

Session Initiation Protocol (SIP) Extension Header Field
for Service Route Discovery During Registration

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 (2003). All Rights Reserved.

Abstract

This document defines a Session Initiation Protocol (SIP) extension header field used in conjunction with responses to REGISTER requests to provide a mechanism by which a registrar may inform a registering user agent (UA) of a service route that the UA may use to request outbound services from the registrar's domain.

Table of Contents

1. Terminology	2
2. Background	2
3. Discussion of Mechanism	4
4. Applicability Statement	5
5. Syntax	5
6. Usage	6
6.1. Procedures at the UA	6
6.2. Procedures at the Proxy	7
6.3. Procedures at the Registrar	8
6.4. Examples of Usage	9
6.4.1. Example of Mechanism in REGISTER Transaction	9
6.4.2. Example of Mechanism in INVITE Transaction	12
7. Security Considerations	14
8. IANA Considerations	15
9. Normative References	15
10. Informative References	15

11. Intellectual Property Statement	16
12. Authors' Addresses	16
13. Full Copyright Statement	17

1. Terminology

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 BCP 14, RFC 2119 [1].

2. Background

The Third Generation Partnership Project (3GPP) established a requirement for discovering home proxies during SIP registration and published this requirement in [6]. The 3GPP network dynamically assigns a home service proxy to each address-of-record (AOR). This assignment may occur in conjunction with a REGISTER operation, or out-of-band as needed to support call services when the address-of-record has no registrations. This home service proxy may provide both inbound (UA terminated) and outbound (UA originated) services.

In the inbound case, the Request-Uniform Resource Identifier (URI) of incoming SIP requests matches the address-of-record of a user associated with the home service proxy. The home service proxy then (in most cases) forwards the request to the registered contact address for that AOR. A mechanism for traversing required proxies between the home service proxy and the registered UA is presented in [4].

Outbound (UA originated) session cases raise another issue. Specifically, "How does the UA know which service proxy to use and how to get there?"

Several mechanisms were proposed in list discussions, including:

1. Configuration data in the UA. This raises questions of UA configuration management and updating, especially if proxy assignment is very dynamic, such as in load-balancing scenarios.
2. Use of some other protocol, such as HTTP, to get configuration data from a configuration server in the home network. While functional, this solution requires additional protocol engines, firewall complexity, operations overhead, and significant additional "over the air" traffic.
3. Use of lookup tables in the home network, as may be done for inbound requests in some 3G networks. This has a relatively high overhead in terms of database operations.

3. Discussion of Mechanism

UAs may include a Route header field in an initial request to force that request to visit and potentially be serviced by one or more proxies. Using such a route (called a "service route" or "preloaded route") allows a UA to request services from a specific home proxy or network of proxies. The open question is, "How may a UA discover what service route to use?"

This document defines a header field called "Service-Route" which can contain a route vector that, if used as discussed above, will direct requests through a specific sequence of proxies. A registrar may use a Service-Route header field to inform a UA of a service route that, if used by the UA, will provide services from a proxy or set of proxies associated with that registrar. The Service-Route header field may be included by a registrar in the response to a REGISTER request. Consequently, a registering UA learns of a service route that may be used to request services from the system it just registered with.

The routing established by the Service-Route mechanism applies only to requests originating in the user agent. That is, it applies only to UA originated requests, and not to requests terminated by that UA.

Simply put, the registrar generates a service route for the registering UA and returns it in the response to each successful REGISTER request. This service route has the form of a Route header field that the registering UA may use to send requests through the service proxy selected by the registrar. The UA would use this route by inserting it as a preloaded Route header field in requests originated by the UA intended for routing through the service proxy.

The mechanism by which the registrar constructs the header field value is specific to the local implementation and outside the scope of this document.

4. Applicability Statement

The Service-Route mechanism is applicable when:

1. The UA registers with a registrar.
2. The registrar has knowledge of a service proxy that should be used by the UA when requesting services from the domain of the registrar. This knowledge may be a result of dynamic assignment or some other mechanism outside the scope of this document.

3. The registrar(s) has/have sufficient knowledge of the network topology, policy, and situation such that a reasonable service route can be constructed.
4. The service route constructed by the registrar is the same for all contacts associated with a single address-of-record. This mechanism does not provide for contact-specific service routes.
5. Other mechanisms for proposing a service route to the UA are not available or are inappropriate for use within the specific environment.

Other methods may also be available by which a UA may be informed of a service route. Such alternative methods are outside the scope of this document. Discussion of why one might wish to assign a service route during registration or when it might be appropriate to do so is outside the scope of this document.

5. Syntax

The syntax for the Service-Route header field is:

```
Service-Route = "Service-Route" HCOLON sr-value *( COMMA sr-value )
sr-value = name-addr *( SEMI rr-param )
```

Note that the Service-Route header field values MUST conform to the syntax of a Route element as defined in [3]. As suggested therein, such values MUST include the loose-routing indicator parameter ";lr" for full compliance with [3].

The allowable usage of header fields is described in Tables 2 and 3 of [3]. The following additions to this table are needed for Service-Route.

Addition of Service-Route to SIP Table 3:

Header field	where	proxy	ACK	BYE	CAN	INV	OPT	REG	PRA
Service-Route	2xx	ar	-	-	-	-	-	o	-

6. Usage

6.1. Procedures at the UA

The UA performs a registration as usual. The REGISTER response may contain a Service-Route header field. If so, the UA MAY store the value of the Service-Route header field in an association with the address-of-record for which the REGISTER transaction had registered a contact. If the UA supports multiple addresses-of-record, it may be able to store multiple service routes, one per address-of-record. If the UA refreshes the registration, the stored value of the Service-Route is updated according to the Service-Route header field of the latest 200 class response. If there is no Service-Route header field in the response, the UA clears any service route for that address-of-record previously stored by the UA. If the re-registration request is refused or if an existing registration expires and the UA chooses not to re-register, the UA SHOULD discard any stored service route for that address-of-record.

The UA MAY choose to exercise a service route for future requests associated with a given address-of-record for which a service route is known. If so, it uses the content of the Service-Route header field as a preloaded Route header field in outgoing initial requests [3]. The UA MUST preserve the order, in case there is more than one Service-Route header field or header field value.

Loose routes may interact with routing policy in interesting ways. The specifics of how the service route set integrates with any locally required default route and local policy are implementation dependent. For example, some devices will use locally-configured explicit loose routing to reach a next-hop proxy, and others will use a default outbound-proxy routing rule. However, for the result to function, the combination MUST provide valid routing in the local environment. In general, the service route set is appended to any locally configured route needed to egress the access proxy chain. Systems designers must match the service routing policy of their nodes with the basic SIP routing policy in order to get a workable system.

6.2. Procedures at the Proxy

The Service-Route header field is generally treated like any other unknown header field by intermediate proxies. They simply forward it on towards the destination. Note that, as usual, intermediate proxies that need to be traversed by future requests within a dialog may record-route. Proxies should not assume that they will be traversed by future requests in a dialog simply because they appear in the Service-Route header field.

There is a question of whether proxies processing a REGISTER response may add themselves to the route set in the Service-Route header field. While this would enable dynamic construction of service routes, it has two significant problems. The first is one of transparency, as seen by the registrar: Intermediate proxies could add themselves without the knowledge or consent of the registrar. The second problem is interaction with end-to-end security. If the registrar uses S/MIME techniques to protect the REGISTER response, such additions would be visible to the UA as "man in the middle" alterations in the response. Consequently, intermediate proxies SHOULD NOT alter the value of Service-Route in REGISTER responses, and if they do, the UA MUST NOT be required to accept the alteration.

Additional considerations apply if a proxy is "dual homed", meaning connected to two (or more) different networks such that requests are received on one interface and proxied out through another network interface. Proxies implementing multi-homing precisely as documented in [3] record-route a request with the sending interface. When processing the reply, they replace the Record-Route header field value that represents the interface onto which they proxied the request with a new value that represents the interface onto which they will proxy the response. Consequently, the route vector seen at the User Agent Server (UAS) is not the exact inverse of the route vector seen at the User Agent Client (UAC). While in itself harmless, this complicates matters for nodes that use the recorded route vector (or recorded Path vector as per [4]) in the determination of a service route for future use.

Instead of following the procedure in [3], proxies used with Service-Route that are inserting Record-Route or Path header field values SHOULD record not one but two route values when processing the request. The first value recorded indicates the receiving interface, and the second indicates the sending interface. When processing the response, no modification of the recorded route is required. This optimization provides for fully invertible routes that can be effectively used in construction of service routes.

6.3. Procedures at the Registrar

When a registrar receives a successful REGISTER request, it MAY choose to return one or more Service-Route header field(s) in the 200 class response. The determination(s) of whether to include these header fields(s) into the 200 class response and what value(s) to insert are a matter of local policy and outside the scope of this document.

Having inserted a Service-Route header field or fields, the registrar returns the 200 class response to the UA in accordance with standard procedures.

A REGISTER operation performing a Fetching Bindings (i.e., no Contact header field is present in the request) SHOULD return the same value of Service-Route as returned in the corresponding previous REGISTER response for the address-of-record in question. In some cases, the Service-Route may be dynamically calculated by the registrar rather than stored, and the decision as to whether this route should be recalculated in the event of a Fetching Bindings operation is left to the implementation.

Note: A Fetching Bindings operation could be used by the UA to recover a lost value of Service-Route. Alternatively, a UA in this situation could just re-REGISTER.

Certain network topologies MAY require a specific proxy (e.g., firewall proxy) to be traversed before the home service proxy. Thus, a registrar with specific knowledge of the network topology MAY return more than one Service-Route header field or element in the 200 class response; the order is specified as top-down, meaning the topmost Service-Route entry will be visited first. Such constructions are implementation specific and outside the scope of this document.

In general, the Service-Route header field contains references to elements strictly within the administrative domain of the registrar and home service proxy. For example, consider a case where a user leaves the "home" network and roams into a "visited" network. The registrar cannot be assumed to have knowledge of the topology of the visited network, so the Service-Route it returns contains elements only within the home network.

Note that the inserted Service-Route element(s) MUST conform to the syntax of a Route element as defined in [3]. As suggested therein, such route elements MUST include the loose-routing indicator parameter ";lr" for full compliance with [3].

F2 Register P1 -> P2

```
REGISTER sip:HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKlJuB1mcr
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKcR1ntRAp
To: Lawyer <sip:UA1@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=981211
Call-ID: 843817637684230@998sdasdh09
CSeq: 1826 REGISTER
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
. . .
```

F3 Register P2 -> R

```
REGISTER sip:HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/UDP P2.HOME.EXAMPLE.COM:5060;branch=z9hG4bKvE0R2l07o2b6T
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKlJuB1mcr
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKcR1ntRAp
To: Lawyer <sip:UA1@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=981211
Call-ID: 843817637684230@998sdasdh09
CSeq: 1826 REGISTER
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
. . .
```

F4 R executes Register

```
R Stores:
For <sip:UA1@HOME.EXAMPLE.COM>
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
```

F5 R calculates Service Route

In this example, R is statically configured to reference HSP as a service route, and R also knows that P2 is used as the provider edge proxy, so:

```
Service-Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
               <sip:HSP.HOME.EXAMPLE.COM;lr>
```

F6 Register Response r -> P2

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP P2.HOME.EXAMPLE.COM:5060;branch=z9hG4bKvE0R2l07o2b6T
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKlJuB1mcr
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKcR1ntRAP
To: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=87654
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=981211
Call-ID: 843817637684230@998sdasdh09
CSeq: 1826 REGISTER
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Service-Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
               <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

F7 Register Response P2 -> P1

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKlJuB1mcr
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKcR1ntRAP
To: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=87654
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=981211
Call-ID: 843817637684230@998sdasdh09
CSeq: 1826 REGISTER
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Service-Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
               <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

F8 Register Response P1 -> UA1

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKcR1ntRAP
To: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=87654
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=981211
Call-ID: 843817637684230@998sdasdh09
CSeq: 1826 REGISTER
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Service-Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
               <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

F9 UA1 stores service route for UA1@HOME.EXAMPLE.COM

6.4.2. Example of Mechanism in INVITE Transaction

This example shows the message sequence for an INVITE transaction originating from UA1 eventually arriving at UA2 using outbound services from HOME.EXAMPLE.COM. UA1 has previously registered with HOME.EXAMPLE.COM and been informed of a service route through HSP.HOME.EXAMPLE.COM. The service being provided by HOME.EXAMPLE.COM is a "logging" service, which provides a record of the call for UA1's use (perhaps the user of UA1 is an attorney who bills for calls to customers).

Note that in this example UA1 and UA2 are assumed to be registered with the same network (HOME.EXAMPLE.COM). This does not generally need to be the case to use the herein described service route mechanism.

Message sequence for INVITE using Service-Route:

F1 Invite UA1 -> P1

```
INVITE sip:UA2@HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKnashds7
To: Customer <sip:UA2@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=456248
Call-ID: 38615183343@s1i1l12j6u
CSeq: 18 INVITE
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
       <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

Note: P1 is selected using the "outbound proxy" rule in UA1.

F2 Invite P1 -> P2

```
INVITE sip:UA2@HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bK34ghi7ab04
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKnashds7
To: Customer <sip:UA2@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=456248
Call-ID: 38615183343@s1i1l12j6u
CSeq: 18 INVITE
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Record-Route: <sip:P1.VISITED.EXAMPLE.ORG;lr>
Route: <sip:P2.HOME.EXAMPLE.COM;lr>,
       <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

Note: P1 has added itself to the Record Route.

F3 Invite P2 -> HSP

```
INVITE sip:UA2@HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/UDP P2.HOME.EXAMPLE.COM:5060;branch=z9hG4bKiokioukju908
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bK34ghi7ab04
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKnashds7
To: Customer <sip:UA2@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=456248
Call-ID: 38615183343@sl1ll2j6u
CSeq: 18 INVITE
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Record-Route: <sip:P2.HOME.EXAMPLE.COM;lr>
Record-Route: <sip:P1.VISITED.EXAMPLE.ORG;lr>
Route: <sip:HSP.HOME.EXAMPLE.COM;lr>
. . .
```

Note: HSP is selected using a DNS lookup for HSP within HOME.EXAMPLE.COM.

P2 has added itself to the Record-Route.

P2 has removed itself from the Route.

F4 HSP executes service

HSP identifies the service to be executed from UA1's stored profile. The specifics of this are outside the scope of this document. For this example HSP writes a record to "Lawyer's log book", then looks up the AOR "sip:UA2@HOME.EXAMPLE.COM" and discovers that the current contact for UA2 is at host UAADDR2.HOME.EXAMPLE.COM. This will be the Request-URI of the next-hop INVITE.

F5 Invite HSP -> P2

```
INVITE sip:UA2@UAADDR2.HOME.EXAMPLE.COM SIP/2.0
Via: SIP/2.0/USP HSP.HOME.EXAMPLE.COM:5060;branch=z9hG4bKHSP10120323
Via: SIP/2.0/UDP P2.HOME.EXAMPLE.COM:5060;branch=z9hG4bKiokioukju908
Via: SIP/2.0/UDP P1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bK34ghi7ab04
Via: SIP/2.0/UDP UADDR1.VISITED.EXAMPLE.ORG:5060;branch=z9hG4bKnashds7
To: Customer <sip:UA2@HOME.EXAMPLE.COM>
From: Lawyer <sip:UA1@HOME.EXAMPLE.COM>;tag=456248
Call-ID: 38615183343@sl1ll2j6u
CSeq: 18 INVITE
Contact: <sip:UA1@UADDR1.VISITED.EXAMPLE.ORG>
Record-Route: <sip:HSP.HOME.EXAMPLE.COM;lr>
Record-Route: <sip:P2.HOME.EXAMPLE.COM;lr>
Record-Route: <sip:P1.VISITED.EXAMPLE.ORG;lr>
```

. . .

Note: P2 selected by outbound proxy rule on HSP.
HSP has removed itself from the Route.

INVITE propagates toward UA2 as usual.

7. Security Considerations

It is possible for proxies between the UA and the registrar during the REGISTER transaction to modify the value of Service-Route returned by the registrar, or to insert a Service-Route even when one was not returned by the registrar. The consequence of such an attack is that future requests made by the UA using the service route might be diverted to or through a node other than would normally be visited. It is also possible for proxies on the INVITE path to execute many different attacks. It is therefore desirable to apply transitive mutual authentication using sips: or other available mechanisms in order to prevent such attacks.

The "sips:" URI as defined in [3] defines a mechanism by which a UA may request transport-level message integrity and mutual authentication. Since there is no requirement for proxies to modify messages, S/MIME signed bodies may be used to provide end-to-end protection for the returned value.

Systems using Service-Route SHOULD provide hop-by-hop message integrity and mutual authentication. UAs SHOULD request this support by using a "sips:" URI. Registrars returning a Service-Route MUST implement end-to-end protection using S/MIME and SHOULD use S/MIME to protect all such responses. UAs receiving Service-Route SHOULD authenticate attached S/MIME bodies if present.

8. IANA Considerations

This document defines the SIP extension header field "Service-Route" which has been included in the registry of SIP header fields defined in [3]. The change process for SIP, [5] mandates that general SIP extension header fields be defined by a standards-track RFC. This document provides the required definition.

The following is the registration for the Service-Route header field:

RFC Number: RFC 3608

Header Field Name: Service-Route

Compact Form: none

9. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Postel, J. and J. Reynolds, "Instructions to RFC Authors", RFC 2223, October 1997.
- [3] 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.
- [4] Willis, D. and B. Hoeneisen, "Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts", RFC 3327, December 2002.
- [5] Mankin, A., Bradner, S., Mahy, R., Willis, D., Ott, J. and B. Rosen, "Change Process for the Session Initiation Protocol (SIP)", BCP 67, RFC 3427, December 2002.

10. Informative References

- [6] Garcia-Martin, M., "3rd-Generation Partnership Project (3GPP) Release 5 requirements on the Session Initiation Protocol (SIP)", Work in Progress, October 2002.

11. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

12. Authors' Addresses

Dean Willis
dynamicsoft Inc.
3100 Independence Parkway
#311-164
Plano, TX 75075
US

Phone: +1 972 473 5455
EMail: dean.willis@softarmor.com

Bernie Hoeneisen
Switch
Limmatquai 138
CH-8001 Zuerich
Switzerland

Phone: +41 1 268 1515
EMail: hoeneisen@switch.ch, b.hoeneisen@ieee.org
URI: <http://www.switch.ch/>

13. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

