

TRAM
Internet-Draft
Intended status: Standards Track
Expires: December 21, 2014

T. Reddy
P. Patil
R. Ravindranath
Cisco
J. Uberti
Google
June 19, 2014

TURN Extension for Third Party Authorization
draft-reddy-tram-turn-third-party-authz-03

Abstract

This document proposes the use of OAuth to obtain and validate ephemeral tokens that can be used for TURN authentication. The usage of ephemeral tokens ensure that access to a TURN server can be controlled even if the tokens are compromised, as is the case in WebRTC where TURN credentials must be specified in Javascript.

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 December 21, 2014.

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	2
2. Terminology	3
3. Solution Overview	3
4. Obtaining a Token Using OAuth	5
4.1. Key Establishment	7
4.1.1. DSKPP	8
4.1.2. HTTP interactions	8
4.1.3. Manual provisioning	9
5. Forming a Request	10
6. STUN Attributes	10
6.1. THIRD-PARTY-AUTHORIZATION	10
6.2. ACCESS-TOKEN	10
7. Receiving a request with ACCESS-TOKEN attribute	12
8. Changes to TURN Client	13
9. Security Considerations	13
10. IANA Considerations	13
11. Acknowledgements	14
12. References	14
12.1. Normative References	14
12.2. Informative References	14
Authors' Addresses	15

1. Introduction

Traversal Using Relay NAT (TURN) TURN [RFC5766] is a protocol that is often used to improve the connectivity of P2P applications. By providing a cloud-based relay service, TURN ensures that a connection can be established even when one or both sides is incapable of a direct P2P connection. However, as a relay service, it imposes a nontrivial cost on the service provider. Therefore, access to a TURN service is almost always access-controlled.

TURN provides a mechanism to control access via "long-term" username/password credentials that are provided as part of the TURN protocol. It is expected that these credentials will be kept secret; if the credentials are discovered, the TURN server could be used by unauthorized users or applications. However, in web applications, ensuring this secrecy is typically impossible. To address this problem and the ones described in [I-D.ietf-tram-auth-problems], this document proposes the use of third party authorization using OAuth for TURN.

To achieve third party authorization, a resource owner e.g. WebRTC server, authorizes a TURN client to access resources on the TURN server.

Using OAuth, a client obtains an ephemeral token from an authorization server e.g. WebRTC server, and the token is presented to the TURN server instead of the traditional mechanism of presenting username/password credentials. The TURN server validates the authenticity of the token and provides required services.

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].

- o WebRTC Server: A web server that supports WebRTC [I-D.ietf-rtcweb-overview].
- o Access Token: OAuth 2.0 access token.
- o mac_key: The session key generated by the authorization server. Note that the lifetime of the session key is equal to the lifetime of the access token.
- o kid: An ephemeral and unique key identifier. The kid also allows the resource server to select the appropriate keying material for decryption.

3. Solution Overview

This specification uses the token type 'Assertion' (aka self-contained token) described in [RFC6819] where all the information necessary to authenticate the validity of the token is contained within the token itself. This approach has the benefit of avoiding a protocol between the TURN server and the authorization server for token validation, thus reducing latency. The exact mechanism used by a client to obtain a token from the OAuth authorization server is outside the scope of this document. For example, a client could make an HTTP request to an authorization server to obtain a token that can be used to avail TURN services. The TURN token is returned in JSON, along with other OAuth Parameters like token type, mac_key, kid, token lifetime etc. The client is oblivious to the content of the token. The token is embedded within a TURN request sent to the TURN server. Once the TURN server has determined the token is valid, TURN services are offered for a determined period of time.




Figure 1: TURN Third Party Authorization

Note : An implementation may choose to contact the WebRTC server to obtain a token even before it makes an allocate request, if it knows the server details before hand. For example, once a client has learnt that a TURN server supports Third Party authorization from a WebRTC server, the client can obtain the token before making subsequent allocate requests.

For example, the client learns the TURN server name "turn1@example.com" from THIRD-PARTY-AUTHORIZATION attribute value and makes the following HTTP request for the access token using transport-layer security (with extra line breaks for display purposes only):

```
POST /o/oauth2/token HTTP/1.1
Audience: turn1@example.com
Content-Type: application/x-www-form-urlencoded
timestamp=1361471629
grant_type=implicit
```

Figure 2: Request

If the client is authorized then the authorization server issues an access token. An example of successful response:

```
HTTP/1.1 200 OK
Content-Type: application/json
Cache-Control: no-store

{
  "access_token":
  "U2FsdGVkXl8qJK/kkWmRcnfHglrVTJSpsS6yU32kmHmOrfGyI3m1gQj1jRPsr0uBb
  HctuycAgsfRX7nJW2BdukGyKMXSiNGNnBzigkAofP6+Z3vkJlQ5pWbfSRroOkWBn",
  "token_type": "mac",
  "expires_in": 1800,
  "kid": "22BIjxU93h/IgwEb",
  "mac_key": "v51N62OM65kyMvfTI080"
}
```

Figure 3: Response

Access token and other attributes issued by the authorization server are explained in Section 6.2.

4. Obtaining a Token Using OAuth


A TURN client should know the authentication capability of the TURN server before deciding to use third party authorization with it. A TURN client initially makes a request without any authorization. If the TURN server supports or mandates third party authorization, it will return an error message indicating support for third party authorization. The TURN server includes an ERROR-CODE attribute with a value of 401 (Unauthorized), a nonce value in a NONCE attribute and a SOFTWARE attribute that gives information about the TURN server's software. The TURN servers also includes additional STUN attribute THIRD-PARTY-AUTHORIZATION signaling the TURN client that the TURN server supports third party authorization.

The following mapping of OAuth concepts to WebRTC is used :

OAuth	WebRTC
Client	WebRTC client
Resource owner	WebRTC server
Authorization server	Authorization server
Resource server	TURN Server

Figure 4: OAuth terminology mapped to WebRTC terminology

Using the OAuth 2.0 authorization framework, a WebRTC client (third-party application) obtains limited access to a TURN (resource server) on behalf of the WebRTC server (resource owner or authorization server). The WebRTC client requests access to resources controlled by the resource owner (WebRTC server) and hosted by the resource server (TURN server). The WebRTC client obtains access token, lifetime, session key (in the mac_key parameter) and key id (kid). The TURN client conveys the access token and other OAuth parameters learnt from the authorization server to the resource server (TURN server). The TURN server obtains the session key from the access token. The TURN server validates the token, computes the message integrity of the request and takes appropriate action i.e permits the TURN client to create allocations. This is shown in an abstract way in Figure 5.



User : Alice

****: Out-of-Band Long-Term Key Establishment

Figure 5: Interactions

OAuth in [RFC6749] defines four grant types. This specification uses the OAuth grant type "Implicit" explained in section 1.3.2 of [RFC6749] where the WebRTC client is issued an access token directly. The scope of the access token explained in section 3.3 of [RFC6749] MUST be TURN.

4.1. Key Establishment

The TURN and authorization servers MUST establish a symmetric key (K), using an out of band mechanism. Symmetric key MUST be chosen to ensure that the size of encrypted token is not large because usage of asymmetric keys will result in large encrypted tokens which may not fit into a single STUN message. The AS-RS, AUTH keys will be derived from K. AS-RS key is used for encrypting the self-contained token and message integrity of the encrypted token is calculated using the AUTH key. The TURN and authorization servers MUST establish the symmetric key over an authenticated secure channel. The establishment of symmetric key is outside the scope of this specification. For example, implementations could use one of the following mechanisms in to establish a symmetric key.

4.1.1.1. DSKPP

The two servers could choose to use Dynamic Symmetric Key Provisioning Protocol (DSKPP) [RFC6063] to establish a symmetric key (K). The encryption and MAC algorithms will be negotiated using the KeyProvClientHello, KeyProvServerHello messages. A unique key identifier (referred to as KeyID) for the symmetric key is generated by the DSKPP server (i.e. Authorization server) and signalled to the DSKPP client (i.e TURN server) which is equivalent to the kid defined in this specification. The AS-RS, AUTH keys would be derived from the symmetric key using (HMAC)-based key derivation function (HKDF) [RFC5869] and the default hash function is SHA-256. For example if the input symmetric key (K) is 32 octets length, encryption algorithm is AES_128_CBC and HMAC algorithm is HMAC-SHA-256-128 then the secondary keys AS-RS, AUTH are generated from the input key K as follows

1. HKDF-Extract(zero, K) -> PRK
2. HKDF-Expand(PRK, zero, 16) -> AS-RS key
3. HKDF-Expand(PRK, zero, 32) -> AUTH key

4.1.1.2. HTTP interactions

The two servers could choose to use REST API to establish a symmetric key. To retrieve a new symmetric key, the TURN server makes an HTTP GET request to the authorization server, specifying TURN as the service to allocate the symmetric keys for, and specifying the name of the TURN server. The response is returned with content-type "application/json", and consists of a JSON object containing the symmetric key.

Request

service - specifies the desired service (turn)
name - TURN server name be associated with the key

example: GET /?service=turn&name=turn1@example.com

Response

key - Long-term key (K)
ttl - the duration for which the key is valid, in seconds.

example:

```
{
  "key" :
  "ESIZRFVmd4iZABEiM0RVZgKn6WjLaTC1FXAghRMVTzkBGNaan496523WIIKerLi",
  "ttl" : 86400,
  "kid" : "22BIjxU93h/IgwEb"
}
```

The AS-RS, AUTH keys are derived from K using HKDF as discussed in Section 4.1.1. Authorization server must also signal a unique key identifier (kid) to the TURN server which will be used to select the appropriate keying material for decryption. The default encryption algorithm to encrypt the self-contained token could be Advanced Encryption Standard (AES) in Cipher Block Chaining (CBC) mode (AES_128_CBC). The default HMAC algorithm to calculate the integrity of the token could be HMAC-SHA-256-128. In this case AS-RS key length must be 128-bit, AUTH key length must be 256-bit (section 2.6 of [RFC4868]).

4.1.3. Manual provisioning

TURN and authorization servers could be manually configured with a symmetric key (K) and kid. The default encryption and HMAC algorithms could be AES_256_CBC, HMAC-SHA-256-128.

Note : The mechanisms specified in Section 4.1.2 Section 4.1.3 are easy to implement and deploy compared to DSKPP but lack encryption and HMAC algorithm agility.

5. Forming a Request

When a TURN server responds that third party authorization is required, a TURN client re-attempts the request, this time including access token and kid values in ACCESS-TOKEN and USERNAME STUN attributes. The TURN client includes a MESSAGE-INTEGRITY attribute as the last attribute in the message over the contents of the TURN message. The HMAC for the MESSAGE-INTEGRITY attribute is computed as described in section 15.4 of [RFC5389] where the mac_key is used as the input key for the HMAC computation. The TURN client and server will use the mac_key to compute the message integrity and doesn't have to perform MD5 hash on the credentials.

6. STUN Attributes

The following new STUN attributes are introduced by this specification to accomplish third party authorization.

6.1. THIRD-PARTY-AUTHORIZATION

This attribute is used by the TURN server to inform the client that it supports third party authorization. This attribute value contains the TURN server name. The TURN server may have tie-up with multiple authorization servers and vice versa, so the client MUST provide the TURN server name to the authorization server so that it can select the appropriate keying material to generate the self-contained token. The THIRD-PARTY-AUTHORIZATION attribute is a comprehension-optional attribute (see Section 15 from [RFC5389]).

6.2. ACCESS-TOKEN

The access token is issued by the authorization server. OAuth does not impose any limitation on the length of the access token but if path MTU is unknown then STUN messages over IPv4 would need to be less than 548 bytes (Section 7.1 of [RFC5389]), access token length needs to be restricted to fit within the maximum STUN message size. Note that the self-contained token is opaque to the client and it MUST NOT examine the ticket. The ACCESS-TOKEN attribute is a comprehension-optional attribute (see Section 15 from [RFC5389]).

The token is structured as follows:

```
struct {  
    opaque {  
        ushort key_length;  
        opaque mac_key[key_length];  
        opaque timestamp[8];  
        long    lifetime;  
    } encrypted_block;  
    opaque mac[mac_length];  
} token;
```

Figure 6: Self-contained token format

The fields are described below:

key_length: Length of the session key. Key length of 160-bits MUST be supported (i.e only 160-bit key is used by HMAC-SHA-1 for message integrity of STUN message). The key length facilitates the hash agility plan discussed in section 16.3 of [RFC5389].

mac_key: The session key generated by the authorization server.

Timestamp: 64-bit unsigned integer field containing a timestamp. The value indicates the time since January 1, 1970, 00:00 UTC, by using a fixed point format. In this format, the integer number of seconds is contained in the first 48 bits of the field, and the remaining 16 bits indicate the number of 1/64K fractions of a second (Native format - Unix).

Lifetime: The lifetime of the access token, in seconds. For example, the value 3600 indicates one hour. The Lifetime value SHOULD be equal to the "expires_in" parameter defined in section 4.2.2 of [RFC6749].

mac: The Hashed Message Authentication Code (HMAC) is calculated with AUTH key over the encrypted portion of the token and the TURN server name (N) conveyed in the THIRD-PARTY-AUTHORIZATION response. Encryption is applied before authentication on the sender side and conversely on the receiver side. The length of the mac field is known to the TURN and authorization server based on the negotiated MAC algorithm.

For example the encryption process can be illustrated as follows. Here C, N denote the ciphertext and TURN server name.

- o C = AES_128_CBC(AS-RS, encrypted_block)
- o mac = HMAC-SHA-256-128(AUTH, C || N)

The token MUST be encoded as defined in Section 4 of [RFC4648] and then encrypted using the symmetric long-term key established between the resource server and the authorization server, as shown in Figure 5 as AS-RS key. HMAC is computed using the encrypted portion of the token and TURN server name to ensure that the client does not use the same token to gain illegal access to other TURN servers provided by the same administrative domain. This attack is possible when multiple TURN servers in a single administrative domain share the same symmetric key with the authorization server. Since the access token is valid for a specific period of time the resource server MUST cache it so that it need not to be provided in every request within an existing allocation. The access token can be re-used for multiple Allocate requests to the same TURN server.

The TURN client MUST include the ACCESS-TOKEN attribute only in Allocate and Refresh requests.

7. Receiving a request with ACCESS-TOKEN attribute

The TURN server, on receiving a request with ACCESS-TOKEN attribute, performs checks listed in section 10.2.2 of [RFC5389] in addition to the following steps to verify that the access token is valid:

- o TURN server selects the keying material based on kid signalled in the USERNAME attribute.
- o It performs the verification of the token message integrity by calculating HMAC over the encrypted portion in the self-contained token and TURN server name using AUTH key and if the resulting value does not match the mac field in the self-contained token then it rejects the request with an error response 401 (Unauthorized).
- o TURN server obtains the mac_key by retrieving the content of the access token (which requires decryption of the self-contained token using the AS-RS key).
- o The TURN server verifies that no replay took place by performing the following check:
 - * The access token is accepted if the timestamp field (TS) in the self-contained token is recent enough to the reception time of the TURN request (RDnew) using the following formula: $\text{Lifetime} + \text{Delta} > \text{abs}(\text{RDnew} - \text{TS})$. The RECOMMENDED value for the allowed Delta is 5 seconds. If the timestamp is NOT within the boundaries then the TURN server discards the request with error response 401 (Unauthorized).

- o The TURN server uses the mac_key to compute the message integrity over the request and if the resulting value does not match the contents of the MESSAGE-INTEGRITY attribute then it rejects the request with an error response 401 (Unauthorized).
- o If all the checks pass, the TURN server continues to process the request. Any response generated by the server MUST include the MESSAGE-INTEGRITY attribute, computed using the mac_key.

The lifetime provided by the TURN server in the Allocate and Refresh responses MUST be less than or equal to the lifetime of the token.

8. Changes to TURN Client

- o A TURN response is discarded by the client if the value computed for message integrity using mac_key does not match the contents of the MESSAGE-INTEGRITY attribute.
- o If the access token expires then the client MUST obtain a new token from the authorization server and use it for new allocations. The client MUST also use the new token to refresh existing allocations. This way client has to maintain only one token per TURN server.

9. Security Considerations

When OAuth is used the interaction between the client and the authorization server requires Transport Layer Security (TLS) with a ciphersuite offering confidentiality protection. The session key MUST NOT be transmitted in clear since this would completely destroy the security benefits of the proposed scheme. If an attacker tries to replay message with ACCESS-TOKEN attribute then the server can detect that the transaction ID as used for an old request and thus prevent the replay attack.

Security considerations discussed in [I-D.ietf-oauth-v2-http-mac] and [RFC5766] are to be taken into account.

10. IANA Considerations

IANA is requested to add the following attributes to the STUN attribute registry [iana-stun],

- o THIRD-PARTY-AUTHORIZATION
- o ACCESS-TOKEN

11. Acknowledgements

Authors would like to thank Dan Wing, Pal Martinsen, Oleg Moskalenko and Charles Eckel for comments and review. The authors would like to give special thanks to Brandon Williams for his help.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", RFC 4868, May 2007.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC6749] Hardt, D., "The OAuth 2.0 Authorization Framework", RFC 6749, October 2012.
- [iana-stun] IANA, , "IANA: STUN Attributes", April 2011, <<http://www.iana.org/assignments/stun-parameters/stun-parameters.xml>>.

12.2. Informative References

- [I-D.ietf-oauth-v2-http-mac] Richer, J., Mills, W., Tschofenig, H., and P. Hunt, "OAuth 2.0 Message Authentication Code (MAC) Tokens", draft-ietf-oauth-v2-http-mac-05 (work in progress), January 2014.
- [I-D.ietf-rtcweb-overview] Alvestrand, H., "Overview: Real Time Protocols for Browser-based Applications", draft-ietf-rtcweb-overview-10 (work in progress), June 2014.
- [I-D.ietf-tram-auth-problems] Reddy, T., R, R., Perumal, M., and A. Yegin, "Problems with STUN long-term Authentication for TURN", draft-ietf-tram-auth-problems-01 (work in progress), May 2014.

- [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.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, May 2010.
- [RFC6063] Doherty, A., Pei, M., Machani, S., and M. Nystrom, "Dynamic Symmetric Key Provisioning Protocol (DSKPP)", RFC 6063, December 2010.
- [RFC6819] Lodderstedt, T., McGloin, M., and P. Hunt, "OAuth 2.0 Threat Model and Security Considerations", RFC 6819, January 2013.

Authors' Addresses

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

Email: tireddy@cisco.com

Prashanth Patil
Cisco Systems, Inc.
Bangalore
India

Email: praspati@cisco.com

Ram Mohan Ravindranath
Cisco Systems, Inc.
Cessna Business Park,
Kadabeesanahalli Village, Varthur Hobli,
Sarjapur-Marathahalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: rmohanr@cisco.com

Justin Uberti
Google
747 6th Ave S
Kirkland, WA
98033
USA

Email: justin@uberti.name