.9. Reduce Reliance on Out-of-Band Trust Mechanisms	11
3.10. Replay Detection	11
3.11. Compromise or Disaster Recovery	12
4. Security Considerations	12
5. References	13
5.1. Normative References	13
5.2. Informative References	13

1. Introduction

Digital signatures are used in many applications. For digital signatures to provide integrity and authentication, the public key used to verify the digital signature must be "trusted", i.e., accepted by a relying party (RP) as appropriate for use in the given context. A public key used to verify a signature must be configured as a trust anchor (TA) or contained in a certificate that can be transitively verified by a certification path terminating at a trust anchor. A trust anchor is a public key and associated data used by a relying party to validate a signature on a signed object where the object is either:

- o a public key certificate that begins a certification path terminated by a signature certificate or encryption certificate
- o an object, other than a public key certificate or certificate revocation list (CRL), that cannot be validated via use of a certification path

Trust anchors have only local significance, i.e., each RP is configured with a set of trust anchors, either by the RP or by an entity that manages TAs in the context in which the RP operates. The associated data defines the scope of a trust anchor by imposing constraints on the signatures that the trust anchor may be used to verify. For example, if a trust anchor is used to verify signatures on X.509 certificates, these constraints may include a combination of name spaces, certificate policies, or application/usage types.

One use of digital signatures is the verification of signatures on firmware packages loaded into hardware modules, such as cryptographic modules, cable boxes, routers, etc. Since such devices are often managed remotely, the devices must be able to authenticate the source of management interactions and can use trust anchors to perform this authentication. However, trust anchors require management as well. Other applications requiring trust anchor management include web browsers (which use trust anchors when authenticating web servers) and email clients (which use trust anchors when validating signed email and when authenticating recipients of encrypted email).

All applications that rely upon digital signatures rely upon some means of managing one or more sets of trust anchors. Each set of trust anchors is referred to in this document as a trust anchor store. Often, the means of managing trust anchor stores are application-specific and rely upon out-of-band means to establish and maintain trustworthiness. An application may use multiple trust

anchor stores, and a given trust anchor store may be used by multiple applications. Each trust anchor store is managed by at least one TA manager; a TA manager may manage multiple TA stores.

The requirements stated in this document were prepared prior to the publication of [RFC5914] and [RFC5934]. The document was not published at that time to allow for changes in requirements during the development of the associated technical specifications. The requirements described below are those that were considered during the development of [RFC5914] and [RFC5934].

This section provides an introduction and defines basic terminology. Section 2 describes problems with current trust anchor management methods. Sections 3 and 4 describe requirements and security considerations for a trust anchor management solution.

1.1. Terminology

The following terms are defined in order to provide a vocabulary for describing requirements for trust anchor management.

Trust Anchor: A trust anchor represents an authoritative entity via a public key and associated data. The public key is used to verify digital signatures, and the associated data is used to constrain the types of information for which the trust anchor is authoritative. A relying party uses trust anchors to determine if a digitally signed object is valid by verifying a digital signature using the trust anchor's public key, and by enforcing the constraints expressed in the associated data for the trust anchor.

Trust Anchor Manager: A trust anchor manager is an entity responsible for managing the contents of a trust anchor store. Throughout this document, each trust anchor manager is assumed to be represented as or delegated by a distinct trust anchor.

Trust Anchor Store: A trust anchor store is a set of one or more trust anchors stored in a device. A trust anchor store may be managed by one or more trust anchor managers. A device may have more than one trust anchor store, each of which may be used by one or more applications.

1.2. Requirements Notation

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

2. Problem Statement

Trust anchors are used to support many application scenarios. Most Internet browsers and email clients use trust anchors when authenticating Transport Layer Security (TLS) sessions, verifying signed email, and generating encrypted email by validating a certification path to a server's certificate, an email originator's certificate, or an email recipient's certificate, respectively. Many software distributions are digitally signed to enable authentication of the software source prior to installation. Trust anchors that support these applications are typically installed as part of the operating system (OS) or application, installed using an enterprise configuration management system, or installed directly by an OS or application user.

Trust anchors are typically stored in application-specific or OS-specific trust anchor stores. Often, a single machine may have a number of different trust anchor stores that may not be synchronized. Reviewing the contents of a particular trust anchor store typically involves use of a proprietary tool that interacts with a particular type of trust store.

The presence of a trust anchor in a particular store often conveys implicit authorization to validate signatures for any contexts from which the store is accessed. For example, the public key of a timestamp authority (TSA) may be installed in a trust anchor store to validate signatures on timestamps [RFC3161]. However, if the store containing this TA is used by multiple applications that serve different purposes, the same key may be used (inappropriately) to validate other types of objects such as certificates or Online Certificate Status Protocol (OCSP) responses. Prior to publication of [RFC5914], there was no standard general-purpose mechanism for limiting the applicability (scope) of a trust anchor. A common practice to address this problem is to place different TAs in different stores and limit the set of applications that access a given TA store.

Trust relationships between Public Key Infrastructures (PKIs) are negotiated by policy authorities. Negotiations frequently require significant time to ensure all participating parties' requirements are satisfied. These requirements are expressed, to some extent, in public key certificates via policy constraints, name constraints, etc. In order for these requirements to be enforced, trust anchor stores must be managed in accord with policy authority intentions. Otherwise, the constraints defined in a cross-certificate could be circumvented by recognizing the subject of the cross certificate as a trust anchor, which would enable path processing implementations to avoid the cross-certificate.

Trust anchors are often represented as self-signed certificates, which provide no useful means of establishing the validity of the information contained in the certificate. Confidence in the integrity of a trust anchor is typically established through out-of-band means, often by checking the "fingerprint" (one-way hash) of the self-signed certificate with an authoritative source. Routine trust anchor rekey operations typically require similar out-of-band checks, though in-band rekey of a trust anchor is supported by the Certificate Management Protocol (CMP) [RFC4210]. Ideally, only the initial set of trust anchors are installed in a particular trust anchor store should require out-of-band verification, particularly when the costs of performing out-of-band checks commensurate with the security requirements of applications using the trust anchor store are high.

Despite the prevalent use of trust anchors, there is neither a standard means for discovering the set of trust anchors installed in a particular trust anchor store nor a standard means of managing those trust anchors. The remainder of this document describes requirements for a solution to this problem along with some security considerations.

3. Requirements

This section describes the requirements for a trust anchor management protocol. Requirements are provided for trust anchor contents as well as for trust anchor store management operations.

3.1. Transport Independence

3.1.1. Functional Requirements

A general-purpose solution for the management of trust anchors MUST be transport independent in order to apply to a range of device communications environments. It MUST work in both session-oriented and store-and-forward communications environments as well as in both push and pull distribution models. To accommodate both communication models in a uniform fashion, connectionless integrity and data origin authentication for TA transactions MUST be provided at the application layer. Confidentiality MAY be provided for such transactions.

3.1.2. Rationale

Not all devices that use trust anchors are available for online management operations; some devices may require manual interaction for trust anchor management. Data origin authentication and integrity are required to ensure that the transaction has not been

modified en route. Only connectionless integrity is required, for compatibility with store-and-forward contexts.

3.2. Basic Management Operations

3.2.1. Functional Requirements

At a minimum, a protocol used for trust anchor management MUST enable a trust anchor manager to perform the following operations:

- o Determine which trust anchors are installed in a particular trust anchor store
- o Add one or more trust anchors to a trust anchor store
- o Remove one or more trust anchors from a trust anchor store
- o Replace an entire trust anchor store

A trust anchor management protocol MUST provide support for these basic operations; however, not all implementations must support each option. For example, some implementations may support only replacement of trust anchor stores.

3.2.2. Rationale

These requirements describe the core operations required to manage the contents of a trust anchor store. An edit operation was omitted for the sake of simplicity, with consecutive remove and add operations used for this purpose. A single add or remove operation can act upon more than one trust anchor to avoid unnecessary round trips and are provided to avoid the need to always replace an entire trust anchor store. Trust anchor store replacement may be useful as a simple, higher-bandwidth alternative to add and remove operations.

3.3. Management Targets

3.3.1. Functional Requirements

A protocol for TA management MUST allow a TA management transaction to be directed to:

All TA stores for which the manager is responsible

An enumerated list of one or more named groups of trust anchor stores

An individual trust anchor store

3.3.2. Rationale

Connections between PKIs can be accomplished using different means. Unilateral or bilateral cross-certification can be performed, or a community may simply elect to explicitly accept a trust anchor from another community. Typically, these decisions occur at the enterprise level. In some scenarios, it can be useful to establish these connections for a small community within an enterprise. Enterprise-wide mechanisms such as cross-certificates are ill-suited for this purpose since certificate revocation or expiration affects the entire enterprise.

A trust anchor management protocol can address this issue by supporting limited installation of trust anchors (i.e., installation of TAs in subsets of the enterprise user community), and by supporting expression of constraints on trust anchor use by relying parties. Limited installation requires the ability to identify the members of the community that are intended to rely upon a particular trust anchor, as well as the ability to query and report on the contents of trust anchor stores. Trust anchor constraints can be used to represent the limitations that might otherwise be expressed in a cross-certificate, and limited installation ensures the recognition of the trust anchor does not necessarily encompass an entire enterprise.

Trust anchor configurations may be uniform across an enterprise, or they may be unique to a single application or small set of applications. Many devices and some applications utilize multiple trust anchor stores. By providing means of addressing a specific store or collections of stores, a trust anchor management protocol can enable efficient management of all stores under a trust anchor manager's control.

3.4. Delegation of TA Manager Authority

3.4.1. Functional Requirements

A trust anchor management protocol MUST enable secure transfer of control of a trust anchor store from one trust anchor manager to another. It also SHOULD enable delegation for specific operations without requiring delegation of the overall trust anchor management capability itself.

3.4.2. Rationale

Trust anchor manager rekey is one type of transfer that must be supported. In this case, the new key will be assigned the same privileges as the old key.

Creation of trust anchors for specific purposes, such as firmware signing, is another example of delegation. For example, a trust anchor manager may delegate only the authority to sign firmware to an entity, but disallow further delegation of that privilege, or the trust anchor manager may allow its delegate to further delegate firmware signing authority to other entities.

3.5. RFC 5280 Support

3.5.1. Functional Requirements

A trust anchor management protocol **MUST** enable management of trust anchors that will be used to validate certification paths and CRLs in accordance with [RFC5280] and [RFC5055]. A trust anchor format **MUST** enable the representation of constraints that influence certification path validation or otherwise establish the scope of usage of the trust anchor public key. Examples of such constraints are name constraints, certificate policies, and key usage.

3.5.2. Rationale

Certification path validation is one of the most common applications of trust anchors. The rules for using trust anchors for path validation are established in [RFC5280]. [RFC5055] describes the use of trust anchors for delegated path validation. Trust anchors used to validate certification paths are responsible for providing, possibly through a delegate, the revocation status information of certificates it issues; this is often accomplished by signing a CRL.

3.6. Support Purposes other than Certification Path Validation

3.6.1. Functional Requirements

A trust anchor management protocol **MUST** enable management of trust anchors that can be used for purposes other than certification path validation, including trust anchors that cannot be used for certification path validation. It **SHOULD** be possible to authorize a trust anchor to delegate authority (to other TAs or certificate holders) and to prevent a trust anchor from delegating authority.

3.6.2. Rationale

Trust anchors are used to validate a variety of signed objects, not just public key certificates and CRLs. For example, a trust anchor may be used to verify firmware packages [RFC4108], OCSP responses [RFC2560], Server-Based Certificate Validation Protocol (SCVP) responses [RFC5055], or timestamps [RFC3161]. TAs that are authorized for use with some or all of these other types of

operations may not be authorized to verify public key certificates or CRLs. Thus, it is important to be able to impose constraints on the ways in which a given TA is employed.

3.7. Trust Anchor Format

3.7.1. Functional Requirements

Minimally, a trust anchor management protocol MUST support management of trust anchors represented as self-signed certificates and trust anchors represented as a distinguished name, public key information, and, optionally, associated data. The definition of a trust anchor MUST include a public key, a public key algorithm, and, if necessary, public key parameters. When the public key is used to validate certification paths or CRLs, a distinguished name also MUST be included per [RFC5280]. A trust anchor format SHOULD enable specification of a public key identifier to enable other applications of the trust anchor, for example, verification of data signed using the Cryptographic Message Syntax (CMS) SignedData structure [RFC5652]. A trust anchor format also SHOULD enable the representation of constraints that can be applied to restrict the use of a trust anchor.

3.7.2. Rationale

Prior to the publication of [RFC5914], there was no standardized format for trust anchors. Self-signed X.509 certificates are typically used, but [RFC5280] does not mandate a particular trust anchor representation. It requires only that a trust anchor's public key information and distinguished name be available during certification path validation. CMS is widely used to protect a variety of types of content using digital signatures, including contents that may be verified directly using a trust anchor, such as firmware packages [RFC4108]. Constraints may include a validity period, constraints on certification path validation, etc.

3.8. Source Authentication

3.8.1. Functional Requirements

An entity receiving trust anchor management data MUST be able to authenticate the identity of the party providing the information and MUST be able to confirm the party is authorized to provide that trust anchor information.

A trust anchor manager **MUST** be able to authenticate which trust anchor store corresponds to a report listing the contents of the trust anchor store and be able to confirm the contents of the report have not been subsequently altered.

3.8.2. Rationale

Data origin authentication and integrity are required to support remote management operations, even when TA management transactions are effected via store-and-forward communications.

3.9. Reduce Reliance on Out-of-Band Trust Mechanisms

3.9.1. Functional Requirements

When performing add operations, a trust anchor management protocol **SHOULD** enable TA integrity to be checked automatically by a relying party without relying on out-of-band trust mechanisms.

3.9.2. Rationale

Traditionally, a trust anchor is distributed out-of-band with its integrity checked manually prior to installation. Installation typically is performed by anyone with sufficient administrative privilege on the system receiving the trust anchor. Reliance on out-of-band trust mechanisms is one problem with current trust anchor management approaches, and reduction of the need to use out-of-band trust mechanisms is a primary motivation for developing a trust anchor management protocol. Ideally, out-of-band trust mechanisms will be required only during trust anchor store initialization.

3.10. Replay Detection

3.10.1. Functional Requirements

A trust anchor management protocol **MUST** enable participants engaged in a trust anchor management protocol exchange to detect replay attacks. A replay detection mechanism that does not introduce a requirement for a reliable source of time **MUST** be available. Mechanisms that do require a reliable source of time **MAY** be available.

3.10.2. Rationale

Detection of replays of trust anchor management transactions is required to support remote management operations. Replay of old trust anchor management transactions could result in the

reintroduction of compromised trust anchors to a trust anchor store, potentially exposing applications to malicious signed objects or certification paths.

Some devices that utilize trust anchors have no access to a reliable source of time, so a replay detection mechanism that requires a reliable time source is insufficient.

3.11. Compromise or Disaster Recovery

3.11.1. Functional Requirements

A trust anchor management protocol **MUST** enable recovery from the compromise or loss of a trust anchor private key, including the private key authorized to serve as a trust anchor manager, without requiring re-initialization of the trust store.

3.11.2. Rationale

Compromise or loss of a private key corresponding to a trust anchor can have significant negative consequences. Currently, in some cases, re-initialization of all affected trust anchor stores is required to recover from a lost or compromised trust anchor key. Due to the costs associated with re-initialization, a trust anchor management protocol should support recovery options that do not require trust anchor store re-initialization.

4. Security Considerations

The public key used to authenticate a TA management transaction may have been placed in the client as the result of an earlier TA management transaction or during an initial bootstrap configuration operation. In most scenarios, at least one public key authorized for trust anchor management must be placed in each trust anchor store to be managed during the initial configuration of the trust anchor store. This public key may be transported and checked using out-of-band means. In all scenarios, regardless of the authentication mechanism, at least one trust anchor manager must be established for each trust anchor store during the initial configuration of the trust anchor store.

Compromise of a trust anchor's private key can result in many security problems including issuance of bogus certificates or installation of rogue trust anchors.

Usage of trust anchor-based constraints requires great care when defining trust anchors. Errors on the part of a trust anchor manager could result in denial of service or have serious security

consequences. For example, if a name constraint for a trust anchor that serves as the root of a PKI includes a typo, denial of service results for certificate holders and relying parties. If a trust anchor manager inadvertently delegates all of its privileges and the delegate subsequently removes the trust anchor manager from trust anchor stores now under its control, recovery may require re-initialization of all effected trust anchor stores.

RFC 5280 requires that certificate path validation be initialized with a TA subject name and public key, but does not require processing of other information, such as name constraints extensions. Inclusion of constraints in trust anchors is optional. When constraints are explicitly included by a trust anchor manager using a trust anchor management protocol, there exists an expectation that the certificate path validation algorithm will make use of the constraints. Application owners must confirm the path processing implementations support the processing of TA-based constraints, where required.

Many of the security considerations from [RFC5280] are also applicable to trust anchor management.

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5055] Freeman, T., Housley, R., Malpani, A., Cooper, D., and W. Polk, "Server-Based Certificate Validation Protocol (SCVP)", RFC 5055, December 2007.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

5.2. Informative References

- [RFC2560] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 2560, June 1999.
- [RFC3161] Adams, C., Cain, P., Pinkas, D., and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", RFC 3161, August 2001.

- [RFC4108] Housley, R., "Using Cryptographic Message Syntax (CMS) to Protect Firmware Packages", RFC 4108, August 2005.
- [RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, September 2005.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, September 2009.
- [RFC5914] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Format", RFC 5914, June 2010.
- [RFC5934] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Management Protocol (TAMP)", RFC 5934, August 2010.

Authors' Addresses

Raksha Reddy
National Security Agency
Suite 6599
9800 Savage Road
Fort Meade, MD 20755

EMail: r.reddy@radium.ncsc.mil

Carl Wallace
Cygnacom Solutions
Suite 5400
7925 Jones Branch Drive
McLean, VA 22102

EMail: cwallace@cygnacom.com

