

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 3, 2015

B. Williams
Akamai
T. Reddy
Cisco
December 30, 2014

Peer-specific Redirection for Traversal Using Relays around NAT (TURN)
draft-williams-peer-redirect-03

Abstract

This specification describes a peer-specific redirection method that allows the TURN server to redirect a client for the purpose of improving communication with a specific peer without negatively affecting communication with other peers.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 3, 2015.

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
1.1. Redirection for Performance	3
1.2. Redirection for Load Balancing	5
2. Terminology	5
3. Peer-specific Server Redirect Mechanism	5
3.1. Attribute Usage	5
3.2. Sending a CreatePermission or ChannelBind Request	7
3.2.1. The CHECK-ALTERNATE Attribute	7
3.2.2. The XOR-OTHER-ADDRESS attribute	8
3.3. Receiving a CreatePermission or ChannelBind Request	8
3.4. Receiving a CreatePermission or ChannelBind Error Response	9
3.5. Receiving a CreatePermission or ChannelBind Success Response	10
4. Security Considerations	10
4.1. CHECK-ALTERNATE Flood	11
4.2. Unsolicited or Invalid ALTERNATE-SERVER	11
5. IANA Considerations	12
6. References	12
6.1. Normative References	12
6.2. Informative References	12
Authors' Addresses	13

1. Introduction

A Traversal Using Relay around NAT (TURN) [RFC5766] service provider may provide multiple candidate TURN servers for use by a host, but it might not be possible to determine which candidate TURN server will provide the best performance until both peers have been identified. This could be true for a variety of reasons, including:

- o Using the selected relay for a specific peer results in a sub-optimal end-to-end Internet path.
- o Load conditions on the selected relay have changed since the allocation was established such that it cannot support the new data flow.


At the same time, the above conditions might apply to one peer but not another, such that it would be best to selectively use the existing relay allocation for peers that will receive reasonable performance and redirect data flows for other peers to an alternate server. These scenarios are discussed in greater detail below.

The Session Traversal Utilities for NAT (STUN) protocol [RFC5389] defines an ALTERNATE-SERVER mechanism with which a server can redirect a client to another server by replying to a request message with an error response with error code 300 (Try Alternate). The TURN protocol describes error code 300 as one of the possible error codes for an Allocate error response.

This specification describes an additional use of the ALTERNATE-SERVER STUN attribute for TURN that allows the TURN server to redirect a client for the purpose of improving communication with a specific peer without negatively affecting communication with other peers. The client application indicates the nature of the desired response, which allows the client to treat the alternate server selection as either a requirement or a suggestion. This flexibility gives the client the option to choose the best way for the Interactive Connectivity Establishment (ICE) protocol [RFC5245] to respond (e.g. discarding the existing relay candidate for communication with this peer versus evaluating the two candidate servers using ICE connectivity checks and selecting the best one).

1.1. Redirection for Performance

Consider the following example:



When Peer B wishes to communicate with either Peer A or Peer C, it performs a DNS lookup and discovers TURN Relay C, the nearest of the candidate TURN servers. Peer B then sends a TURN Allocate request to TURN Relay C to determine the reflexive and relay candidates to offer. After the reflexive candidate has been chosen, Peer B sends a ChannelBind request to TURN Relay C to establish a channel for communication with the peer. If Peer C is the remote peer, the existing allocation will perform reasonably well, but if Peer A is the remote peer, the latency for relayed packets will be nearly twice as long as if TURN Relay A had been selected as the relay candidate. The problem is worse if Peer B wishes to communicate with both Peer A and Peer C, since there is no single relay candidate that would provide optimum performance for both peers.

If TURN Relay C and TURN Relay A are part of a common TURN service, it would be possible for TURN Relay C to determine that TURN Relay A will provide optimal service for communication between Peer B and Peer A. This allows the TURN service to redirect just the data channel between Peer A and Peer B to TURN relay A, thus providing optimal performance for both relay channels.

The above example describes the problem in terms of physical geography instead of network geography in order to help clarify the discussion. However, readers should note that the problem of selecting a relay server to achieve optimal end-to-end routing is much more complicated than the above description suggests, requiring a detailed real-time view of network connectivity characteristics and the peering relationships between autonomous systems. A naive approach based solely on the physical location of the hosts involved is just as likely to produce negative results as positive ones.

That said, a relay service provider with a broadly distributed system for actively monitoring network performance across the relevant parts of the Internet could make use of the resulting data set to select the optimal relay for each peer pair.

1.2. Redirection for Load Balancing

At the point when a relay allocation is first established, it can be difficult to determine how much aggregate concurrent load could eventually be associated with that allocation. The initiating peer could attempt to use that allocation for any number of peer-to-peer data flows over an extended period of time, during which time load conditions on the relay could change substantially, such that quality of service for already established flows would degrade if the relay were to accept additional flows.

Under these conditions, a TURN service provider with multiple relay hosts and distributed capacity could improve service quality by redirecting data flows to a different host that has more available capacity. At the same time, it is desirable to avoid disrupting established data flows by continuing to handle established flows on the current relay and only redirecting new flows elsewhere.

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 [RFC2119].

3. Peer-specific Server Redirect Mechanism

This specification describes two new uses of the existing STUN ALTERNATE-SERVER attribute. In the first case, the ALTERNATE-SERVER attribute is included with either a CreatePermission error response or a ChannelBind error response. In the second case, the ALTERNATE-SERVER attribute is included with either a CreatePermission success response or a ChannelBind success response.

This specification also defines two new comprehension-optional STUN attributes: CHECK-ALTERNATE and XOR-OTHER-ADDRESS. The CHECK-ALTERNATE attribute is used by the client to request that the server perform peer-specific redirection. The XOR-OTHER-ADDRESS is used by the client to provide an alternate peer address for location identification in the event that the XOR-PEER-ADDRESS attribute in the CreatePermission or ChannelBind request is not expected to reliably serve this purpose.

3.1. Attribute Usage

When sending a CreatePermission or a ChannelBind request, the CHECK-ALTERNATE STUN attribute allows a TURN client to indicate support for

peer-specific server redirection. To maintain backward compatibility with [RFC5766] compliant TURN servers that do not support peer-specific redirection, this attribute is defined as comprehension-optional, which allows a TURN server that does not support peer-specific redirection to ignore the attribute. To maintain backward compatibility with [RFC5766] compliant TURN clients that do not support peer-specific redirection, a TURN server only sends the ALTERNATE-SERVER attribute in CreatePermission and ChannelBind responses when the CHECK-ALTERNATE STUN attribute was present in the request. This prevents transmission of the ALTERNATE-SERVER attribute in cases where the receiving client might not consider the usage legitimate.

The CHECK-ALTERNATE STUN attribute's value indicates the expected server response type: error or success. This capability to declare the expected response type allows TURN client implementers greater flexibility during session establishment. For example, a TURN client implementer may wish to maintain the smallest number of permissions possible during session establishment in order keep the internal client implementation simple, in which case an error response would be desirable. On the other hand, a TURN client implementer may wish to optimize for faster session establishment by continuing to use a sub-optimal allocation while setting up the new one, in which case a success response would be desirable. This second case could be achieved with an error response if the client were to send a second request without the CHECK-ALTERNATE attribute, but such an approach would require an extra RTT.

The XOR-OTHER-ADDRESS STUN attribute allows the TURN client to provide an alternate peer address that can be used by the server to identify the network geographic location of the peer when performing the peer-specific redirection check. Use of this attribute is only necessary if the XOR-PEER-ADDRESS already contained in the CreatePermission or ChannelBind request does not adequately serve this purpose, which should only be true when both peers require a TURN relay for end-to-end data flow. In this case, the TURN CreatePermission or ChannelBind request will provide the peer's TURN relay address as the XOR-PEER-ADDRESS value. If the RTT between the peer and its TURN relay server is very small, the TURN relay address might still be an appropriate address to use for the peer-specific redirection check. As the RTT grows, the TURN relay address will become less suitable for this purpose. For this reason, it is generally the case that the peer's public address (i.e. its host or reflexive address) is a better indication of its network geographic location than its TURN relay address.

Even in cases where both peers require a TURN relay, a typical ICE protocol implementation will give higher candidate priority to the

peer's host and reflexive addresses, which means that the first CreatePermission or ChannelBind request will provide the peer's public address as the XOR-PEER-ADDRESS value and no XOR-OTHER-ADDRESS attribute is necessary. However, although ICE recommends this priority, it does not require it, and so the first request may contain the peer's TURN relay address. With such an implementation, the XOR-OTHER-ADDRESS attribute allows the client to provide the peer's reflexive address in a request that populates the XOR-PEER-ADDRESS attribute with the peer's relay address.

3.2. Sending a CreatePermission or ChannelBind Request

A client that supports peer-specific server redirection and desires such redirection to be performed MUST include the CHECK-ALTERNATE attribute in the first CreatePermission or ChannelBind request when that request is expected to form a new permission or binding. A client MUST NOT include the CHECK-ALTERNATE attribute in a CreatePermission or ChannelBind request that is intended to extend the lifetime of an existing permission or binding.

Peer-specific server redirection is only supported for requests that include a single XOR-PEER-ADDRESS attribute. When forming a CreatePermission request with multiple XOR-PEER-ADDRESS attributes, the client MUST NOT include the CHECK-ALTERNATE attribute.

When the CreatePermission or ChannelBind request includes the CHECK-ALTERNATE attribute, the client MAY also include an XOR-OTHER-ADDRESS attribute with a value appropriate for the above described purpose. The XOR-OTHER-ADDRESS attribute SHOULD NOT be included in the request if its value will be identical to the request's XOR-PEER-ADDRESS attribute.

3.2.1. The CHECK-ALTERNATE Attribute

When forming a CHECK-ALTERNATE attribute, the STUN Type is TBD-CA. This type is in the comprehension-optional range, which means that STUN agents can safely ignore the attribute if they do not understand it.

The CHECK-ALTERNATE attribute takes a 1-byte Value, which means that the Length is 1 and 3 bytes of padding are required after the Value. The format of the Value is:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
|E|      RFFU      |
+---+---+---+---+---+---+
```

The Value contains a single 1-bit flag:

E: If 1, the server is requested to send a Try Alternate (300) error response when redirection is expected. If 0, the server is request to include an ALTERNATE-SERVER attribute in the success response for the request.

The other 7 bits of the attribute's value must be set to zero on transmission and ignored on reception.

3.2.2. The XOR-OTHER-ADDRESS attribute

When forming an XOR-OTHER-ADDRESS attribute, the STUN Type is TBD-XOA. This type is in the comprehension-optional range, which means that STUN agents can safely ignore the attribute if they do not understand it.

The XOR-OTHER-ADDRESS value specifies an address and port suitable for identification of the peer's network geographic location. It is encoded in the same way as XOR-MAPPED-ADDRESS [RFC5389].

3.3. Receiving a CreatePermission or ChannelBind Request

When a server receives a CreatePermission or ChannelBind request that includes a CHECK-ALTERNATE attribute, it processes as per the TURN specification [RFC5766] plus the specific rules mentioned here.

The server checks the following:

- o If the CHECK-ALTERNATE attribute is not recognized, ignore the attribute because its type indicates that it is comprehension-optional. This should be the existing behavior.
- o If the message is a CreatePermission request with multiple XOR-PEER-ADDRESS attributes, ignore the CHECK-ALTERNATE attribute if present.
- o If peer-specific redirection is not supported by the server, ignore the attribute.
- o If the associated permission or binding already exists, ignore the attribute.

If none of the above causes the attribute to be ignored and no other cause for sending an error response has been found, the server attempts to identify an alternate server that will provide better performance for the session based on the criteria supported by the TURN service (e.g. optimal data path and/or load balancing). When an

XOR-OTHER-ADDRESS attribute is found in the request message, the server SHOULD use this address for peer location identification. Otherwise, the server SHOULD use the address provided in the XOR-PEER-ADDRESS attribute.

If no alternate server is identified, the server replies with a success response that does not include an ALTERNATE-SERVER attribute.

If an alternate server is identified and the client requested an error response for redirection, the server rejects the request with a 300 (Try Alternate) error. No new permission or binding is generated on the server in this case.

If an alternate server is identified and the client did not request an error response for redirection, the server creates the permission or binding. The server then replies to the request with a success response, including an ALTERNATE-SERVER attribute in the message.

3.4. Receiving a CreatePermission or ChannelBind Error Response

If the client receives a CreatePermission or ChannelBind error response with error code 420 (Unknown Attribute) and CHECK-ALTERNATE is listed in the UNKNOWN-ATTRIBUTE attribute of the message, the client SHOULD retransmit the original request without the CHECK-ALTERNATE attribute. This case is not expected due to the use of a comprehension-optional attribute type.

If the client receives a CreatePermission or ChannelBind error response with error code 300 (Try Alternate), the client SHOULD attempt to form an allocation to the TURN server indicated in the ALTERNATE-SERVER attribute.

If the alternate server responds to the Allocate request with a success response, the client SHOULD attempt to form a new permission or binding using the new allocation from the alternate server. The CreatePermission or ChannelBind request to the alternate server MAY include a CHECK-ALTERNATE attribute but SHOULD NOT request redirection via an error response. This helps to avoid the possibility of redirection loops.

If the alternate server responds to the Allocate request with an error response, the client MAY resend the original CreatePermission or ChannelBind request, either without the CHECK-ALTERNATE attribute or with a CHECK-ALTERNATE attribute that does not request an error response.

See Section 4 below for discussion of how the client should respond when receiving a Try Alternate error response that was not requested.

3.5. Receiving a CreatePermission or ChannelBind Success Response

If the client receives a CreatePermission or ChannelBind success response, it proceeds with processing according to the TURN specification [RFC5766]. If the message does not include an ALTERNATE-SERVER attribute, no additional processing is required.

If the success response includes an ALTERNATE-SERVER attribute, the client SHOULD attempt to form an allocation to the TURN server indicated in the ALTERNATE-SERVER attribute.

If the alternate server responds to the Allocate request with a success response, the client SHOULD attempt to form a new permission or binding using the new allocation from the alternate server. The CreatePermission or ChannelBind request to the alternate server MAY include a CHECK-ALTERNATE attribute with either attribute value. If this is done, care should be taken in the client implementation to recognize and avoid redirection loops.

While waiting for the new allocation and permission or binding to form via the indicated alternate server, the client SHOULD use the original permission or binding from the request that included the CHECK-ALTERNATE attribute. In this way, peer-specific redirection without an error response can be considered a "hint" that allows the client to establish an alternate path and test its quality before switching to it.

See Section 4 below for discussion of how the client should respond when receiving an ALTERNATE-SERVER attribute that was not requested.

4. Security Considerations

This section considers attacks that are possible in a TURN deployment through the specified protocol extension, and discusses how they are mitigated by mechanisms in the protocol or recommended practices in the implementation.

The specified mechanism affects the use of TURN CreatePermission request messages, ChannelBind request messages, and their respective success and error response messages. Each of these TURN message types requires the MESSAGE-INTEGRITY STUN attribute, which limits attacks that attempt to make use of the specified mechanism to authenticated clients and servers.

4.1. CHECK-ALTERNATE Flood

A compromised TURN client could send a large number of CreatePermission or ChannelBind request messages, which would drive increased load on the TURN server. The CHECK-ALTERNATE attribute does not make such an attack more likely, though it could make it possible to increase the impact of such an attack due to the additional load associated with determining whether an alternate server should be used by the client. The TURN server MAY be configured to ignore the CHECK-ALTERNATE attribute under some conditions in order to limit the associated load. The conditions under which it is appropriate for a TURN server to ignore the CHECK-ALTERNATE attribute are implementation dependent.

4.2. Unsolicited or Invalid ALTERNATE-SERVER

A compromised TURN server could send the "Try Alternate" error code in response to a request message that did not contain the CHECK-ALTERNATE attribute or where the value of the attribute did not request an error response. For client connectivity, this is no worse than any other error response code that could be sent. No matter what the error response code may be, the client is unable to relay data to the remote peer. The client MUST ignore the ALTERNATE-SERVER attribute in error responses when the CHECK-ALTERNATE attribute was not included in the associated request. The client SHOULD ignore the ALTERNATE-SERVER attribute in error responses when the CHECK-ALTERNATE attribute was included in the associated request if the attribute value did not request an error response. The client MAY discontinue use of the associated TURN allocation when an unsolicited Try Alternate error is received.

A compromised TURN server could send an ALTERNATE-SERVER attribute in a success response message for a request message that did not contain the CHECK-ALTERNATE attribute. The client MUST ignore the ALTERNATE-SERVER attribute in success responses when the CHECK-ALTERNATE attribute was not included in the associated request message. The client SHOULD ignore the ALTERNATE-SERVER attribute in success responses when the CHECK-ALTERNATE attribute was included in the associated request if the attribute value requested an error response. The client MAY discontinue use of the associated TURN allocation when an unsolicited ALTERNATE-SERVER attribute is received.

A compromised TURN server could send an invalid ALTERNATE-SERVER attribute value in either an error or a success response message, where the value refers to an unaffiliated TURN server to which the sending TURN server is not allowed to redirect traffic. Such an attack is already allowed by the use of Try Alternate errors in

response to Allocate request messages. Use of the ALTERNATE-SERVER attribute in the context of peer-specific redirection does not make such an attack more likely, though it could make it possible to increase the scale of such an attack by allowing multiple ALTERNATE-SERVER attributes to each client, one per requested permission or binding. A client SHOULD ignore all future ALTERNATE-SERVER attributes received from the TURN server after an authentication failure with any server identified via an ALTERNATE-SERVER attribute. A client MAY discontinue use of the associated TURN allocation after an authentication failure with any server identified via an ALTERNATE-SERVER attribute.

5. IANA Considerations

[Paragraphs below in braces should be removed by the RFC Editor upon publication]

[The CHECK-ALTERNATE attribute requires that IANA allocate a value in the "STUN attributes Registry" from the comprehension-optional range (0x8000-0xFFFF), to be replaced for TBD-CA throughout this document]

This document defines the CHECK-ALTERNATE STUN attribute, described in Section 3.2.1. IANA has allocated the comprehension-optional codepoint TBD-CA for this attribute.

[The XOR-OTHER-ADDRESS attribute requires that IANA allocate a value in the "STUN attributes Registry" from the comprehension-optional range (0x8000-0xFFFF), to be replaced for TBD-XOA throughout this document]

This document defines the XOR-OTHER-ADDRESS STUN attribute, described in Section 3.2.2. IANA has allocated the comprehension-optional codepoint TBD-XOA for this attribute.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

6.2. Informative References

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245,

April 2010.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing,
"Session Traversal Utilities for NAT (STUN)", RFC 5389,
October 2008.

[RFC5766] Mahy, R., Matthews, P., and J. Rosenberg, "Traversal Using
Relays around NAT (TURN): Relay Extensions to Session
Traversal Utilities for NAT (STUN)", RFC 5766, April 2010.

Authors' Addresses

Brandon Williams
Akamai, Inc.
8 Cambridge Center
Cambridge, MA 02142
USA

Email: brandon.williams@akamai.com

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tiredy@cisco.com

