Internet Engineering Task Force (IETF)
Request for Comments: 6125
Category: Standards Track
ISSN: 2070-1721

P. Saint-Andre
Cisco
J. Hodges
PayPal

March 2011

Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)

Abstract

Many application technologies enable secure communication between two entities by means of Internet Public Key Infrastructure Using X.509 (PKIX) certificates in the context of Transport Layer Security (TLS). This document specifies procedures for representing and verifying the identity of application services in such interactions.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc6125.

Table of Contents

1. Introduction

1.1. Motivation

The visible face of the Internet largely consists of services that employ a client-server architecture in which an interactive or automated client communicates with an application service in order to retrieve or upload information, communicate with other entities, or access a broader network of services. When a client communicates with an application service using Transport Layer Security [TLS] or Datagram Transport Layer Security [DTLS], it references some notion of the server's identity (e.g., "the website at example.com") while attempting to establish secure communication. Likewise, during TLS negotiation, the server presents its notion of the service's identity in the form of a public-key certificate that was issued by a certification authority (CA) in the context of the Internet Public Key Infrastructure using X.509 [PKIX]. Informally, we can think of these identities as the client's "reference identity" and the server's "presented identity" (these rough ideas are defined more precisely later in this document through the concept of particular identifiers). In general, a client needs to verify that the server's presented identity matches its reference identity so it can authenticate the communication.

Many application technologies adhere to the pattern just outlined. Such protocols have traditionally specified their own rules for representing and verifying application service identity. Unfortunately, this divergence of approaches has caused some confusion among certification authorities, application developers, and protocol designers.

Therefore, to codify secure procedures for the implementation and deployment of PKIX-based authentication, this document specifies recommended procedures for representing and verifying application service identity in certificates intended for use in application protocols employing TLS.

1.2. Audience

The primary audience for this document consists of application protocol designers, who can reference this document instead of defining their own rules for the representation and verification of application service identity. Secondarily, the audience consists of certification authorities, service providers, and client developers from technology communities that might reuse the recommendations in this document when defining certificate issuance policies, generating certificate signing requests, or writing software algorithms for identity matching.

1.3. How to Read This Document

This document is longer than the authors would have liked because it was necessary to carefully define terminology, explain the underlying concepts, define the scope, and specify recommended behavior for both certification authorities and application software implementations. The following sections are of special interest to various audiences:

- Protocol designers might want to first read the checklist in Section 3.
- Certification authorities might want to first read the recommendations for representation of server identity in Section 4.
- Service providers might want to first read the recommendations for requesting of server certificates in Section 5.
- Software implementers might want to first read the recommendations for verification of server identity in Section 6.

The sections on terminology (Section 1.8), naming of application services (Section 2), document scope (Section 1.7), and the like provide useful background information regarding the recommendations and guidelines that are contained in the above-referenced sections, but are not absolutely necessary for a first reading of this document.

1.4. Applicability

This document does not supersede the rules for certificate issuance or validation provided in [PKIX]. Therefore, [PKIX] is authoritative on any point that might also be discussed in this document. Furthermore, [PKIX] also governs any certificate-related topic on which this document is silent, including but not limited to certificate syntax, certificate extensions such as name constraints and extended key usage, and handling of certification paths.

This document addresses only name forms in the leaf "end entity" server certificate, not any name forms in the chain of certificates used to validate the server certificate. Therefore, in order to ensure proper authentication, application clients need to verify the entire certification path per [PKIX].

This document also does not supersede the rules for verifying service identity provided in specifications for existing application protocols published prior to this document, such as those excerpted under Appendix B. However, the procedures described here can be referenced by future specifications, including updates to specifications for existing application protocols if the relevant technology communities agree to do so.

1.5. Overview of Recommendations

To orient the reader, this section provides an informational overview of the recommendations contained in this document.

For the primary audience of application protocol designers, this document provides recommended procedures for the representation and verification of application service identity within PKIX certificates used in the context of TLS.

For the secondary audiences, in essence this document encourages certification authorities, application service providers, and application client developers to coalesce on the following practices:

- Move away from including and checking strings that look like domain names in the subject’s Common Name.

- Move toward including and checking DNS domain names via the subjectAlternativeName extension designed for that purpose: dNSName.

- Move toward including and checking even more specific subjectAlternativeName extensions where appropriate for using the protocol (e.g., uniformResourceIdentifier and the otherName form SRVName).

- Move away from the issuance of so-called wildcard certificates (e.g., a certificate containing an identifier for "*.example.com").

These suggestions are not entirely consistent with all practices that are currently followed by certification authorities, client developers, and service providers. However, they reflect the best aspects of current practices and are expected to become more widely adopted in the coming years.

1.6. Generalization from Current Technologies

This document attempts to generalize best practices from the many application technologies that currently use PKIX certificates with TLS. Such technologies include, but are not limited to:

- The Internet Message Access Protocol [IMAP] and the Post Office Protocol [POP3]; see also [USINGTLS]

- The Hypertext Transfer Protocol [HTTP]; see also [HTTP-TLS]

- The Lightweight Directory Access Protocol [LDAP]; see also [LDAP-AUTH] and its predecessor [LDAP-TLS]

- The Simple Mail Transfer Protocol [SMTP]; see also [SMTP-AUTH] and [SMTP-TLS]

- The Extensible Messaging and Presence Protocol [XMPP]; see also [XMPP-OLD]

- The Network News Transfer Protocol [NNTP]; see also [NNTP-TLS]

- The NETCONF Configuration Protocol [NETCONF]; see also [NETCONF-SSH] and [NETCONF-TLS]

- The Syslog Protocol [SYSLOG]; see also [SYSLOG-TLS] and [SYSLOG-DTLS]

- The Session Initiation Protocol [SIP]; see also [SIP-CERTS]

- The Simple Network Management Protocol [SNMP]; see also [SNMP-TLS]

- The General Internet Signalling Transport [GIST]

However, as noted, this document does not supersede the rules for verifying service identity provided in specifications for those application protocols.

1.7. Scope

1.7.1. In Scope

This document applies only to service identities associated with fully qualified DNS domain names, only to TLS and DTLS (or the older Secure Sockets Layer (SSL) technology), and only to PKIX-based systems. As a result, the scenarios described in the following section are out of scope for this specification (although they might be addressed by future specifications).

1.7.2. Out of Scope

The following topics are out of scope for this specification:

- Client or end-user identities.

  Certificates representing client or end-user identities (e.g., the rfc822Name identifier) can be used for mutual authentication between a client and server or between two clients, thus enabling stronger client-server security or end-to-end security. However, certification authorities, application developers, and service operators have less experience with client certificates than with server certificates, thus giving us fewer models from which to generalize and a less solid basis for defining best practices.

- Identifiers other than fully qualified DNS domain names.

  Some certification authorities issue server certificates based on IP addresses, but preliminary evidence indicates that such certificates are a very small percentage (less than 1%) of issued certificates. Furthermore, IP addresses are not necessarily reliable identifiers for application services because of the existence of private internets [PRIVATE], host mobility, multiple interfaces on a given host, Network Address Translators (NATs) resulting in different addresses for a host from different locations on the network, the practice of grouping many hosts together behind a single IP address, etc. Most fundamentally, most users find DNS domain names much easier to work with than IP addresses, which is why the domain name system was designed in the first place. We prefer to define best practices for the much more common use case and not to complicate the rules in this specification.

Furthermore, we focus here on application service identities, not specific resources located at such services. Therefore this document discusses Uniform Resource Identifiers [URI] only as a way to communicate a DNS domain name (via the URI "host" component or its equivalent), not as a way to communicate other aspects of a service such as a specific resource (via the URI "path" component) or parameters (via the URI "query" component).

We also do not discuss attributes unrelated to DNS domain names, such as those defined in [X.520] and other such specifications (e.g., organizational attributes, geographical attributes, company logos, and the like).

- Security protocols other than [TLS], [DTLS], or the older Secure Sockets Layer (SSL) technology.

  Although other secure, lower-layer protocols exist and even employ PKIX certificates at times (e.g., IPsec [IPSEC]), their use cases can differ from those of TLS-based and DTLS-based application technologies. Furthermore, application technologies have less experience with IPsec than with TLS, thus making it more difficult to gather feedback on proposed best practices.

- Keys or certificates employed outside the context of PKIX-based systems.

  Some deployed application technologies use a web of trust model based on or similar to OpenPGP [OPENPGP], or use self-signed certificates, or are deployed on networks that are not directly connected to the public Internet and therefore cannot depend on Certificate Revocation Lists (CRLs) or the Online Certificate Status Protocol [OCSP] to check CA-issued certificates. However, the method for binding a public key to an identifier in OpenPGP differs essentially from the method in X.509, the data in self-signed certificates has not been certified by a third party in any way, and checking of CA-issued certificates via CRLs or OCSP is critically important to maintaining the security of PKIX-based systems. Attempting to define best practices for such technologies would unduly complicate the rules defined in this specification.

- Certification authority policies, such as:

  * What types or "classes" of certificates to issue and whether to apply different policies for them (e.g., allow the wildcard character in certificates issued to individuals who have provided proof of identity but do not allow the wildcard character in "Extended Validation" certificates [EV-CERTS]).

* Whether to issue certificates based on IP addresses (or some other form, such as relative domain names) in addition to fully qualified DNS domain names.

* Which identifiers to include (e.g., whether to include SRV-IDs or URI-IDs as defined in the body of this specification).

* How to certify or validate fully qualified DNS domain names and application service types.

* How to certify or validate other kinds of information that might be included in a certificate (e.g., organization name).

o Resolution of DNS domain names.

  Although the process whereby a client resolves the DNS domain name of an application service can involve several steps (e.g., this is true of resolutions that depend on DNS SRV resource records, Naming Authority Pointer (NAPTR) DNS resource records [NAPTR], and related technologies such as [S-NAPTR]), for our purposes we care only about the fact that the client needs to verify the identity of the entity with which it communicates as a result of the resolution process. Thus the resolution process itself is out of scope for this specification.

o User interface issues.

  In general, such issues are properly the responsibility of client software developers and standards development organizations dedicated to particular application technologies (see, for example, [WSC-UI]).

1.8. Terminology

Because many concepts related to "identity" are often too vague to be actionable in application protocols, we define a set of more concrete terms for use in this specification.

application service: A service on the Internet that enables interactive and automated clients to connect for the purpose of retrieving or uploading information, communicating with other entities, or connecting to a broader network of services.

application service provider: An organization or individual that hosts or deploys an application service.

application service type: A formal identifier for the application protocol used to provide a particular kind of application service at a domain; the application service type typically takes the form of a Uniform Resource Identifier scheme [URI] or a DNS SRV Service [DNS-SRV].

attribute-type-and-value pair: A colloquial name for the ASN.1-based construction comprising a Relative Distinguished Name (RDN), which itself is a building-block component of Distinguished Names. See Section 2 of [LDAP-DN].

automated client: A software agent or device that is not directly controlled by a human user.

delegated domain: A domain name or host name that is explicitly configured for communicating with the source domain, by either (a) the human user controlling an interactive client or (b) a trusted administrator. In case (a), one example of delegation is an account setup that specifies the domain name of a particular host to be used for retrieving information or connecting to a network, which might be different from the server portion of the user's account name (e.g., a server at mailhost.example.com for connecting to an IMAP server hosting an email address of juliet@example.com). In case (b), one example of delegation is an admin-configured host-to-address/address-to-host lookup table.

derived domain: A domain name or host name that a client has derived from the source domain in an automated fashion (e.g., by means of a [DNS-SRV] lookup).

identifier: A particular instance of an identifier type that is either presented by a server in a certificate or referenced by a client for matching purposes.

identifier type: A formally defined category of identifier that can be included in a certificate and therefore that can also be used for matching purposes. For conciseness and convenience, we define the following identifier types of interest, which are based on those found in the PKIX specification [PKIX] and various PKIX extensions.

* CN-ID = a Relative Distinguished Name (RDN) in the certificate subject field that contains one and only one attribute-type-and-value pair of type Common Name (CN), where the value matches the overall form of a domain name (informally, dot-separated letter-digit-hyphen labels); see [PKIX] and also [LDAP-SCHEMA]

* DNS-ID = a subjectAltName entry of type dNSName; see [PKIX]

* SRV-ID = a subjectAltName entry of type otherName whose name form is SRVName; see [SRVNAME]

* URI-ID = a subjectAltName entry of type uniformResourceIdentifier whose value includes both (i) a "scheme" and (ii) a "host" component (or its equivalent) that matches the "reg-name" rule (where the quoted terms represent the associated [ABNF] productions from [URI]); see [PKIX] and [URI]

interactive client: A software agent or device that is directly controlled by a human user. (Other specifications related to security and application protocols, such as [WSC-UI], often refer to this entity as a "user agent".)

pinning: The act of establishing a cached name association between the application service's certificate and one of the client's reference identifiers, despite the fact that none of the presented identifiers matches the given reference identifier. Pinning is accomplished by allowing a human user to positively accept the mismatch during an attempt to communicate with the application service. Once a cached name association is established, the certificate is said to be pinned to the reference identifier and in future communication attempts the client simply verifies that the service's presented certificate matches the pinned certificate, as described under Section 6.6.2. (A similar definition of "pinning" is provided in [WSC-UI].)

PKIX: PKIX is a short name for the Internet Public Key Infrastructure using X.509 defined in RFC 5280 [PKIX], which comprises a profile of the X.509v3 certificate specifications and X.509v2 certificate revocation list (CRL) specifications for use in the Internet.

PKIX-based system: A software implementation or deployed service that makes use of X.509v3 certificates and X.509v2 certificate revocation lists (CRLs).

PKIX certificate: An X.509v3 certificate generated and employed in the context of PKIX.

presented identifier: An identifier that is presented by a server to a client within a PKIX certificate when the client attempts to establish secure communication with the server; the certificate can include one or more presented identifiers of different types,

and if the server hosts more than one domain then the certificate might present distinct identifiers for each domain.

reference identifier: An identifier, constructed from a source domain and optionally an application service type, used by the client for matching purposes when examining presented identifiers.

source domain: The fully qualified DNS domain name that a client expects an application service to present in the certificate (e.g., "www.example.com"), typically input by a human user, configured into a client, or provided by reference such as in a hyperlink. The combination of a source domain and, optionally, an application service type enables a client to construct one or more reference identifiers.

subjectAltName entry: An identifier placed in a subjectAltName extension.

subjectAltName extension: A standard PKIX certificate extension [PKIX] enabling identifiers of various types to be bound to the certificate subject -- in addition to, or in place of, identifiers that may be embedded within or provided as a certificate's subject field.

subject field: The subject field of a PKIX certificate identifies the entity associated with the public key stored in the subject public key field (see Section 4.1.2.6 of [PKIX]).

subject name: In an overall sense, a subject's name(s) can be represented by or in the subject field, the subjectAltName extension, or both (see [PKIX] for details). More specifically, the term often refers to the name of a PKIX certificate's subject, encoded as the X.501 type Name and conveyed in a certificate's subject field (see Section 4.1.2.6 of [PKIX]).

TLS client: An entity that assumes the role of a client in a Transport Layer Security [TLS] negotiation. In this specification we generally assume that the TLS client is an (interactive or automated) application client; however, in application protocols that enable server-to-server communication, the TLS client could be a peer application service.

TLS server: An entity that assumes the role of a server in a Transport Layer Security [TLS] negotiation; in this specification we assume that the TLS server is an application service.

Most security-related terms in this document are to be understood in the sense defined in [SECTERMS]; such terms include, but are not limited to, "attack", "authentication", "authorization", "certification authority", "certification path", "certificate", "credential", "identity", "self-signed certificate", "trust", "trust anchor", "trust chain", "validate", and "verify".

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

2. Naming of Application Services

This section discusses naming of application services on the Internet, followed by a brief tutorial about subject naming in PKIX.

2.1. Naming Application Services

This specification assumes that the name of an application service is based on a DNS domain name (e.g., "example.com") -- supplemented in some circumstances by an application service type (e.g., "the IMAP server at example.com").

From the perspective of the application client or user, some names are direct because they are provided directly by a human user (e.g., via runtime input, prior configuration, or explicit acceptance of a client communication attempt), whereas other names are indirect because they are automatically resolved by the client based on user input (e.g., a target name resolved from a source name using DNS SRV or NAPTR records). This dimension matters most for certificate consumption, specifically verification as discussed in this document.

From the perspective of the application service, some names are unrestricted because they can be used in any type of service (e.g., a certificate might be reused for both the HTTP service and the IMAP service at example.com), whereas other names are restricted because they can be used in only one type of service (e.g., a special-purpose certificate that can be used only for an IMAP service). This dimension matters most for certificate issuance.

Therefore, we can categorize the identifier types of interest as follows:

- A CN-ID is direct and unrestricted.
- A DNS-ID is direct and unrestricted.

- An SRV-ID can be either direct or (more typically) indirect, and is restricted.

- A URI-ID is direct and restricted.

We summarize this taxonomy in the following table.

	Direct	Restricted
CN-ID	Yes	No
DNS-ID	Yes	No
SRV-ID	Either	Yes
URI-ID	Yes	Yes

When implementing software, deploying services, and issuing certificates for secure PKIX-based authentication, it is important to keep these distinctions in mind. In particular, best practices differ somewhat for application server implementations, application client implementations, application service providers, and certification authorities. Ideally, protocol specifications that reference this document will specify which identifiers are mandatory-to-implement by servers and clients, which identifiers ought to be supported by certificate issuers, and which identifiers ought to be requested by application service providers. Because these requirements differ across applications, it is impossible to categorically stipulate universal rules (e.g., that all software implementations, service providers, and certification authorities for all application protocols need to use or support DNS-IDs as a baseline for the purpose of interoperability).

However, it is preferable that each application protocol will at least define a baseline that applies to the community of software developers, application service providers, and CAs actively using or supporting that technology (one such community, the CA/Browser Forum, has codified such a baseline for "Extended Validation Certificates" in [EV-CERTS]).

2.2. DNS Domain Names

For the purposes of this specification, the name of an application service is (or is based on) a DNS domain name that conforms to one of the following forms:

1. A "traditional domain name", i.e., a fully qualified DNS domain name or "FQDN" (see [DNS-CONCEPTS]) all of whose labels are "LDH labels" as described in [IDNA-DEFS]. Informally, such labels are constrained to [US-ASCII] letters, digits, and the hyphen, with the hyphen prohibited in the first character position. Additional qualifications apply (please refer to the above-referenced specifications for details), but they are not relevant to this specification.

2. An "internationalized domain name", i.e., a DNS domain name that conforms to the overall form of a domain name (informally, dot-separated letter-digit-hyphen labels) but includes at least one label containing appropriately encoded Unicode code points outside the traditional US-ASCII range. That is, it contains at least one U-label or A-label, but otherwise may contain any mixture of NR-LDH labels, A-labels, or U-labels, as described in [IDNA-DEFS] and the associated documents.

2.3. Subject Naming in PKIX Certificates

In theory, the Internet Public Key Infrastructure using X.509 [PKIX] employs the global directory service model defined in [X.500] and [X.501]. Under that model, information is held in a directory information base (DIB) and entries in the DIB are organized in a hierarchy called the directory information tree (DIT). An object or alias entry in that hierarchy consists of a set of attributes (each of which has a defined type and one or more values) and is uniquely identified by a Distinguished Name (DN). The DN of an entry is constructed by combining the Relative Distinguished Names of its superior entries in the tree (all the way down to the root of the DIT) with one or more specially nominated attributes of the entry itself (which together comprise the Relative Distinguished Name (RDN) of the entry, so-called because it is relative to the Distinguished Names of the superior entries in the tree). The entry closest to the root is sometimes referred to as the "most significant" entry, and the entry farthest from the root is sometimes referred to as the "least significant" entry. An RDN is a set (i.e., an unordered group) of attribute-type-and-value pairs (see also [LDAP-DN]), each of which asserts some attribute about the entry.

In practice, the certificates used in [X.509] and [PKIX] borrow key concepts from X.500 and X.501 (e.g., DNs and RDNs) to identify entities, but such certificates are not necessarily part of a global directory information base. Specifically, the subject field of a PKIX certificate is an X.501 type Name that "identifies the entity associated with the public key stored in the subject public key field" (see Section 4.1.2.6 of [PKIX]). However, it is perfectly acceptable for the subject field to be empty, as long as the

certificate contains a subject alternative name ("subjectAltName") extension that includes at least one subjectAltName entry, because the subjectAltName extension allows various identities to be bound to the subject (see Section 4.2.1.6 of [PKIX]). The subjectAltName extension itself is a sequence of typed entries, where each type is a distinct kind of identifier.

For our purposes, an application service can be identified by a name or names carried in the subject field (i.e., a CN-ID) and/or in one of the following identifier types within subjectAltName entries:

- DNS-ID
- SRV-ID
- URI-ID

Existing certificates often use a CN-ID in the subject field to represent a fully qualified DNS domain name; for example, consider the following three subject names, where the attribute of type Common Name contains a string whose form matches that of a fully qualified DNS domain name ("im.example.org", "mail.example.net", and "www.example.com", respectively):

```
CN=im.example.org,O=Example Org,C=GB

C=CA,O=Example Internetworking,CN=mail.example.net

O=Examples-R-Us,CN=www.example.com,C=US
```

However, the Common Name is not strongly typed because a Common Name might contain a human-friendly string for the service, rather than a string whose form matches that of a fully qualified DNS domain name (a certificate with such a single Common Name will typically have at least one subjectAltName entry containing the fully qualified DNS domain name):

```
CN=A Free Chat Service,O=Example Org,C=GB
```

Or, a certificate’s subject might contain both a CN-ID as well as another common name attribute containing a human-friendly string:

```
CN=A Free Chat Service,CN=im.example.org,O=Example Org,C=GB
```

In general, this specification recommends and prefers use of subjectAltName entries (DNS-ID, SRV-ID, URI-ID, etc.) over use of the subject field (CN-ID) where possible, as more completely described in the following sections. However, specifications that reuse this one

can legitimately encourage continued support for the CN-ID identifier type if they have good reasons to do so, such as backward compatibility with deployed infrastructure (see, for example, [EV-CERTS]).

2.3.1. Implementation Notes

Confusion sometimes arises from different renderings or encodings of the hierarchical information contained in a certificate.

Certificates are binary objects and are encoded using the Distinguished Encoding Rules (DER) specified in [X.690]. However, some implementations generate displayable (a.k.a. printable) renderings of the certificate issuer, subject field, and subjectAltName extension, and these renderings convert the DER-encoded sequences into a "string representation" before being displayed. Because a certificate subject field (of type Name [X.509], the same as for a Distinguished Name (DN) [X.501]) is an ordered sequence, order is typically preserved in subject string representations, although the two most prevalent subject (and DN) string representations differ in employing left-to-right vs. right-to-left ordering. However, because a Relative Distinguished Name (RDN) is an unordered group of attribute-type-and-value pairs, the string representation of an RDN can differ from the canonical DER encoding (and the order of attribute-type-and-value pairs can differ in the RDN string representations or display orders provided by various implementations). Furthermore, various specifications refer to the order of RDNs in DNs or certificate subject fields using terminology that is implicitly related to an information hierarchy (which may or may not actually exist), such as "most specific" vs. "least specific", "left-most" vs. "right-most", "first" vs. "last", or "most significant" vs. "least significant" (see, for example, [LDAP-DN]).

To reduce confusion, in this specification we avoid such terms and instead use the terms provided under Section 1.8; in particular, we do not use the term "(most specific) Common Name field in the subject field" from [HTTP-TLS] and instead state that a CN-ID is a Relative Distinguished Name (RDN) in the certificate subject containing one and only one attribute-type-and-value pair of type Common Name (thus removing the possibility that an RDN might contain multiple AVAs (Attribute Value Assertions) of type CN, one of which could be considered "most specific").

Finally, although theoretically some consider the order of RDNs within a subject field to have meaning, in practice that rule is often not observed. An AVA of type CN is considered to be valid at any position within the subject field.

3. Designing Application Protocols

This section provides guidelines for designers of application protocols, in the form of a checklist to follow when reusing the recommendations provided in this document.

- Does your technology use DNS SRV records to resolve the DNS domain names of application services? If so, consider recommending or requiring support for the SRV-ID identifier type in PKIX certificates issued and used in your technology community. (Note that many existing application technologies use DNS SRV records to resolve the DNS domain names of application services, but do not rely on representations of those records in PKIX certificates by means of SRV-IDs as defined in [SRVNAME].)

- Does your technology use URIs to identify application services? If so, consider recommending or requiring support for the URI-ID identifier type. (Note that many existing application technologies use URIs to identify application services, but do not rely on representation of those URIs in PKIX certificates by means of URI-IDs.)

- Does your technology need to use DNS domain names in the Common Name of certificates for the sake of backward compatibility? If so, consider recommending support for the CN-ID identifier type as a fallback.

- Does your technology need to allow the wildcard character in DNS domain names? If so, consider recommending support for wildcard certificates, and specify exactly where the wildcard character is allowed to occur (e.g., only the complete left-most label of a DNS domain name).

Sample text is provided under Appendix A.

4. Representing Server Identity

This section provides rules and guidelines for issuers of certificates.

4.1. Rules

When a certification authority issues a certificate based on the fully qualified DNS domain name at which the application service provider will provide the relevant application, the following rules apply to the representation of application service identities. The

reader needs to be aware that some of these rules are cumulative and can interact in important ways that are illustrated later in this document.

1. The certificate SHOULD include a "DNS-ID" if possible as a baseline for interoperability.

2. If the service using the certificate deploys a technology for which the relevant specification stipulates that certificates ought to include identifiers of type SRV-ID (e.g., this is true of [XMPP]), then the certificate SHOULD include an SRV-ID.

3. If the service using the certificate deploys a technology for which the relevant specification stipulates that certificates ought to include identifiers of type URI-ID (e.g., this is true of [SIP] as specified by [SIP-CERTS], but not true of [HTTP] since [HTTP-TLS] does not describe usage of a URI-ID for HTTP services), then the certificate SHOULD include a URI-ID. The scheme SHALL be that of the protocol associated with the application service type and the "host" component (or its equivalent) SHALL be the fully qualified DNS domain name of the service. A specification that reuses this one MUST specify which URI schemes are to be considered acceptable in URI-IDs contained in PKIX certificates used for the application protocol (e.g., "sip" but not "sips" or "tel" for SIP as described in [SIP-SIPS], or perhaps http and https for HTTP as might be described in a future specification).

4. The certificate MAY include other application-specific identifiers for types that were defined before publication of [SRVNAME] (e.g., XmppAddr for [XMPP]) or for which service names or URI schemes do not exist; however, such application-specific identifiers are not applicable to all application technologies and therefore are out of scope for this specification.

5. Even though many deployed clients still check for the CN-ID within the certificate subject field, certification authorities are encouraged to migrate away from issuing certificates that represent the server’s fully qualified DNS domain name in a CN-ID. Therefore, the certificate SHOULD NOT include a CN-ID unless the certification authority issues the certificate in accordance with a specification that reuses this one and that explicitly encourages continued support for the CN-ID identifier type in the context of a given application technology.

6. The certificate MAY contain more than one DNS-ID, SRV-ID, or URI-ID but SHOULD NOT contain more than one CN-ID, as further explained under Section 7.4.

7. Unless a specification that reuses this one allows continued support for the wildcard character '*', the DNS domain name portion of a presented identifier SHOULD NOT contain the wildcard character, whether as the complete left-most label within the identifier (following the description of labels and domain names in [DNS-CONCEPTS], e.g., "*.example.com") or as a fragment thereof (e.g., *oo.example.com, f*o.example.com, or fo*.example.com). A more detailed discussion of so-called "wildcard certificates" is provided under Section 7.2.

4.2. Examples

Consider a simple website at "www.example.com", which is not discoverable via DNS SRV lookups. Because HTTP does not specify the use of URIs in server certificates, a certificate for this service might include only a DNS-ID of "www.example.com". It might also include a CN-ID of "www.example.com" for backward compatibility with deployed infrastructure.

Consider an IMAP-accessible email server at the host "mail.example.net" servicing email addresses of the form "user@example.net" and discoverable via DNS SRV lookups on the application service name of "example.net". A certificate for this service might include SRV-IDs of "_imap.example.net" and "_imaps.example.net" (see [EMAIL-SRV]) along with DNS-IDs of "example.net" and "mail.example.net". It might also include CN-IDs of "example.net" and "mail.example.net" for backward compatibility with deployed infrastructure.

Consider a SIP-accessible voice-over-IP (VoIP) server at the host "voice.example.edu" servicing SIP addresses of the form "user@voice.example.edu" and identified by a URI of <sip: voice.example.edu>. A certificate for this service would include a URI-ID of "sip:voice.example.edu" (see [SIP-CERTS]) along with a DNS-ID of "voice.example.edu". It might also include a CN-ID of "voice.example.edu" for backward compatibility with deployed infrastructure.

Consider an XMPP-compatible instant messaging (IM) server at the host "im.example.org" servicing IM addresses of the form "user@im.example.org" and discoverable via DNS SRV lookups on the "im.example.org" domain. A certificate for this service might include SRV-IDs of "_xmpp-client.im.example.org" and "_xmpp-server.im.example.org" (see [XMPP]), a DNS-ID of "im.example.org", and an XMPP-specific "XmppAddr" of "im.example.org" (see [XMPP]). It might also include a CN-ID of "im.example.org" for backward compatibility with deployed infrastructure.

5. Requesting Server Certificates

This section provides rules and guidelines for service providers regarding the information to include in certificate signing requests (CSRs).

In general, service providers are encouraged to request certificates that include all of the identifier types that are required or recommended for the application service type that will be secured using the certificate to be issued.

If the certificate might be used for any type of application service, then the service provider is encouraged to request a certificate that includes only a DNS-ID.

If the certificate will be used for only a single type of application service, then the service provider is encouraged to request a certificate that includes a DNS-ID and, if appropriate for the application service type, an SRV-ID or URI-ID that limits the deployment scope of the certificate to only the defined application service type.

If a service provider offering multiple application service types (e.g., a World Wide Web service, an email service, and an instant messaging service) wishes to limit the applicability of certificates using SRV-IDs or URI-IDs, then the service provider is encouraged to request multiple certificates, i.e., one certificate per application service type. Conversely, the service provider is discouraged from requesting a single certificate containing multiple SRV-IDs or URI-IDs identifying each different application service type. This guideline does not apply to application service type "bundles" that are used to identify manifold distinct access methods to the same underlying application (e.g., an email application with access methods denoted by the application service types of "imap", "imaps", "pop3", "pop3s", and "submission" as described in [EMAIL-SRV]).

6. Verifying Service Identity

This section provides rules and guidelines for implementers of application client software regarding algorithms for verification of application service identity.

6.1. Overview

At a high level, the client verifies the application service's identity by performing the actions listed below (which are defined in the following subsections of this document):

1. The client constructs a list of acceptable reference identifiers based on the source domain and, optionally, the type of service to which the client is connecting.

2. The server provides its identifiers in the form of a PKIX certificate.

3. The client checks each of its reference identifiers against the presented identifiers for the purpose of finding a match.

4. When checking a reference identifier against a presented identifier, the client matches the source domain of the identifiers and, optionally, their application service type.

Naturally, in addition to checking identifiers, a client might complete further checks to ensure that the server is authorized to provide the requested service. However, such checking is not a matter of verifying the application service identity presented in a certificate, and therefore methods for doing so (e.g., consulting local policy information) are out of scope for this document.

6.2. Constructing a List of Reference Identifiers

6.2.1. Rules

The client MUST construct a list of acceptable reference identifiers, and MUST do so independently of the identifiers presented by the service.

The inputs used by the client to construct its list of reference identifiers might be a URI that a user has typed into an interface (e.g., an HTTPS URL for a website), configured account information (e.g., the domain name of a particular host or URI used for retrieving information or connecting to a network, which might be different from the DNS domain name portion of a username), a hyperlink in a web page that triggers a browser to retrieve a media object or script, or some other combination of information that can yield a source domain and an application service type.

The client might need to extract the source domain and application service type from the input(s) it has received. The extracted data MUST include only information that can be securely parsed out of the inputs (e.g., parsing the fully qualified DNS domain name out of the "host" component (or its equivalent) of a URI or deriving the application service type from the scheme of a URI) or information that is derived in a manner not subject to subversion by network attackers (e.g., pulling the data from a delegated domain that is explicitly established via client or system configuration, resolving

the data via [DNSSEC], or obtaining the data from a third-party domain mapping service in which a human user has explicitly placed trust and with which the client communicates over a connection or association that provides both mutual authentication and integrity checking). These considerations apply only to extraction of the source domain from the inputs; naturally, if the inputs themselves are invalid or corrupt (e.g., a user has clicked a link provided by a malicious entity in a phishing attack), then the client might end up communicating with an unexpected application service.

> Example: Given an input URI of <sips:alice@example.net>, a client would derive the application service type "sip" from the "scheme" and parse the domain name "example.net" from the "host" component (or its equivalent).

Each reference identifier in the list SHOULD be based on the source domain and SHOULD NOT be based on a derived domain (e.g., a host name or domain name discovered through DNS resolution of the source domain). This rule is important because only a match between the user inputs and a presented identifier enables the client to be sure that the certificate can legitimately be used to secure the client's communication with the server. There is only one scenario in which it is acceptable for an interactive client to override the recommendation in this rule and therefore communicate with a domain name other than the source domain: because a human user has "pinned" the application service's certificate to the alternative domain name as further discussed under Section 6.6.4 and Section 7.1. In this case, the inputs used by the client to construct its list of reference identifiers might include more than one fully qualified DNS domain name, i.e., both (a) the source domain and (b) the alternative domain contained in the pinned certificate.

Using the combination of fully qualified DNS domain name(s) and application service type, the client constructs a list of reference identifiers in accordance with the following rules:

- The list SHOULD include a DNS-ID. A reference identifier of type DNS-ID can be directly constructed from a fully qualified DNS domain name that is (a) contained in or securely derived from the inputs (i.e., the source domain), or (b) explicitly associated with the source domain by means of user configuration (i.e., a derived domain).

- If a server for the application service type is typically discovered by means of DNS SRV records, then the list SHOULD include an SRV-ID.

- If a server for the application service type is typically associated with a URI for security purposes (i.e., a formal protocol document specifies the use of URIs in server certificates), then the list SHOULD include a URI-ID.

- The list MAY include a CN-ID, mainly for the sake of backward compatibility with deployed infrastructure.

Which identifier types a client includes in its list of reference identifiers is a matter of local policy. For example, in certain deployment environments, a client that is built to connect only to a particular kind of service (e.g., only IM services) might be configured to accept as valid only certificates that include an SRV-ID for that application service type; in this case, the client would include only SRV-IDs matching the application service type in its list of reference identifiers (not, for example, DNS-IDs). By contrast, a more lenient client (even one built to connect only to a particular kind of service) might include both SRV-IDs and DNS-IDs in its list of reference identifiers.

> Implementation Note: It is highly likely that implementers of client software will need to support CN-IDs for the foreseeable future, because certificates containing CN-IDs are so widely deployed. Implementers are advised to monitor the state of the art with regard to certificate issuance policies and migrate away from support CN-IDs in the future if possible.
>
> Implementation Note: The client does not need to construct the foregoing identifiers in the actual formats found in a certificate (e.g., as ASN.1 types); it only needs to construct the functional equivalent of such identifiers for matching purposes.
>
> Security Warning: A client MUST NOT construct a reference identifier corresponding to Relative Distinguished Names (RDNs) other than those of type Common Name and MUST NOT check for RDNs other than those of type Common Name in the presented identifiers.

6.2.2. Examples

A web browser that is connecting via HTTPS to the website at "www.example.com" might have two reference identifiers: a DNS-ID of "www.example.com" and, as a fallback, a CN-ID of "www.example.com".

A mail user agent that is connecting via IMAPS to the email service at "example.net" (resolved as "mail.example.net") might have five reference identifiers: an SRV-ID of "_imaps.example.net" (see [EMAIL-SRV]), DNS-IDs of "example.net" and "mail.example.net", and, as a fallback, CN-IDs of "example.net" and "mail.example.net". (A

legacy email user agent would not support [EMAIL-SRV] and therefore would probably be explicitly configured to connect to "mail.example.net", whereas an SRV-aware user agent would derive "example.net" from an email address of the form "user@example.net" but might also accept "mail.example.net" as the DNS domain name portion of reference identifiers for the service.)

A voice-over-IP (VoIP) user agent that is connecting via SIP to the voice service at "voice.example.edu" might have only one reference identifier: a URI-ID of "sip:voice.example.edu" (see [SIP-CERTS]).

An instant messaging (IM) client that is connecting via XMPP to the IM service at "im.example.org" might have three reference identifiers: an SRV-ID of "_xmpp-client.im.example.org" (see [XMPP]), a DNS-ID of "im.example.org", and an XMPP-specific "XmppAddr" of "im.example.org" (see [XMPP]).

6.3. Preparing to Seek a Match

Once the client has constructed its list of reference identifiers and has received the server's presented identifiers in the form of a PKIX certificate, the client checks its reference identifiers against the presented identifiers for the purpose of finding a match. The search fails if the client exhausts its list of reference identifiers without finding a match. The search succeeds if any presented identifier matches one of the reference identifiers, at which point the client SHOULD stop the search.

> Implementation Note: A client might be configured to perform multiple searches, i.e., to match more than one reference identifier. Although such behavior is not forbidden by this specification, rules for matching multiple reference identifiers are a matter for implementation or future specification.
>
> Security Warning: A client MUST NOT seek a match for a reference identifier of CN-ID if the presented identifiers include a DNS-ID, SRV-ID, URI-ID, or any application-specific identifier types supported by the client.

Before applying the comparison rules provided in the following sections, the client might need to split the reference identifier into its DNS domain name portion and its application service type portion, as follows:

- A reference identifier of type DNS-ID does not include an application service type portion and thus can be used directly as the DNS domain name for comparison purposes. As an example, a

DNS-ID of "www.example.com" would result in a DNS domain name portion of "www.example.com".

- o A reference identifier of type CN-ID also does not include an application service type portion and thus can be used directly as the DNS domain name for comparison purposes. As previously mentioned, this document specifies that a CN-ID always contains a string whose form matches that of a DNS domain name (thus differentiating a CN-ID from a Common Name containing a human-friendly name).

- o For a reference identifier of type SRV-ID, the DNS domain name portion is the Name and the application service type portion is the Service. As an example, an SRV-ID of "_imaps.example.net" would be split into a DNS domain name portion of "example.net" and an application service type portion of "imaps" (mapping to an application protocol of IMAP as explained in [EMAIL-SRV]).

- o For a reference identifier of type URI-ID, the DNS domain name portion is the "reg-name" part of the "host" component (or its equivalent) and the application service type portion is the application service type associated with the scheme name matching the [ABNF] "scheme" rule from [URI] (not including the ':' separator). As previously mentioned, this document specifies that a URI-ID always contains a "host" component (or its equivalent) containing a "reg-name". (Matching only the "reg-name" rule from [URI] limits verification to DNS domain names, thereby differentiating a URI-ID from a uniformResourceIdentifier entry that contains an IP address or a mere host name, or that does not contain a "host" component at all.) Furthermore, note that extraction of the "reg-name" might necessitate normalization of the URI (as explained in [URI]). As an example, a URI-ID of "sip: voice.example.edu" would be split into a DNS domain name portion of "voice.example.edu" and an application service type of "sip" (associated with an application protocol of SIP as explained in [SIP-CERTS]).

Detailed comparison rules for matching the DNS domain name portion and application service type portion of the reference identifier are provided in the following sections.

6.4. Matching the DNS Domain Name Portion

The client MUST match the DNS domain name portion of a reference identifier according to the following rules (and SHOULD also check the application service type as described under Section 6.5). The rules differ depending on whether the domain to be checked is a "traditional domain name" or an "internationalized domain name" (as

defined under Section 2.2). Furthermore, to meet the needs of clients that support presented identifiers containing the wildcard character '*', we define a supplemental rule for so-called "wildcard certificates". Finally, we also specify the circumstances under which it is acceptable to check the "CN-ID" identifier type.

6.4.1. Checking of Traditional Domain Names

If the DNS domain name portion of a reference identifier is a "traditional domain name", then matching of the reference identifier against the presented identifier is performed by comparing the set of domain name labels using a case-insensitive ASCII comparison, as clarified by [DNS-CASE] (e.g., "WWW.Example.Com" would be lower-cased to "www.example.com" for comparison purposes). Each label MUST match in order for the names to be considered to match, except as supplemented by the rule about checking of wildcard labels (Section 6.4.3).

6.4.2. Checking of Internationalized Domain Names

If the DNS domain name portion of a reference identifier is an internationalized domain name, then an implementation MUST convert any U-labels [IDNA-DEFS] in the domain name to A-labels before checking the domain name. In accordance with [IDNA-PROTO], A-labels MUST be compared as case-insensitive ASCII. Each label MUST match in order for the domain names to be considered to match, except as supplemented by the rule about checking of wildcard labels (Section 6.4.3; but see also Section 7.2 regarding wildcards in internationalized domain names).

6.4.3. Checking of Wildcard Certificates

A client employing this specification's rules MAY match the reference identifier against a presented identifier whose DNS domain name portion contains the wildcard character '*' as part or all of a label (following the description of labels and domain names in [DNS-CONCEPTS]).

For information regarding the security characteristics of wildcard certificates, see Section 7.2.

If a client matches the reference identifier against a presented identifier whose DNS domain name portion contains the wildcard character '*', the following rules apply:

1. The client SHOULD NOT attempt to match a presented identifier in which the wildcard character comprises a label other than the left-most label (e.g., do not match bar.*.example.net).

2. If the wildcard character is the only character of the left-most label in the presented identifier, the client SHOULD NOT compare against anything but the left-most label of the reference identifier (e.g., *.example.com would match foo.example.com but not bar.foo.example.com or example.com).

3. The client MAY match a presented identifier in which the wildcard character is not the only character of the label (e.g., baz*.example.net and *baz.example.net and b*z.example.net would be taken to match baz1.example.net and foobaz.example.net and buzz.example.net, respectively). However, the client SHOULD NOT attempt to match a presented identifier where the wildcard character is embedded within an A-label or U-label [IDNA-DEFS] of an internationalized domain name [IDNA-PROTO].

6.4.4. Checking of Common Names

As noted, a client MUST NOT seek a match for a reference identifier of CN-ID if the presented identifiers include a DNS-ID, SRV-ID, URI-ID, or any application-specific identifier types supported by the client.

Therefore, if and only if the presented identifiers do not include a DNS-ID, SRV-ID, URI-ID, or any application-specific identifier types supported by the client, then the client MAY as a last resort check for a string whose form matches that of a fully qualified DNS domain name in a Common Name field of the subject field (i.e., a CN-ID). If the client chooses to compare a reference identifier of type CN-ID against that string, it MUST follow the comparison rules for the DNS domain name portion of an identifier of type DNS-ID, SRV-ID, or URI-ID, as described under Section 6.4.1, Section 6.4.2, and Section 6.4.3.

6.5. Matching the Application Service Type Portion

When a client checks identifiers of type SRV-ID and URI-ID, it MUST check not only the DNS domain name portion of the identifier but also the application service type portion. The client does this by splitting the identifier into the DNS domain name portion and the application service type portion (as described under Section 6.3), then checking both the DNS domain name portion (as described under Section 6.4) and the application service type portion as described in the following subsections.

> Implementation Note: An identifier of type SRV-ID or URI-ID provides an application service type portion to be checked, but that portion is combined only with the DNS domain name portion of the SRV-ID or URI-ID itself. For example, if a client's list of

reference identifiers includes an SRV-ID of "_xmpp-client.im.example.org" and a DNS-ID of "apps.example.net", the client would check (a) the combination of an application service type of "xmpp-client" and a DNS domain name of "im.example.org" and (b) a DNS domain name of "apps.example.net". However, the client would not check (c) the combination of an application service type of "xmpp-client" and a DNS domain name of "apps.example.net" because it does not have an SRV-ID of "_xmpp-client.apps.example.net" in its list of reference identifiers.

6.5.1. SRV-ID

The application service name portion of an SRV-ID (e.g., "imaps") MUST be matched in a case-insensitive manner, in accordance with [DNS-SRV]. Note that the "_" character is prepended to the service identifier in DNS SRV records and in SRV-IDs (per [SRVNAME]), and thus does not need to be included in any comparison.

6.5.2. URI-ID

The scheme name portion of a URI-ID (e.g., "sip") MUST be matched in a case-insensitive manner, in accordance with [URI]. Note that the ":" character is a separator between the scheme name and the rest of the URI, and thus does not need to be included in any comparison.

6.6. Outcome

The outcome of the matching procedure is one of the following cases.

6.6.1. Case #1: Match Found

If the client has found a presented identifier that matches a reference identifier, then the service identity check has succeeded. In this case, the client MUST use the matched reference identifier as the validated identity of the application service.

6.6.2. Case #2: No Match Found, Pinned Certificate

If the client does not find a presented identifier matching any of the reference identifiers but the client has previously pinned the application service's certificate to one of the reference identifiers in the list it constructed for this communication attempt (as "pinning" is explained under Section 1.8), and the presented certificate matches the pinned certificate (including the context as described under Section 7.1), then the service identity check has succeeded.

6.6.3. Case #3: No Match Found, No Pinned Certificate

If the client does not find a presented identifier matching any of the reference identifiers and the client has not previously pinned the certificate to one of the reference identifiers in the list it constructed for this communication attempt, then the client MUST proceed as described under Section 6.6.4.

6.6.4. Fallback

If the client is an interactive client that is directly controlled by a human user, then it SHOULD inform the user of the identity mismatch and automatically terminate the communication attempt with a bad certificate error; this behavior is preferable because it prevents users from inadvertently bypassing security protections in hostile situations.

> Security Warning: Some interactive clients give advanced users the option of proceeding with acceptance despite the identity mismatch, thereby "pinning" the certificate to one of the reference identifiers in the list constructed by the client for this communication attempt. Although this behavior can be appropriate in certain specialized circumstances, in general it ought to be exposed only to advanced users. Even then it needs to be handled with extreme caution, for example by first encouraging even an advanced user to terminate the communication attempt and, if the advanced user chooses to proceed anyway, by forcing the user to view the entire certification path and only then allowing the user to pin the certificate (on a temporary or permanent basis, at the user’s option).

Otherwise, if the client is an automated application not directly controlled by a human user, then it SHOULD terminate the communication attempt with a bad certificate error and log the error appropriately. An automated application MAY provide a configuration setting that disables this behavior, but MUST enable the behavior by default.

7. Security Considerations

7.1. Pinned Certificates

As defined under Section 1.8, a certificate is said to be "pinned" to a DNS domain name when a user has explicitly chosen to associate a service’s certificate with that DNS domain name despite the fact that the certificate contains some other DNS domain name (e.g., the user has explicitly approved "apps.example.net" as a domain associated with a source domain of "example.com"). The cached name association

MUST take account of both the certificate presented and the context in which it was accepted or configured (where the "context" includes the chain of certificates from the presented certificate to the trust anchor, the source domain, the application service type, the service's derived domain and port number, and any other relevant information provided by the user or associated by the client).

7.2. Wildcard Certificates

This document states that the wildcard character '*' SHOULD NOT be included in presented identifiers but MAY be checked by application clients (mainly for the sake of backward compatibility with deployed infrastructure). As a result, the rules provided in this document are more restrictive than the rules for many existing application technologies (such as those excerpted under Appendix B). Several security considerations justify tightening the rules:

- Wildcard certificates automatically vouch for any and all host names within their domain. This can be convenient for administrators but also poses the risk of vouching for rogue or buggy hosts. See for example [Defeating-SSL] (beginning at slide 91) and [HTTPSbytes] (slides 38-40).

- Specifications for existing application technologies are not clear or consistent about the allowable location of the wildcard character, such as whether it can be:

  * only the complete left-most label (e.g., *.example.com)

  * some fragment of the left-most label (e.g., fo*.example.com, f*o.example.com, or *oo.example.com)

  * all or part of a label other than the left-most label (e.g., www.*.example.com or www.foo*.example.com)

  * all or part of a label that identifies a so-called "public suffix" (e.g., *.co.uk or *.com)

  * included more than once in a given label (e.g., f*b*r.example.com

  * included as all or part of more than one label (e.g., *.*.example.com)

  These ambiguities might introduce exploitable differences in identity checking behavior among client implementations and necessitate overly complex and inefficient identity checking algorithms.

- o There is no specification that defines how the wildcard character may be embedded within the A-labels or U-labels [IDNA-DEFS] of an internationalized domain name [IDNA-PROTO]; as a result, implementations are strongly discouraged from including or attempting to check for the wildcard character embedded within the A-labels or U-labels of an internationalized domain name (e.g., "xn--kcry6tjko*.example.org"). Note, however, that a presented domain name identifier MAY contain the wildcard character as long as that character occupies the entire left-most label position, where all of the remaining labels are valid NR-LDH labels, A-labels, or U-labels (e.g., "*.xn--kcry6tjko.example.org").

Notwithstanding the foregoing security considerations, specifications that reuse this one can legitimately encourage continued support for the wildcard character if they have good reasons to do so, such as backward compatibility with deployed infrastructure (see, for example, [EV-CERTS]).

7.3. Internationalized Domain Names

Allowing internationalized domain names can lead to the inclusion of visually similar (so-called "confusable") characters in certificates; for discussion, see for example [IDNA-DEFS].

7.4. Multiple Identifiers

A given application service might be addressed by multiple DNS domain names for a variety of reasons, and a given deployment might service multiple domains (e.g., in so-called "virtual hosting" environments). In the default TLS handshake exchange, the client is not able to indicate the DNS domain name with which it wants to communicate, and the TLS server returns only one certificate for itself. Absent an extension to TLS, a typical workaround used to facilitate mapping an application service to multiple DNS domain names is to embed all of the domain names into a single certificate.

A more recent approach, formally specified in [TLS-EXT], is for the client to use the TLS "Server Name Indication" (SNI) extension when sending the client_hello message, stipulating the DNS domain name it desires or expects of the service. The service can then return the appropriate certificate in its Certificate message, and that certificate can represent a single DNS domain name.

To accommodate the workaround that was needed before the development of the SNI extension, this specification allows multiple DNS-IDs, SRV-IDs, or URI-IDs in a certificate; however, it explicitly discourages multiple CN-IDs. Although it would be preferable to forbid multiple CN-IDs entirely, there are several reasons at this

time why this specification states that they SHOULD NOT (instead of MUST NOT) be included:

- At least one significant technology community of interest explicitly allows multiple CN-IDs [EV-CERTS].
- At least one significant certification authority is known to issue certificates containing multiple CN-IDs.
- Many service providers often deem inclusion of multiple CN-IDs necessary in virtual hosting environments because at least one widely deployed operating system does not yet support the SNI extension.

It is hoped that the recommendation regarding multiple CN-IDs can be further tightened in the future.

8. Contributors

The following individuals made important contributions to the text of this document: Shumon Huque, RL 'Bob' Morgan, and Kurt Zeilenga.

9. Acknowledgements

The editors and contributors wish to thank the following individuals for their feedback and suggestions: Bernard Aboba, Richard Barnes, Uri Blumenthal, Nelson Bolyard, Kaspar Brand, Anthony Bryan, Scott Cantor, Wan-Teh Chang, Bil Corry, Dave Cridland, Dave Crocker, Cyrus Daboo, Charles Gardiner, Philip Guenther, Phillip Hallam-Baker, Bruno Harbulot, Wes Hardaker, David Harrington, Paul Hoffman, Love Hornquist Astrand, Henry Hotz, Russ Housley, Jeffrey Hutzelman, Cullen Jennings, Simon Josefsson, Geoff Keating, John Klensin, Scott Lawrence, Matt McCutchen, Alexey Melnikov, Subramanian Moonesamy, Eddy Nigg, Ludwig Nussel, Joe Orton, Tom Petch, Yngve N. Pettersen, Tim Polk, Robert Relyea, Eric Rescorla, Pete Resnick, Martin Rex, Joe Salowey, Stefan Santesson, Jim Schaad, Rob Stradling, Michael Stroeder, Andrew Sullivan, Peter Sylvester, Martin Thomson, Paul Tiemann, Sean Turner, Nicolas Williams, Dan Wing, Dan Winship, and Stefan Winter.

Thanks also to Barry Leiba and Ben Campbell for their reviews on behalf of the Security Directorate and the General Area Review Team, respectively.

The responsible Area Director was Alexey Melnikov.

10. References

10.1. Normative References

[DNS-CONCEPTS] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[DNS-SRV] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

[IDNA-DEFS] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010.

[IDNA-PROTO] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", RFC 5891, August 2010.

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

[LDAP-DN] Zeilenga, K., Ed., "Lightweight Directory Access Protocol (LDAP): String Representation of Distinguished Names", RFC 4514, June 2006.

[PKIX] 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.

[SRVNAME] Santesson, S., "Internet X.509 Public Key Infrastructure Subject Alternative Name for Expression of Service Name", RFC 4985, August 2007.

[URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

10.2. Informative References

[ABNF] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

[DNS-CASE] Eastlake 3rd, D., "Domain Name System (DNS) Case Insensitivity Clarification", RFC 4343, January 2006.

[DNSSEC] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[DTLS] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, April 2006.

[Defeating-SSL] Marlinspike, M., "New Tricks for Defeating SSL in Practice", BlackHat DC, February 2009, <http://www.blackhat.com/presentations/bh-dc-09/Marlinspike/ BlackHat-DC-09-Marlinspike-Defeating-SSL.pdf>.

[EMAIL-SRV] Daboo, C., "Use of SRV Records for Locating Email Submission/Access Services", RFC 6186, March 2011.

[EV-CERTS] CA/Browser Forum, "Guidelines For The Issuance And Management Of Extended Validation Certificates", October 2009, <http://www.cabforum.org/Guidelines_v1_2.pdf>.

[GIST] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.

[HTTP] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[HTTP-TLS] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.

[HTTPSbytes] Sokol, J. and R. Hansen, "HTTPS Can Byte Me", BlackHat Abu Dhabi, November 2010, <https://media.blackhat.com/bh-ad-10/Hansen/Blackhat-AD-2010-Hansen-Sokol-HTTPS-Can-Byte-Me-slides.pdf>.

[IDNA2003] Faltstrom, P., Hoffman, P., and A. Costello, "Internationalizing Domain Names in Applications (IDNA)", RFC 3490, March 2003.

[IMAP] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, March 2003.

[IP] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[IPSEC] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[LDAP] Sermersheim, J., "Lightweight Directory Access Protocol (LDAP): The Protocol", RFC 4511, June 2006.

[LDAP-AUTH] Harrison, R., "Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms", RFC 4513, June 2006.

[LDAP-SCHEMA] Sciberras, A., Ed., "Lightweight Directory Access Protocol (LDAP): Schema for User Applications", RFC 4519, June 2006.

[LDAP-TLS] Hodges, J., Morgan, R., and M. Wahl, "Lightweight Directory Access Protocol (v3): Extension for Transport Layer Security", RFC 2830, May 2000.

[NAPTR] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", RFC 3403, October 2002.

[NETCONF] Enns, R., Ed., "NETCONF Configuration Protocol", RFC 4741, December 2006.

[NETCONF-SSH] Wasserman, M. and T. Goddard, "Using the NETCONF Configuration Protocol over Secure SHell (SSH)", RFC 4742, December 2006.

[NETCONF-TLS] Badra, M., "NETCONF over Transport Layer Security (TLS)", RFC 5539, May 2009.

[NNTP] Feather, C., "Network News Transfer Protocol (NNTP)", RFC 3977, October 2006.

[NNTP-TLS] Murchison, K., Vinocur, J., and C. Newman, "Using Transport Layer Security (TLS) with Network News Transfer Protocol (NNTP)", RFC 4642, October 2006.

[OCSP] 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.

[OPENPGP] Callas, J., Donnerhacke, L., Finney, H., Shaw, D., and R. Thayer, "OpenPGP Message Format", RFC 4880, November 2007.

[PKIX-OLD] Housley, R., Ford, W., Polk, T., and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and CRL Profile", RFC 2459, January 1999.

[POP3] Myers, J. and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, May 1996.

[PRIVATE] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[S-NAPTR] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", RFC 3958, January 2005.

[SECTERMS] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.

[SIP] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[SIP-CERTS] Gurbani, V., Lawrence, S., and A. Jeffrey, "Domain Certificates in the Session Initiation Protocol (SIP)", RFC 5922, June 2010.

[SIP-SIPS] Audet, F., "The Use of the SIPS URI Scheme in the Session Initiation Protocol (SIP)", RFC 5630, October 2009.

[SMTP] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.

[SMTP-AUTH] Siemborski, R., Ed. and A. Melnikov, Ed., "SMTP Service Extension for Authentication", RFC 4954, July 2007.

[SMTP-TLS] Hoffman, P., "SMTP Service Extension for Secure SMTP over Transport Layer Security", RFC 3207, February 2002.

[SNMP] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.

[SNMP-TLS] Hardaker, W., "Transport Layer Security (TLS) Transport Model for the Simple Network Management Protocol (SNMP)", RFC 5953, August 2010.

[SYSLOG] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.

[SYSLOG-DTLS] Salowey, J., Petch, T., Gerhards, R., and H. Feng, "Datagram Transport Layer Security (DTLS) Transport Mapping for Syslog", RFC 6012, October 2010.

[SYSLOG-TLS] Miao, F., Ed., Ma, Y., Ed., and J. Salowey, Ed., "Transport Layer Security (TLS) Transport Mapping for Syslog", RFC 5425, March 2009.

[TLS] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[TLS-EXT] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.

[US-ASCII] American National Standards Institute, "Coded Character Set - 7-bit American Standard Code for Information Interchange", ANSI X3.4, 1986.

[USINGTLS] Newman, C., "Using TLS with IMAP, POP3 and ACAP", RFC 2595, June 1999.

[WSC-UI] Saldhana, A. and T. Roessler, "Web Security Context: User Interface Guidelines", World Wide Web Consortium LastCall WD-wsc-ui-20100309, March 2010, <http://www.w3.org/TR/2010/WD-wsc-ui-20100309>.

[X.500] International Telecommunications Union, "Information Technology - Open Systems Interconnection - The Directory: Overview of concepts, models and services", ITU-T Recommendation X.500, ISO Standard 9594-1, August 2005.

[X.501] International Telecommunications Union, "Information Technology - Open Systems Interconnection - The Directory: Models", ITU-T Recommendation X.501, ISO Standard 9594-2, August 2005.

[X.509] International Telecommunications Union, "Information Technology - Open Systems Interconnection - The Directory: Public-key and attribute certificate frameworks", ITU-T Recommendation X.509, ISO Standard 9594-8, August 2005.

[X.520] International Telecommunications Union, "Information Technology - Open Systems Interconnection - The Directory: Selected attribute types", ITU-T Recommendation X.509, ISO Standard 9594-6, August 2005.

[X.690] International Telecommunications Union, "Information Technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, ISO Standard 8825-1, August 2008.

[XMPP] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, March 2011.

[XMPP-OLD] Saint-Andre, P., Ed., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 3920, October 2004.

Appendix A. Sample Text

At the time of this writing, two application technologies reuse the recommendations in this specification: email [EMAIL-SRV] and XMPP [XMPP]. Here we include the text from [XMPP] to illustrate the thought process that might be followed by protocol designers for other application technologies. Specifically, because XMPP uses DNS SRV records for resolution of the DNS domain names for application services, the XMPP specification recommends the use of SRV-IDs.

The text regarding certificate issuance is as follows:

######

In a PKIX certificate to be presented by an XMPP server (i.e., a "server certificate"), the certificate MUST include one or more XMPP addresses (i.e., domainparts) associated with XMPP services hosted at the server. The rules and guidelines defined in [this specification] apply to XMPP server certificates, with the following XMPP-specific considerations:

- Support for the DNS-ID identifier type [PKIX] is REQUIRED in XMPP client and server software implementations. Certification authorities that issue XMPP-specific certificates MUST support the DNS-ID identifier type. XMPP service providers SHOULD include the DNS-ID identifier type in certificate requests.

- Support for the SRV-ID identifier type [SRVNAME] is REQUIRED for XMPP client and server software implementations (for verification purposes XMPP client implementations need to support only the "_xmpp-client" application service type, whereas XMPP server implementations need to support both the "_xmpp-client" and "_xmpp-server" application service types). Certification authorities that issue XMPP-specific certificates SHOULD support the SRV-ID identifier type. XMPP service providers SHOULD include the SRV-ID identifier type in certificate requests.

- Support for the XmppAddr identifier type is encouraged in XMPP client and server software implementations for the sake of backward-compatibility, but is no longer encouraged in certificates issued by certification authorities or requested by XMPP service providers.

- DNS domain names in server certificates MAY contain the wildcard character '*' as the complete left-most label within the identifier.

######

The text regarding certificate verification is as follows:

######

For server certificates, the rules and guidelines defined in [this specification] apply, with the proviso that the XmppAddr identifier is allowed as a reference identifier.

The identities to be checked are set as follows:

- The initiating entity sets its reference identifier to the 'to' address it communicates in the initial stream header; i.e., this is the identity it expects the receiving entity to provide in a PKIX certificate.

- The receiving entity sets its reference identifier to the 'from' address communicated by the initiating entity in the initial stream header; i.e., this is the identity that the initiating entity is trying to assert.

######

Appendix B. Prior Art

(This section is non-normative.)

The recommendations in this document are an abstraction from recommendations in specifications for a wide range of application protocols. For the purpose of comparison and to delineate the history of thinking about application service identity verification within the IETF, this informative section gathers together prior art by including the exact text from various RFCs (the only modifications are changes to the names of several references to maintain coherence with the main body of this document, and the elision of irrelevant text as marked by the characters "[...]").

B.1. IMAP, POP3, and ACAP (1999)

In 1999, [USINGTLS] specified the following text regarding application service identity verification in IMAP, POP3, and ACAP:

######

2.4. Server Identity Check

During the TLS negotiation, the client MUST check its understanding of the server hostname against the server's identity as presented in the server Certificate message, in order to prevent man-in-the-middle attacks. Matching is performed according to these rules:

- The client MUST use the server hostname it used to open the connection as the value to compare against the server name as expressed in the server certificate. The client MUST NOT use any form of the server hostname derived from an insecure remote source (e.g., insecure DNS lookup). CNAME canonicalization is not done.

- If a subjectAltName extension of type dNSName is present in the certificate, it SHOULD be used as the source of the server's identity.

- Matching is case-insensitive.

- A "*" wildcard character MAY be used as the left-most name component in the certificate. For example, *.example.com would match a.example.com, foo.example.com, etc. but would not match example.com.

- If the certificate contains multiple names (e.g. more than one dNSName field), then a match with any one of the fields is considered acceptable.

If the match fails, the client SHOULD either ask for explicit user confirmation, or terminate the connection and indicate the server's identity is suspect.

######

B.2. HTTP (2000)

In 2000, [HTTP-TLS] specified the following text regarding application service identity verification in HTTP:

######

3.1. Server Identity

In general, HTTP/TLS requests are generated by dereferencing a URI. As a consequence, the hostname for the server is known to the client. If the hostname is available, the client MUST check it against the server's identity as presented in the server's Certificate message, in order to prevent man-in-the-middle attacks.

If the client has external information as to the expected identity of the server, the hostname check MAY be omitted. (For instance, a client may be connecting to a machine whose address and hostname are dynamic but the client knows the certificate that the server will present.) In such cases, it is important to narrow the scope of acceptable certificates as much as possible in order to prevent man in the middle attacks. In special cases, it may be appropriate for the client to simply ignore the server's identity, but it must be understood that this leaves the connection open to active attack.

If a subjectAltName extension of type dNSName is present, that MUST be used as the identity. Otherwise, the (most specific) Common Name field in the Subject field of the certificate MUST be used. Although the use of the Common Name is existing practice, it is deprecated and Certification Authorities are encouraged to use the dNSName instead.

Matching is performed using the matching rules specified by [PKIX-OLD]. If more than one identity of a given type is present in the certificate (e.g., more than one dNSName name, a match in any one of the set is considered acceptable.) Names may contain the wildcard character * which is considered to match any single domain name component or component fragment. E.g., *.a.com matches foo.a.com but not bar.foo.a.com. f*.com matches foo.com but not bar.com.

In some cases, the URI is specified as an IP address rather than a hostname. In this case, the iPAddress subjectAltName must be present in the certificate and must exactly match the IP in the URI.

If the hostname does not match the identity in the certificate, user oriented clients MUST either notify the user (clients MAY give the user the opportunity to continue with the connection in any case) or terminate the connection with a bad certificate error. Automated clients MUST log the error to an appropriate audit log (if available) and SHOULD terminate the connection (with a bad certificate error). Automated clients MAY provide a configuration setting that disables this check, but MUST provide a setting which enables it.

Note that in many cases the URI itself comes from an untrusted source. The above-described check provides no protection against attacks where this source is compromised. For example, if the URI was obtained by clicking on an HTML page which was itself obtained without using HTTP/TLS, a man in the middle could have replaced the URI. In order to prevent this form of attack, users should carefully examine the certificate presented by the server to determine if it meets their expectations.

######

B.3. LDAP (2000/2006)

In 2000, [LDAP-TLS] specified the following text regarding application service identity verification in LDAP:

######

3.6. Server Identity Check

The client MUST check its understanding of the server's hostname against the server's identity as presented in the server's Certificate message, in order to prevent man-in-the-middle attacks.

Matching is performed according to these rules:

- The client MUST use the server hostname it used to open the LDAP connection as the value to compare against the server name as expressed in the server's certificate. The client MUST NOT use the server's canonical DNS name or any other derived form of name.

- If a subjectAltName extension of type dNSName is present in the certificate, it SHOULD be used as the source of the server's identity.

- Matching is case-insensitive.

- The "*" wildcard character is allowed. If present, it applies only to the left-most name component.

E.g. *.bar.com would match a.bar.com, b.bar.com, etc. but not bar.com. If more than one identity of a given type is present in the certificate (e.g. more than one dNSName name), a match in any one of the set is considered acceptable.

If the hostname does not match the dNSName-based identity in the certificate per the above check, user-oriented clients SHOULD either notify the user (clients MAY give the user the opportunity to continue with the connection in any case) or terminate the connection and indicate that the server's identity is suspect. Automated clients SHOULD close the connection, returning and/or logging an error indicating that the server's identity is suspect.

Beyond the server identity checks described in this section, clients SHOULD be prepared to do further checking to ensure that the server is authorized to provide the service it is observed to provide. The client MAY need to make use of local policy information.

######

In 2006, [LDAP-AUTH] specified the following text regarding application service identity verification in LDAP:

######

3.1.3. Server Identity Check

In order to prevent man-in-the-middle attacks, the client MUST verify the server's identity (as presented in the server's Certificate message). In this section, the client's understanding of the server's identity (typically the identity used to establish the transport connection) is called the "reference identity".

The client determines the type (e.g., DNS name or IP address) of the reference identity and performs a comparison between the reference identity and each subjectAltName value of the corresponding type until a match is produced. Once a match is produced, the server's identity has been verified, and the server identity check is complete. Different subjectAltName types are matched in different ways. Sections 3.1.3.1 - 3.1.3.3 explain how to compare values of various subjectAltName types.

The client may map the reference identity to a different type prior to performing a comparison. Mappings may be performed for all available subjectAltName types to which the reference identity can be mapped; however, the reference identity should only be mapped to types for which the mapping is either inherently secure (e.g., extracting the DNS name from a URI to compare with a subjectAltName

of type dNSName) or for which the mapping is performed in a secure manner (e.g., using [DNSSEC], or using user- or admin-configured host-to-address/address-to-host lookup tables).

The server's identity may also be verified by comparing the reference identity to the Common Name (CN) [LDAP-SCHEMA] value in the last Relative Distinguished Name (RDN) of the subject field of the server's certificate (where "last" refers to the DER-encoded order, not the order of presentation in a string representation of DER-encoded data). This comparison is performed using the rules for comparison of DNS names in Section 3.1.3.1, below, with the exception that no wildcard matching is allowed. Although the use of the Common Name value is existing practice, it is deprecated, and Certification Authorities are encouraged to provide subjectAltName values instead. Note that the TLS implementation may represent DNs in certificates according to X.500 or other conventions. For example, some X.500 implementations order the RDNs in a DN using a left-to-right (most significant to least significant) convention instead of LDAP's right-to-left convention.

If the server identity check fails, user-oriented clients SHOULD either notify the user (clients may give the user the opportunity to continue with the LDAP session in this case) or close the transport connection and indicate that the server's identity is suspect. Automated clients SHOULD close the transport connection and then return or log an error indicating that the server's identity is suspect or both.

Beyond the server identity check described in this section, clients should be prepared to do further checking to ensure that the server is authorized to provide the service it is requested to provide. The client may need to make use of local policy information in making this determination.

3.1.3.1. Comparison of DNS Names

If the reference identity is an internationalized domain name, conforming implementations MUST convert it to the ASCII Compatible Encoding (ACE) format as specified in Section 4 of RFC 3490 [IDNA2003] before comparison with subjectAltName values of type dNSName. Specifically, conforming implementations MUST perform the conversion operation specified in Section 4 of RFC 3490 as follows:

- in step 1, the domain name SHALL be considered a "stored string";
- in step 3, set the flag called "UseSTD3ASCIIRules";
- in step 4, process each label with the "ToASCII" operation; and

o in step 5, change all label separators to U+002E (full stop).

After performing the "to-ASCII" conversion, the DNS labels and names MUST be compared for equality according to the rules specified in Section 3 of RFC3490.

The '*' (ASCII 42) wildcard character is allowed in subjectAltName values of type dNSName, and then only as the left-most (least significant) DNS label in that value. This wildcard matches any left-most DNS label in the server name. That is, the subject *.example.com matches the server names a.example.com and b.example.com, but does not match example.com or a.b.example.com.

3.1.3.2. Comparison of IP Addresses

When the reference identity is an IP address, the identity MUST be converted to the "network byte order" octet string representation [IP] [IPv6]. For IP Version 4, as specified in RFC 791, the octet string will contain exactly four octets. For IP Version 6, as specified in RFC 2460, the octet string will contain exactly sixteen octets. This octet string is then compared against subjectAltName values of type iPAddress. A match occurs if the reference identity octet string and value octet strings are identical.

3.1.3.3. Comparison of Other subjectName Types

Client implementations MAY support matching against subjectAltName values of other types as described in other documents.

######

B.4. SMTP (2002/2007)

In 2002, [SMTP-TLS] specified the following text regarding application service identity verification in SMTP:

######

4.1 Processing After the STARTTLS Command

[...]

The decision of whether or not to believe the authenticity of the other party in a TLS negotiation is a local matter. However, some general rules for the decisions are:

o A SMTP client would probably only want to authenticate an SMTP server whose server certificate has a domain name that is the domain name that the client thought it was connecting to.

[...]

######

In 2006, [SMTP-AUTH] specified the following text regarding application service identity verification in SMTP:

######

14. Additional Requirements When Using SASL PLAIN over TLS

[...]

After a successful [TLS] negotiation, the client MUST check its understanding of the server hostname against the server's identity as presented in the server Certificate message, in order to prevent man-in-the-middle attacks. If the match fails, the client MUST NOT attempt to authenticate using the SASL PLAIN mechanism. Matching is performed according to the following rules:

The client MUST use the server hostname it used to open the connection as the value to compare against the server name as expressed in the server certificate. The client MUST NOT use any form of the server hostname derived from an insecure remote source (e.g., insecure DNS lookup). CNAME canonicalization is not done.

If a subjectAltName extension of type dNSName is present in the certificate, it SHOULD be used as the source of the server's identity.

Matching is case-insensitive.

A "*" wildcard character MAY be used as the leftmost name component in the certificate. For example, *.example.com would match a.example.com, foo.example.com, etc., but would not match example.com.

If the certificate contains multiple names (e.g., more than one dNSName field), then a match with any one of the fields is considered acceptable.

######

B.5. XMPP (2004)

In 2004, [XMPP-OLD] specified the following text regarding application service identity verification in XMPP:

######

14.2. Certificate Validation

When an XMPP peer communicates with another peer securely, it MUST validate the peer's certificate. There are three possible cases:

Case #1: The peer contains an End Entity certificate which appears to be certified by a certification path terminating in a trust anchor (as described in Section 6.1 of [PKIX]).

Case #2: The peer certificate is certified by a Certificate Authority not known to the validating peer.

Case #3: The peer certificate is self-signed.

In Case #1, the validating peer MUST do one of two things:

1. Verify the peer certificate according to the rules of [PKIX]. The certificate SHOULD then be checked against the expected identity of the peer following the rules described in [HTTP-TLS], except that a subjectAltName extension of type "xmpp" MUST be used as the identity if present. If one of these checks fails, user-oriented clients MUST either notify the user (clients MAY give the user the opportunity to continue with the connection in any case) or terminate the connection with a bad certificate error. Automated clients SHOULD terminate the connection (with a bad certificate error) and log the error to an appropriate audit log. Automated clients MAY provide a configuration setting that disables this check, but MUST provide a setting that enables it.

2. The peer SHOULD show the certificate to a user for approval, including the entire certification path. The peer MUST cache the certificate (or some non-forgeable representation such as a hash). In future connections, the peer MUST verify that the same certificate was presented and MUST notify the user if it has changed.

In Case #2 and Case #3, implementations SHOULD act as in (2) above.

######

Although [XMPP-OLD] defined its own rules, [XMPP] reuses the rules in this document regarding application service identity verification in XMPP.

B.6. NNTP (2006)

In 2006, [NNTP-TLS] specified the following text regarding application service identity verification in NNTP:

######

5. Security Considerations

[...]

During the TLS negotiation, the client MUST check its understanding of the server hostname against the server's identity as presented in the server Certificate message, in order to prevent man-in-the-middle attacks. Matching is performed according to these rules:

- The client MUST use the server hostname it used to open the connection (or the hostname specified in TLS "server_name" extension [TLS]) as the value to compare against the server name as expressed in the server certificate. The client MUST NOT use any form of the server hostname derived from an insecure remote source (e.g., insecure DNS lookup). CNAME canonicalization is not done.

- If a subjectAltName extension of type dNSName is present in the certificate, it SHOULD be used as the source of the server's identity.

- Matching is case-insensitive.

- A "*" wildcard character MAY be used as the left-most name component in the certificate. For example, *.example.com would match a.example.com, foo.example.com, etc., but would not match example.com.

- If the certificate contains multiple names (e.g., more than one dNSName field), then a match with any one of the fields is considered acceptable.

If the match fails, the client SHOULD either ask for explicit user confirmation or terminate the connection with a QUIT command and indicate the server's identity is suspect.

Additionally, clients MUST verify the binding between the identity of the servers to which they connect and the public keys presented by those servers. Clients SHOULD implement the algorithm in Section 6 of [PKIX] for general certificate validation, but MAY supplement that algorithm with other validation methods that achieve equivalent levels of verification (such as comparing the server certificate against a local store of already-verified certificates and identity bindings).

######

B.7. NETCONF (2006/2009)

In 2006, [NETCONF-SSH] specified the following text regarding application service identity verification in NETCONF:

######

6. Security Considerations

The identity of the server MUST be verified and authenticated by the client according to local policy before password-based authentication data or any configuration or state data is sent to or received from the server. The identity of the client MUST also be verified and authenticated by the server according to local policy to ensure that the incoming client request is legitimate before any configuration or state data is sent to or received from the client. Neither side should establish a NETCONF over SSH connection with an unknown, unexpected, or incorrect identity on the opposite side.

######

In 2009, [NETCONF-TLS] specified the following text regarding application service identity verification in NETCONF:

######

3.1. Server Identity

During the TLS negotiation, the client MUST carefully examine the certificate presented by the server to determine if it meets the client's expectations. Particularly, the client MUST check its understanding of the server hostname against the server's identity as presented in the server Certificate message, in order to prevent man-in-the-middle attacks.

Matching is performed according to the rules below (following the example of [NNTP-TLS]):

- The client MUST use the server hostname it used to open the connection (or the hostname specified in the TLS "server_name" extension [TLS]) as the value to compare against the server name as expressed in the server certificate. The client MUST NOT use any form of the server hostname derived from an insecure remote source (e.g., insecure DNS lookup). CNAME canonicalization is not done.

- If a subjectAltName extension of type dNSName is present in the certificate, it MUST be used as the source of the server's identity.

- Matching is case-insensitive.

- A "*" wildcard character MAY be used as the leftmost name component in the certificate. For example, *.example.com would match a.example.com, foo.example.com, etc., but would not match example.com.

- If the certificate contains multiple names (e.g., more than one dNSName field), then a match with any one of the fields is considered acceptable.

If the match fails, the client MUST either ask for explicit user confirmation or terminate the connection and indicate the server's identity is suspect.

Additionally, clients MUST verify the binding between the identity of the servers to which they connect and the public keys presented by those servers. Clients SHOULD implement the algorithm in Section 6 of [PKIX] for general certificate validation, but MAY supplement that algorithm with other validation methods that achieve equivalent levels of verification (such as comparing the server certificate against a local store of already-verified certificates and identity bindings).

If the client has external information as to the expected identity of the server, the hostname check MAY be omitted.

######

B.8. Syslog (2009)

In 2009, [SYSLOG-TLS] specified the following text regarding application service identity verification in Syslog:

######

5.2. Subject Name Authorization

Implementations MUST support certification path validation [PKIX]. In addition, they MUST support specifying the authorized peers using locally configured host names and matching the name against the certificate as follows.

- Implementations MUST support matching the locally configured host name against a dNSName in the subjectAltName extension field and SHOULD support checking the name against the common name portion of the subject distinguished name.

- The '*' (ASCII 42) wildcard character is allowed in the dNSName of the subjectAltName extension (and in common name, if used to store the host name), but only as the left-most (least significant) DNS label in that value. This wildcard matches any left-most DNS label in the server name. That is, the subject *.example.com matches the server names a.example.com and b.example.com, but does not match example.com or a.b.example.com. Implementations MUST support wildcards in certificates as specified above, but MAY provide a configuration option to disable them.

- Locally configured names MAY contain the wildcard character to match a range of values. The types of wildcards supported MAY be more flexible than those allowed in subject names, making it possible to support various policies for different environments. For example, a policy could allow for a trust-root-based authorization where all credentials issued by a particular CA trust root are authorized.

- If the locally configured name is an internationalized domain name, conforming implementations MUST convert it to the ASCII Compatible Encoding (ACE) format for performing comparisons, as specified in Section 7 of [PKIX].

- Implementations MAY support matching a locally configured IP address against an iPAddress stored in the subjectAltName extension. In this case, the locally configured IP address is converted to an octet string as specified in [PKIX], Section 4.2.1.6. A match occurs if this octet string is equal to the value of iPAddress in the subjectAltName extension.

######

B.9. SIP (2010)

In 2010, [SIP-CERTS] specified the following text regarding application service identity verification in SIP:

######

7.2. Comparing SIP Identities

When an implementation (either client or server) compares two values as SIP domain identities:

Implementations MUST compare only the DNS name component of each SIP domain identifier; an implementation MUST NOT use any scheme or parameters in the comparison.

Implementations MUST compare the values as DNS names, which means that the comparison is case insensitive as specified by [DNS-CASE]. Implementations MUST handle Internationalized Domain Names (IDNs) in accordance with Section 7.2 of [PKIX].

Implementations MUST match the values in their entirety:

Implementations MUST NOT match suffixes. For example, "foo.example.com" does not match "example.com".

Implementations MUST NOT match any form of wildcard, such as a leading "." or "*." with any other DNS label or sequence of labels. For example, "*.example.com" matches only "*.example.com" but not "foo.example.com". Similarly, ".example.com" matches only ".example.com", and does not match "foo.example.com."

[HTTP-TLS] allows the dNSName component to contain a wildcard; e.g., "DNS:*.example.com". [PKIX], while not disallowing this explicitly, leaves the interpretation of wildcards to the individual specification. [SIP] does not provide any guidelines on the presence of wildcards in certificates. Through the rule above, this document prohibits such wildcards in certificates for SIP domains.

######

B.10. SNMP (2010)

In 2010, [SNMP-TLS] specified the following text regarding application service identity verification in SNMP:

######

If the server's presented certificate has passed certification path validation [PKIX] to a configured trust anchor, and an active row exists with a zero-length snmpTlstmAddrServerFingerprint value, then the snmpTlstmAddrServerIdentity column contains the expected host name. This expected host name is then compared against the server's certificate as follows:

- Implementations MUST support matching the expected host name against a dNSName in the subjectAltName extension field and MAY support checking the name against the CommonName portion of the subject distinguished name.

- The '*' (ASCII 0x2a) wildcard character is allowed in the dNSName of the subjectAltName extension (and in common name, if used to store the host name), but only as the left-most (least significant) DNS label in that value. This wildcard matches any left-most DNS label in the server name. That is, the subject *.example.com matches the server names a.example.com and b.example.com, but does not match example.com or a.b.example.com. Implementations MUST support wildcards in certificates as specified above, but MAY provide a configuration option to disable them.

- If the locally configured name is an internationalized domain name, conforming implementations MUST convert it to the ASCII Compatible Encoding (ACE) format for performing comparisons, as specified in Section 7 of [PKIX].

If the expected host name fails these conditions then the connection MUST be closed.

######

B.11. GIST (2010)

In 2010, [GIST] specified the following text regarding application service identity verification in the General Internet Signalling Transport:

######

5.7.3.1. Identity Checking in TLS

After TLS authentication, a node MUST check the identity presented by the peer in order to avoid man-in-the-middle attacks, and verify that the peer is authorised to take part in signalling at the GIST layer. The authorisation check is carried out by comparing the presented identity with each Authorised Peer Database (APD) entry in turn, as discussed in Section 4.4.2. This section defines the identity comparison algorithm for a single APD entry.

For TLS authentication with X.509 certificates, an identity from the DNS namespace MUST be checked against each subjectAltName extension of type dNSName present in the certificate. If no such extension is present, then the identity MUST be compared to the (most specific) Common Name in the Subject field of the certificate. When matching DNS names against dNSName or Common Name fields, matching is case-insensitive. Also, a "*" wildcard character MAY be used as the left-most name component in the certificate or identity in the APD. For example, *.example.com in the APD would match certificates for a.example.com, foo.example.com, *.example.com, etc., but would not match example.com. Similarly, a certificate for *.example.com would be valid for APD identities of a.example.com, foo.example.com, *.example.com, etc., but not example.com.

Additionally, a node MUST verify the binding between the identity of the peer to which it connects and the public key presented by that peer. Nodes SHOULD implement the algorithm in Section 6 of [PKIX] for general certificate validation, but MAY supplement that algorithm with other validation methods that achieve equivalent levels of verification (such as comparing the server certificate against a local store of already-verified certificates and identity bindings).

For TLS authentication with pre-shared keys, the identity in the psk_identity_hint (for the server identity, i.e. the Responding node) or psk_identity (for the client identity, i.e. the Querying node) MUST be compared to the identities in the APD.

######

Authors' Addresses

Peter Saint-Andre
Cisco
1899 Wyknoop Street, Suite 600
Denver, CO 80202
USA

Phone: +1-303-308-3282
EMail: psaintan@cisco.com

Jeff Hodges
PayPal
2211 North First Street
San Jose, California 95131
US

EMail: Jeff.Hodges@PayPal.com