

Network Working Group
Request for Comments: 4398
Obsoletes: 2538
Category: Standards Track

S. Josefsson
March 2006

Storing Certificates in the Domain Name System (DNS)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Cryptographic public keys are frequently published, and their authenticity is demonstrated by certificates. A CERT resource record (RR) is defined so that such certificates and related certificate revocation lists can be stored in the Domain Name System (DNS).

This document obsoletes RFC 2538.

Table of Contents

1. Introduction	3
2. The CERT Resource Record	3
2.1. Certificate Type Values	4
2.2. Text Representation of CERT RRs	6
2.3. X.509 OIDs	6
3. Appropriate Owner Names for CERT RRs	7
3.1. Content-Based X.509 CERT RR Names	8
3.2. Purpose-Based X.509 CERT RR Names	9
3.3. Content-Based OpenPGP CERT RR Names	9
3.4. Purpose-Based OpenPGP CERT RR Names	10
3.5. Owner Names for IPKIX, ISPKI, IPGP, and IACPKIX	10
4. Performance Considerations	11
5. Contributors	11
6. Acknowledgements	11
7. Security Considerations	12
8. IANA Considerations	12
9. Changes since RFC 2538	13
10. References	14
10.1. Normative References	14
10.2. Informative References	15
Appendix A. Copying Conditions	16

1. Introduction

Public keys are frequently published in the form of a certificate, and their authenticity is commonly demonstrated by certificates and related certificate revocation lists (CRLs). A certificate is a binding, through a cryptographic digital signature, of a public key, a validity interval and/or conditions, and identity, authorization, or other information. A certificate revocation list is a list of certificates that are revoked, and of incidental information, all signed by the signer (issuer) of the revoked certificates. Examples are X.509 certificates/CRLs in the X.500 directory system or OpenPGP certificates/revocations used by OpenPGP software.

Section 2 specifies a CERT resource record (RR) for the storage of certificates in the Domain Name System [1] [2].

Section 3 discusses appropriate owner names for CERT RRs.


Sections 4, 7, and 8 cover performance, security, and IANA considerations, respectively.

Section 9 explains the changes in this document compared to RFC 2538.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [3].

2. The CERT Resource Record

The CERT resource record (RR) has the structure given below. Its RR type code is 37.



The type field is the certificate type as defined in Section 2.1 below.

The key tag field is the 16-bit value computed for the key embedded in the certificate, using the RRSIG Key Tag algorithm described in Appendix B of [12]. This field is used as an efficiency measure to

pick which CERT RRs may be applicable to a particular key. The key tag can be calculated for the key in question, and then only CERT RRs with the same key tag need to be examined. Note that two different keys can have the same key tag. However, the key MUST be transformed to the format it would have as the public key portion of a DNSKEY RR before the key tag is computed. This is only possible if the key is applicable to an algorithm and complies to limits (such as key size) defined for DNS security. If it is not, the algorithm field MUST be zero and the tag field is meaningless and SHOULD be zero.

The algorithm field has the same meaning as the algorithm field in DNSKEY and RRSIG RRs [12], except that a zero algorithm field indicates that the algorithm is unknown to a secure DNS, which may simply be the result of the algorithm not having been standardized for DNSSEC [11].

2.1. Certificate Type Values

The following values are defined or reserved:

Value	Mnemonic	Certificate Type
----	-----	-----
0		Reserved
1	PKIX	X.509 as per PKIX
2	SPKI	SPKI certificate
3	PGP	OpenPGP packet
4	IPKIX	The URL of an X.509 data object
5	ISPKI	The URL of an SPKI certificate
6	IPGP	The fingerprint and URL of an OpenPGP packet
7	ACPKIX	Attribute Certificate
8	IACPKIX	The URL of an Attribute Certificate
9-252		Available for IANA assignment
253	URI	URI private
254	OID	OID private
255		Reserved
256-65279		Available for IANA assignment
65280-65534		Experimental
65535		Reserved

These values represent the initial content of the IANA registry; see Section 8.

The PKIX type is reserved to indicate an X.509 certificate conforming to the profile defined by the IETF PKIX working group [8]. The certificate section will start with a one-octet unsigned OID length and then an X.500 OID indicating the nature of the remainder of the

certificate section (see Section 2.3, below). (NOTE: X.509 certificates do not include their X.500 directory-type-designating OID as a prefix.)

The SPKI and ISPKI types are reserved to indicate the SPKI certificate format [15], for use when the SPKI documents are moved from experimental status. The format for these two CERT RR types will need to be specified later.

The PGP type indicates an OpenPGP packet as described in [5] and its extensions and successors. This is used to transfer public key material and revocation signatures. The data is binary and MUST NOT be encoded into an ASCII armor. An implementation SHOULD process transferable public keys as described in Section 10.1 of [5], but it MAY handle additional OpenPGP packets.

The ACPKIX type indicates an Attribute Certificate format [9].

The IPKIX and IACPKIX types indicate a URL that will serve the content that would have been in the "certificate, CRL, or URL" field of the corresponding type (PKIX or ACPKIX, respectively).

The IPGP type contains both an OpenPGP fingerprint for the key in question, as well as a URL. The certificate portion of the IPGP CERT RR is defined as a one-octet fingerprint length, followed by the OpenPGP fingerprint, followed by the URL. The OpenPGP fingerprint is calculated as defined in RFC 2440 [5]. A zero-length fingerprint or a zero-length URL are legal, and indicate URL-only IPGP data or fingerprint-only IPGP data, respectively. A zero-length fingerprint and a zero-length URL are meaningless and invalid.

The IPKIX, ISPKI, IPGP, and IACPKIX types are known as "indirect". These types MUST be used when the content is too large to fit in the CERT RR and MAY be used at the implementer's discretion. They SHOULD NOT be used where the DNS message is 512 octets or smaller and could thus be expected to fit a UDP packet.

The URI private type indicates a certificate format defined by an absolute URI. The certificate portion of the CERT RR MUST begin with a null-terminated URI [10], and the data after the null is the private format certificate itself. The URI SHOULD be such that a retrieval from it will lead to documentation on the format of the certificate. Recognition of private certificate types need not be based on URI equality but can use various forms of pattern matching so that, for example, subtype or version information can also be encoded into the URI.

The OID private type indicates a private format certificate specified by an ISO OID prefix. The certificate section will start with a one-octet unsigned OID length and then a BER-encoded OID indicating the nature of the remainder of the certificate section. This can be an X.509 certificate format or some other format. X.509 certificates that conform to the IETF PKIX profile SHOULD be indicated by the PKIX type, not the OID private type. Recognition of private certificate types need not be based on OID equality but can use various forms of pattern matching such as OID prefix.

2.2. Text Representation of CERT RRs

The RDATA portion of a CERT RR has the type field as an unsigned decimal integer or as a mnemonic symbol as listed in Section 2.1, above.

The key tag field is represented as an unsigned decimal integer.

The algorithm field is represented as an unsigned decimal integer or a mnemonic symbol as listed in [12].

The certificate/CRL portion is represented in base 64 [16] and may be divided into any number of white-space-separated substrings, down to single base-64 digits, which are concatenated to obtain the full signature. These substrings can span lines using the standard parenthesis.

Note that the certificate/CRL portion may have internal sub-fields, but these do not appear in the master file representation. For example, with type 254, there will be an OID size, an OID, and then the certificate/CRL proper. However, only a single logical base-64 string will appear in the text representation.

2.3. X.509 OIDs

OIDs have been defined in connection with the X.500 directory for user certificates, certification authority certificates, revocations of certification authority, and revocations of user certificates. The following table lists the OIDs, their BER encoding, and their length-prefixed hex format for use in CERT RRs:

```
id-at-userCertificate
  = { joint-iso-ccitt(2) ds(5) at(4) 36 }
    == 0x 03 55 04 24
id-at-cACertificate
  = { joint-iso-ccitt(2) ds(5) at(4) 37 }
    == 0x 03 55 04 25
id-at-authorityRevocationList
  = { joint-iso-ccitt(2) ds(5) at(4) 38 }
    == 0x 03 55 04 26
id-at-certificateRevocationList
  = { joint-iso-ccitt(2) ds(5) at(4) 39 }
    == 0x 03 55 04 27
```

3. Appropriate Owner Names for CERT RRs

It is recommended that certificate CERT RRs be stored under a domain name related to their subject, i.e., the name of the entity intended to control the private key corresponding to the public key being certified. It is recommended that certificate revocation list CERT RRs be stored under a domain name related to their issuer.

Following some of the guidelines below may result in DNS names with characters that require DNS quoting as per Section 5.1 of RFC 1035 [2].

The choice of name under which CERT RRs are stored is important to clients that perform CERT queries. In some situations, the clients may not know all information about the CERT RR object it wishes to retrieve. For example, a client may not know the subject name of an X.509 certificate, or the email address of the owner of an OpenPGP key. Further, the client might only know the hostname of a service that uses X.509 certificates or the Key ID of an OpenPGP key.

Therefore, two owner name guidelines are defined: content-based owner names and purpose-based owner names. A content-based owner name is derived from the content of the CERT RR data; for example, the Subject field in an X.509 certificate or the User ID field in OpenPGP keys. A purpose-based owner name is a name that a client retrieving CERT RRs ought to know already; for example, the host name of an X.509 protected service or the Key ID of an OpenPGP key. The content-based and purpose-based owner name may be the same; for example, when a client looks up a key based on the From: address of an incoming email.

Implementations SHOULD use the purpose-based owner name guidelines described in this document and MAY use CNAME RRs at content-based owner names (or other names), pointing to the purpose-based owner name.

Note that this section describes an application-based mapping from the name space used in a certificate to the name space used by DNS. The DNS does not infer any relationship amongst CERT resource records based on similarities or differences of the DNS owner name(s) of CERT resource records. For example, if multiple labels are used when mapping from a CERT identifier to a domain name, then care must be taken in understanding wildcard record synthesis.

3.1. Content-Based X.509 CERT RR Names

Some X.509 versions, such as the PKIX profile of X.509 [8], permit multiple names to be associated with subjects and issuers under "Subject Alternative Name" and "Issuer Alternative Name". For example, the PKIX profile has such Alternate Names with an ASN.1 specification as follows:

```
GeneralName ::= CHOICE {
    otherName          [0]    OtherName,
    rfc822Name         [1]    IA5String,
    dNSName            [2]    IA5String,
    x400Address        [3]    ORAddress,
    directoryName      [4]    Name,
    ediPartyName       [5]    EDIPartyName,
    uniformResourceIdentifier [6] IA5String,
    iPAddress          [7]    OCTET STRING,
    registeredID       [8]    OBJECT IDENTIFIER }
```

The recommended locations of CERT storage are as follows, in priority order:

1. If a domain name is included in the identification in the certificate or CRL, that ought to be used.
2. If a domain name is not included but an IP address is included, then the translation of that IP address into the appropriate inverse domain name ought to be used.
3. If neither of the above is used, but a URI containing a domain name is present, that domain name ought to be used.
4. If none of the above is included but a character string name is included, then it ought to be treated as described below for OpenPGP names.
5. If none of the above apply, then the distinguished name (DN) ought to be mapped into a domain name as specified in [4].

Example 1: An X.509v3 certificate is issued to /CN=John Doe /DC=Doe/DC=com/DC=xy/O=Doe Inc/C=XY/ with Subject Alternative Names of (a) string "John (the Man) Doe", (b) domain name john-doe.com, and (c) URI <https://www.secure.john-doe.com:8080/>. The storage locations recommended, in priority order, would be

1. john-doe.com,
2. www.secure.john-doe.com, and
3. Doe.com.xy.

Example 2: An X.509v3 certificate is issued to /CN=James Hacker/L=Basingstoke/O=Widget Inc/C=GB/ with Subject Alternate names of (a) domain name widget.foo.example, (b) IPv4 address 10.251.13.201, and (c) string "James Hacker <hacker@mail.widget.foo.example>". The storage locations recommended, in priority order, would be

1. widget.foo.example,
2. 201.13.251.10.in-addr.arpa, and
3. hacker.mail.widget.foo.example.

3.2. Purpose-Based X.509 CERT RR Names

Due to the difficulty for clients that do not already possess a certificate to reconstruct the content-based owner name, purpose-based owner names are recommended in this section. Recommendations for purpose-based owner names vary per scenario. The following table summarizes the purpose-based X.509 CERT RR owner name guidelines for use with S/MIME [17], SSL/TLS [13], and IPsec [14]:

Scenario	Owner name
-----	-----
S/MIME Certificate	Standard translation of an RFC 2822 email address. Example: An S/MIME certificate for "postmaster@example.org" will use a standard hostname translation of the owner name, "postmaster.example.org".
TLS Certificate	Hostname of the TLS server.
IPsec Certificate	Hostname of the IPsec machine and/or, for IPv4 or IPv6 addresses, the fully qualified domain name in the appropriate reverse domain.

An alternate approach for IPsec is to store raw public keys [18].

3.3. Content-Based OpenPGP CERT RR Names

OpenPGP signed keys (certificates) use a general character string User ID [5]. However, it is recommended by OpenPGP that such names include the RFC 2822 [7] email address of the party, as in "Leslie Example <Leslie@host.example>". If such a format is used, the CERT ought to be under the standard translation of the email address into

a domain name, which would be `leslie.host.example` in this case. If no RFC 2822 name can be extracted from the string name, no specific domain name is recommended.

If a user has more than one email address, the CNAME type can be used to reduce the amount of data stored in the DNS. For example:

```
$ORIGIN example.org.  
smith          IN CERT PGP 0 0 <OpenPGP binary>  
john.smith     IN CNAME smith  
js             IN CNAME smith
```

3.4. Purpose-Based OpenPGP CERT RR Names

Applications that receive an OpenPGP packet containing encrypted or signed data but do not know the email address of the sender will have difficulties constructing the correct owner name and cannot use the content-based owner name guidelines. However, these clients commonly know the key fingerprint or the Key ID. The key ID is found in OpenPGP packets, and the key fingerprint is commonly found in auxiliary data that may be available. In this case, use of an owner name identical to the key fingerprint and the key ID expressed in hexadecimal [16] is recommended. For example:

```
$ORIGIN example.org.  
0424D4EE81A0E3D119C6F835EDA21E94B565716F IN CERT PGP ...  
F835EDA21E94B565716F                       IN CERT PGP ...  
B565716F                                     IN CERT PGP ...
```

If the same key material is stored for several owner names, the use of CNAME may help avoid data duplication. Note that CNAME is not always applicable, because it maps one owner name to the other for all purposes, which may be sub-optimal when two keys with the same Key ID are stored.

3.5. Owner Names for IPKIX, ISPKI, IPGP, and IACPKIX

These types are stored under the same owner names, both purpose- and content-based, as the PKIX, SPKI, PGP, and ACPKIX types.

4. Performance Considerations

The Domain Name System (DNS) protocol was designed for small transfers, typically below 512 octets. While larger transfers will perform correctly and work is underway to make larger transfers more efficient, it is still advisable at this time that every reasonable effort be made to minimize the size of certificates stored within the DNS. Steps that can be taken may include using the fewest possible optional or extension fields and using short field values for necessary variable-length fields.

The RDATA field in the DNS protocol may only hold data of size 65535 octets (64kb) or less. This means that each CERT RR MUST NOT contain more than 64kb of payload, even if the corresponding certificate or certificate revocation list is larger. This document addresses this by defining "indirect" data types for each normal type.

Deploying CERT RRs to support digitally signed email changes the access patterns of DNS lookups from per-domain to per-user. If digitally signed email and a key/certificate lookup based on CERT RRs are deployed on a wide scale, this may lead to an increased DNS load, with potential performance and cache effectiveness consequences. Whether or not this load increase will be noticeable is not known.

5. Contributors

The majority of this document is copied verbatim from RFC 2538, by Donald Eastlake 3rd and Olafur Gudmundsson.

6. Acknowledgements

Thanks to David Shaw and Michael Graff for their contributions to earlier works that motivated, and served as inspiration for, this document.

This document was improved by suggestions and comments from Olivier Dubuisson, Scott Hollenbeck, Russ Housley, Peter Koch, Olaf M. Kolkman, Ben Laurie, Edward Lewis, John Loughney, Allison Mankin, Douglas Otis, Marcos Sanz, Pekka Savola, Jason Sloderbeck, Samuel Weiler, and Florian Weimer. No doubt the list is incomplete. We apologize to anyone we left out.

7. Security Considerations

By definition, certificates contain their own authenticating signatures. Thus, it is reasonable to store certificates in non-secure DNS zones or to retrieve certificates from DNS with DNS security checking not implemented or deferred for efficiency. The results may be trusted if the certificate chain is verified back to a known trusted key and this conforms with the user's security policy.

Alternatively, if certificates are retrieved from a secure DNS zone with DNS security checking enabled and are verified by DNS security, the key within the retrieved certificate may be trusted without verifying the certificate chain if this conforms with the user's security policy.

If an organization chooses to issue certificates for its employees, placing CERT RRs in the DNS by owner name, and if DNSSEC (with NSEC) is in use, it is possible for someone to enumerate all employees of the organization. This is usually not considered desirable, for the same reason that enterprise phone listings are not often publicly published and are even marked confidential.

Using the URI type introduces another level of indirection that may open a new vulnerability. One method of securing that indirection is to include a hash of the certificate in the URI itself.

If DNSSEC is used, then the non-existence of a CERT RR and, consequently, certificates or revocation lists can be securely asserted. Without DNSSEC, this is not possible.

8. IANA Considerations

The IANA has created a new registry for CERT RR: certificate types. The initial contents of this registry is:

Decimal	Type	Meaning	Reference
-----	----	-----	-----
0		Reserved	RFC 4398
1	PKIX	X.509 as per PKIX	RFC 4398
2	SPKI	SPKI certificate	RFC 4398
3	PGP	OpenPGP packet	RFC 4398
4	IPKIX	The URL of an X.509 data object	RFC 4398
5	ISPKI	The URL of an SPKI certificate	RFC 4398
6	IPGP	The fingerprint and URL of an OpenPGP packet	RFC 4398
7	ACPKIX	Attribute Certificate	RFC 4398
8	IACPKIX	The URL of an Attribute Certificate	RFC 4398

9-252		Available for IANA assignment by IETF Standards action	
253	URI	URI private	RFC 4398
254	OID	OID private	RFC 4398
255		Reserved	RFC 4398
256-65279		Available for IANA assignment by IETF Consensus	
65280-65534		Experimental	RFC 4398
65535		Reserved	RFC 4398

Certificate types 0x0000 through 0x00FF and 0xFF00 through 0xFFFF can only be assigned by an IETF standards action [6]. This document assigns 0x0001 through 0x0008 and 0x00FD and 0x00FE. Certificate types 0x0100 through 0xFEFF are assigned through IETF Consensus [6] based on RFC documentation of the certificate type. The availability of private types under 0x00FD and 0x00FE ought to satisfy most requirements for proprietary or private types.

The CERT RR reuses the DNS Security Algorithm Numbers registry. In particular, the CERT RR requires that algorithm number 0 remain reserved, as described in Section 2. The IANA will reference the CERT RR as a user of this registry and value 0, in particular.

9. Changes since RFC 2538

1. Editorial changes to conform with new document requirements, including splitting reference section into two parts and updating the references to point at latest versions, and to add some additional references.
2. Improve terminology. For example replace "PGP" with "OpenPGP", to align with RFC 2440.
3. In Section 2.1, clarify that OpenPGP public key data are binary, not the ASCII armored format, and reference 10.1 in RFC 2440 on how to deal with OpenPGP keys, and acknowledge that implementations may handle additional packet types.
4. Clarify that integers in the representation format are decimal.
5. Replace KEY/SIG with DNSKEY/RRSIG etc, to align with DNSSECbis terminology. Improve reference for Key Tag Algorithm calculations.
6. Add examples that suggest use of CNAME to reduce bandwidth.
7. In Section 3, appended the last paragraphs that discuss "content-based" vs "purpose-based" owner names. Add Section 3.2 for purpose-based X.509 CERT owner names, and Section 3.4 for purpose-based OpenPGP CERT owner names.
8. Added size considerations.
9. The SPKI types has been reserved, until RFC 2692/2693 is moved from the experimental status.
10. Added indirect types IPKIX, ISPKI, IPGP, and IACPKIX.

11. An IANA registry of CERT type values was created.

10. References

10.1. Normative References

- [1] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [2] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [4] Kille, S., Wahl, M., Grimstad, A., Huber, R., and S. Sataluri, "Using Domains in LDAP/X.500 Distinguished Names", RFC 2247, January 1998.
- [5] Callas, J., Donnerhacke, L., Finney, H., and R. Thayer, "OpenPGP Message Format", RFC 2440, November 1998.
- [6] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [7] Resnick, P., "Internet Message Format", RFC 2822, April 2001.
- [8] Housley, R., Polk, W., Ford, W., and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3280, April 2002.
- [9] Farrell, S. and R. Housley, "An Internet Attribute Certificate Profile for Authorization", RFC 3281, April 2002.
- [10] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [11] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [12] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

10.2. Informative References

- [13] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.
- [14] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [15] Ellison, C., Frantz, B., Lampson, B., Rivest, R., Thomas, B., and T. Ylonen, "SPKI Certificate Theory", RFC 2693, September 1999.
- [16] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 3548, July 2003.
- [17] Ramsdell, B., "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification", RFC 3851, July 2004.
- [18] Richardson, M., "A Method for Storing IPsec Keying Material in DNS", RFC 4025, March 2005.

Appendix A. Copying Conditions

Regarding the portion of this document that was written by Simon Josefsson ("the author", for the remainder of this section), the author makes no guarantees and is not responsible for any damage resulting from its use. The author grants irrevocable permission to anyone to use, modify, and distribute it in any way that does not diminish the rights of anyone else to use, modify, and distribute it, provided that redistributed derivative works do not contain misleading author or version information. Derivative works need not be licensed under similar terms.

Author's Address

Simon Josefsson

E-Mail: simon@josefsson.org

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

Intellectual Property

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

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

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

Acknowledgement

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