INTERNET-DRAFT
NGTRANS Working Group
Obsoletes draft-ietf-ngtrans-dstm-03.txt
Expires July 2001

Jim Bound
Nokia Networks
Laurent Toutain
Francis Dupont
ENST Bretagne
Hossam Afifi
INT
Alain Durand
Sun Microsystems

Dual Stack Transition Mechanism (DSTM)

<draft-ietf-ngtrans-dstm-04.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 Global IPv4 Addresses to IPv6/IPv4 nodes over a native IPv6 Network, use 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.

Table of Contents:

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 Global IPv4 Addresses to IPv6/IPv4 nodes over a native IPv6 Network, use 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 we manage one routing plan for IPv6 only.

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 their network. In addition the IPv4 globally routable address space available to the network is a scarce resource, and the user may not want to deploy DHCPv4[15] to assign temporary IPv4 addresses to IPv6 nodes, and would rather require those nodes to use IPv6 to obtain or be given the IPv4 temporary addresses from DHCPv6. Also, to reduce the IPv4 applications a user has to support and to obtain a temporary IPv6 IPv4-Mapped Address (see Section 6), the client only has to run a DHCPv6 client process with the DTI mechanisms in this specification.

The DSTM architecture is composed of a DHCPv6 server, that provides for the assignment of IPv4 Global Addresses to IPv6 Hosts. The DHCPv6 server will allocate temporary IPv4 Global Addresses to IPv6 nodes. The DHCPv6 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. Also a DSTM daemon exists working with a DHCPv6 client to resolve the address space mechanics, between IPv4 and IPv6.

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 DHCPv6 option requirements. Section 7 provides the DSTM Applicability

Statement.

2. Terminology

2.1 IPv6 DSTM Terminology

DSTM Domain	The network areas on an Intranet where a DHCPv6 Server has access to IPv6 nodes participating in DSTM for that network, and IPv4 routing access is not necessary within a DSTM domain.
DSTM Border Router	A border router within a DSTM domain and access to an external IPv4-ONLY domain.
DSTM Host	A Host that supports a dual IP layer IPv4 and IPv6 stack, DTI, and a DHCPv6 Client process.
IPv6 Protocol Terms:	See [3]
IPv6 Transition Terms:	See [14]
DHCPv6 Terms:	See [4]
DTI:	Dynamic Tunneling Interface. An interface encapsulating IPv4 packets into IPv6 packets.
IPv4 Global Address:	An IPv4 address that is globally routable on the Internet.
Tunnel End Point (TEP)	Destination of the IPv6 packet containing an IPv4 packet. In most cases this will be a dual stack 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 [8]. 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 which uses existing protocols. DSTM does not specify a protocol. However, DSTM defines a new DHCPv6 Option for transition.

The motivation for DSTM is to provide IPv6 nodes a means to acquire an IPv4 Global 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 which carry 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), which 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.

- Standard DHCPv6 is used to support the extension to provide and accept from DHCPv6 Servers Global IPv4 Addresses.

- 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. DHCPv6 is used to provide TEPs to IPv6 nodes supporting DTI, as part of the new DHCPv6 Option.

- 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 DHCPv6 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 Global IPv4 Addresses.

  o Software within an IPv6 node to support the dynamic tunneling mechanisms in this specification to encapsulate IPv4 packets within IPv6 on an IPv6 node. In addition a daemon must exist to access a DHCPv6 client for Global IPv4 Mapped Addresses and TEPs. How this daemon communicates with a DHCPv6 Client implementation is implementation defined, and left as an exercise for implementors of this transition mechanism.

  o Software in DSTM Border Routers to recall or be able to cache the association of IPv6 and IPv4 addresses of nodes during decapsulation and encapsulation.

A simplistic overview of DSTM is as follows:

```
-----------------------------------------------
                                              |    IPv4 Internet or Intranet
        DSTM Domain Intranet                  |      IPv4 Applications
                                              |         Domain
                _____________________         |
               |                     |        |
               | DHCPv6 Server       |        |
               |_____________________|        |
                             ^                |
                             |                |
   __________________        |               _|_______
  |                  |       |              |        |
  | IPv6/IPv4 Node   |       |              |  DSTM  |
  |------------------|       |              | Border |
  |  DSTM Daemon     |       |              | Router |
  |  DHCPv6 client   |<-------              |  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 IPv6 node is configured with the IPv6 address of a TEP, where an IPv4 encapsulated packet will be sent.

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

```
        DHCPv6
        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 DHCPv6 server for an
    |            |            |      IPv4 address using DHCPv6
    |<====       |            |    - The DHCPv6 server locates the client
    |            |            |      and provides a temporary IPv4
    |            |            |      global address.
    |+++++++++++>|            |    - The DTI sends the IPv6 packet to the
    |            |            |      TEP.
    |            |----------->|    - Y sends the packet to the destination Z4
    |            |            |    - Y caches the association between
    |            |            |      the IPv4 and IPv6 addresses of X.
```

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) that will decapsulate them and inject 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 it has to send the first IPv4 packet, a request is sent to the DHCPv6 client. The DHCPv6 client will send a DHCPv6 request to the DHCPv6 server to get the temporary IPv4 Global Address and a TEP.

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.

5.2 DTI Encapsulation of IPv4 packets

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.

5.3 DTI IPv6 destination address

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 can be either statically configured or dynamically acquired when the IPv6 node acquires an IPv4 Compatible Address from a DHCPv6 Server.

The TEP SHOULD be provided by the DHCPv6 server when the DSTM node receives an IPv4-Mapped IPv6 Address (section 6). But, a DSTM node MAY manually configure the TEP during early deployment of IPv6, this will not scale and is not recommended as a long term transition solution.

6. DHCPv6 Requirements

The DSTM processes will use the DHCPv6 services [4] 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.

6.1 DHCPv6 Global IPv4 Address Option

The DHCPv6 IPv4 Address Option informs a DHCPv6 Client or Server that the Identity Association Option (IA) [4] following this option will contain an IPv4-Mapped IPv6 Address [19] 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.

6.1.1 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.

6.1.2 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 [4].

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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|              TBD              |           variable            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            IA UUID                            |
|                          (8 octets)                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              T1                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              T2                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   num-addrs   |              IPv6 address                     |
+-+-+-+-+-+-+-+-+              (16 octets)                      |
|                                                               |
|                                                               |
+               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               |   pref. len   |      preferred lifetime       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| pref. lifetime (cont.)        |        valid lifetime         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| valid lifetime (cont.)        |         IPv6 address          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
|                              ...                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

option-code

TBD

option-len
Variable; equal to 17 + num-addrs*25

IA UUID
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.

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

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).

6.1.3 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.

6.2 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.

6.3 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 [18]. 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 [12, 14].

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.

DSTM requires the use of DHCPv6 to obtain IPv4 addresses and TEPs for a DSTM node. 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 Hosts 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 [4], and for communications between the IPv6 node, once it has an IPv4 address, and the remote IPv4 node, IPsec [7] can be used as DSTM does not break secure end-to-end communications at any point in the mechanism.

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.

References

[1] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, USC/Information Sciences Institute, November 1987.

[2] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, USC/Information Sciences Institute, November 1987.

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

[4] J. Bound, M. Carney, C. Perkins, and R. Droms. Dynamic Host Configuration Protocol for IPv6. draft-ietf-dhc-dhcpv6-16.txt November 2000 (work in progress).

[5] P. Vixie, S. Thomson, Y. Rekhter, and J. Bound. Dynamic Updates to the Domain Name System (DNS). RFC 2136, April 1997.

[6] William R. Cheswick and Steven Bellovin. Firewalls and Internet

Security. Addison-Wesley, Reading, MA 1994 (ISBN:

0-201-63357-4).

[7] 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.

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

[9] D. Eastlake and C. Kaufman. Domain Name System Security Extensions. RFC 2065, January 1997.

[10] D. Eastlake. Secure Domain Name System Dynamic Update. RFC 2137, April 1997.

[11] R. Callon and D. Haskins. Routing Aspects Of IPv6 Transition RFC 2185, September 1997.

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

[13] E. Nordmark. Stateless IP/ICMP Translator (SIIT) RFC 2765, February 2000.

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

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

[16] Rekhter, Moskowitz, Karrenburg, Groot. Address Allocation for Private Networks. RFC 1918. February 1996.

[17] M. Crawford, C. Huitema. DNS Extensions to Support IPv6 Address Aggregation and Renumbering. RFC 2874, July 2000.

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

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

Authors' Address

Jim Bound
Nokia Networks
5 Wayside Road
Burlington, MA 01803
Phone: +1 781 492 0613
Email: Jim.Bound@nokia.com

Laurent Toutain
ENST Bretagne
BP 78
35 512 Cesson
Phone : +33 2 99 12 70 26
Email : Laurent.Toutain@enst-bretagne.fr

Hossam Afifi
INT
91 011 M-Ivry
Phone : +33 1 60 76 40 40
Email : Hossam.Afifi@int-evry.fr

Francis Dupont
ENST Bretagne
BP 78
35 512 Cesson
Phone : +33 2 99 12 70 36
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