

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: October 18, 2013

U. Herberg, Ed.
Fujitsu
A. Cardenas
University of Texas at Dallas
T. Iwao
Fujitsu
M. Dow
Freescale
S. Cespedes
U. Icesi
April 16, 2013

Depth-First Forwarding in Unreliable Networks (DFF)
draft-cardenas-dff-11

Abstract

This document specifies the "Depth-First Forwarding" (DFF) protocol for IPv6 networks, a data forwarding mechanism that can increase reliability of data delivery in networks with dynamic topology and/or lossy links. The protocol operates entirely on the forwarding plane, but may interact with the routing plane. DFF forwards data packets using a mechanism similar to a "depth-first search" for the destination of a packet. The routing plane may be informed of failures to deliver a packet or loops. This document specifies the DFF mechanism both for IPv6 networks (as specified in [RFC2460](#)) and in addition also for LoWPAN "mesh-under" networks (as specified in [RFC4944](#)). DFF assumes that the underlying link layer provides means to detect if a packet has been successfully delivered to the next hop or not, is designed for networks with little traffic, and is used for unicast transmissions only.

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 October 18, 2013.

Copyright Notice

Copyright (c) 2013 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	5
1.1.	Motivation	5
1.2.	Experiments to be conducted	6
2.	Notation and Terminology	7
2.1.	Notation	7
2.2.	Terminology	8
3.	Applicability Statement	9
4.	Protocol Overview and Functioning	11
4.1.	Information Sets Overview	12
4.2.	Signaling Overview	12
5.	Protocol Dependencies	13
6.	Information Sets	13
6.1.	Bidirectional Neighbor List	13
6.2.	Processed Set	14
7.	Packet Header Fields	14
8.	Protocol Parameters	15
9.	Data Packet Generation and Processing	16
9.1.	Data Packets Entering the DFF Routing Domain	16
9.2.	Data Packet Processing	17
10.	Unsuccessful Packet Transmission	19
11.	Determining the Next Hop for a Packet	20
12.	Informing the Routing Protocol	21
13.	Sequence Numbers	22
14.	Modes of Operation	22
14.1.	Route-Over	22
14.1.1.	Mapping of DFF Terminology to IPv6 Terminology	22
14.1.2.	Packet Format	23
14.2.	Mesh-Under	25
14.2.1.	Mapping of DFF Terminology to LowPAN Terminology	25
14.2.2.	Packet Format	26
15.	Scope Limitation of DFF	27
15.1.	Route-Over MoP	28
15.2.	Mesh-Under MoP	30
16.	MTU Exceedance	31
17.	Security Considerations	31
17.1.	Attacks Out of Scope	32
17.2.	Protection Mechanisms of DFF	32
17.3.	Attacks In Scope	32
17.3.1.	Denial of Service	32
17.3.2.	Packet Header Modification	33
17.3.2.1.	Return Flag Tampering	33
17.3.2.2.	Duplicate Flag Tampering	33
17.3.2.3.	Sequence Number Tampering	34
18.	IANA Considerations	34
19.	Acknowledgements	34
20.	References	35

20.1.	Normative References	35
20.2.	Informative References	35
Appendix A.	Examples	36
A.1.	Example 1: Normal Delivery	36
A.2.	Example 2: Forwarding with Link Failure	37
A.3.	Example 3: Forwarding with Missed Link Layer Acknowledgment	38
A.4.	Example 4: Forwarding with a Loop	39
Appendix B.	Deployment Experience	39
B.1.	Deployments in Japan	39
B.2.	Kit Carson Electric Cooperative	40
B.3.	Simulations	40
B.4.	Open Source Implementation	40
	Authors' Addresses	40

1. Introduction

This document specifies the Depth-First Forwarding (DFF) protocol for IPv6 networks, both for IPv6 forwarding ([\[RFC2460\]](#), henceforth denoted "route-over"), and also for "mesh-under" forwarding using the LoWPAN adaptation layer ([\[RFC4944\]](#)). The protocol operates entirely on the forwarding plane, but may interact with the routing plane. The purpose of DFF is to increase reliability of data delivery in networks with dynamic topologies and/or lossy links.

DFF forwards data packets using a "depth-first search" for the destination of the packets. DFF relies on an external neighborhood discovery mechanism which lists neighbors of a router that may be attempted as next hops for a data packet. In addition, DFF may use information from the Routing Information Base (RIB) for deciding in which order to try to send the packet to the neighboring routers.

If the packet makes no forward progress using the first selected next hop, DFF will successively try all neighbors of the router. If none of the next hops successfully receives or forwards the packet, DFF returns the packet to the previous hop, which in turn tries to send it to alternate neighbors.

As network topologies do not necessarily form trees, loops can occur. Therefore, DFF contains a loop detection and avoidance mechanism.

DFF may provide information, which may - by a mechanism outside of this specification - be used for updating cost of routes in the RIB based on failed or successful delivery of packets through alternative next hops. Such information may also be used by a routing protocol.

DFF assumes that the underlying link layer provides means to detect if a packet has been successfully delivered to the next hop or not, is designed for networks with little traffic, and is used for unicast transmissions only.

1.1. Motivation

In networks with dynamic topologies and/or lossy links, even frequent exchanges of control messages between routers for updating the routing tables cannot guarantee that the routes correspond to the effective topology of the network at all times. Packets may not be delivered to their destination because the topology has changed since the last routing protocol update.

More frequent routing protocol updates can mitigate that problem to a certain extent, however this requires additional signaling, consuming channel and router resources (e.g., when flooding control messages

through the network). This is problematic in networks with lossy links, where further control traffic exchange can worsen the network stability because of collisions. Moreover, additional control traffic exchange may drain energy from battery-driven routers.

The data-forwarding mechanism specified in this document allows for forwarding data packets along alternate paths for increasing reliability of data delivery, using a depth-first search. The objective is to decrease the necessary control traffic overhead in the network, and at the same time to increase delivery success rates.

As this specification is intended for experimentation, the mechanism is also specified for forwarding on the LOWPAN adaption layer (according to [Section 11 of \[RFC4944\]](#)), in addition to IPv6 forwarding as specified in [\[RFC2460\]](#). Other than different header formats, the DFF mechanism for route-over and mesh-under is similar, and is therefore first defined in general, and then more specifically for both IPv6 route-over forwarding (as specified in [Section 14.1](#)), and for LOWPAN adaption layer mesh-under (as specified in [Section 14.2](#)).

1.2. Experiments to be conducted

This document is presented as an Experimental specification that can increase reliability of data delivery in networks with dynamic topology and/or lossy links. It is anticipated that, once sufficient operational experience has been gained, this specification will be revised to progress it on to the Standards Track. Experience reports regarding DFF implementation and deployment are encouraged particularly with respect to:

- o Optimal values for the parameter P_HOLD_TIME, depending on the size of the network, the topology and the amount of traffic originated per router. The longer a Processed Tuple is hold, the more memory is consumed on a router. Moreover, if a tuple is hold too long, a sequence number wrap-around may occur, and a new packet may have the same sequence number as one indicated in an old Processed Tuple. However, if the tuple is expired too soon (before the packet has been completed its path to the destination), it may be mistakenly detected as new packet instead of one already seen.
- o Optimal values for the parameter MAX_HOP_LIMIT, depending on the size of the network, the topology, and the lossyness of the link layer. MAX_HOP_LIMIT makes sure that packets do not unnecessarily traverse in the network; it may be used to limit the "detour" of packets that is acceptable. The value may also be based on a per-packet-basis if hop-count information is available from the RIB or

routing protocol. In such a case, the hop-limit for the packet may be a percentage (e.g., 200%) of the hop-count value indicated in the routing table.

- o Optimal methods to increase cost of a route when a loop or lost L2 ACK is detected by DFF. While this is not specified as a normative part of this document, it may be of interest in an experiment to find good values of how much to increase link cost in the RIB or routing protocol.
- o Performance of using DFF in combination with different routing protocols, such as reactive and proactive protocols. This also implies how routes are updated by the RIB / routing protocol when informed by DFF about loops or broken links.

2. Notation and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

Additionally, this document uses the notation in [Section 2.1](#) and the terminology in [Section 2.2](#).

2.1. Notation

The following notations are used in this document:

List: A list of elements is defined as [] for an empty list, [element] for a list with one element, and [element1, element2, ...] for a list with multiple elements.

Concatenation of lists: If L1 and L2 are lists, then L1@L2 is a new list with first all elements of L1, followed by all elements of L2 in that order.

Byte order: All packet formats in this specification use network byte order (most significant octet first) for all fields. The most significant bit in an octet is numbered bit 0, and the least significant bit of an octet is numbered bit 7.

Assignment: a := b

An assignment operator, whereby the left side (a) is assigned the value of the right side (b).

Comparison: `c = d`

A comparison operator, returning true if the value of the left side (c) is equal to the value of the right side (d).

Flags: This specification uses multiple 1-bit flags. A value of '0' of a flag means 'false', a value of '1' means 'true'.

2.2. Terminology

The terms "route-over" and "mesh-under", introduced in [[RFC6775](#)] are used in this document, where "route-over" is not only limited to 6LoWPANs but applies to general IPv6 networks:

Mesh-under: A topology where nodes are connected to a [6LoWPAN Border Router] 6LBR through a mesh using link-layer forwarding. Thus, in a mesh-under configuration, all IPv6 hosts in a LoWPAN are only one IP hop away from the 6LBR. This topology simulates the typical IP-subnet topology with one router with multiple nodes in the same subnet.

Route-over: A topology where hosts are connected to the 6LBR through the use of intermediate layer-3 (IP) routing. Here, hosts are typically multiple IP hops away from a [6LoWPAN Router] 6LBR. The route-over topology typically consists of a 6LBR, a set of 6LRs, and hosts.

The following terms are used in this document. As the DFF mechanism is specified both for route-over IPv6 and for mesh-under LoWPAN adaptation layer, the terms are generally defined in this section, and then specifically mapped for each of the different modes of operation in [Section 14](#).

Depth-first search: "Depth-first search (DFS) is an algorithm for traversing or searching a tree, tree structure, or graph. One starts at the root (selecting some node as the root in the graph case) and explores as far as possible along each branch before backtracking" [[DFS wikipedia](#)]. In this document, the algorithm for traversing a graph is applied to forwarding packets in a computer network, with nodes being routers.

Routing Information Base (RIB): A table stored in the user-space of an operating system of a router or host. The table lists routes to network destinations, as well as associated metrics with these routes.

Mode of Operation (MoP): The DFF mechanism specified in this document can either be used as "route-over" IPv6 forwarding mechanism (Mode of Operation: "route-over"), or as "mesh-under" LoWPAN adaptation layer (Mode of Operation: "mesh-under").

Packet: An IPv6 Packet (for "route-over" MoP) or a "LoWPAN encapsulated packet" (for "mesh-under" MoP) containing an IPv6 Packet as payload.

Packet Header: An IPv6 extension header (for "route-over" MoP) or a LoWPAN header (for "mesh-under" MoP).

Address: An IPv6 address (for "route-over" MoP), or a 16-bit short or EUI-64 link layer address (for "mesh-under" MoP).

Originator: The router which added the DFF header (specified in [Section 7](#)) to a Packet.

Originator Address: An Address of the Originator. This Address SHOULD be an Address configured on the interface which transmits the Packet, selected according to [[RFC6724](#)].

Destination: The router or host to which a Packet is finally destined. In case this router or host is outside of the routing domain in which DFF is used, the Destination is the router that removes the DFF header (specified in [Section 7](#)) from the Packet. This case is described in [Section 15.1](#).

Destination Address: An Address to which the Packet is sent.

Next Hop: An Address of the next hop router to which the Packet is sent along the path to the Destination.

Previous Hop: The Address of the previous hop router from which a Packet has been received. In case the Packet has been received by a router from outside of the routing domain where DFF is used (i.e., no DFF header is contained in the Packet), the Originator Address of the router adding the DFF header to the Packet is used as the Previous Hop.

Hop Limit: An upper bound how many times the Packet may be forwarded.

[3.](#) Applicability Statement

This document specifies DFF, a packet forwarding mechanism intended for use in networks with dynamic topology and/or lossy links with the

purpose of increasing reliability of data delivery. The protocol's applicability is determined by its characteristics, which are that this protocol:

- o Is applicable for use in IPv6 networks, either as "route-over" forwarding mechanism using IPv6 ([[RFC2460](#)]), or as "mesh-under" forwarding mechanism using the frame format for transmission of IPv6 packets defined in [[RFC4944](#)].
- o Assumes addresses used in the network are either IPv6 addresses (if the protocol is used as "route-over"), or 16-bit short or EUI-64 link layer addresses, as specified in [[RFC4944](#)] if the protocol is used as "mesh-under". In "mesh-under" mode, mixed 16-bit and EUI-64 addresses within one DFF routing domain are allowed (if conform with [[RFC4944](#)]), as long as DFF is limited to be used within one PAN. It is assumed that the "route-over" mode and "mesh-under" mode are mutually exclusive in the same routing domain.
- o Assumes that the underlying link layer provides means to detect if a Packet has been successfully delivered to the Next Hop or not (e.g., by L2 ACK messages).
- o Is designed to work in networks with lossy links and/or with a dynamic topology. In networks with very stable links (e.g. Ethernet) and fixed topology, DFF will not bring any benefit (but also not be harmful, other than the additional overhead for the Packet header).
- o Is designed to work in a completely distributed manner, and does not depend on any central entity.
- o Is designed for networks with little traffic in terms of numbers of Packets per second, since each recently forwarded Packet increases the state on a router. The amount of traffic per time that is supported by DFF depends on the memory resources of the router running DFF, on the density of the network, on the loss rate of the channel, and the maximum hop limit for each Packet: for each recently seen Packet, a list of Next Hops that the Packet has been sent to is stored in memory. The stored entries can be deleted after an expiration time, so that only recently received Packets require storage on the router.
- o Is designed for dense topologies with multiple paths between each source and each destination. Certain topologies are less suitable for DFF: topologies that can be partitioned by the removal of a single router or link, topologies with multiple stub routers that each have a single link to the network, topologies with only a

single path to a destination, or topologies where the "detour" that a Packet makes during the depth-first search in order to reach the destination would be too long. Note that the number of retransmissions of a Packet that stipulate a "too long" path depends on the underlying link layer (capacity and probability of Packet loss), as well as how much bandwidth is required for data traffic by applications running in the network. In such topologies, the Packet may never reach the Destination, and therefore unnecessary transmissions of data Packets may occur until the Hop Limit of the Packet reaches zero and the Packet is dropped. This may consume channel and router resources.

- o Is used for unicast transmissions only (not for anycast or multicast).
- o Is for use within stub networks, and for traffic between a router inside the routing domain in which DFF is used and a known border router. Examples of such networks are LOWPANs. Scope limitations are described in [Section 15](#).

4. Protocol Overview and Functioning

When a Packet is to be forwarded by a router using DFF, the router creates a list of candidate Next Hops for that Packet. This list (created per packet) is ordered, and [Section 11](#) provides recommendations of how to order the list, e.g., first listing Next Hops listed in the RIB, if available, ordered in increasing cost, followed by other neighbors provided by an external neighborhood discovery. DFF proceeds to forward the Packet to the Next Hop listed first in the list. If the transmission was not successful (as determined by the underlying link layer) or if the Packet was "returned" by a Next Hop to which it had been sent before, the router will try to forward the Packet to the next Next Hop on the list. A router "returns" a Packet to the router from which it was originally received, once it has unsuccessfully tried to forward the Packet to all elements in the candidate Next Hop list. If the Packet is eventually returned to the Originator of the Packet, it is dropped.

For each recently forwarded Packet, a router running DFF stores the list of Next Hops to which a Packet has been sent. Packets are identified by a sequence number that is included in the Packet Header (refer to [Section 13](#) for further details on the sequence number). This list of recently forwarded Packets also allows for avoiding loops when forwarding a Packet. Entries of the list (identified by a sequence number of a Packet) expire after a given expiration timeout, and are removed.

4.1. Information Sets Overview

This specification requires a single set on each router, the Processed Set. Moreover, a list of bidirectional neighbors must be provided by an external neighborhood discovery mechanism, or may be determined from the RIB (e.g., if the RIB provides routes to adjacent routers, and if these one-hop routes are verified to be bidirectional). The Processed Set stores the sequence number, the Originator Address, the Previous Hop and a list of Next Hops, to which the Packet has been sent, for each recently seen Packet. Entries in the set are removed after a predefined time-out. Each time a Packet is forwarded to a Next Hop, that Next Hop is added to the list of Next Hops of the entry for the Packet.

Note that an implementation of this protocol may maintain the information of the Processed Set in the indicated form, or in any other organization which offers access to this information. In particular, it is not necessary to remove Tuples from a Set at the exact time indicated, only to behave as if the Tuples were removed at that time.

4.2. Signaling Overview

DFF requires additional header information in each data Packet by a router using this specification. This information is stored in a Packet Header that is specified in this document as LOWPAN header and as IPv6 Hop-by-Hop Options extension header respectively, for the intended "route-over" and "mesh-under" Modes of Operations. This DFF header contains a sequence number used for uniquely identifying a Packet, and two flags: RET (for "return") and DUP (for "duplicated").

While a router successively tries sending a data Packet to one or more of its neighbors, $RET = 0$. If none of the transmissions of the Packet to the neighbors of a router have succeeded, the Packet is returned to the Previous Hop, indicated by setting the return flag ($RET := 1$). The RET flag is required to discern between a deliberately returned Packet and a looping Packet: if a router receives a Packet with $RET = 1$ (and $DUP = 0$ or $DUP = 1$) that it has already forwarded, the Packet was deliberately returned, and the router will continue to successively send the Packet to routers from the candidate Next Hop list. If that Packet has $RET = 0$, the router assumes that the Packet is looping and returns it to the Previous Hop. An external mechanism may use this information for increasing the route cost of the route to the Destination using the Next Hop which resulted in the loop in the RIB. Alternatively, or in addition, the routing protocol may be informed. Note that once DUP is set to 1, loop detection is not possible any more as the flag is not reset any more. Therefore, a Packet may loop if the RIBs of

routers in the domain are inconsistent, until hop limit has reached 0.

Whenever a Packet transmission to a neighbor has failed (as determined by the underlying link layer, e.g., using L2 ACKs), the duplicate (DUP) flag is set in the Packet Header for the following transmissions. The rationale is that the Packet may have been successfully received by the neighbor and only the L2 ACK has been lost, resulting in possible duplicates of the Packet in the network. The DUP flag tags such a possible duplicate. The DUP flag is required to discern between a duplicated Packet and a looping Packet: if a router receives a Packet with DUP = 1 (and RET = 0) that it has already forwarded, the Packet is not considered looping, and successively forwarded to the next router from the candidate Next Hop list. If the received Packet has DUP = 0 (and RET = 0), the router assumes that the Packet is looping, sets RET := 1, and returns it to the Previous Hop. Again, an external mechanism may use this information for increasing route costs and/or informing the routing protocol.

5. Protocol Dependencies

DFF MAY use information from the Routing Information Base (RIB), specifically for determining an order of preference for to which next hops a packet should be forwarded (e.g., the packet may be forwarded first to neighbors that are listed in the RIB as next hops to the destination, preferring those with the lowest route cost). [Section 11](#) provides recommendations about the order of preference for the next hops of a packet.

DFF MUST have access to a list of bidirectional neighbors for each router, provided by a mechanism such as, e.g., NHDP [[RFC6130](#)]. That neighborhood discovery protocol is not specified in this document.

6. Information Sets

This section specifies the information sets used by DFF.

6.1. Bidirectional Neighbor List

DFF MUST have access to a list of Addresses of bidirectional neighbors of the router. This list can be provided by an external neighborhood discovery mechanism, or alternatively may be determined from the RIB (e.g., if the RIB provides routes to adjacent routers, and if these one-hop routes are verified to be bidirectional). The list of Addresses of bidirectional neighbors is not specified within

this document. The Addresses in the list are used to construct a list of candidate Next Hops for a Packet, as specified in [Section 11](#).

6.2. Processed Set

Each router maintains a Processed Set in order to support the loop detection functionality. The Processed Set lists sequence numbers of previously received Packets, as well as a list of Next Hops to which the Packet has been sent successively as part of the depth-first forwarding mechanism. The Processed Set SHOULD be stored in non-volatile memory and restored after a reboot of the router.

The set consists of Processed Tuples

```
(P_orig_address, P_seq_number, P_prev_hop,  
P_next_hop_neighbor_list, P_time)
```

where

P_orig_address is the Originator Address of the received Packet;

P_seq_number is the sequence number of the received Packet;

P_prev_hop is the Address of the Previous Hop of the Packet;

P_next_hop_neighbor_list is a list of Addresses of Next Hops to which the Packet has been sent previously, as part of the depth-first forwarding mechanism, as specified in [Section 9.2](#);

P_time specifies when this Tuple expires and MUST be removed.

The consequences when no or not enough non-volatile storage is available on a router (e.g., because of limited resources) are that Packets that are already in a loop caused by the routing protocol may continue to loop until the hop limit is exhausted. Non-looping Packets may be sent to Next Hops that have already received the Packet previously and will return the Packet, leading to some unnecessary retransmissions. This effect is only temporary and applies only for Packets already traversing the network.

7. Packet Header Fields

This section specifies the information required by DFF in the Packet Header. Note that, depending on whether DFF is used in the "route-over" MoP or in the "mesh-under" MoP, the DFF header is either an IPv6 Hop-by-Hop Options extension header (as specified in [Section 14.1.2](#)) or a LowPAN header (as specified in [Section 14.2.2](#)).

[Section 14.1.2](#) and [Section 14.2.2](#) specify the precise order, format and encoding of the fields that are listed in this section.

Version (VER) - This 2-bit value indicates the version of DFF that is used. This specification defines value 00. Packets with other values of the version MUST be forwarded using forwarding as defined in [[RFC2460](#)] and [[RFC4944](#)] for route-over and mesh-under MoP respectively.

Duplicate Packet Flag (DUP) - This 1-bit flag is set in the DFF header of a Packet, when that Packet is being re-transmitted due to a signal from the link-layer that the original transmission failed, as specified in [Section 9.2](#). Once the flag is set to 1, it MUST NOT be modified by routers forwarding the Packet.

Return Packet Flag (RET) - The 1-bit flag MUST be set to 1 prior to sending the Packet back to the Previous Hop. Upon receiving a packet with RET = 1, and before sending it to a new Candidate Next Hop, that flag MUST be set to 0 as specified in [Section 9.2](#).

Sequence Number - A 16-bit field, containing an unsigned integer sequence number generated by the Originator, unique to each router for each Packet to which the DFF has been added, as specified in [Section 13](#). The Originator Address concatenated with the sequence number represents an identifier of previously seen data Packets. Refer to [Section 13](#) for further information about sequence numbers.

8. Protocol Parameters

The parameters used in this specification are listed in this section. These parameters are configurable and do not need to be stored in non-volatile memory. Default values for the parameters depend on the network size, topology, link layer and traffic patterns. Part of the experimentation described in [Section 1.2](#) is to determine suitable default values.

P_HOLD_TIME - Is the time period after which a newly created or modified Processed Tuple expires and MUST be deleted. An implementation SHOULD use a value for P_HOLD_TIME that is high enough that the Processed Tuple for a Packet is still in memory on all forwarding routers while the Packet is transiting the routing domain. The value SHOULD at least be MAX_HOP_LIMIT times the expected time to send a Packet to a router on the same link. The value MUST be lower than the time it takes until the same sequence number is reached again after a wrap-around on the router identified by P_orig_address of the Processed Tuple.

MAX_HOP_LIMIT - Is the initial value of Hop Limit, and therefore the maximum number of times that a Packet is forwarded in the routing domain. When choosing the value of MAX_HOP_LIMIT, the size of the network, the distance between source and destination in number of hops, and the maximum possible "detour" of a Packet SHOULD be considered (compared to the shortest path). Such information MAY be used from the RIB, if provided.

9. Data Packet Generation and Processing

The following sections specify the process of handling a Packet entering the DFF routing domain (i.e., without DFF header) in [Section 9.1](#), as well as forwarding a data Packet from another router running DFF in [Section 9.2](#).

9.1. Data Packets Entering the DFF Routing Domain

This section applies for any data Packets upon their first entry into a routing domain, in which DFF is used. This occurs when a new data Packet is generated on this router, or when a data Packet is forwarded from outside the routing domain (i.e., from a host attached to this router or from a router outside the routing domain in which DFF is used). Before such a data Packet (henceforth denoted "current Packet") is transmitted, the following steps MUST be executed:

1. If required, encapsulate the Packet as specified in [Section 15](#).
2. Add the DFF header to the current Packet (to the outer header if the Packet has been encapsulated), with:
 - * DUP := 0;
 - * RET := 0;
 - * Sequence Number := a new sequence number of the Packet (as specified in [Section 13](#)).
3. Check that the Packet does not exceed the MTU as specified in [Section 16](#). In case it does, execute the procedures listed in [Section 16](#) and do not further process the Packet.
4. Select the Next Hop (henceforth denoted "next_hop") for the current Packet, as specified in [Section 11](#).
5. Add a Processed Tuple to the Processed Set with:

- * P_orig_address := the Originator Address of the current Packet;
 - * P_seq_number := the sequence number of the current Packet;
 - * P_prev_hop := the Originator Address of the current Packet;
 - * P_next_hop_neighbor_list := [next_hop];
 - * P_time := current time + P_HOLD_TIME.
6. Pass the current Packet to the underlying link layer for transmission to next_hop. If the transmission fails (as determined by the link layer), the procedures in [Section 10](#) MUST be executed.

9.2. Data Packet Processing

When a Packet (henceforth denoted the "current Packet") is received by a router, then the following tasks MUST be performed:

1. If the Packet Header is malformed (i.e., the header format is not as expected by this specification), drop the Packet.
2. Otherwise, if the Destination Address of the Packet matches an Address of an interface of this router, deliver the Packet to upper layers and do not further process the Packet as specified below.
3. Decrement the value of the Hop Limit field by one (1).
4. Drop the Packet if Hop Limit is decremented to zero and do not further process the Packet as specified below.
5. If no Processed Tuple (henceforth denoted the "current tuple") exists in the Processed Set, with:
 - + P_orig_address = the Originator Address of the current Packet, AND;
 - + P_seq_number = the sequence number of the current Packet.

Then:

1. Add a Processed Tuple (henceforth denoted the "current tuple") with:

- + P_orig_address := the Originator Address of the current Packet;
 - + P_seq_number := the sequence number of the current Packet;
 - + P_prev_hop := the Previous Hop Address of the current Packet;
 - + P_next_hop_neighbor_list := [];
 - + P_time := current time + P_HOLD_TIME.
2. Set RET to 0 in the DFF header.
 3. Select the Next Hop (henceforth denoted "next_hop") for the current Packet, as specified in [Section 11](#).
 4. P_next_hop_neighbor_list := P_next_hop_neighbor_list@[next_hop].
 5. Pass the current Packet to the underlying link layer for transmission to next_hop. If the transmission fails (as determined by the link layer), the procedures in [Section 10](#) MUST be executed.
 6. Otherwise, if a tuple exists:
 1. If the return flag of the current Packet is not set (RET = 0) (i.e., a loop has been detected):
 1. Set RET := 1.
 2. Pass the current Packet to the underlying link layer for transmission to the Previous Hop.
 2. Otherwise, if the return flag of the current Packet is set (RET = 1):
 1. If the Previous Hop of the Packet is not contained in P_next_hop_neighbor_list of the current tuple, drop the Packet.
 2. If the Previous Hop of the Packet (i.e., the address of the router from which the current Packet has just been received) is equal to P_prev_hop of current tuple (i.e., the address of the router from which the current Packet has been first received), drop the Packet.

3. Set `RET := 0`.
4. Select the Next Hop (henceforth denoted "next_hop") for the current Packet, as specified in [Section 11](#).
5. Modify the current tuple:
 - `P_next_hop_neighbor_list := P_next_hop_neighbor_list@ [next_hop]`;
 - `P_time := current time + P_HOLD_TIME`.
6. If the selected Next Hop is equal to `P_prev_hop` of the current tuple, as specified in [Section 11](#), (i.e., all Candidate Next Hops have been unsuccessfully tried), set `RET := 1`. If this router (i.e., the router receiving the current packet) has the same Address as the Originator Address of the current Packet, drop the Packet.
7. Pass the current Packet to the underlying link layer for transmission to next_hop. If transmission fails (as determined by the link layer), the procedures in [Section 10](#) MUST be executed.

[10](#). Unsuccessful Packet Transmission

DFF requires that the underlying link layer provides information as to whether a Packet is successfully received by the Next Hop. Absence of such a signal is interpreted as delivery failure of the Packet (henceforth denoted the "current Packet"). Note that the underlying link layer MAY retry sending the Packet multiple times (e.g., using exponential back-off) before determining that the Packet has not been successfully received by the Next Hop. Whenever [Section 9](#) explicitly requests it in case of such a delivery failure, the following steps are executed:

1. Set the duplicate flag (DUP) of the DFF header of the current Packet to 1.
2. Select the Next Hop (henceforth denoted "next_hop") for the current Packet, as specified in [Section 11](#).
3. Find the Processed Tuple (the "current tuple") in the Processed Set, with:

- + P_orig_address = the Originator Address of the current Packet,
AND;
- + P_seq_number = the sequence number of the current Packet,
- 4. If no current tuple is found, drop the Packet.
- 5. Otherwise, modify the current tuple:
 - * P_next_hop_neighbor_list := P_next_hop_neighbor_list@[next_hop];
 - * P_time := current time + P_HOLD_TIME.
- 6. If the selected next_hop is equal to P_prev_hop of the current tuple, as specified in [Section 11](#) (i.e., all neighbors have been unsuccessfully tried):
 - * RET := 1
 - * Decrement the value of the Hop Limit field by one (1). Drop the Packet if Hop Limit is decremented to zero.
- 7. Otherwise
 - * RET := 0
- 8. Transmit the current Packet to next_hop. If transmission fails (determined by the link layer), and if the next_hop does not equal P_prev_hop from the current tuple, the procedures in [Section 10](#) MUST be executed.

[11.](#) Determining the Next Hop for a Packet

When forwarding a Packet, a router determines a valid Next Hop for that Packet as specified in this section. As a Processed Tuple was either existing when receiving the Packet (henceforth denoted the "current Packet"), or otherwise was created, it can be assumed the a Processed Tuple for that Packet (henceforth denoted the "current tuple") is available.

The Next Hop is chosen from a list of candidate Next Hops in order of decreasing priority. This list is created per Packet. The maximum candidate Next Hop List for a Packet contains all the neighbors of the router (as determined from an external neighborhood discovery process), except for the Previous Hop of the current Packet. A smaller list MAY be used, if desired, and the exact selection of the

size of the candidate Next Hop List is a local decision in each router, which does not affect interoperability. If information from the RIB is used, then the candidate Next Hop list MUST contain at least the Next Hop, indicated in the RIB as the Next Hop on the shortest path to the destination, and SHOULD contain all Next Hops, indicated to the RIB as Next Hops on paths to the destination. If a Next Hop from the RIB equals the Previous Hop of the current Packet, it MUST NOT be added to the candidate Next Hop list, and the RIB MAY be informed about a potential loop as specified in [Section 12](#).

The list MUST NOT contain Addresses which are listed in `P_next_hop_neighbor_list` of the current tuple, in order to avoid sending the Packet to the same neighbor multiple times. Moreover, an Address MUST NOT appear more than once in the list, for the same reason. Also, Addresses of an interface of this router MUST NOT be added to the list.

The list has an order of preference, where Next Hops at the top of the list being the ones that Packets are sent to first in the depth-first processing specified in [Section 9.1](#) and [Section 9.2](#). The following order is RECOMMENDED, with the elements listed on top having the highest preference:

1. The neighbor that is indicated in the RIB as the Next Hop on the shortest path to the destination of the current Packet;
2. Other neighbors indicated in the RIB as Next Hops on path to the destination of the current Packet;
3. All other bidirectional neighbors (except the Previous Hop of the current Packet).

Additional information from the RIB or the list of bidirectional neighbors MAY be used for determining the order, such as route cost or link quality.

If the candidate Next Hop list created as specified in this section is empty, the selected Next Hop MUST be `P_prev_hop` of the current tuple; this case applies when returning the Packet to the Previous Hop.

[12](#). Informing the Routing Protocol

When a Packet is returned (i.e., a Packet with `RET = 1` is received by a router) or a link layer acknowledgment (ACK) has not been received for a forwarded Packet, an external mechanism (not specified in this document) MAY use this information to increase the cost for the route

in the RIB, and/or to inform the routing protocol. In particular, DFF can inform a routing protocol if a Packet is received by a router that has been received before (as indicated by an existing Processed Tuple), and DUP = 0 and RET = 0, which indicates a loop in the routing topology. Care has to be taken by this external mechanism not to create loops. The rationale for such a mechanism is to update routes based on information from DFF, so that future packet transmissions will take better routes.

13. Sequence Numbers

Whenever a router generates a Packet or forwards a Packet on behalf of a host or a router outside the routing domain where DFF is used, a sequence number MUST be created and included in the DFF header. This sequence number MUST be unique locally on each router where it is created. A sequence number MUST start at 0 for the first Packet to which the DFF header is added, and then increase in steps of 1 for each new Packet. The sequence number MUST NOT be greater than 65535 and MUST wrap around to 0.

14. Modes of Operation

DFF can be used either as "route-over" IPv6 forwarding protocol, or alternatively as "mesh-under" data forwarding protocol for the LowPAN adaptation layer ([RFC4944]). Previous sections have specified the DFF mechanism in general; specific differences for each MoP are specified in this section.

14.1. Route-Over

This section maps the general terminology from [Section 2.2](#) to the specific terminology when using the "route-over" MoP.

14.1.1. Mapping of DFF Terminology to IPv6 Terminology

The following terms are those listed in [Section 2.2](#), and their meaning is explicitly defined when DFF is used in the "route-over" MoP:

Packet - An IPv6 packet, as specified in [[RFC2460](#)].

Packet Header - An IPv6 extension header, as specified in [[RFC2460](#)].

Address - An IPv6 address, as specified in [[RFC4291](#)].

Originator Address - The Originator Address corresponds to the Source address field of the IPv6 header as specified in [[RFC2460](#)].

Destination Address - The Destination Address corresponds to the Destination field of the IPv6 header as specified in [[RFC2460](#)].

Next Hop - The Next Hop is the IPv6 address of the next hop to which the Packet is sent; the link layer address from that IP address is resolved by a mechanism such as ND [[RFC4861](#)]. The link layer address is then used by L2 as destination.

Previous Hop - The Previous Hop is the IPv6 address from the interface of the previous hop from which the Packet has been received.

Hop Limit - The Hop Limit corresponds to the Hop Limit field in the IPv6 header as specified in [[RFC2460](#)].

[14.1.2](#). Packet Format

In the "route-over" MoP, all IPv6 Packets MUST conform with the format specified in [[RFC2460](#)].

The DFF header, as specified below, is an IPv6 Extension Hop-by-Hop Options header, and is depicted in Figure 1 (where DUP is abbreviated to D, and RET is abbreviated to R because of the limited space in the figure). This document specifies a new option to be used inside the Hop-by-Hop Options header, which contains the DFF fields (DUP and RET flags and sequence number, as specified in [Section 7](#)).

[RFC6564] specifies:

New options for the existing Hop-by-Hop Header SHOULD NOT be created or specified unless no alternative solution is feasible. Any proposal to create a new option for the existing Hop-by-Hop Header MUST include a detailed explanation of why the hop-by-hop behavior is absolutely essential in the document proposing the new option with hop-by-hop behavior.

[RFC6564] recommends to use Destination Headers instead of Hop-by-Hop Option headers. Destination Headers are only read by the destination of an IPv6 packet, not by intermediate routers. However, the mechanism specified in this document relies on intermediate routers reading and editing the header. Specifically, the sequence number and the DUP and RET flags are read by each router running the DFF protocol. Modifying the DUP flag and RET flag is essential for this

protocol to tag duplicate or returned Packets. Without the DUP flag, a duplicate Packet cannot be discerned from a looping Packet, and without the RET flag, a returned Packet cannot be discerned from a looping Packet.




Figure 1: IPv6 DFF Header

Field definitions of the DFF header are as follows:

Next Header - 8-bit selector. Identifies the type of header immediately following the Hop-by-Hop Options header. As specified in [\[RFC2460\]](#).

Hdr Ext Len - 8-bit unsigned integer. Length of the Hop-by-Hop Options header in 8-octet units, not including the first 8 octets. As specified in [\[RFC2460\]](#). This value is set to 0 (zero).

OptTypeDFF - 8-bit identifier of the type of option as specified in [\[RFC2460\]](#). This value is set to IP_DFF. The two high order bits of the option type MUST be set to '11' and the third bit is equal to '1'. With these bits, according to [\[RFC2460\]](#), routers that do not understand this option on a received Packet discard the packet and, only if the packet's Destination Address was not a multicast address, send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type. Also, according to [\[RFC2460\]](#), the values within the option are expected to change en route.

OptDataLenDFF - 8-bit unsigned integer. Length of the Option Data field of this option, in octets as specified in [\[RFC2460\]](#). This value is set to 2 (two).

DFF fields - A 2-bit version field (abbreviated as VER), the DUP (abbreviated as D) and RET (abbreviated as R) flags follow after Mesh Forw, as specified in [Section 7](#). The version specified in this document is 00. All other bits (besides VER, DUP, and RET) of this octet are reserved and MUST be set to 0.

Sequence Number - A 16-bit field, containing an unsigned integer sequence number, as specified in [Section 7](#).

Pad1 - Since the Hop-by-Hop Options header must have a length of multiples of 8 octets, a Pad1 option is used, as specified in [\[RFC2460\]](#). All bits of this octet are 0.

[14.2](#). Mesh-Under

This section maps the general terminology from [Section 2.2](#) to the specific terminology when using the "mesh-under" MoP.

[14.2.1](#). Mapping of DFF Terminology to LoWPAN Terminology

The following terms are those listed in [Section 2.2](#) (besides "Mode of Operation"), and their meaning is explicitly defined when DFF is used in the "mesh-under" MoP:

Packet - A "LoWPAN encapsulated packet" (as specified in [\[RFC4944\]](#), which contains an IPv6 packet as payload.

Packet Header - A LoWPAN header, as specified in [\[RFC4944\]](#).

Address - A 16-bit short or EUI-64 link layer address, as specified in [\[RFC4944\]](#).

Originator Address - The Originator Address corresponds to the Originator Address field of the Mesh Addressing header as specified in [\[RFC4944\]](#).

Destination Address - The Destination Address corresponds to the Final Destination field of the Mesh Addressing header as specified in [\[RFC4944\]](#).

Next Hop - The Next Hop is the destination address of a frame containing a LoWPAN encapsulated packet, as specified in [\[RFC4944\]](#).

Previous Hop - The Previous Hop is the source address of the frame containing a LoWPAN encapsulated packet, as specified in [\[RFC4944\]](#).

Hop Limit - The Hop Limit corresponds to the Deep Hops Left field in the Mesh Addressing header as specified in [\[RFC4944\]](#).

14.2.2. Packet Format

In the "mesh-under" MoP, all IPv6 Packets MUST conform with the format specified in [\[RFC4944\]](#). All data Packets exchanged by routers using this specification MUST contain the Mesh Addressing header as part of the LowPAN encapsulation, as specified in [\[RFC4944\]](#).

The DFF header, as specified below, **MUST** follow the Mesh Addressing header. After these two headers, any other LoWPAN header, e.g., header compression or fragmentation headers, **MAY** also be added before the actual payload. Figure 2 depicts the Mesh Addressing header defined in [RFC4944], and Figure 3 depicts the DFF header.




Figure 2: Mesh Addressing Header




Figure 3: Header for DFF data Packets

Field definitions of the Mesh Addressing header are as specified in [\[RFC4944\]](#). When adding that header to the LoWPAN encapsulation on the Originator, the fields of the Mesh Addressing header MUST be set to the following values:

- o $V := 0$ if the Originator Address is an IEEE extended 64-bit address (EUI-64); otherwise, $V := 1$ if it is a short 16-bit address.
- o $F := 0$ if the Final Destination Address is an IEEE extended 64-bit address (EUI-64); otherwise, $F := 1$ if it is a short 16-bit address.
- o Hops Left $:= 0xF$ (i.e., reserved value indicating that the Deep Hops Left field is following);

- o Deep Hops Left := MAX_HOP_LIMIT.

Field definitions of the DFF header are as follows:

Mesh Forw - A 6-bit identifier that allows for the use of different mesh forwarding mechanisms. As specified in [\[RFC4944\]](#), additional mesh forwarding mechanisms should use the reserved dispatch byte values following LOWPAN_BC0; therefore, 0 1 MUST precede Mesh Forw. The value of Mesh Forw is LOWPAN_DFF.

DFF fields - A 2-bit version field (abbreviated as VER), the DUP (abbreviated as D) and RET (abbreviated as R) flags follow after Mesh Forw, as specified in [Section 7](#). The version specified in this document is 00. All other bits (besides VER, DUP, and RET) of this octet are reserved and MUST be set to 0.

Sequence Number - A 16-bit field, containing an unsigned integer sequence number, as specified in [Section 7](#).

15. Scope Limitation of DFF

The forwarding mechanism specified in this document MUST be limited in scope to the routing domain in which DFF is used. That also implies that any headers specific to DFF do not traverse the boundaries of the routing domain. This section specifies, both for the "route-over" MoP and the "mesh-under" MoP, how to limit the scope of DFF to the routing domain in which it is used.

Figure 4 to Figure 7 depict four different cases for source and destination of traffic with regards to the scope of the routing domain in which DFF is used. [Section 15.2](#) and [Section 15.1](#) specify how routers limit the scope of DFF for the "route-over" MoP and the "mesh-under" MoP respectively for these cases. In these sections, all routers "inside the routing domain" use DFF, and all sources or destinations "outside the routing domain" are either hosts attached to a router running DFF or routers not running DFF.




Figure 4: Traffic within the routing domain (from S to D)




Figure 5: Traffic from within the routing domain to outside of the domain (from S to D)




Figure 6: Traffic from outside the routing domain to inside the domain (from S to D)




Figure 7: Traffic from outside the routing domain, traversing the domain and then to the outside of the domain (from S to D)

15.1. Route-Over MoP

In Figure 4, both the source and destination of the traffic are routers within the routing domain. If traffic is originated at S, the DFF header is added to the IPv6 header (as specified in [Section 14.1.2](#)). The Originator Address is set to S and the Destination Address is set to D. The Packet is forwarded to D using this specification. When router D receives the Packet, it processes the payload of the IPv6 Packet in upper layers. This case assumes that S has knowledge that D is in the routing domain, e.g., because of administrative setting based on the IP address of the Destination. If S has no knowledge about whether D is in the routing domain, IPv6-in-IPv6 tunnels as specified in [\[RFC2473\]](#) MUST be used. These cases

are described in the following paragraphs.

In Figure 5, the source of the traffic (S) is within the routing domain, and the destination (D) is outside of the routing domain. The IPv6 Packet, originated at S, MUST be encapsulated according to [\[RFC2473\]](#) (IPv6-in-IPv6 tunnels), and the DFF header added to the outer IPv6 header. S chooses the next router that should process the Packet as the tunnel exit-point. In some cases, the tunnel exit-point will be the final router along a path towards the Packet's Destination, and the Packet will only traverse a single tunnel (e.g., if R is a known border router then S can choose R as tunnel exit-point). In other cases, the tunnel exit-point will not be the final router along the path to D, and the Packet may traverse multiple tunnels to reach the Destination; note that in this case, the DFF mechanism is only used inside each IPv6-in-IPv6 tunnel. If no information is available which router to choose as tunnel exit-point, the Next Hop MUST be used as tunnel exit-point. The Originator Address of the Packet is set to S and the Destination Address is set to the tunnel exit-point (in the outer IPv6 header). The Packet is forwarded to the tunnel exit-point using this specification (potentially using multiple consecutive IPv6-in-IPv6 tunnels). When router R receives the Packet, it decapsulates the IPv6 Packet and forwards the inner IPv6 Packet to D, using normal IPv6 forwarding as specified in [\[RFC2460\]](#).

In Figure 6, the source of the traffic (S) is outside of the routing domain, and the destination (D) is inside of the routing domain. The IPv6 Packet, originated at S, is forwarded to R using normal IPv6 forwarding as specified in [\[RFC2460\]](#). Router R MUST encapsulate the IPv6 Packet according to [\[RFC2473\]](#), and add the DFF header (as specified in [Section 14.1.2](#)) to the outer IPv6 header. Like in the previous case, R has to select a tunnel exit-point; if it knows that D is in the routing domain (e.g., based on administrative settings), it SHOULD select D as the tunnel exit-point. In case it does not have any information which exit-point to select, it MUST use the Next Hop as tunnel exit-point, limiting the effectiveness of DFF to inside each IPv6-in-IPv6 tunnel. The Originator Address of the Packet is set to R, the Destination Address to the tunnel exit-point (both in the outer IPv6 header), the sequence number in the DFF header is generated locally on R. The Packet is forwarded to D using this specification. When router D receives the Packet, it decapsulates the inner IPv6 Packet and processes the payload of the inner IPv6 Packet in upper layers.

This mechanism is typically not used in transit networks; therefore, this case is discouraged, but described nevertheless for completeness: In Figure 7, both the source of the traffic (S) and the destination (D) are outside of the routing domain. The IPv6 Packet,

originated at S, is forwarded to R1 using normal IPv6 forwarding as specified in [\[RFC2460\]](#). Router R1 MUST encapsulate the IPv6 Packet according to [\[RFC2473\]](#) and add the DFF header (as specified in [Section 14.1.2](#)). R1 selects a tunnel exit-point like in the previous cases; if R2 is, e.g., a known border router, then R1 can select R2 as tunnel exit-point. The Originator Address is set to R1, the Destination Address to the tunnel exit-point (both in the outer IPv6 header), the sequence number in the DFF header is generated locally on R1. The Packet is forwarded to the tunnel exit-point using this specification (potentially traversing multiple consecutive IPv6-in-IPv6 tunnels). When router R2 receives the Packet, it decapsulates the inner IPv6 Packet and forwards the inner IPv6 Packet to D, using normal IPv6 forwarding as specified in [\[RFC2460\]](#).

[15.2.](#) Mesh-Under MoP

In Figure 4, both the source and destination of the traffic are routers within the routing domain. If traffic is originated at router S, the LoWPAN encapsulated Packet is created from the IPv6 packet as specified in [\[RFC4944\]](#). Then, the Mesh Addressing header and the DFF header (as specified in [Section 14.2.2](#)) are added to the LoWPAN encapsulation on router S. The Originator Address is set to S and the Destination Address is set to D. The Packet is then forwarded using this specification. When router D receives the Packet, it processes the payload of the Packet in upper layers.

In Figure 5, the source of the traffic (S) is within the routing domain, and the destination (D) is outside of the routing domain (which is known by S to be outside the routing domain because D uses a different IP prefix from the PAN). The LoWPAN encapsulated Packet, originated at router S, is created from the IPv6 packet as specified in [\[RFC4944\]](#). Then, the Mesh Addressing header and the DFF header (as specified in [Section 14.2.2](#)) are added to the LoWPAN encapsulation on router S. The Originator Address is set to S and the Destination Address is set to R, which is a known border router of the PAN. The Packet is then forwarded using this specification. When router R receives the Packet, it restores the IPv6 packet from the LoWPAN encapsulated Packet and forwards it to D, using normal IPv6 forwarding as specified in [\[RFC2460\]](#).

In Figure 6, the source of the traffic (S) is outside of the routing domain, and the destination (D) is inside of the routing domain. The IPv6 packet, originated at S, is forwarded to R using normal IPv6 forwarding as specified in [\[RFC2460\]](#). Router R (which is a known border router to the PAN) creates the LoWPAN encapsulated Packet from the IPv6 packet as specified in [\[RFC4944\]](#). Then, R adds the Mesh Addressing header and the DFF header (as specified in [Section 14.2.2](#)). The Originator Address is set to R, the Destination

Address to D, the sequence number in the DFF header is generated locally on R. The Packet is forwarded to D using this specification. When router D receives the Packet, it restores the IPv6 packet from the LoWPAN encapsulated Packet and processes the payload in upper layers.

As LoWPANs are typically no transit networks, this case is discouraged, but described nevertheless for completeness: In Figure 7, both the source of the traffic (S) and the destination (D) are outside of the routing domain. The IPv6 packet, originated at S, is forwarded to R1 using normal IPv6 forwarding as specified in [RFC2460]. Router R1 (which is a known border router of the PAN) creates the LoWPAN encapsulated Packet from the IPv6 Packet as specified in [RFC4944]. Then, it adds the Mesh Addressing header and the DFF header (as specified in [Section 14.2.2](#)). The Originator Address is set to R1, the Destination Address to R2 (which is another border router towards the Destination), the sequence number in the DFF header is generated locally on R1. The Packet is forwarded to R2 using this specification. When router R2 receives the Packet, it restores the IPv6 packet from the LoWPAN encapsulated Packet and forwards the IPv6 packet to D, using normal IPv6 forwarding as specified in [RFC2460].

16. MTU Exceedance

When adding the DFF header as specified in [Section 9.1](#) or when encapsulating the Packet as specified in [Section 15](#), the Packet size may exceed the MTU. This is described in [Section 5 of \[RFC2460\]](#). When the Packet size of a Packet to be forwarded by DFF exceeds the MTU, the following steps are executed:

1. The router MUST discard the Packet.
2. The router MAY log the event locally (depending on the storage capabilities of the router).
3. The router MUST send back an ICMP Packet Too Big to the source of the Packet reporting back the Next Hop MTU considering the additional overhead of adding the headers.

17. Security Considerations

Based on the recommendations in [RFC3552], this section describes security threats to DFF, lists which attacks are out of scope, which attacks DFF is susceptible to, and which attacks DFF protects against.

17.1. Attacks Out of Scope

As DFF is a data forwarding protocol, any security issues concerning the payload of the Packets are not considered in this section.

It is the responsibility of upper layers to use appropriate security mechanisms (IPsec, TLS, ...) according to application requirements. As DFF does not modify the contents of IP datagrams, other than the DFF header (which is a Hop-by-Hop Options extension header in the "route-over" MoP, and therefore not protected by IPsec), no special considerations for IPsec have to be addressed.

Any attack that is not specific to DFF, but that applies in general to the link layer (e.g., wireless, PLC), is out of scope. In particular, these attacks are: Eavesdropping, Packet insertion, Packet replaying, Packet deletion, and man-in-the-middle attacks. Appropriate link-layer encryption can mitigate part of these attacks and is therefore RECOMMENDED.

17.2. Protection Mechanisms of DFF

DFF itself does not provide any additional integrity, confidentiality or authentication. Therefore, the level of protection of DFF depends on the underlying link layer security as well as protection of the payload by upper layer security (e.g., IPsec).

In the following sections, whenever encrypting or digitally signing Packets is suggested for protecting DFF, it is assumed that routers are not compromised.

17.3. Attacks In Scope

This section discusses security threats to DFF, and for each describes whether (and how) DFF is affected by the threat. DFF is designed to be used in lossy and unreliable networks. Predominant examples of lossy networks are wireless networks, where routers send Packets via broadcast. The attacks listed below are easier to exploit in wireless media, but can also be observed in wired networks.

17.3.1. Denial of Service

Denial of Service attacks are possible when using DFF by either exceeding the storage on a router, or by exceeding the available bandwidth of the channel. As DFF does not contain any algorithms with high complexity, it is unlikely that the processing power of the router could be exhausted by an attack on DFF.

The storage of a router can be exhausted by increasing the size of the Processed Set, i.e., by adding new tuples, or by increasing the size of each tuple. New tuples can be added by injecting new Packets in the network, or by forwarding overheard Packets.

Another possible DoS attack is to send Packets to a non-existing Address in the network. DFF would perform a depth-first search until the Hop Limit has reached zero. It is therefore RECOMMENDED to set the Hop Limit to a value that limits the path length.

If security provided by the link layer is used, this attack can be mitigated if the malicious router does not possess valid credentials, since other routers would not forward data through the malicious router.

17.3.2. Packet Header Modification

The following attacks can be exploited by modifying the Packet Header information, unless additional security (such as link layer security) is used:

17.3.2.1. Return Flag Tampering

A malicious router may tamper the "return" flag of a DFF Packet, and send it back to the previous hop, but only if that router had been selected as next hop by the receiving router before (as specified in [Section 9.2](#)). If the malicious router had not been selected as next hop, then a returned Packet is dropped by the receiving router. If, otherwise, the malicious router had been selected as next hop by the receiving router, and the malicious router has set the return flag, the receiving router would then try alternative neighbors. This may lead to Packets never reaching their Destination, as well as unnecessary depth-first search in the network (bandwidth exhaustion / energy drain).

This attack can be mitigated by using appropriate security of the underlying link layer.

17.3.2.2. Duplicate Flag Tampering

A malicious router may modify the Duplicate Flag of a Packet that it forwards.

If it changes the flag from 0 to 1, the Packet would be detected as duplicate by other routers in the network and not as looping packet. This may have an impact on route repair mechanisms, if an external mechanism as described in [Section 12](#) is used.

If the Duplicate Flag is set from 1 to 0, and a router receives that Packet for the second time (i.e., it has already received a Packet with the same Originator Address and sequence number before), it will wrongly detect a loop. This may have an impact on route repair mechanisms, if an external mechanism as described in [Section 12](#) is used.

This attack can be mitigated by using appropriate security of the underlying link layer.

[17.3.2.3](#). Sequence Number Tampering

A malicious router may modify the sequence number of a Packet that it forwards.

In particular, if the sequence number is modified to a number of another, previously sent, Packet of the same Originator, this Packet may wrongly be perceived as looping packet. This may have an impact on route repair mechanisms, if an external mechanism as described in [Section 12](#) is used.

This attack can be mitigated by using appropriate security of the underlying link layer.

[18](#). IANA Considerations

IANA is requested to allocate a value from the Dispatch Type Field registry for LOWPAN_DFF.

IANA is requested to allocate a value from the Destination Options and Hop-by-Hop Options registry for IP_DFF. The first three bits of that value MUST be 001.

[19](#). Acknowledgements

Jari Arkko (Ericsson), Antonin Bas (Ecole Polytechnique), Thomas Clausen (Ecole Polytechnique), Yuichi Igarashi (Hitachi), Kazuya Monden (Hitachi), Geoff Mulligan (Proto6), Hiroki Satoh (Hitachi), Ganesh Venkatesh (Mobelitix), and Jiazi Yi (Ecole Polytechnique) provided useful reviews of the draft and discussions which helped to improve this document.

The authors also would like to thank Ralph Droms, Adrian Farrel, Ted Lemon, Alvaro Retana, Dan Romascanu, and Martin Stiemerling for their reviews during IETF LC and IESG evaluation.

20. References

20.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", [RFC 2473](#), December 1998.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.
- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", [RFC 6130](#), April 2011.
- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", [RFC 6564](#), April 2012.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#), September 2012.

20.2. Informative References

- [DFF_paper1] Cespedes, S., Cardenas, A., and T. Iwao, "Comparison of Data Forwarding Mechanisms for AMI Networks", 2012 IEEE Innovative Smart Grid Technologies Conference (ISGT), January 2012.
- [DFF_paper2] Iwao, T., Iwao, T., Yura, M., Nakaya, Y., Cardenas, A., Lee, S., and R. Masuoka, "Dynamic Data Forwarding in Wireless Mesh Networks", First IEEE International Conference on Smart Grid Communications (SmartGridComm), October 2010.
- [DFS_wikipedia]

"Dynamic Data Forwarding in Wireless Mesh Networks",
http://en.wikipedia.org/wiki/Depth-first_search,
March 2013.

[KCEC_press_release]

Kit Carson Electric Cooperative (KCEC), "DFF deployed by
KCEC (Press Release)", [http://www.kitcarson.com/
index.php?option=com_content&view=article&id=45&Itemid=1](http://www.kitcarson.com/index.php?option=com_content&view=article&id=45&Itemid=1),
2011.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC
Text on Security Considerations", [BCP 72](#), [RFC 3552](#),
July 2003.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
"Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#),
September 2007.

[RFC6775] Shelby, Z., Chakrabarti, S., Nordmark, E., and C. Bormann,
"Neighbor Discovery Optimization for IPv6 over Low-Power
Wireless Personal Area Networks (6LoWPANs)", [RFC 6775](#),
November 2012.

[Appendix A](#). Examples

In this section, some example network topologies are depicted, using
the DFF mechanism for data forwarding. In these examples, it is
assumed that a routing protocol is running which adds or inserts
entries into the RIB.

[A.1](#). Example 1: Normal Delivery

Figure 8 depicts a network topology with seven routers A to G, with
links between them as indicated by lines. It is assumed that router
A sends a Packet to G, through B and D, according to the routing
protocol.




Figure 8: Example 1: normal delivery

If no link fails in this topology, and no loop occurs, then DFF forward the Packet along the Next Hops listed in each of the routers RIB along the path towards the destination. Each router adds a Processed Tuple for the incoming Packet, and selects the Next Hop as specified in [Section 11](#), i.e., it will first select the next hop for router G as determined by the routing protocol.

A.2. Example 2: Forwarding with Link Failure

Figure 9 depicts the same topology as the Example 1, but both links between B and D and between B and E are unavailable (e.g., because of wireless link characteristics).




Figure 9: Example 2: link failure

When B receives the Packet from router A, it adds a Processed Tuple, and then tries to forward the Packet to D. Once B detects that the Packet cannot be successfully delivered to D because it does not receive link layer ACKs, it will follow the procedures listed in [Section 10](#), by setting the DUP flag to 1, selecting E as new next hop, adding E to the list of next hops in the Processed Tuple, and then forwarding the Packet to E.

As the link to E also fails, B will again follow the procedure in [Section 10](#). As all possible next hops (D and E) are listed in the Processed Tuple, B will set the RET flag in the Packet and return it to A.

A determines that it already has a Processed Tuple for the returned Packet, reset the RET flag of the Packet and select a new next hop for the Packet. As B is already in the list of next hops in the Processed Tuple, it will select C as next hop and forward the Packet to it. C will then forward the Packet to F, and F delivers the Packet to its Destination G.

A.3. Example 3: Forwarding with Missed Link Layer Acknowledgment

Figure 10 depicts the same topology as the Example 1, but the link layer acknowledgments from C to A are lost (e.g., because the link is uni-directional). It is assumed that A prefers a path to G through C and F.




Figure 10: Example 3: missed link layer acknowledgment

While C successfully receives the Packet from A, A does not receive the L2 ACK and assumes the Packet has not been delivered to C. Therefore, it sets the DUP flag of the Packet to 1, in order to

indicate that this Packet may be a duplicate. Then, it forwards the Packet to B.

A.4. Example 4: Forwarding with a Loop

Figure 11 depicts the same topology as the Example 1, but there is a loop from D to A, and A sends the Packet to G through B and D.




Figure 11: Example 4: loop

When A receives the Packet through the loop from D, it will find a Processed Tuple for the Packet. Router A will set the RET flag and return the Packet to D, which in turn will return it to B. B will then select E as next hop, which will then forward it to G.

Appendix B. Deployment Experience

DFF has been deployed and experimented with both in real deployments and in network simulations, as described in the following.

B.1. Deployments in Japan

The majority of the large Advanced Metering Infrastructure (AMI) deployments using DFF are located in Japan, but the data of these networks is property of Japanese utilities and cannot be disclosed.

B.2. Kit Carson Electric Cooperative

DFF has been deployed at Kit Carson Electric Cooperative (KCEC), a non-profit organization distributing electricity to about 30,000 customers in New Mexico. As described in a press release [[KCEC press release](#)], DFF is running on currently about 2000 electric meters. All meters are connected through a mesh network using an unreliable, wireless medium. DFF is used together with a distance vector routing protocol. Metering data from each meter is sent towards a gateway periodically every 15 minutes. The data delivery reliability is over 99%.

B.3. Simulations

DFF has been evaluated in Ns2 and OMNEST simulations, in conjunction with a distance vector routing protocol. The performance of DFF has been compared to using only the routing protocol without DFF. The results published in peer-reviewed academic papers ([[DFF_paper1](#)][[DFF_paper2](#)]) show significant improvements of the Packet delivery ratio compared to using only the distance vector protocol.

B.4. Open Source Implementation

Fujitsu Laboratories of America is currently working on an open source implementation of DFF, which is to be released in early 2013, and which allows for interoperability testings of different DFF implementations. The implementation is written in Java, and can be used both on real machines and in the Ns2 simulator.

Authors' Addresses

Ulrich Herberg (editor)
Fujitsu
1240 E. Arques Avenue, M/S 345
Sunnyvale, CA 94085
US

Phone: +1 408 530-4528
Email: ulrich.herberg@us.fujitsu.com

Alvaro A. Cardenas
University of Texas at Dallas
School of Computer Science, 800 West Campbell Rd, EC 31
Richardson, TX 75080-3021
US

Email: alvaro.cardenas@me.com

Tadashige Iwao
Fujitsu
Shiodome City Center, 5-2, Higashi-shimbashi 1-chome, Minato-ku
Tokyo,
JP

Phone: +81-44-754-3343
Email: smartnetpro-iwao_std@ml.css.fujitsu.com

Michael L. Dow
Freescale
6501 William Cannon Drive West
Austin, TX 78735
USA

Phone: +1 512 895 4944
Email: m.dow@freescale.com

Sandra L. Cespedes
U. Icesi
Calle 18 No. 122-135 Pance
Cali, Valle
Colombia

Phone:
Email: scespedes@icesi.edu.co

