Network Working Group
Internet-Draft
Expires: December 30, 2002

M. Blanchet
Viagenie
O. Medina
ENST Bretagne
F. Parent
Viagenie

July 1, 2002

DSTM IPv4 over IPv6 tunnel profile for Tunnel Setup Protocol(TSP)
draft-blanchet-ngtrans-tsp-dstm-profile-01

Status of this Memo

Copyright Notice

Abstract

Based on the actions they perform, The network model presented in DSTM [1] defines three types of equipments: a DSTM server, DSTM nodes and a Tunnel End Point (TEPs). Within this model, a protocol is required for configuration data exchange among these equipments. This document presents a method to perform these actions based on TSP [2].

Table of Contents

1. Introduction

Based on the actions they perform, The network model presented in DSTM [1] defines three types of equipments: a DSTM server, DSTM nodes and a Tunnel End Point (TEPs). Within this model, a protocol is required for configuration data exchange among these equipments. This document presents a method to perform these actions based on TSP [2].

The Tunnel Setup Protocol, TSP, is a protocol designed to negotiate tunnel information, such as IP addresses, network prefixes and routing information. TSP provides optional authentication, transport over IPv6 and redundancy of the service. Other protocols, such as DHCPv6 [4], can be used to deploy DSTM but, in the short term, such protocols may be more complex to implement.

The use of TSP for DSTM address allocation and tunnel set up demands the definition of four types of messages:

- 'Tunnel Create' messages are used to request the establishment of a 4over6 tunnel between a node and a given TEP. For first-time requests, tunnel creation implies the allocation of a temporary IPv4 address to the requesting node. In addition, this type of message is also used to ask for extension of the validity of an already allocated address.

- 'Tunnel Delete' messages are sent by the server to destroy an existing 4over6 tunnel. The server MUST send this type of message to the client (and to the TEP, if server and TEP are not co-located) when the allocation timer for a given address expires.

- 'Tunnel Info' messages are sent as a reply to Tunnel Create or Tunnel Delete requests. This type of message may contain configuration data to be used by a node, or simply confirm the creation/deletion of a 4over6 tunnel.

- Finally, Error Messages inform about the impossibility to allocate a temporary address or establish a 4over6 tunnel.

TSP provides authentication services using SASL [5]. If DSTM client authentication is required, the DSTM server can be configured to negotiate with the client the authentication scheme that will be used. In this mode, only authenticated clients are authorized to receive an IPv4 address. If no authentication is required, the ANONYMOUS authentication scheme can be used to allow any client to receive a temporary IPv4 address.

2. General Description of the Protocol

Figure 2.1 presents the message exchanges required by DSTM when the allocation process is started by a DSTM node. In this document we do not explain the different mechanisms that can be put in place to detect the need of an IPv4 address in a Dual Stack node. As required by DSTM, all TSP message exchanges take place in IPv6 using TCP transport. Remark that exchanges between DSTM Servers and TEPs are not required if both functionalities are implemented in the same host.

The allocation process greatly depends on a parameter called "Lifetime". It specifies the time (in seconds) over which an IPv4 address is assigned to a node, defining implicitly how often requests for allocation renewals are to be sent.

TSP message exchange starts whenever a DSTM node requires an IPv4 address. The node may start the exchange, but it may also be possible that DSTM servers send Unsolicited Allocation messages to nodes. This would be useful for implementations where it is allowed to originate connections from outside the DSTM domain (probably using a DNS-ALG). The exact description of this possibility is outside the scope of this document.

Address Allocation Process using SAAP

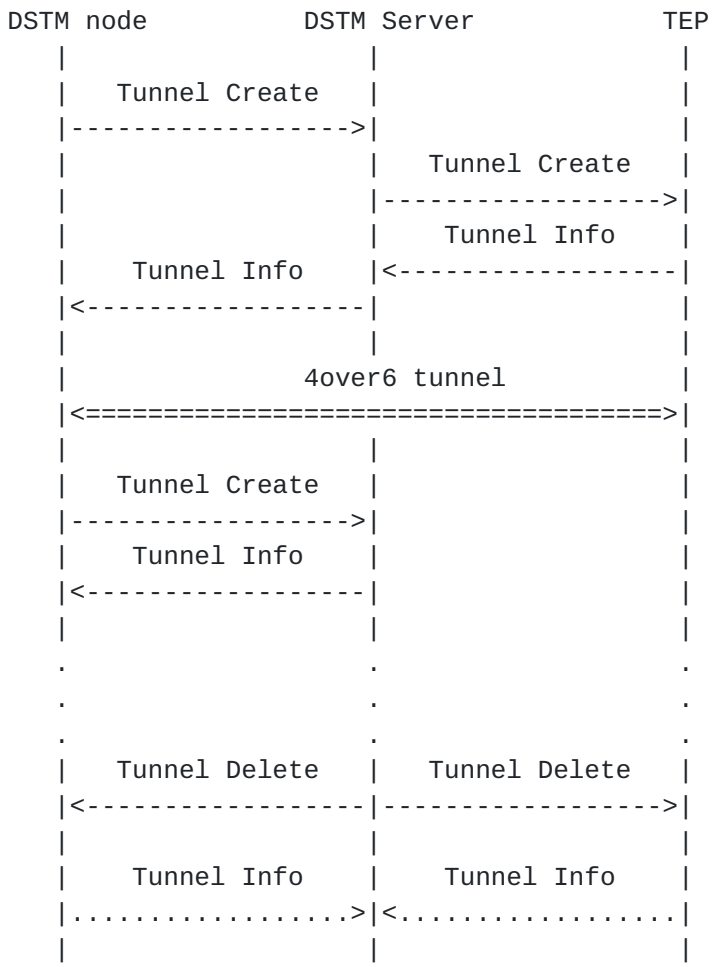

As shown in the figure, DSTM makes use of three types of TSP message: Create, Delete and Info. 'Tunnel Create' messages are sent by a DSTM node to ask for 4over6 Tunnel Configuration Parameters (implicitly including the request for a temporary IPv4 address). The same type of message is used by the DSTM server to configure the TEP and by the DSTM node to ask for renewal of the allocation. 'Tunnel Info' messages are usually sent as a reply to a previous 'Tunnel Create' request. Such a message may also be used to acknowledge the reception of a 'Tunnel Delete' command. Finally, DSTM servers send 'Tunnel Delete' messages to destroy 4over6 tunnels when the allocation time for an address expires.

2.1 Initial Address Allocation

As described in TSP [2], the first phase in TSP involves authentication (which can be ANONYMOUS). If authentication fails, an

'Authentication Failure' error message (type 300) is generated and no

address is allocated to the requesting node. If authentication succeeds, TSP enters into command phase and the allocation process can take place.

As shown on figure 2.1, the address allocation process starts when a DSTM node sends a 'Tunnel Create' request to the DSTM Server. This message contains the Link-Local address of the node and the Global IPv6 address that the node would use to establish the 4over6 tunnel. No other information is needed.

Next, the DSTM server processes the request. It may result in an error due to Address Pool exhaustion (error type 306). If an IPv4 address is available, the server configures the TEP using another 'Tunnel Create' message. The message includes the global IPv6 and the allocated IPv4 addresses of the requesting node.

The TEP MUST be configured to accept TSP messages only from a valid DSTM server. At the arrival of a 'Tunnel Create' Request, the TEP updates its IPv4/IPv6 mapping table and sets up the 4over6 tunnel as requested. If, for some reason, it is not possible to update the table, or the 4over6 tunnel cannot be set up, the TEP replies with an error message (error type 307). In that case, the DSTM server SHOULD forward the error message to the requesting node.

If tunnel configuration succeeds, the DSTM server receives a 'Tunnel Info' message from the TEP. This message contains the IPv6 and IPv4 addresses of the TEP for the new tunnel.

At this point, the server updates its own tables and sends a 'Tunnel Info' message to the requesting node. This message contains the temporary IPv4 address of the node, its period of validity (the 'Life Time') and address information of the TEP. TEP information MUST be the same that the TEP provided.

Finally, the IPv4 stack of the node is configured. A 4over6 tunnel is established between the node and the TEP. An IPv4 default route is added pointing to the 4over6 tunnel. Communication in IPv4 can take place. A timer configured with the 'Life Time' parameter informs the node when to ask for renewal of allocation, if needed.

2.2 Allocation Renewal

As long as an IPv4 address is needed at the node, 'Tunnel Create' messages are sent to the DSTM Server as a request for allocation renewal. The frequency of such requests depends on the 'Life Time' parameter. The temporary IPv4 address for which allocation renewal is requested MUST be included in the messages.

Based on the contents of the message and local policy, the server may reply with a 'Tunnel Info' message. At the node, the reception of such a message means that allocation time has been extended: the timer is reset to the value contained in the 'Life Time' field. No modification is needed in the IPv4 stack nor in the TEP. If allocation cannot be extended, an error message MUST be sent to the node (error type 308) and tunnel information MUST be deleted at the TEP.

2.3 End of Allocation

If properly configured, there will be a time where the node will no longer need an IPv4 address. At this time, it will stop sending 'Tunnel Create' requests for renewal. At the server, when allocation time expires, a 'Tunnel Delete' message MUST be sent to both the node and the corresponding TEP. The server SHOULD NOT wait for an acknowledge from the node before updating its own tables and deleting the configuration at the TEP. However, implementations may wait until the server receives a reply before releasing the address.

A 'Tunnel Delete' message contains the IPv4 and IPv6 addresses of the node for which the entry in the mapping table is to be deleted. The TEP MUST stop forwarding packets for that node as a reaction to this type of message. Depending on implementation, TEPs may acknowledge tunnel deletion using a 'Tunnel Info' message.

3. TSP Profile for DSTM

This section describes the TSP profile for IPv4 over IPv6 tunnels in DSTM.

3.1 Overview

The TSP profile uses the included DTD for the XML format of the message. The DTD (c.f. Annexe) contains the description of the tunnel XML message. This message is used by a TSP-DSTM compliant server to provide the necessary information to DSTM nodes and the TEP in order to establish 4over6 tunnels. Three types of action are defined in a 'tunnel' message: Create, Delete and Info.

The 'Create' action is used to request a new tunnel or to renew an address allocation.

The 'Delete' action is used by the server to remove an existing tunnel from a node and the TEP.

The 'Info' action is used by the server to send tunnel configuration data. It is also used by nodes and the TEP to acknowledge a previous

command (Create or Delete).

The 'tunnel' message may have one or two elements:

- client: Client's information
- server: Server's information

Server is used in the context of the other party in the TSP connection. It can be the DSTM server if the client is the DSTM node, or the TEP if the client is the DSTM server.

3.2 Client element

The client element contains 'address' elements. The 'address' element is used to identify the client IPv6 endpoint of the 4over6 tunnel. The client MUST send its link- local and global IPv6 addresses to the server. The server will then return a temporary IPv4 address inside the 'client' element when the tunnel is created or allocation is renewed.

3.3 Server element

The 'server' element contains 'address' elements. This element is used to identify the addresses at the TEP. The 'address' element provides both IPv4 and IPv6 addresses of the TEP.

4. DSTM protocol using TSP

TSP message exchanges are done using TCP over IPv6 transport. Once the TCP session is established between the DSTM node and server, it MAY be kept connected for the duration of the address allocation lease time. This TCP connection can be used by the server to send requests to the client on a communication channel already established (and potentially authenticated) by the client.

This section presents an example of a DSTM host requesting an IPv4 address allocation to a DSTM server. As described in TSP[ref], the first TSP phase involves authentication (which can be ANONYMOUS) followed by a command phase that takes care of the allocation negotiation.

4.1 Initial Address Allocation

Allocation Requests coming from a node consist of a 'tunnel' element using the attributes action set to 'create' and type set to 'v4v6'. The 'tunnel' element contains one 'client' element.

Simple tunnel request made by a client.

```
          -- Successful TCP Connection --
          C:VERSION=1.0 CR LF
          S:CAPABILITY TUNNEL=V4V6 AUTH=DIGEST-MD5 AUTH=ANONYMOUS CR LF
          C:AUTHENTICATE ANONYMOUS CR LF
          S:OK Authentication successful CR LF
          C:Content-length: 228 CR LF
            <tunnel action="create" type="v4v6">
               <client>
                   <address
type="ipv6">fe80:0000:0000:0000:0000:0000:0000:0001</address>
                   <address
type="ipv6">3ffe:0b00:0c18:ffff:0000:0000:0000:0001</address>
               </client>
            </tunnel> CR LF
```

If the allocation request is accepted, the DSTM server will acknowledge the allocation to the client by sending a 'tunnel' element with the attribute 'action' set to 'info', 'type' set to 'v4v6' and the 'lifetime' attribute set to the period of validity or lease time of the allocation. The 'tunnel' element contains 'server' and 'client' elements.

Server response

```
          S: Content-length: 370 CR LF
             200 OK CR LF
             <tunnel action="info" type="v4v6" lifetime="1440">
               <server>
                  <address type="ipv4" length="30">206.123.31.2</address>
                  <address type="ipv6">3ffe:b00:c18:ffff:
0000:0000:0000:0002</address>
               </server>
               <client>
                  <address type="ipv4" length="30">206.123.31.1</address>
                  <address
type="ipv6">3ffe:b00:c18:ffff::0000:0000:0000:0001</address>
               </client>
            </tunnel> CR LF
```

4.2 Allocation Renewal

A DSTM host asks for renewal of an IPv4 address allocation by sending a 'Tunnel Create' message to a DSTM server. The request consists of

a 'tunnel' element using the attributes action set to 'create' and

type set to 'v4v6'. The 'tunnel' element contains one 'client' element. The temporary IPv4 address for which allocation renewal is requested MUST be included in the messages.

Renewal of the same client

```
          C:Content-length: 228 CR LF
            <tunnel action="create" type="v4v6">
               <client>
                   <address
type="ipv6">fe80:0000:0000:0000:0000:0000:0000:0001</address>
                   <address
type="ipv6">3ffe:0b00:0c18:ffff:0000:0000:0000:0001</address>
                   <address type="ipv4" length="30">206.123.31.1</address>
               </client>
            </tunnel> CR LF
```

If the allocation request is accepted, the DSTM server will acknowledge the renewal to the client by sending a 'tunnel' element with the attribute 'action' set to 'info', 'type' set to 'v4v6' and the 'lifetime' attribute set to the period of validity or lease time of the allocation. No message is sent to the TEP in this case. At the node, the reception of such a message means that allocation time has been extended; the timer is reset to the value contained in the 'lifetime' field.

Server's response to the renewal

```
          S: Content-length: 370 CR LF
             200 OK CR LF
             <tunnel action="info" type="v4v6" lifetime="1440">
               <server>
                  <address type="ipv4" length="30">206.123.31.2</address>
                  <address type="ipv6"
length="64">3ffe:b00:c18:ffff:0000:0000:0000:0002</address>
               </server>
               <client>
                  <address type="ipv4" length="30">206.123.31.1</address>
                  <address type="ipv6"
length="64">3ffe:b00:c18:ffff::0000:0000:0000:0001</address>
               </client>
            </tunnel> CR LF
```

4.3 End of Allocation

A DSTM server uses a 'Tunnel Delete' message to end the IPv4 address allocation of a client. The release request consists of a 'tunnel' element using the attributes action set to 'delete' and type set to 'v4v6'. The 'tunnel' element contains 'server' and 'client' elements representing the address allocation that is released.

Server sending a release request

```
          S: Content-length: 370 CR LF
             200 OK CR LF
             <tunnel action="delete" type="v4v6">
               <server>
                  <address type="ipv4" length="30">206.123.31.2</address>
                  <address type="ipv6"
length="64">3ffe:b00:c18:ffff:0000:0000:0000:0002</address>
               </server>
               <client>
                  <address type="ipv4" length="30">206.123.31.1</address>
                  <address type="ipv6"
length="64">3ffe:b00:c18:ffff::0000:0000:0000:0001</address>
               </client>
            </tunnel> CR LF
```

5. Error Codes

This list describes the error codes used in this document.

300 Authentication failed

306 Address Pool Exhausted

307 Configuration Error at TEP

308 Requested Address Unavailable

309 Invalid IPv6 address

310 IPv4 Invalid Address

6. IANA Considerations

The TUNNELTYPE "v4v6" is registered for this document.

7. Security

TSP provides authentication services using SASL [RFC2222]. If DSTM client authentication is required, TSP can be configured at the server to negotiate with the client the authentication scheme that will be used.

In the context where the server sends a request to the client, some form of authentication is required so that the client can be sure that the request comes from a trusted DSTM server.

This document proposes that in the case where the client initially authenticates to the DSTM server, this TCP session MAY be kept connected for the duration of the address allocation lease time. This TCP connection can be used by the server to send requests to the client on a communication channel already established by the client. A more secure solution would be to provide mutual authentication between the parties.

References

[1] Bound, J., "Dual Stack Transition Mechanism (DSTM)", draft-ietf-ngtrans-dstm-05 (work in progress), November 2001.

[2] Blanchet, M., "Tunnel Setup Protocol", July 2001.

[3] Hagino, J., "Possible abuse against IPv6 transition technologies", July 2000.

[4] Droms, R., Perkins, C., Bound, J. and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", draft-ietf-dhc-dhcpv6-21 (work in progress), November 2001.

[5] Myers, J., "Simple Authentication and Security Layer (SASL)", RFC 2222, October 1997.

Authors' Addresses

Marc Blanchet
Viagenie
2875 boul. Laurier, bureau 300
Sainte-Foy, QC G1V 2M2
Canada

Phone: +1 418 656 9254
EMail: Marc.Blanchet@viagenie.qc.ca
URI: http://www.viagenie.qc.ca/

Octavio Medina
ENST Bretagne
BP 78
Cesson Sevigne, Cedex 35512
France

Phone: +33 2 99 12 70 23
EMail: Octavio.Medina@enst-bretagne.fr
URI: http://www.enst-bretagne.fr

Florent Parent
Viagenie
2875 boul. Laurier, bureau 300
Sainte-Foy, QC G1V 2M2
Canada

Phone: +1 418 656 9254
EMail: Florent.Parent@viagenie.qc.ca
URI: http://www.viagenie.qc.ca/

Appendix A. Appendix A. IPv4 over IPv6 tunnel DTD

DTD

```
<?xml version="1.0"?>

<!DOCTYPE tunnel  [

<!ELEMENT tunnel        (server?,client?,broker?)>

  <!ATTLIST tunnel action   (create|info|list) #REQUIRED >
  <!ATTLIST tunnel type     (v4v6|broker)      #REQUIRED >
  <!ATTLIST tunnel lifetime CDATA              "1440"    >

<!ELEMENT server        (address+,router?)>

<!ELEMENT client        (address+,router?)>

<!ELEMENT broker        (address+)>

<!ELEMENT router        (prefix?,dns_server?,as?)>
  <!ATTLIST router protocol (rip|bgp) "">

<!ELEMENT dns_server    (address+)>

<!ELEMENT as EMPTY>
  <!ATTLIST as number CDATA #REQUIRED>

<!ELEMENT prefix        (#PCDATA)>
  <!ATTLIST prefix length CDATA #REQUIRED>

<!ELEMENT address       (#PCDATA)>
  <!ATTLIST address type (ipv4|ipv6|dn) #REQUIRED>
  <!ATTLIST address length CDATA "">

]>
```

Full Copyright Statement

Acknowledgement

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