INTERNET-DRAFT
NGTRANS Working Group
Obsoletes draft-ietf-ngtrans-dstm-04.txt
Expires May 2002

Jim Bound
Compaq
Laurent Toutain
Francis Dupont
Octavio Medina
ENST Bretagne
Hossam Afifi
INT
Alain Durand
Sun Microsystems

Dual Stack Transition Mechanism (DSTM)

<draft-ietf-ngtrans-dstm-05.txt>

Status of this Memo

Abstract

The initial deployment of IPv6 will require a tightly coupled use of IPv4 addresses to support the interoperation of IPv6 and IPv4, within an IPv6 Network. Nodes will still need to communicate with IPv4 nodes that do not have a dual IP layer supporting both IPv4 and IPv6. The Dual Stack Transition Mechanism (DSTM) provides a method to assign temporary IPv4 Addresses to IPv6/IPv4 nodes over a native IPv6 Network. DSTM uses dynamic tunnels within an IPv6 Network to carry IPv4 traffic and a defined set of processes and architecture for the supporting infrastructure required for this transition mechanism.

Table of Contents:

1. Introduction 3
2. Terminology 4
2.1 IPv6 DSTM Terminology 4
2.2 Specification Language 5
3. DSTM Overview and Assumptions 5
4. DSTM Deployment Example 7
4.1 DSTM Client/Server Example 8
5 DTI Architecture 9
5.1 Assignment of the IPv4 address to the DTI 10
5.2 DTI proceeding of IPv4 packets 10
5.3 DTI IPv6 packet 10
6. DSTM Server Requirements 11
7. Applicability Statement 11
8. Security Considerations 12
Annexe A: Static Configuration 12
Annexe B: RPC 13
Annexe C: DHCPv6 13
C.1 DHCPv6 Global IPv4 Address Option 14
C.1.1 Client Request of IPv4 Global Address 15
C.1.2 Server Reply of IPv4 Global Address Option 15
C.1.3 Client Processing of IPv4 Address Option 17
C.2 Server Processing of an IPv4 Address Option 17
C.3 Client Processing of an IPv4 Address Option 18
Acknowledgments 19
Normative References 20
Informative References 20
Authors' Address 20

1. Introduction

The initial deployment of IPv6 will require a tightly coupled use of IPv4 addresses to support the interoperation of IPv6 and IPv4, within an IPv6 Network. Nodes will still need to communicate with IPv4 nodes that do not have a dual IP layer supporting both IPv4 and IPv6. The Dual Stack Transition Mechanism (DSTM) provides a method to assign temporary IPv4 Addresses to IPv6/IPv4 nodes over a native IPv6 Network. DSTM uses of dynamic tunnels within an IPv6 Network to carry IPv4 traffic, and a defined set of processes and architecture for the supporting infrastructure required for this transition mechanism.

The DSTM assigns, when needed an IPv4 address to a dual IP layer node. This will allow either IPv6 nodes to communicate with IPv4-only nodes, or for IPv4-only applications to run without modification on an IPv6 node. This allocation mechanism is coupled with the ability to perform dynamic tunneling of an IPv4 packet inside an IPv6 packet, to hide IPv4 packets in the native IPv6 domain. This will simplify the network management of IPv6 deployment, since routers need only IPv6 routing tables to move IPv4 packets across an IPv6 network. This means that only the IPv6 routing plan is managed inside the network.

DSTM is targeted to help the interoperation of IPv6 newly deployed networks with existing IPv4 networks. DSTM assumes that a user will deploy an IPv6 network to reduce the need and reliability on IPv4 within a portion of his network. In addition the IPv4 globally routable address space available to the network is a scarce resource, and DHCPv4[7] may not be directly used to assign temporary IPv4 addresses to IPv6 nodes, since no IPv4 connectivity is maintained into the network. Also, to reduce the IPv4 applications a user has to support and to obtain a temporary IPv4 Address (see Section 6), the client only has to run a client process with the DTI mechanisms in this specification.

The DSTM architecture is composed of an addresses server, that can provides the assignment of IPv4 addresses to IPv6 Nodes. The server will also be used to maintain the mapping between the allocated IPv4 address and the permanent IPv6 address of the node. Each IPv6 DSTM will have an IPv4 interface called the Dynamic Tunneling Interface (DTI) designed to encapsulate IPv4 packets into IPv6 packets. A DSTM client on the node SHOULD be used for IPv4 address allocation and MAY be used to solve the mapping between IPv4 and IPv6 addresses.

The specification will begin by defining the terminology (section 2), then section 3 provides a technical overview of the DSTM methodology as a transition mechanism. Then in section 4 we provide a DSTM example. Section 5 describes the DTI Architecture and Section 6 discusses the properties of the IPv4 allocation mechanisms that can be used. Section 7 provides the DSTM Applicability Statement.

Annexes give ways to use this specification in different situations. Annexe A gives a simple user configuration. This simplifies a lot the DTI interface but require more IPv4 addresses. Annexe B documents the use of RPCv6 to allocate addresses. This is a simple but limited protocol. Annexe C describes the DHCPv6 mechanism for DSTM. This is the most appropriate and generic mechanism, but due to some standardisation delay, it could not be deployed as fast as RPCv6.

2. Terminology

2.1 IPv6 DSTM Terminology

DSTM Domain	The network areas on an Intranet where a temporary IPv4 allocation Server has access to IPv6 nodes participating in DSTM for that network, and IPv4 routing access is not necessary within a DSTM domain.
DSTM Node	A Node that supports a dual IP layer IPv4 and IPv6 stack, DTI, and an IPv4 allocation Client. The DSTM node generate only IPv6 packets on the network.
DSTM Border Router	A border router within a DSTM domain and access to an external IPv4-ONLY domain.
DSTM client	A process on the DSTM Node that managed the temporary IPv4 address assigned by the DSTM Server.
DSTM Server	A process in charge of managing the IPv4 address space that will be allocated to DSTM Nodes.
IPv6 Protocol Terms:	See [1]
IPv6 Transition Terms:	See [6]
DHCPv6 Terms:	See [2]
DTI	Dynamic Tunneling Interface. An interface encapsulating IPv4 packets into IPv6 packets.
Tunnel End Point (TEP)	Destination of the IPv6 packet containing an

IPv4 packet. In most cases this will be a DSTM border router.

2.2 Specification Language

In this document, several words are used to signify the requirements of the specification, in accordance with RFC 2119 [4]. These words are often capitalized.

MUST — This word, or the adjective "required", means that the definition is an absolute requirement of the specification.

MUST NOT — This phrase means that the definition is an absolute prohibition of the specification.

SHOULD — This word, or the adjective "recommended", means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications must be understood and carefully weighed before choosing a different course. Unexpected results may result otherwise.

MAY — This word, or the adjective "optional", means that this item is one of an allowed set of alternatives. An implementation which does not include this option MUST be prepared to interoperate with another implementation which does include the option.

silently discard — The implementation discards the packet without further processing, and without indicating an error to the sender. The implementation SHOULD provide the capability of logging the error, including the contents of the discarded packet, and SHOULD record the event in a statistics counter.

3. DSTM Overview and Assumptions

DSTM as discussed in the introduction is a method that uses existing protocols. DSTM does not specify a protocol. However, DSTM defines client, server and TEP behaviour and the properties of the temporary addresses allocation mechanisms.

The motivation for DSTM is to provide IPv6 nodes a means to acquire an IPv4 address, for communications with IPv4-only nodes or IPv4 applications.

The core assumption within this mechanism is that it is totally transparent to applications, which can continue to work with IPv4 addresses. It is also transparent to the network carring only IPv6 packets. It is the authors' viewpoint that the user in this case, has deployed IPv6 to support end to end computing, without translation. This aspect is fundamental during a transition process to guarantee that every existing application will continue to work (e.g. IPsec, H.323), with embed IPv4 addresses in the payload of a packet.

The DSTM model and assumptions are as follows:

- The DSTM domain is within an Intranet not on the Internet.

- IPv6 nodes do not maintain IPv4 addresses except on a temporary basis, to communicate with IPv4-only and IPv4 Applications.

- The temporary IPv4 address allocation is done by the DSTM server, different protocols such as DHCPv6 or RPCv6 can be used to assign the IPv4 address.

- The DSTM domain for the IPv6 nodes will keep IPv4 routing tables to a minimum and use IPv6 routing, hence, reducing the network management required for IPv4 during transition.

- Once IPv6 nodes have obtained IPv4 addresses Dynamic Tunneling is used to encapsulate the IPv4 packet within IPv6 and then forward that packet to an IPv6 TEP, where the packet will be decapulated and forwarded using IPv4. The IPv4 allocation mechanism may also provide the TEP IPv6 address.

- Existing IPv4 applications or nodes do not have to be modified to communicate with DSTM.

- Implementation defined software will have to exist to support DSTM:

  o Ability within a DSTM Server implementation to maintain configuration information about TEPs for encapsulating IPv4 packets between IPv6 nodes that can forward IPv4 packets to an IPv4 routing realm, and to maintain a pool of IPv4 Addresses.

  o an IPv6 node MUST support the dynamic tunneling mechanisms in this specification to encapsulate IPv4 packets within IPv6 on an IPv6 node. In addition a DSTM client SHOULD be present on the node for IPv4

Mapped Addresses and TEPs management.

- o DSTM Border Routers MAY recall or be able to cache the association of IPv6 and IPv4 addresses of nodes during the forwarding process.

A schematic overview of DSTM is as follows:

```
   -----------------------------------------------
                                                 |    IPv4 Internet or Intranet
           DSTM Domain Intranet                  |      IPv4 Applications
                                                 |         Domain
                   _____________________         |
                  |                     |        |
                  |   DSTM Server       |        |
                  |_____________________|        |
                                ^                |
                                |                |
      __________________        |               _|_______
     |                  |       |              |        |
     | IPv6/IPv4 Node   |       |              |  DSTM  |
     |------------------|       |              | Border |
     |  DSTM client     |       |              | Router |
     |                  |<-------              |  IPv6  |
     |------------------|                      |    &   |
     |    DTI/Route     |<-------------------->|  IPv4  |
      -------------------                       ---------
                                                |
   ----------------------------------------------
```

For an IPv6 node to participate in DSTM it MUST have a dual IP layer, supporting both an IPv4 and an IPv6 stack. DSTM is not a solution for IPv6 ONLY nodes.

4. DSTM Deployment Example

In the example below, the following notation will be used:

```
X    will designate an IPv6 node with a dual stack, X6 will be the IPv6
     address of this node and X4 the IPv4 address
Y    will designate a DSTM border router at the boundary between an
     IPv6 DSTM domain and an IPv4-only domain.
Z    will designate an IPv4-only node and Z4 its address.
==>  means an IPv6 packet
-->  means an IPv4 packet
++>  means a tunneled IPv4 packet is encapsulated in an IPv6 packet
```

..> means a DNS query or response. The path taken by this

packet does not matter in the examples
"a" means the DNS name of a node

4.1 DSTM Client/Server Example

This example describes the case where an application (either compiled for the IPv6 or IPv4 API) running on an IPv6 node (X6) wants to establish a session with an IPv4 application on an IPv4-only node (Z4).

The IPv4 routing table of node X is configured to send IPv4 packets to the DTI interface.

```
            DSTM
           Server          DNS
      X6          Y6/Y4         Z4
       |            |            |
       |. . . . . . . .>    Z    |     - X6 asks the DNS for the A RR for "Z"
       |<. . . . . . . .    Z4   |     - the answer is Z4
       |            |            |
       |            |            |     - The application sends its first IPv4
       |            |            |       packet which arrives to the DTI
       |            |            |       interface.
       |            |            |       (If the application is compiled for
IPv6
       |            |            |       this can be done through an IPv4-mapped
       |            |            |       address).
       |            |            |
       |            |            |     - X6 needs an IPv4 address (first use)
       |====>       |            |     - X6 queries the DSTM server for an
       |            |            |       IPv4 address
       |<====       |            |     - The DSTM server locates the client
       |            |            |       and provides a temporary IPv4
       |            |            |       global address and the IPv6 TEP
address.
       |+++++++++++>|            |     - The DTI sends the IPv6 packet to the
       |            |            |       TEP.
       |            |----------->|     - Y sends the packet to the destination
Z4
       |            |            |     - Y caches the association between
       |            |<-----------|     - Z4 answers.
       |            |            |
       |<+++++++++++|            |     - Y uses the mapping between X4 and X6
       |            |            |       to tunnel the packet to the destination
       |            |            |
```

When Z responds the packet returns back through Y. Y having cached the association between the IPv4 and the IPv6 address of X, is able to send the packet encapsulating the IPv4 packet within IPv6 back to X.

5 DTI Architecture

In the absence of an IPv4 routing infrastructure, a DSTM node can not directly send IPv4 packets on the network. It has to encapsulate them into IPv6 packets and send them to a tunnel end point (TEP), which is a particular DSTM node, that will decapsulate the packet and forward them in the IPv4 network.

On a DSTM node, this encapsulation is done by the DTI interface. An IPv4 packet can be directed to that interface by an IPv4 routing table entry.

The exact details of the DTI interface and the associated routing table entries are implementation dependant.

5.1 Assignment of the IPv4 address to the DTI

When the DTI interface is activated, an IPv4 address is not given to that interface. When the first IPv4 packet has to be sent by the DTI interface, a request is sent to the DSTM serveur to get the temporary IPv4 address and the TEP IPv6 address.

An IPv6 node can know it needs an IPv4 address if the DNS resolver on the node knows that the destination address will be an IPv4 address. This can also be trigged by the DTI interface if no IPv4 addresss is associated to that interface (see next paragraph).

5.2 DTI proceeding of IPv4 packets

The DSTM server allocates the source address of the IPv4 packet. If the DTI interface does not have an IPv4 address, the process using this interface SHOULD be blocked and the request mechanism for a temporary IPv4 address SHOULD be started.

The other fields of the IPv4 packet are normally filled.

5.3 DTI IPv6 packet

When a DTI has to encapsulate an IPv4 packet into an IPv6 packet, the DTI has to determine the TEP IPv6 address for the destination. The TEP can be the node destination or, if the destination node is IPv4-only, the IPv6 address of an IPv4/IPv6 DSTM Border Router.

The TEP IPv6 address can be either statically configured or dynamically acquired when the IPv6 node acquires an IPv4 address from a DSTM Server.

The TEP IPv6 address SHOULD be provided by the DSTM server when the DSTM node receives an temporary IPv4 Address (section 6). But, a DSTM node MAY manually configure the TEP during early deployment of DSTM, this will not scale and is not recommended as a long term transition solution.

The next header type for IPv4 encapsulation is 4 (as for IPv4 tunneling over IPv4). When a tunneled packet arrives to the IPv6 destination, the IPv6 header is removed and the packet is processed by the IPv4 layer. The DSTM Border Router SHOULD cache the association between the IPv4 and IPv6 source addresses. The IPv4 packet will then be forwarded by the DSTM border router using the IPv4 infrastructure.

The IPv6 source address of an encapsulated packet will be the IPv6 address of the interface on which the IPv6 packet will be sent.

6. DSTM Server Requirements

The DSTM server is mostly in charge of the temporary IPv4 address allocation. This allocation is very simple since there is no localisation purpose in this address. The DSTM server has just to guaranty the uniqueness of the IPv4 address for a period of time. The DTSM server MUST also memorize the mapping between the IPv6 address of the node requesting a temporary address and the allocated IPv4 address.

The temporary IPv4 address is allocated by the server for a fixed among of time. This duration MUST be included in the response. If the client needs the IPv4 address for a longer period of time, the client MUST renew the lease.

The pool of IPv4 global addresses MUST be routed to one or more TEP in the DSTM domain.

The response SHOULD include the TEP IPv6 address in charge of the temporary IPv4 address.

The communication between the DSTM client and the server MUST be in IPv6.

The DSTM server MAY allocate a temporary IPv4 address without a request from he client.

The DSTM server SHOULD be able to authenticate the DSTM client.

7. Applicability Statement

DSTM is applicable for use from within the DSTM Domain to IPv4 nodes or applications on a user Intranet or over the Internet.

DSTM's motivation is to support dual IP layer DSTM node to communicate using global IPv4 addresses across an Intranet or

Internet, where global addresses are required. But, DSTM has been defined to also permit the use of Private IPv4 address space to permit the Intranet use of DSTM where users require temporary access to IPv4 services within their Intranet.

if DSTM requires the use of DHCPv6 to obtain IPv4 addresses and TEPs for a DSTM node, the Communications between the DSTM Daemon and the DHCPv6 client is implementation defined. The DTI mechanism is also implementation defined. DSTM does permit optionally for DSTM node to manually configure TEPs for DTI for early deployment of DSTM but highly recommends not doing this and configuring DHCPv6 servers with this information is really the way to execute DSTM on an IPv6 Network.

DSTM also assumes that all packets returning from an IPv4 node to a DSTM dual IP layer node return through the orginating DSTM Border Router which has cached the association of the DSTM's IPv4+IPv6 addresses. At this time it is beyond the scope of DSTM to permit IPv4 packets destined for DSTM node to return packets through a non-orginating DSTM border router.

DSTM also through the new DHCPv6 extension permits Network Operators to inform DSTM Nodes they will need IPv4 addresses for communications using the DHCPv6 Reconfigure-Init message.

DSTM as future work can be extended to support multiple border routers for returning IPv4 packets, and for the discovery of DSTM node using IPv4 DNS queries as future work for DSTM.

8. Security Considerations

The DSTM mechanism can use all the defined security specifications for each functional part of the operation. For DNS the DNS Security Extensions/Update can be used [9, 10], for DHCPv6 the DHCPv6 Authentication Message can be used [2], and for communications between the IPv6 node, once it has an IPv4 address, and the remote IPv4 node,

IPsec [3] can be used as DSTM does not break secure end-to-end communications at any point in the mechanism.

Annexe A: Static Configuration

The DTI interface in a DSTM client can be a static tunnel such as a gif interface in the KAME stack. An IPv4 address can be manually assigned to the interface. In that case, no DSTM server is needed, the TEP will maintain the mapping between the v4 and the v6 addresses of the DSTM client.

The following listing gives a configuration example for the KAME stack:

```
gifconfig gif0 inet6 3ffe:ffe:1002:1::1 3ffe:3ffe:1002:1:260:caff:fe85:abcd
ifconfig gif0 inet 193.109.121.195 193.109.121.199 netmask 255.255.255.255
up
route change default 193.109.121.199
```

Annexe B: RPC

RPC is a simple method that can be used for communication between the DSTM client and the DSTM server. This method is efficient when the address request is triggered by the DSTM client. The following listing gives the structure used by RPC in this case:

```
const REQUEST_REQ = 1;
const REQUEST_REP = 2;
const RELEASE_REQ = 3;
const RELEASE_REP = 4;

struct packet {
        int type;                       /* Message opcode/type */
        opaque local[16];               /* Link-local address */
        opaque mask[4];                 /* Netmask */
        opaque i4addr[4];               /* IPv4 address */
        opaque tep4[4];                 /* TEP IPv4 address */
        opaque tep6[16];                /* TEP IPv6 address */
        unsigned long ends;             /* When the lease ends */
        unsigned long starts;           /* When the lease starts */
        unsigned long extend;           /* How long to extend */
        unsigned long keep;             /* How long to keep */
};

program RPC {
        version RPC_ONE {
                struct packet REQUEST(struct packet) = 1;
                struct packet RELEASE(struct packet) = 2;
        } = 1;
} = [to be assigned];
```

Annexe C: DHCPv6

The DSTM processes will use the DHCPv6 services [2] to communicate between the DHCPv6 Server and the DHCPv6 Client. A new option is required for DHCPv6 to support DSTM. But there are some additional requirements placed on the DSTM processes that are not specific to the DHCPv6 protocol as a transition and interoperation set of mechanisms for the IPv6 node.

DHCPv6 clients solicit servers, and servers advertise their availability. Then DHCPv6 clients request configuration parameters, and a server sends those parameters back in a reply message. The client requests parameters by specifying options with the DHCPv6 request messaqge. This new DSTM option will request that the server return an IPv4-Mapped IPv6 address to the client.

DHCPv6 servers also support a Reconfigure message sent to clients to ask clients to initiate a request message for a specific option. This permits DHCPv6 servers to offer clients IPv4-Mapped IPv6 addresses.

C.1 DHCPv6 Global IPv4 Address Option

The DHCPv6 IPv4 Address Option informs a DHCPv6 Client or Server that the Identity Association Option (IA) [2] following this option will contain an IPv4-Mapped IPv6 Address [9] in the case of a DHCPv6 Client receiving the option, or is a Request for an IPv4-Mapped IPv6 Address from a client in the case of a DHCPv6 Server receiving the option. The option can also provide an IPv6 address to be used as the TEP to encapsulate an IPv4 packet within IPv6.

This option can be used with the DHCPv6 Request, Reply, and Reconfigure- Init Messages for cases where a DHCPv6 Server wants to assign to clients IPv4-Mapped IPv6 Addresses, thru the Option Request Option (ORO) in DHCPv6.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              option-code      |             option-length     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Tunnel End Point (TEP)               |
|                           (If Present)                        |
|                            (16 octets)                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
option-code:            TBD
option-length:          Variable: 0 or 16
Tunnel End Point:       IPv6 Address if Present
```

An IPv4 Global Address Option MUST only apply to the IA following it this option.

C.1.1 Client Request of IPv4 Global Address

When the client requests an IPv4 address from the DHCPv6 Server the TEP field MUST not be present in the Global IPv4 Address Option.

C.1.2 Server Reply of IPv4 Global Address Option

The server will reply to the client with a Global IPv4 Address Option, that can contain an IPv6 Address Tunnel End Point, and an IA Option which MUST include an IPv4 IPv6-Mapped Address. The IA Option is provided as a reference in this document [2].

The format of the IA option is:

The identity association option is used to carry an identity association, the parameters associated with the IA and the addresses assigned to the IA.

The format of the IA option is:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           OPTION IA           |          option-len           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         IAID (4 octets)                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               T1                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               T2                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   IA status   |   num-addrs   |T| addr status | prefix length |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                          IPv6 address                         |
|                           (16 octets)                         |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       preferred lifetime                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         valid lifetime                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|T| addr status | prefix length |                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
|                          IPv6 address                         |
```

```
   |                          (16 octets)                           |
```

```
 |                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                               |      preferred lifetime       |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 | pref. lifetime (cont.)        |        valid lifetime         |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 | valid lifetime (cont.)        |T| addr status | prefix length |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                         IPv6 address                          |
 |                          (16 octets)                          |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                              ...                              |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

option-code   OPTION_IA (TBD)

option-len   Variable; equal to 24 + num-addrs*26

IA ID   The unique identifier for this IA; chosen by the client

T1   The time at which the client contacts the server from which the addresses in the IA were obtained to extend the lifetimes of the addresses assigned to the IA.

T2   The time at which the client contacts any available server to extend the lifetimes of the addresses assigned to the IA.

T   When set to 1, indicates that this address is a "temporary address" [7]; when set to 0, the address is not a temporary address.

IA status   Status of the IA in this option.

num-addrs   An unsigned integer giving the number of addresses carried in this IA option (MAY be zero).

addr status   Status of the addresses in this IA.

prefix length   Prefix length for this address.

IPv6 address   An IPv6 address assigned to this IA.

preferred lifetime   The preferred lifetime for the associated

IPv6 address.

valid lifetime — The valid lifetime for the associated IPv6 address.

The ``IPv6 address'', ``preferred lifetime'' and ``valid lifetime'' fields are repeated for each address in the IA option (as determined by the ``num-addrs'' field).

C.1.3 Client Processing of IPv4 Address Option

The processing of the IPv4 Global Address Option on the client is implementation defined but here are some guidelines for developers.

When processing the IA Option following the IPv4 Global Address Option, an IP Address provided will be an IPv4-Mapped IPv6 Address. A conceptual implementation model would be to add this address to the nodes IPv6 mechanisms that maintain timing procedures for IPv6 addresses on the IPv6 stack, and then configure the IPv4 interface for DTI, as a procedure called from the DHCPv6 client.

As the IPv4 IPv6-Mapped Address is an IPv6 address all other processing for DHCPv6 is as specified in that document, the IPv4 Global Address Option just informs the client that an address within the IA option will be an IPv4 IPv6-Mapped Address.

C.2 Server Processing of an IPv4 Address Option

When a DHCPv6 Server receives an IPv4 Global Address Option in a DHCPv6 Request message, the client is requesting an IPv4 IPv6-Mapped Address.

A DHCPv6 Server can send a Client a Reconfigure-Init message using the IPv4 Global Address Option to ask the Client to request an IPv4 Global Address thru an ORO. The Client will then send a request to the server for an IPv4 IPv6-Mapped Address.

The Server will know a priori the Clients IPv6 routable address, when sending a Reconfiguration-Init message.

The Server will look in its implementation defined IPv4 Address configuration to determine if a TEP is available for a specific IPv6 Address Prefix. If that is the case the Server will put the address for the TEP in the Global IPv4 Address Option.

C.3 Client Processing of an IPv4 Address Option

When the Server supplies an IPv4 Global Address in a Reply.

The Client MUST not update the DNS with this new address.

A conceptual model to configure an IPv4 IPv6-Mapped address on a client is as follows:

1. In an implementation defined manner the Client MUST assign the address to an interface, supporting the Client's IPv4 stack implementation.

2. In an implementation defined manner the Client MUST create an entry as an IPv4-Mapped IPv6 Address supporting the processing required for an IPv6 address regarding the valid and preferred lifetimes as specified in IPv6 Addrconf [8]. Once the IPv4-Mapped IPv6 Address valid lifetime expires the IPv4 address MUST be deleted from the respective interface and a DHCPv6 Release Message MUST be sent to the DHCPv6 Server to delete the IPv4 IPv6-Mapped Address from the Servers bindings.

3. If a TEP address is provided in the Global IPv4 Address Option, the Client MUST create a configured tunnel to the TEP address, in an implementation defined manner. These encapsulation mechanisms are defined in other IPv6 specifications [5, 6].

Changes from draft 04 to draft 05

1. Give in the normative part only DSTM server requierments

2. Create 3 annexes for different way to configure DTSM client

Changes from draft 03 to draft 04

1. Changed DHCPv6 options and processing to comply with draft-ietf-dhc-dhcpv6-16.txt

Changes from draft 02 to draft 03

1. Working Group Edits

Changes from draft 01 to draft 02

1. Added futher clarifications to DSTM components.

2. Added client/server details for DHCPv6 ngtrans extension.

3. Removed optional scenarios to simplify this mechanism.

4. Removed AIIH concepts and changed to be DSTM components.

5. Add Applicability Statement

6. Added acknowledgment section and new coauthors Francis Dupont and Alain Durand.

Changes from draft 00 to draft 01

1. Added text explaining why the draft does not use DHCPv4 to assign IPv4 compatible addresses to the "Introduction".

2. Defined what is mandatory and what is optional and added relative text in various places to clarify this change. And added RFC 2119 adjectives to the spec where appropriate.

3. Scenario 1 where IPv6 node wants to communicate with IPv4 node is mandatory.

4. Scenarios 2 and 3 are now optional where an IPv6 node is assigned an IPv4 compatible address because an external IPv4 node is attempting communications with the IPv6 node.

5. For scenario 1 DHCPv6 is only needed for DSTM and not the tightly coupled paradigm of a co-existent DHCPv6 and DNS server. Also added mandatory and optional to the DSTM AIIH/NODE/ROUTER Diagram.

6. Made Static Tunnel Endpoints mandatory and Dyanmic Tunnel End Points optional.

7. Fixed DHCPv6 Reconfigure statements to take into account changes to the Reconfigure message in the DHCPv6 working group, to support AIIH processing.

Acknowledgments

The authors would like to acknowledge the implementation contributions by Stephane Atheo at ENST Bretagne who has implemented a DSTM prototype on FreeBSD and input to this specification. We would also like to thank the NGTRANS Working Group for their input.

Normative References

[1] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Architecture", RFC 2460, December 1998.

[3] IPSEC -
S. Kent, R. Atkinson. Security Architecture for the Internet Protocol. RFC 2401, November 1998.
S. Kent, R. Atkinson. IP Authentication Header. RFC 2402, November 1998.
S. Kent, R. Atkinson. IP Encapsulating Security Payload RFC 2406, November 1998.

[4] S. Bradner. Key words for use in RFCs to indicate Requirement Levels. RFC 2119, March 1997.

[5] A. Conta and S. Deering. Generic Packet Tunneling in IPv6. RFC 2473, December 1998.

[6] R. Gilligan and E. Nordmark. Transition Mechanisms for IPv6 Hosts and Routers. RFC 2893, August 2000.

[7] R. Droms. Dynamic Host Configuration Protocol. RFC 2131, March 1997.

[8] Thomson, Narten. IPv6 Stateless Address Configuration. RFC 2462, December 1998.

[9] Hinden, Deering. IP Version 6 Addressing Architecture. RFC 2373, July 1998.

Informative References

[2] J. Bound, M. Carney, C. Perkins, and R. Droms. Dynamic Host Configuration Protocol for IPv6. draft-ietf-dhc-dhcpv6-16.txt October 2001 (work in progress).

Authors' Address

Jim Bound
Compaq Computer Corporation
110 Spitbrook Road
Nashua, NH 003062
USA

Laurent Toutain

ENST Bretagne
BP 78
35512 Cesson Sevigne Cedex
Phone : +33 2 99 12 70 26
Email : Laurent.Toutain@enst-bretagne.fr

Octavio Medina
ENST Bretagne
BP 78
35512 Cesson Sevigne Cedex
Phone : +33 2 99 12 70 23
Email / Octavio.Medina@enst-bretagne.fr

Hossam Afifi
INT
91011 Evry
Phone : +33 1 60 76 40 40
Email : Hossam.Afifi@int-evry.fr

Francis Dupont
ENST Bretagne
BP 78
35 512 Cesson Sevigne Cedex
Phone : +33 2 99 12 70 33
Email : Francis.Dupont@enst-bretagne.fr

Alain Durand
Sun Microsystems
901 San Antonio Road
UMPK 17-202
Palo Alto, CA 94303-4900
Tel: +1 650 786 7503
Fax: +1 650 786 5896
Email: Alain.Durand@sun.com