

Internet Engineering Task Force (IETF)  
Request for Comments: 7264  
Category: Standards Track  
ISSN: 2070-1721

N. Zong  
X. Jiang  
R. Even  
Huawei Technologies  
Y. Zhang  
CoolPad / China Mobile  
June 2014

An Extension to the REsource LOcation And Discovery (RELOAD) Protocol  
to Support Relay Peer Routing

Abstract

This document defines an optional extension to the REsource LOcation And Discovery (RELOAD) protocol to support the relay peer routing mode. RELOAD recommends symmetric recursive routing for routing messages. The new optional extension provides a shorter route for responses, thereby reducing overhead on intermediate peers. This document also describes potential cases where this extension can be used.

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/rfc7264>.

## Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction .....	3
2. Terminology .....	4
3. Overview .....	5
3.1. RPR .....	5
3.2. Scenarios Where RPR Can Be Used .....	6
3.2.1. Managed or Closed P2P Systems .....	6
3.2.2. Using Bootstrap Nodes as Relay Peers .....	7
3.2.3. Wireless Scenarios .....	7
4. Relationship between SRR and RPR .....	7
4.1. How RPR Works .....	7
4.2. How SRR and RPR Work Together .....	7
5. RPR Extensions to RELOAD .....	8
5.1. Basic Requirements .....	8
5.2. Modification to RELOAD Message Structure .....	8
5.2.1. Extensive Routing Mode .....	8
5.3. Creating a Request .....	9
5.3.1. Creating a Request for RPR .....	9
5.4. Request and Response Processing .....	9
5.4.1. Destination Peer: Receiving a Request and Sending a Response .....	9
5.4.2. Sending Peer: Receiving a Response .....	10
5.4.3. Relay Peer Processing .....	10
6. Overlay Configuration Extension .....	10
7. Discovery of Relay Peers .....	11
8. Security Considerations .....	11
9. IANA Considerations .....	11
9.1. A New RELOAD Forwarding Option .....	11
10. Acknowledgments .....	11
11. References .....	12
11.1. Normative References .....	12
11.2. Informative References .....	12
Appendix A. Optional Methods to Investigate Peer Connectivity .....	13
Appendix B. Comparison of Cost of SRR and RPR .....	14
B.1. Closed or Managed Networks .....	14
B.2. Open Networks .....	15

## 1. Introduction

The REsource LOcation And Discovery (RELOAD) protocol [RFC6940] recommends symmetric recursive routing (SRR) for routing messages and describes the extensions that would be required to support additional routing algorithms. In addition to SRR, two other routing options -- direct response routing (DRR) and relay peer routing (RPR) -- are also discussed in Appendix A of [RFC6940]. As we show in Section 3, RPR is advantageous over SRR in some scenarios in that RPR can reduce load (CPU and link bandwidth) on intermediate peers. RPR works better in a network where relay peers are provisioned in advance so

that relay peers are publicly reachable in the P2P system. In other scenarios, using a combination of RPR and SRR together is more likely to provide benefits than if SRR is used alone.

Note that in this document we focus on the RPR mode and its extensions to RELOAD to produce a standalone solution. Please refer to [RFC7263] for details on the DRR mode.

We first discuss the problem statement in Section 3. How to combine RPR and SRR is presented in Section 4. An extension to RELOAD to support RPR is defined in Section 5. Discovery of relay peers is introduced in Section 7. Some optional methods to check peer connectivity are introduced in Appendix A. In Appendix B, we give a comparison of the cost of SRR and RPR in both managed and open networks.

## 2. 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 RFC 2119 [RFC2119].

We use terminology and definitions from the base RELOAD specification [RFC6940] extensively in this document. We also use terms defined in the NAT behavior discovery document [RFC5780]. Other terms used in this document are defined inline when used and are also defined below for reference.

**Publicly Reachable:** A peer is publicly reachable if it can receive unsolicited messages from any other peer in the same overlay.  
**Note:** "Publicly" does not mean that the peers must be on the public Internet, because the RELOAD protocol may be used in a closed network.

**Relay Peer:** A relay peer is a type of publicly reachable peer that can receive unsolicited messages from all other peers in the overlay and forward the responses from destination peers towards the sender of the request.

**Relay Peer Routing (RPR):** "RPR" refers to a routing mode in which responses to Peer-to-Peer SIP (P2PSIP) requests are sent by the destination peer to a relay peer transport address that will forward the responses towards the sending peer. For simplicity, the abbreviation "RPR" is used in the rest of this document.

Symmetric Recursive Routing (SRR): "SRR" refers to a routing mode in which responses follow the reverse path of the request to get to the sending peer. For simplicity, the abbreviation "SRR" is used in the rest of this document.

Direct Response Routing (DRR): "DRR" refers to a routing mode in which responses to P2PSIP requests are returned to the sending peer directly from the destination peer based on the sending peer's own local transport address(es). For simplicity, the abbreviation "DRR" is used in the rest of this document.

### 3. Overview

RELOAD is expected to work under a great number of application scenarios. The situations where RELOAD is to be deployed differ greatly. For instance, some deployments are global, such as a Skype-like system intended to provide public service, while others run in small-scale closed networks. SRR works in any situation, but RPR may work better in some specific scenarios.

#### 3.1. RPR

RELOAD is a simple request-response protocol. After sending a request, a peer waits for a response from a destination peer. There are several ways for the destination peer to send a response back to the source peer. In this section, we will provide detailed information on RPR. Note that the same types of illustrative settings can be found in Appendix B.1 of [RFC7263].

If peer A knows it is behind a NAT or NATs and knows one or more relay peers with whom they have had prior connections, peer A can try RPR. Assume that peer A is associated with relay peer R. When sending the request, peer A includes information describing peer R's transport address in the request. When peer X receives the request, peer X sends the response to peer R, which forwards it directly to peer A on the existing connection. Figure 1 illustrates RPR. Note that RPR also allows a shorter route for responses compared to SRR; this means less overhead on intermediate peers. Establishing a connection to the relay with Transport Layer Security (TLS) requires multiple round trips. Please refer to Appendix B for a cost comparison between SRR and RPR.

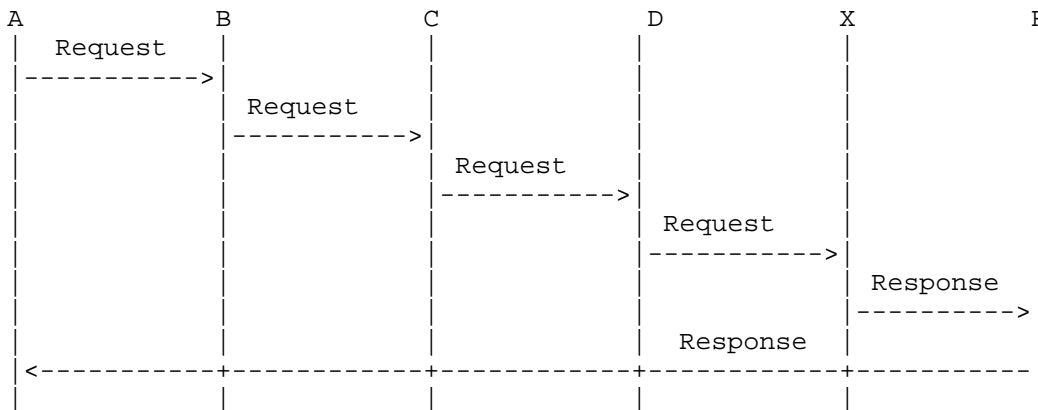


Figure 1: RPR Mode

This technique relies on the relative population of peers such as peer A that require relay peers, and peers such as peer R that are capable of serving as relay peers. It also requires a mechanism to enable peers to know which peers can be used as their relays. This mechanism may be based on configuration -- for example, as part of the overlay configuration, an initial list of relay peers can be supplied. Another option is a response message in which the responding peer can announce that it can serve as a relay peer.

### 3.2. Scenarios Where RPR Can Be Used

In this section, we will list several scenarios where using RPR would improve performance.

#### 3.2.1. Managed or Closed P2P Systems

As described in Section 3.2.1 of [RFC7263], many P2P systems run in a closed or managed environment so that network administrators can better manage their system. For example, the network administrator can deploy several relay peers that are publicly reachable in the system and indicate their presence in the configuration file. After learning where these relay peers are, peers behind NATs can use RPR with help from these relay peers. Peers MUST also support SRR in case RPR fails.

Another usage is to install relay peers on the managed network boundary, allowing external peers to send responses to peers inside the managed network.

### 3.2.2. Using Bootstrap Nodes as Relay Peers

Bootstrap nodes are typically publicly reachable in a RELOAD architecture. As a result, one possible scenario would be to use the bootstrap nodes as relay peers for use with RPR. A relay peer SHOULD be publicly accessible and maintain a direct connection with its client. As such, bootstrap nodes are well suited to play the role of relay peers.

### 3.2.3. Wireless Scenarios

In some mobile deployments, using RPR may help reduce radio battery usage and bandwidth by the intermediate peers. The service provider may recommend using RPR based on his knowledge of the topology.

## 4. Relationship between SRR and RPR

### 4.1. How RPR Works

Peers using RPR MUST maintain a connection with their relay peer(s). This can be done in the same way as establishing a neighbor connection between peers using the Attach method [RFC6940].

A requirement for RPR is that the source peer convey its relay peer's (or peers') transport address(es) in the request so the destination peer knows where the relay peers are and will send the response to a relay peer first. The request MUST also include the requesting peer's Node-ID or IP address, which enables the relay peer to route the response back to the right peer.

Note that being a relay peer does not require that the relay peer have more functionality than an ordinary peer. Relay peers comply with the same procedure as an ordinary peer to forward messages. The only difference is that there may be a larger traffic burden on relay peers. Relay peers can decide whether to accept a new connection based on their current burden.

### 4.2. How SRR and RPR Work Together

RPR is not intended to replace SRR. It is better to use these two modes together to adapt to each peer's specific situation. Note that the informative suggestions for how to transition between SRR and RPR are the same as those for DRR. Please refer to Section 4.2 of [RFC7263] for more details. If a relay peer is provided by the service provider, peers SHOULD prefer RPR over SRR. However, RPR SHOULD NOT be used in the open Internet or if the administrator does

not feel he has enough information about the overlay network topology. A new overlay configuration element specifying the usage of RPR is defined in Section 6.

## 5. RPR Extensions to RELOAD

Adding support for RPR requires extensions to the current RELOAD protocol. In this section, we define the required extensions, including extensions to message structure and message processing.

### 5.1. Basic Requirements

All peers MUST be able to process requests for routing in SRR and MAY support RPR routing requests.

### 5.2. Modification to RELOAD Message Structure

RELOAD provides an extensible framework to accommodate future extensions. In this section, we define an RPR routing option for the extensive routing mode specified in [RFC7263]. The state-keeping flag [RFC7263] is needed to support the RPR mode.

#### 5.2.1. Extensive Routing Mode

The new RouteMode value for RPR is defined below for the ExtensiveRoutingModeOption structure:

```
enum {(0),DRR(1),RPR(2),(255)} RouteMode;
struct {
    RouteMode          routemode;
    OverlayLinkType    transport;
    IpAddressPort      ipaddressport;
    Destination        destinations<1..2^8-1>;
} ExtensiveRoutingModeOption;
```

Note that the DRR value in RouteMode is defined in [RFC7263].

RouteMode: refers to which type of routing mode is indicated to the destination peer.

OverlayLinkType: refers to the transport type that is used to deliver responses from the destination peer to the relay peer.

IpAddressPort: refers to the transport address that the destination peer should use for sending responses. This will be a relay peer address for RPR.



Destination: refers to the relay peer itself. If the routing mode is RPR, then the destination contains two items: the relay peer's Node-ID and the sending peer's Node-ID.

### 5.3. Creating a Request

#### 5.3.1. Creating a Request for RPR

When using RPR for a transaction, the sending peer MUST set the IGNORE-STATE-KEEPING flag in the ForwardingHeader. Additionally, the peer MUST construct and include a ForwardingOption structure in the ForwardingHeader. When constructing the ForwardingOption structure, the fields MUST be set as follows:

- 1) The type MUST be set to extensive\_routing\_mode.
- 2) The ExtensiveRoutingModeOption structure MUST be used for the option field within the ForwardingOption structure. The fields MUST be defined as follows:
  - 2.1) routemode set to 0x02 (RPR).
  - 2.2) transport set as appropriate for the relay peer.
  - 2.3) ipaddressport set to the transport address of the relay peer through which the sender wishes the message relayed.
  - 2.4) The destination structure MUST contain two values. The first MUST be defined as type "node" and set with the values for the relay peer. The second MUST be defined as type "node" and set with the sending peer's own values.

### 5.4. Request and Response Processing

This section gives normative text for message processing after RPR is introduced. Here, we only describe the additional procedures for supporting RPR. Please refer to [RFC6940] for RELOAD base procedures.

#### 5.4.1. Destination Peer: Receiving a Request and Sending a Response

When the destination peer receives a request, it will check the options in the forwarding header. If the destination peer cannot understand the extensive\_routing\_mode option in the request, it MUST attempt to use SRR to return an "Error\_Unknown\_Extension" response (defined in Sections 6.3.3.1 and 14.9 of [RFC6940]) to the sending peer.

If the routing mode is RPR, the destination peer MUST construct a `destination_list` for the response with two entries as defined in [RFC6940]. The first entry MUST be set to the relay peer's Node-ID from the option in the request, and the second entry MUST be the sending peer's Node-ID from the option in the request.

In the event that the routing mode is set to RPR and there are not exactly two destinations, the destination peer MUST try to send an "Error\_Unknown\_Extension" response (defined in Sections 6.3.3.1 and 14.9 of [RFC6940]) to the sending peer using SRR.

After the peer constructs the `destination_list` for the response, it sends the response to the transport address, which is indicated in the `ipaddressport` field in the option using the specific transport mode in the `ForwardingOption`. If the destination peer receives a retransmit with SRR preference on the message it is trying to respond to now, the responding peer SHOULD abort the RPR response and use SRR.

#### 5.4.2. Sending Peer: Receiving a Response

Upon receiving a response, the peer follows the rules in [RFC6940]. If the sender used RPR and did not get a response until the timeout, it MAY resend the message using either RPR (but with a different relay peer, if available) or SRR.

#### 5.4.3. Relay Peer Processing

Relay peers are designed to forward responses to peers who are not publicly reachable. For the routing of the response, this document still uses the `destination_list`. The only difference from SRR is that the `destination_list` is not the reverse of the `via_list`. Instead, it is constructed from the forwarding option as described below.

When a relay peer receives a response, it MUST follow the rules in [RFC6940]. It receives the response, validates the message, readjusts the `destination_list`, and forwards the response to the next hop in the `destination_list` based on the connection table. There is no added requirement for the relay peer.

## 6. Overlay Configuration Extension

This document uses the new RELOAD overlay configuration element, "route-mode", inside each "configuration" element, as defined in Section 6 of [RFC7263]. The route mode MUST be "RPR".

## 7. Discovery of Relay Peers

There are several ways to distribute information about relay peers throughout the overlay. P2P network providers can deploy some relay peers and advertise them in the configuration file. With the configuration file at hand, peers can get relay peers to try RPR. Another way is to consider the relay peer as a service; some service advertisement and discovery mechanism can then also be used for discovering relay peers -- for example, using the same mechanism as that used in Traversal Using Relays around NAT (TURN) server discovery as discussed in [RFC6940]. Another option is to let a peer advertise its capability to be a relay in the response to an Attach or Join [RFC6940].

## 8. Security Considerations

The normative security recommendations of Section 13 of [RFC6940] are applicable to this document. As a routing alternative, the security part of RPR conforms to Section 13.6 of [RFC6940], which describes routing security. RPR behaves like a DRR requesting node towards the destination node. The RPR relay peer is not necessarily an arbitrary node -- for example, a managed network, a bootstrap node, or a configured relay peer; it should be a trusted node, because a trusted node will be less of a risk, as outlined in Section 13 of [RFC6940].

In order to address possible DoS attacks, the relay peer SHOULD also limit the number of maximum connections; this is required in order to also reduce load on the relay peer, as explained in Section 4.1.

## 9. IANA Considerations

### 9.1. A New RELOAD Forwarding Option

A new RELOAD Forwarding Option type has been added to the "RELOAD Forwarding Option Registry" defined in [RFC6940].

Code: 2  
Forwarding Option: extensive\_routing\_mode

## 10. Acknowledgments

David Bryan helped extensively with this document and helped provide some of the text, analysis, and ideas contained here. The authors would like to thank Ted Hardie, Narayanan Vidya, Dondeti Lakshminath, Bruce Lowekamp, Stephane Bryant, Marc Petit-Huguenin, and Carlos Jesus Bernardos Cano for their constructive comments.

## 11. References

### 11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC6940] Jennings, C., Lowekamp, B., Rescorla, E., Baset, S., and H. Schulzrinne, "REsource LOcation And Discovery (RELOAD) Base Protocol", RFC 6940, January 2014.
- [RFC7263] Zong, N., Jiang, X., Even, R., and Y. Zhang, "An Extension to the REsource LOcation And Discovery (RELOAD) Protocol to Support Direct Response Routing", RFC 7263, June 2014.

### 11.2. Informative References

- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [RFC5780] MacDonald, D. and B. Lowekamp, "NAT Behavior Discovery Using Session Traversal Utilities for NAT (STUN)", RFC 5780, May 2010.

## Appendix A. Optional Methods to Investigate Peer Connectivity

This section is for informational purposes only and provides some mechanisms that can be used when the configuration information does not specify if RPR can be used. It summarizes some methods that can be used by a peer to determine its own network location compared with NAT. These methods may help a peer to decide which routing mode it may wish to try. Note that there is no foolproof way to determine whether a peer is publicly reachable, other than via out-of-band mechanisms. This document addresses UNilateral Self-Address Fixing (UNSAF) [RFC3424] considerations by specifying a fallback plan to SRR [RFC6940]. SRR is not an UNSAF mechanism. This document does not define any new UNSAF mechanisms.

For RPR to function correctly, a peer may attempt to determine whether it is publicly reachable. If it is not, RPR may be chosen to route the response with help from relay peers, or the peers should fall back to SRR. NATs and firewalls are two major contributors to preventing RPR from functioning properly. There are a number of techniques by which a peer can get its reflexive address on the public side of the NAT. After obtaining the reflexive address, a peer can perform further tests to learn whether the reflexive address is publicly reachable. If the address appears to be publicly reachable, the peer to which the address belongs can be a candidate to serve as a relay peer. Peers that are not publicly reachable may still use RPR to shorten the response path, with help from relay peers.

Some conditions that are unique in P2PSIP architecture could be leveraged to facilitate the tests. In a P2P overlay network, each peer has only a partial view of the whole network and knows of a few peers in the overlay. P2P routing algorithms can easily deliver a request from a sending peer to a peer with whom the sending peer has no direct connection. This makes it easy for a peer to ask other peers to send unsolicited messages back to the requester.

The approaches for a peer to get the addresses needed for further tests, as well as the test for learning whether a peer may be publicly reachable, are the same as those for DRR. Please refer to Appendix A of [RFC7263] for more details.

## Appendix B. Comparison of Cost of SRR and RPR

The major advantage of using RPR is that it reduces the number of intermediate peers traversed by the response. This reduces the load, such as processing and communication bandwidth, on those peers' resources.

### B.1. Closed or Managed Networks

As described in Section 3, many P2P systems run in a closed or managed environment (e.g., carrier networks), so network administrators would know that they could safely use RPR.

The number of hops for a response in SRR and in RPR are listed in the following table. Note that the same types of illustrative settings can be found in Appendix B.1 of [RFC7263].

Mode	Success	No. of Hops	No. of Msgs
SRR	Yes	$\log(N)$	$\log(N)$
RPR	Yes	2	2
RPR (DTLS)	Yes	2	7+2

Table 1: Comparison of SRR and RPR in Closed Networks

From the above comparison, it is clear that:

- 1) In most cases when the number of peers  $(N) > 4$  ( $2^2$ ), RPR uses fewer hops than SRR. Using a shorter route means less overhead and resource usage on intermediate peers, which is an important consideration for adopting RPR in the cases where such resources as CPU and bandwidth are limited, e.g., the case of mobile, wireless networks.
- 2) In the cases when  $N > 512$  ( $2^9$ ), RPR also uses fewer messages than SRR.
- 3) In the cases when  $N < 512$ , RPR uses more messages than SRR (but still uses fewer hops than SRR), so the consideration of whether to use RPR or SRR depends on other factors such as using less resources (bandwidth and processing) from the intermediate peers. Section 4 provides use cases where RPR has a better chance of working or where the considerations of intermediary resources are important.

## B.2. Open Networks

In open networks (e.g., the Internet) where RPR is not guaranteed to work, RPR can fall back to SRR if it fails after trial, as described in Section 4.2. Based on the same settings as those listed in Appendix B.1, the number of hops, as well as the number of messages for a response in SRR and RPR, are listed in the following table:

Mode	Success	No. of Hops	No. of Msgs
SRR	Yes	$\log(N)$	$\log(N)$
RPR	Yes	2	2
	Fail & fall back to SRR	$2+\log(N)$	$2+\log(N)$
RPR (DTLS)	Yes	2	7+2
	Fail & fall back to SRR	$2+\log(N)$	$9+\log(N)$

Table 2: Comparison of SRR and RPR in Open Networks

From the above comparison, it can be observed that trying to first use RPR could still provide an overall number of hops lower than directly using SRR. The detailed analysis is the same as that for DRR and can be found in [RFC7263].

## Authors' Addresses

Ning Zong  
Huawei Technologies

EEmail: zongning@huawei.com

Xingfeng Jiang  
Huawei Technologies

EEmail: jiang.x.f@huawei.com

Roni Even  
Huawei Technologies

EEmail: roni.even@mail01.huawei.com

Yunfei Zhang  
CoolPad / China Mobile

EEmail: hishigh@gmail.com