

Simplified Multicast Forwarding

Abstract

This document describes a Simplified Multicast Forwarding (SMF) mechanism that provides basic Internet Protocol (IP) multicast forwarding suitable for limited wireless mesh and mobile ad hoc network (MANET) use. It is mainly applicable in situations where efficient flooding represents an acceptable engineering design trade-off. It defines techniques for multicast duplicate packet detection (DPD), to be applied in the forwarding process, for both IPv4 and IPv6 protocol use. This document also specifies optional mechanisms for using reduced relay sets to achieve more efficient multicast data distribution within a mesh topology as compared to Classic Flooding. Interactions with other protocols, such as use of information provided by concurrently running unicast routing protocols or interaction with other multicast protocols, as well as multiple deployment approaches are also described. Distributed algorithms for selecting reduced relay sets and related discussion are provided in the appendices. Basic issues relating to the operation of multicast MANET border routers are discussed, but ongoing work remains in this area and is beyond the scope of this document.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6621>.

Copyright Notice

Copyright (c) 2012 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction and Scope	4
2. Terminology	4
3. Applicability Statement	5
4. Overview and Functioning	6
5. SMF Packet Processing and Forwarding	8
6. SMF Duplicate Packet Detection	10
6.1. IPv6 Duplicate Packet Detection	11
6.1.1. IPv6 SMF_DPD Option Header	12
6.1.2. IPv6 Identification-Based DPD	14
6.1.3. IPv6 Hash-Based DPD	16
6.2. IPv4 Duplicate Packet Detection	17
6.2.1. IPv4 Identification-Based DPD	18
6.2.2. IPv4 Hash-Based DPD	19
7. Relay Set Selection	20
7.1. Non-Reduced Relay Set Forwarding	20
7.2. Reduced Relay Set Forwarding	20
8. SMF Neighborhood Discovery Requirements	23
8.1. SMF Relay Algorithm TLV Types	24
8.1.1. SMF Message TLV Type	24

8.1.2. SMF Address Block TLV Type	25
9. SMF Border Gateway Considerations	26
9.1. Forwarded Multicast Groups	27
9.2. Multicast Group Scoping	28
9.3. Interface with Exterior Multicast Routing Protocols	29
9.4. Multiple Border Routers	29
10. Security Considerations	31
11. IANA Considerations	32
11.1. IPv6 SMF_DPD Header Extension Option Type	33
11.2. TaggerId Types (TidTy)	33
11.3. Well-Known Multicast Address	34
11.4. SMF TLVs	34
11.4.1. Expert Review for Created Type Extension Registries	34
11.4.2. SMF Message TLV Type (SMF_TYPE)	34
11.4.3. SMF Address Block TLV Type (SMF_NBR_TYPE)	35
11.4.4. SMF Relay Algorithm ID Registry	35
12. Acknowledgments	36
13. References	37
13.1. Normative References	37
13.2. Informative References	38
Appendix A. Essential Connecting Dominating Set (E-CDS) Algorithm	40
A.1. E-CDS Relay Set Selection Overview	40
A.2. E-CDS Forwarding Rules	41
A.3. E-CDS Neighborhood Discovery Requirements	41
A.4. E-CDS Selection Algorithm	44
Appendix B. Source-Based Multipoint Relay (S-MPR) Algorithm	46
B.1. S-MPR Relay Set Selection Overview	46
B.2. S-MPR Forwarding Rules	47
B.3. S-MPR Neighborhood Discovery Requirements	48
B.4. S-MPR Selection Algorithm	50
Appendix C. Multipoint Relay Connected Dominating Set (MPR-CDS) Algorithm	52
C.1. MPR-CDS Relay Set Selection Overview	52
C.2. MPR-CDS Forwarding Rules	53
C.3. MPR-CDS Neighborhood Discovery Requirements	53
C.4. MPR-CDS Selection Algorithm	54

1. Introduction and Scope

Unicast routing protocols, designed for MANET and wireless mesh use, often apply distributed algorithms to flood routing control plane messages within a MANET routing domain. For example, algorithms specified within [RFC3626] and [RFC3684] provide distributed methods of dynamically electing reduced relay sets that attempt to efficiently flood routing control messages while maintaining a connected set under dynamic topological conditions.

Simplified Multicast Forwarding (SMF) extends the efficient flooding concept to the data forwarding plane, providing an appropriate multicast forwarding capability for use cases where localized, efficient flooding is considered an effective design approach. The baseline design is intended to provide a basic, best-effort multicast forwarding capability that is constrained to operate within a single MANET routing domain.

An SMF routing domain is an instance of an SMF routing protocol with common policies, under a single network administration authority. The main design goals of this document are to:

- o adapt efficient relay sets in MANET environments [RFC2501], and
- o define the needed IPv4 and IPv6 multicast duplicate packet detection (DPD) mechanisms to support multi-hop packet forwarding.

2. 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].

The terms introduced in [RFC5444], including "packet", "message", "TLV Block", "TLV", and "address", are to be interpreted as described therein.

The following abbreviations are used throughout this document:

Abbreviation	Definition
MANET	Mobile Ad Hoc Network
SMF	Simplified Multicast Forwarding
CF	Classic Flooding
CDS	Connected Dominating Set
MPR	Multipoint Relay
S-MPR	Source-based MPR
MPR-CDS	MPR-based CDS
E-CDS	Essential CDS
NHDP	Neighborhood Discovery Protocol
DPD	Duplicate Packet Detection
I-DPD	Identification-based DPD
H-DPD	Hash-based DPD
HAV	Hash assist value
FIB	Forwarding Information Base
TLV	type-length-value encoding
DoS	Denial of Service
SMF Router	A MANET Router implementing the protocol specified in this document
SMF Routing Domain	A MANET Routing Domain wherein the protocol specified in this document is operating

3. Applicability Statement

Within dynamic wireless routing topologies, maintaining traditional forwarding trees to support a multicast routing protocol is often not as effective as in wired networks due to the reduced reliability and increased dynamics of mesh topologies [MGL04][GM99]. A basic packet forwarding service reaching all connected routers running the SMF protocol within a MANET routing domain may provide a useful group communication paradigm for various classes of applications, for example, multimedia streaming, interactive group-based messaging and applications, peer-to-peer middleware multicasting, and multi-hop mobile discovery or registration services. SMF is likely only appropriate for deployment in limited dynamic MANET routing domains (further defined as administratively scoped multicast forwarding domains in Section 9.2) so that the flooding process can be contained.

A design goal is that hosts may also participate in multicast traffic transmission and reception with standard IP network-layer semantics (e.g., special or unnecessary encapsulation of IP packets should be avoided in this case). SMF deployments are able to connect and

interoperate with existing standard multicast protocols operating within more conventional Internet infrastructures. To this end, a multicast border router or proxy mechanism MUST be used when deployed alongside more fixed-infrastructure IP multicast routing such Protocol Independent Multicast (PIM) variants [RFC3973][RFC4601]. Experimental SMF implementations and deployments have demonstrated gateway functionality at MANET border routers operating with existing external IP multicast routing protocols [CDHM07][DHS08][DHG09]. SMF may be extended or combined with other mechanisms to provide increased reliability and group-specific filtering; the details for this are out of the scope of this document.

This document does not presently support forwarding of packets with directed broadcast addresses as a destination [RFC2644].

4. Overview and Functioning

Figure 1 provides an overview of the logical SMF router architecture, consisting of "Neighborhood Discovery", "Relay Set Selection", and "Forwarding Process" components. Typically, relay set selection (or self-election) occurs based on dynamic input from a neighborhood discovery process. SMF supports the case where neighborhood discovery and/or relay set selection information is obtained from a coexistent process (e.g., a lower-layer mechanism or a unicast routing protocol using relay sets). In some algorithm designs, the forwarding decision for a packet can also depend on previous hop or incoming interface information. The asterisks (*) in Figure 1 mark the primitives and relationships needed by relay set algorithms requiring previous-hop packet-forwarding knowledge.

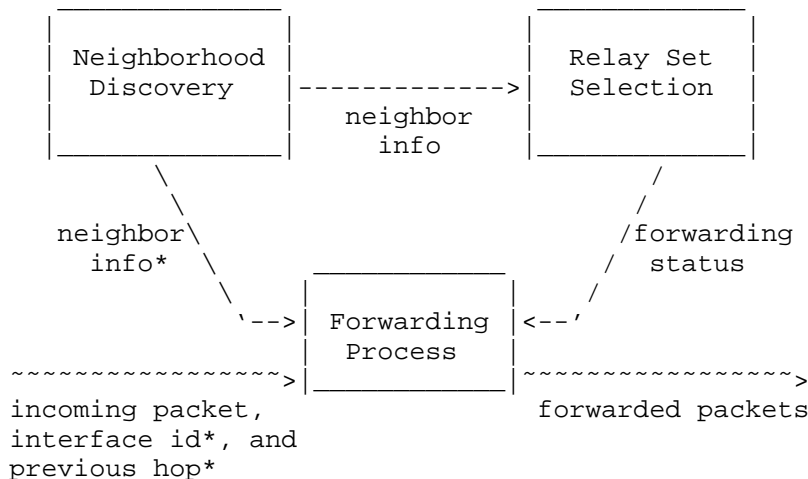


Figure 1: SMF Router Architecture

Certain IP multicast packets, defined in Sections 9.2 and 5, are "non-forwardable". These multicast packets MUST be ignored by the SMF forwarding engine. The SMF forwarding engine MAY also work with policies and management interfaces to allow additional filtering control over which multicast packets are considered for potential SMF forwarding. This interface would allow more refined dynamic forwarding control once such techniques are matured for MANET operation. At present, further discussion of dynamic control is left to future work.

Interoperable SMF implementations MUST use compatible DPD approaches and be able to process the header options defined in this document for IPv6 operation.

Classic Flooding (CF) is defined as the simplest case of SMF multicast forwarding. With CF, all SMF routers forward each received multicast packet exactly once. In this case, the need for any relay set selection or neighborhood topology information is eliminated, at the expense of additional network overhead incurred from unnecessary packet retransmissions. With CF, the SMF DPD functionality is still required. While SMF supports CF as a mode of operation, the use of more efficient relay set modes is RECOMMENDED in order to reduce contention and congestion caused by unnecessary packet retransmissions [NTSC99].

An efficient reduced relay set is constructed by selecting and updating, as needed, a subset of all possible routers in a MANET routing domain to act as SMF forwarders. Known distributed relay set selection algorithms have demonstrated the ability to provide and maintain a dynamic connected set for forwarding multicast IP packets [MDC04]. A few such relay set selection algorithms are described in the appendices of this document, and the basic designs borrow directly from previously documented IETF work. SMF relay set configuration is extensible, and additional relay set algorithms beyond those specified here can be accommodated in future work.

Determining and maintaining an optimized set of relays generally requires dynamic neighborhood topology information. Neighborhood topology discovery functions MAY be provided by a MANET unicast routing protocol or by using the MANET Neighborhood Discovery Protocol (NHDP) [RFC6130], operating concurrently with SMF. This specification also allows alternative lower-layer interfaces (e.g., radio router interfaces) to provide the necessary neighborhood information to aid in supporting more effective relay set selection. An SMF implementation SHOULD provide the ability for multicast forwarding state to be dynamically managed per operating network interface. The relay state maintenance options and interactions are outlined in Section 7. This document states specific requirements

for neighborhood discovery with respect to the forwarding process and the relay set selection algorithms described herein. For determining dynamic relay sets in the absence of other control protocols, SMF relies on NHDP to assist in IP-layer 2-hop neighborhood discovery and maintenance for relay set selection. "SMF_TYPE" and "SMF_NBR_TYPE" Message and Address Block TLV structures (per [RFC5444]) are defined by this document for use with NHDP to carry SMF-specific information. It is RECOMMENDED that all routers performing SMF operation in conjunction with NHDP include these TLV types in any NHDP HELLO messages generated. This capability allows for routers participating in SMF to be explicitly identified along with their respective dynamic relay set algorithm configuration.

5. SMF Packet Processing and Forwarding

The SMF packet processing and forwarding actions are conducted with the following packet handling activities:

1. Processing of outbound, locally generated multicast packets.
2. Reception and processing of inbound packets on specific network interfaces.

The purpose of intercepting outbound, locally generated multicast packets is to apply any added packet marking needed to satisfy the DPD requirements so that proper forwarding may be conducted. Note that for some system configurations, the interception of outbound packets for this purpose is not necessary.

Inbound multicast packets are received by the SMF implementation and processed for possible forwarding. SMF implementations MUST be capable of forwarding IP multicast packets with destination addresses that are not router-local and link-local for IPv6, as defined in [RFC4291], and that are not within the local network control block as defined by [RFC5771].

This will support generic multi-hop multicast application needs or distribute designated multicast traffic ingressing the SMF routing domain via border routers. The multicast addresses to be forwarded should be maintained by an a priori list or a dynamic forwarding information base (FIB) that MAY interact with future MANET dynamic group membership extensions or management functions.

The SL-MANET-ROUTERS multicast group is defined to contain all routers within an SMF routing domain, so that packets transmitted to the multicast address associated with the group will be attempted to be delivered to all connected routers running SMF. Due to the mobile nature of a MANET, routers running SMF may not be topologically

connected at particular times. For IPv6, SL-MANET-ROUTERS is specified to be "site-local". Minimally, SMF MUST forward, as instructed by the relay set selection algorithm, unique (non-duplicate) packets received for SL-MANET-ROUTERS when the Time to Live (TTL) / hop limit or hop limit value in the IP header is greater than 1. SMF MUST forward all additional global-scope multicast addresses specified within the dynamic FIB or configured list as well. In all cases, the following rules MUST be observed for SMF multicast forwarding:

1. Any IP packets not addressed to an IP multicast address MUST NOT be forwarded by the SMF forwarding engine.
2. IP multicast packets with TTL/hop limit ≤ 1 MUST NOT be forwarded.
3. Link local IP multicast packets MUST NOT be forwarded.
4. Incoming IP multicast packets with an IP source address matching one of those of the local SMF router interface(s) MUST NOT be forwarded.
5. Received frames with the Media Access Control (MAC) source address matching any MAC address of the router's interfaces MUST NOT be forwarded.
6. Received packets for which SMF cannot reasonably ensure temporal DPD uniqueness MUST NOT be forwarded.
7. Prior to being forwarded, the TTL/hop limit of the forwarded packet MUST be decremented by one.

Note that rule #4 is important because over some types of wireless interfaces, the originating SMF router may receive retransmissions of its own packets when they are forwarded by adjacent routers. This rule avoids unnecessary retransmission of locally generated packets even when other forwarding decision rules would apply.

An additional processing rule also needs to be considered based upon a potential security threat. As discussed in Section 10, there is a potential DoS attack that can be conducted by remotely "previewing" (e.g., via a directional receive antenna) packets that an SMF router would be forwarding and conducting a "pre-play" attack by transmitting the packet before the SMF router would otherwise receive it, but with a reduced TTL/hop limit field value. This form of attack can cause an SMF router to create a DPD entry that would block the proper forwarding of the valid packet (with correct TTL/hop limit) through the SMF routing domain. A RECOMMENDED approach to

prevent this attack, when it is a concern, would be to cache temporal packet TTL/hop limit values along with the per-packet DPD state (hash value(s) and/or identifier as described in Section 6). Then, if a subsequent matching (with respect to DPD) packet arrives with a larger TTL/hop limit value than the packet that was previously forwarded, SMF should forward the new packet and update the TTL/hop limit value cached with corresponding DPD state to the new, larger TTL/hop limit value. There may be temporal cases where SMF would unnecessarily forward some duplicate packets using this approach, but those cases are expected to be minimal and acceptable when compared with the potential threat of denied service.

Once the SMF multicast forwarding rules have been applied, an SMF implementation **MUST** make a forwarding decision dependent upon the relay set selection algorithm in use. If the SMF implementation is using Classic Flooding (CF), the forwarding decision is implicit once DPD uniqueness is determined. Otherwise, a forwarding decision depends upon the current interface-specific relay set state. The descriptions of the relay set selection algorithms in the appendices to this document specify the respective heuristics for multicast packet forwarding and specific DPD or other processing required to achieve correct SMF behavior in each case. For example, one class of forwarding is based upon relay set selection status and the packet's previous hop, while other classes designate the local SMF router as a forwarder for all neighboring routers.

6. SMF Duplicate Packet Detection

Duplicate packet detection (DPD) is often a requirement in MANET or wireless mesh packet forwarding mechanisms because packets may be transmitted out via the same physical interface as the one over which they were received. Routers may also receive multiple copies of the same packets from different neighbors or interfaces. SMF operation requires DPD, and implementations **MUST** provide mechanisms to detect and reduce the likelihood of forwarding duplicate multicast packets using temporal packet identification. It is **RECOMMENDED** this be implemented by keeping a history of recently processed multicast packets for comparison with incoming packets. A DPD packet cache history **SHOULD** be kept long enough so as to span the maximum network traversal lifetime, `MAX_PACKET_LIFETIME`, of multicast packets being forwarded within an SMF routing domain. The DPD mechanism **SHOULD** avoid keeping unnecessary state for packet flows such as those that are locally generated or link-local destinations that would not be considered for forwarding, as presented in Section 5.

For both IPv4 and IPv6, this document describes two basic multicast duplicate packet detection mechanisms: header content identification-based (I-DPD) and hash-based (H-DPD) duplicate packet detection.

I-DPD is a mechanism using specific packet headers, and option headers in the case of IPv6, in combination with flow state to estimate the temporal uniqueness of a packet. H-DPD uses hashing over header fields and payload of a multicast packet to provide an estimation of temporal uniqueness.

Trade-offs of the two approaches to DPD merit different considerations dependent upon the specific SMF deployment scenario.

Because of the potential addition of a hop-by-hop option header with IPv6, all SMF routers in the same SMF deployments MUST be configured so as to use a common mechanism and DPD algorithm. The main difference between IPv4 and IPv6 SMF DPD specifications is the avoidance of any additional header options for IPv4.

For each network interface, SMF implementations MUST maintain DPD packet state as needed to support the forwarding heuristics of the relay set algorithm used. In general, this involves keeping track of previously forwarded packets so that duplicates are not forwarded, but some relay techniques have additional considerations, such as those discussed in Appendix B.2.

Additional details of I-DPD and H-DPD processing and maintenance for different classes of packets are described in the following subsections.

6.1. IPv6 Duplicate Packet Detection

This section describes the mechanisms and options for SMF IPv6 DPD. The base IPv6 packet header does not provide an explicit packet identification header field that can be exploited for I-DPD. The following options are therefore described to support IPv6 DPD:

1. a hop-by-hop SMF_DPD option header, defined in this document (Section 6.1.1),
2. the use of IPv6 fragment header fields for I-DPD, if one is present (Section 6.1.2),
3. the use of IPsec sequencing for I-DPD when a non-fragmented, IPsec header is detected (Section 6.1.2), and
4. an H-DPD approach assisted, as needed, by the SMF_DPD option header (Section 6.1.3).

SMF MUST provide a DPD marking module that can insert the hop-by-hop IPv6 header option, defined in Section 6.1.1. This module MUST be invoked after any source-based fragmentation that may occur with

IPv6, so as to ensure that all fragments are suitably marked. SMF IPv6 DPD is presently specified to allow either a packet hash or header identification method for DPD. An SMF implementation MUST be configured to operate either in I-DPD or H-DPD mode and perform the corresponding tasks, outlined in Sections 6.1.2 and 6.1.3.

6.1.1. IPv6 SMF_DPD Option Header

This section defines an IPv6 Hop-by-Hop Option [RFC2460], SMF_DPD, to serve the purpose of unique packet identification for IPv6 I-DPD. Additionally, the SMF_DPD option header provides a mechanism to guarantee non-collision of hash values for different packets when H-DPD is used.

If this is the only hop-by-hop option present, the optional TaggerId field (see below) is not included, and the size of the DPD packet identifier (sequence number) or hash token is 24 bits or less, this will result in the addition of 8 bytes to the IPv6 packet header including the "Next Header", "Header Extension Length", SMF_DPD option fields, and padding.




Figure 2: IPv6 SMF_DPD Hop-by-Hop Option Header

"Option Type" = 00001000. The highest order three bits are 000 because this specification requires that routers not recognizing this option type skip over this option and continue processing the header and that the option must not change en route [RFC2460].

"Opt. Data Len" = Length of option content (i.e., 1 + (<IdType> ? (<IdLen> + 1): 0) + Length(DPD ID)).

"H-bit" = a hash indicator bit value identifying DPD marking type. 0 == sequence-based approach with optional TaggerId and a tuple-based sequence number. 1 == indicates a hash assist value (HAV) field follows to aid in avoiding hash-based DPD collisions.

When the "H-bit" is cleared (zero value), the SMF_DPD format to support I-DPD operation is specified as shown in Figure 3 and defines the extension header in accordance with [RFC2460].




Figure 3: IPv6 SMF_DPD Option Header in I-DPD mode

"TidTy" = a 3-bit field indicating the presence and type of the optional TaggerId field.

"TidLen" = a 4-bit field indicating the length (in octets) of the following TaggerId field.

"TaggerId" = a field, is used to differentiate multiple ingressing border gateways that may commonly apply the SMF_DPD option header to packets from a particular source. Table 1 lists the TaggerId types used in this document:

Name	Purpose
NULL	Indicates no TaggerId field is present. "TidLen" MUST also be set to ZERO.
DEFAULT	A TaggerId of non-specific context is present. "TidLen + 1" defines the length of the TaggerId field in bytes.
IPv4	A TaggerId representing an IPv4 address is present. The "TidLen" MUST be set to 3.
IPv6	A TaggerId representing an IPv6 address is present. The "TidLen" MUST be set to 15.

Table 1: TaggerId Types

This format allows a quick check of the "TidTy" field to determine if a TaggerId field is present. If "TidTy" is NULL, then the length of the DPD packet <Identifier> field corresponds to (<Opt. Data Len> - 1). If the <TidTy> is non-NULL, then the length of the TaggerId field is equal to (<TidLen> - 1), and the remainder of the option data comprises the DPD packet <Identifier> field. When the TaggerId field is present, the <Identifier> field can be considered a unique packet identifier in the context of the <TaggerId:srcAddr:dstAddr> tuple. When the TaggerId field is not present, then it is assumed that the source applied the SMF_DPD option and the <Identifier> can

be considered unique in the context of the IPv6 packet header <srcAddr:dstAddr> tuple. IPv6 I-DPD operation details are in Section 6.1.2.

When the "H-bit" in the SMF_DPD option data is set, the data content value is interpreted as a hash assist value (HAV) used to facilitate H-DPD operation. In this case, the source or ingressing gateways apply the SMF_DPD with an HAV only when required to differentiate the hash value of a new packet with respect to hash values in the DPD cache. This situation can be detected locally on the router by running the hash algorithm and checking the DPD cache, prior to ingressing a previously unmarked packet or a locally sourced packet. This helps to guarantee the uniqueness of generated hash values when H-DPD is used. Additionally, this avoids the added overhead of applying the SMF_DPD option header to every packet. For many hash algorithms, it is expected that only sparse use of the SMF_DPD option may be required. The format of the SMF_DPD option header for H-DPD operation is given in Figure 4.




Figure 4: IPv6 SMF_DPD Option Header in H-DPD Mode

The SMF_DPD option should be applied with an HAV to produce a unique hash digest for packets within the context of the IPv6 packet header <srcAddr>. The size of the HAV field is implied by "Opt. Data Len". The appropriate size of the field depends upon the collision properties of the specific hash algorithm used. More details on IPv6 H-DPD operation are provided in Section 6.1.3.

6.1.2. IPv6 Identification-Based DPD

Table 2 summarizes the IPv6 I-DPD processing and forwarding decision approach. Within the table, '*' indicates an ignore field condition.

IPv6 Fragment Header	IPv6 IPsec Header	IPv6 I-DPD Header	SMF IPv6 I-DPD Mode Action
Present	*	Not Present	Use Fragment Header I-DPD Check and Process for Forwarding
Not Present	Present	Not Present	Use IPsec Header I-DPD Check and Process for Forwarding
Present	*	Present	Invalid; do not forward.
Not Present	Present	Present	Invalid; do not forward.
Not Present	Not Present	Not Present	Add I-DPD Header, and Process for Forwarding
Not Present	Present	Present	Use I-DPD Header Check and Process for Forwarding
Not Present	Not Present	Present	Use I-DPD Header Check and Process for Forwarding
Not Present	Present	Present	Use I-DPD Header Check and Process for Forwarding

Table 2: IPv6 I-DPD Processing Rules

1. If a received IPv6 multicast packet is an IPv6 fragment, SMF MUST use the fragment extension header fields for packet identification. This identifier can be considered unique in the context of the <srcAddr:dstAddr> of the IP packet.
2. If the packet is an unfragmented IPv6 IPsec packet, SMF MUST use IPsec fields for packet identification. The IPsec header <sequence> field can be considered a unique identifier in the context of the <IPsecType:srcAddr:dstAddr:SPI> where "IPsecType" is either Authentication Header (AH) or Encapsulating Security Payload (ESP) [RFC4302].
3. For unfragmented, non-IPsec IPv6 packets, the use of the SMF_DPD option header is necessary to support I-DPD operation. The SMF_DPD option header is applied in the context of the <srcAddr> of the IP packet. Hosts or ingressing SMF gateways are responsible for applying this option to support DPD. Table 3 summarizes these packet identification types:

IPv6 Packet Type	Packet DPD ID Context	Packet DPD ID
Fragment IPsec Packet Regular Packet	<srcAddr:dstAddr> <IPsecType:srcAddr:dstAddr:SPI> <[TaggerId:]srcAddr:dstAddr>	<fragmentOffset:id> <sequence> <SMF_DPD option header id>

Table 3: IPv6 I-DPD Packet Identification Types

"IPsecType" is either Authentication Header (AH) or Encapsulating Security Payload (ESP).

The "TaggerId" is an optional field of the IPv6 SMF_DPD option header.

6.1.3. IPv6 Hash-Based DPD

A default hash-based DPD approach (H-DPD) for use by SMF is specified as follows. An SHA-1 [RFC3174] hash of the non-mutable header fields, options fields, and data content of the IPv6 multicast packet is used to produce a 160-bit digest. The approach for calculating this hash value SHOULD follow the same guidelines described for calculating the Integrity Check Value (ICV) described in [RFC4302] with respect to non-mutable fields. This approach should have a reasonably low probability of digest collision when packet headers and content are varying. SHA-1 is being applied in SMF only to provide a low probability of collision and is not being used for cryptographic or authentication purposes. A history of the packet hash values SHOULD be maintained within the context of the IPv6 packet header <srcAddr>. SMF ingress points (i.e., source hosts or gateways) use this history to confirm that new packets are unique with respect to their hash value. The hash assist value (HAV) field described in Section 6.1.1 is provided as a differentiating field when a digest collision would otherwise occur. Note that the HAV is an immutable option field, and SMF MUST process any included HAV values (see Section 6.1.1) in its hash calculation.

If a packet results in a digest collision (i.e., by checking the H-DPD digest history) within the DPD cache kept by SMF forwarders, the packet SHOULD be silently dropped. If a digest collision is detected at an SMF ingress point, the H-DPD option header is constructed with a randomly generated HAV. An HAV is recalculated as needed to produce a non-colliding hash value prior to forwarding.

The multicast packet is then forwarded with the added IPv6 SMF_DPD option header. A common hash approach **MUST** be used by SMF routers for the applied HAV to consistently avoid hash collision and thus inadvertent packet drops.

The SHA-1 indexing and IPv6 HAV approaches are specified at present for consistency and robustness to suit experimental uses. Future approaches and experimentation may discover design trade-offs in hash robustness and efficiency worth considering. Enhancements **MAY** include reducing the maximum payload length that is processed, determining shorter indexes, or applying more efficient hashing algorithms. Use of the HAV functionality may allow for application of "lighter-weight" hashing techniques that might not have been initially considered otherwise due to poor collision properties. Such techniques could reduce packet-processing overhead and memory requirements.

6.2. IPv4 Duplicate Packet Detection

This section describes the mechanisms and options for IPv4 DPD. The following areas are described to support IPv4 DPD:

1. the use of IPv4 fragment header fields for I-DPD when they exist (Section 6.2.1),
2. the use of IPsec sequencing for I-DPD when a non-fragmented IPv4 IPsec packet is detected (Section 6.2.1), and
3. an H-DPD approach (Section 6.2.2) when neither of the above cases can be applied.

Although the IPv4 datagram has a 16-bit Identification (ID) field as specified in [RFC0791], it cannot be relied upon for DPD purposes due to common computer operating system implementation practices and the reasons described in the updated specification of the IPv4 ID Field [IPV4-ID-UPDATE]. An SMF IPv4 DPD marking option like the IPv6 SMF_DPD option header is not specified since IPv4 header options are not as tractable for hosts as they are for IPv6. However, when IPsec is applied or IPv4 packets have been fragmented, the I-DPD approach can be applied reliably using the corresponding packet identifier fields described in Section 6.2.1. For the general IPv4 case (non-IPsec and non-fragmented packets), the H-DPD approach of Section 6.2.2 is **RECOMMENDED**.

Since IPv4 SMF does not specify an option header, the interoperability constraints are looser than in the IPv6 version, and forwarders may operate with mixed H-DPD and I-DPD modes as long as they consistently perform the appropriate DPD routines outlined in the following sections. However, it is RECOMMENDED that a deployment be configured with a common mode for operational consistency.

6.2.1. IPv4 Identification-Based DPD

Table 4 summarizes the IPv4 I-DPD processing approach once a packet has passed the basic forwardable criteria described in Section 5. To summarize, for IPv4, I-DPD is applicable only for packets that have been fragmented or have IPsec applied. In Table 4, '*' indicates an ignore field condition. DF, MF, and Fragment offset correspond to related fields and flags defined in [RFC0791].

DF flag	MF flag	Fragment offset	IPsec	IPv4 I-DPD Action
1	1	*	*	Invalid; do not forward.
1	0	nonzero	*	Invalid; do not forward.
*	0	zero	not Present	Use H-DPD check instead
*	0	zero	Present	IPsec enhanced Tuple I-DPD Check and Process for Forwarding
0	0	nonzero	*	Extended Fragment Offset Tuple I-DPD Check and Process for Forwarding
0	1	zero or nonzero	*	Extended Fragment Offset Tuple I-DPD Check and Process for Forwarding

Table 4: IPv4 I-DPD Processing Rules

For performance reasons, IPv4 network fragmentation and reassembly of multicast packets within wireless MANET networks should be minimized, yet SMF provides the forwarding of fragments when they occur. If the IPv4 multicast packet is a fragment, SMF MUST use the fragmentation header fields for packet identification. This identification can be considered temporally unique in the context of the <protocol:srcAddr:dstAddr> of the IPv4 packet. If the packet is an unfragmented IPv4 IPsec packet, SMF MUST use IPsec fields for packet identification. The IPsec header <sequence> field can be considered a unique

identifier in the context of the <IPsecType:srcAddr:dstAddr:SPI> where "IPsecType" is either AH or ESP [RFC4302]. Table 5 summarizes these packet identification types:

IPv4 Packet Type	Packet Identification Context	Packet Identifier
Fragment IPsec Packet	<protocol:srcAddr:dstAddr> <IPsecType:srcAddr:dstAddr:SPI>	<fragmentOffset:id> <sequence>

Table 5: IPv4 I-DPD Packet Identification Types

"IPsecType" is either Authentication Header (AH) or Encapsulating Security Payload (ESP).

6.2.2. IPv4 Hash-Based DPD

The hashing technique here is similar to that specified for IPv6 in Section 6.1.3, but the H-DPD header option with HAV is not considered. To ensure consistent IPv4 H-DPD operation among SMF routers, a default hashing approach is specified. A common DPD hashing algorithm for an SMF routing area is RECOMMENDED because colliding hash values for different packets result in "false positive" duplicate packet detection, and there is small probability that valid packets may be dropped based on the hashing technique used. Since the "hash assist value" approach is not available for IPv4, use of a common hashing approach minimizes the probability of hash collision packet drops over multiple hops of forwarding.

SMF MUST perform a SHA-1 [RFC3174] hash of the immutable header fields, option fields, and data content of the IPv4 multicast packet resulting in a 160-bit digest. The approach for calculating the hash value SHOULD follow the same guidelines described for calculating the Integrity Check Value (ICV) described in [RFC4302] with respect to non-mutable fields. A history of the packet hash values SHOULD be maintained in the context of <protocol:srcAddr:dstAddr>. The context for IPv4 is more specific than that of IPv6 since the SMF_DPD HAV cannot be employed to mitigate hash collisions. A RECOMMENDED implementation detail for IPv4 H-DPD is to concatenate the 16-bit IPv4 ID value with the computed hash value as an extended DPD hash value that may provide reduced hash collisions in the cases where the IPv4 ID field is being set by host operating systems or gateways.

When this approach is taken, the use of the supplemental "internal hash" technique as described in Section 10 is RECOMMENDED as a security measure.

The SHA-1 hash is specified at present for consistency and robustness. Future approaches and experimentation may discover design trade-offs in hash robustness and efficiency worth considering for future revisions of SMF. This MAY include reducing the packet payload length that is processed, determining shorter indexes, or applying a more efficient hashing algorithm.

7. Relay Set Selection

SMF is flexible in its support of different reduced relay set mechanisms for efficient flooding, the constraints imposed herein being detailed in this section.

7.1. Non-Reduced Relay Set Forwarding

SMF implementations MUST support CF as a basic forwarding mechanism when reduced relay set information is not available or not selected for operation. In CF mode, each router transmits a packet once that has passed the SMF forwarding rules. The DPD techniques described in Section 6 are critical to proper operation and prevention of duplicate packet retransmissions by the same relays.

7.2. Reduced Relay Set Forwarding

MANET reduced relay sets are often achieved by distributed algorithms that can dynamically calculate a topological connected dominating set (CDS).

A goal of SMF is to apply reduced relay sets for more efficient multicast dissemination within dynamic topologies. To accomplish this, an SMF implementation MUST support the ability to modify its multicast packet forwarding rules based upon relay set state received dynamically during operation. In this way, SMF operates effectively as neighbor adjacencies or multicast forwarding policies within the topology change.

In early SMF experimental prototyping, the relay set information was derived from coexistent unicast routing control plane traffic flooding processes [MDC04]. From this experience, extra pruning considerations were sometimes required when utilizing a relay set from a separate routing protocol process. As an example, relay sets formed for the unicast control plane flooding MAY include additional redundancy that may not be desired for multicast forwarding use (e.g., biconnected relay set).

Here is a recommended criteria list for SMF relay set selection algorithm candidates:

1. Robustness to topological dynamics and mobility
2. Localized election or coordination of any relay sets
3. Reasonable minimization of CDS relay set size given the above constraints
4. Heuristic support for preference or election metrics

Some relay set algorithms meeting these criteria are described in the appendices of this document. Additional relay set selection algorithms may be specified in separate specifications in the future. Each appendix subsection in this document can serve as a template for specifying additional relay algorithms.

Figure 5 depicts an information flow diagram of possible relay set control options. The SMF Relay Set State represents the information base that is used by SMF in the forwarding decision process. The diagram demonstrates that the SMF Relay Set State may be determined by three fundamentally different methods:

- o Independent operation with NHDP [RFC6130] input providing dynamic network neighborhood adjacency information, used by a particular relay set selection algorithm.
- o Slave operation with an existing unicast MANET routing protocol, capable of providing CDS election information for use by SMF.
- o Cross-layer operation that may involve L2 triggers or information describing neighbors or links.

Other heuristics to influence and control election can come from network management or other interfaces as shown on the right of Figure 5. CF mode simplifies the control and does not require other input but relies solely on DPD.

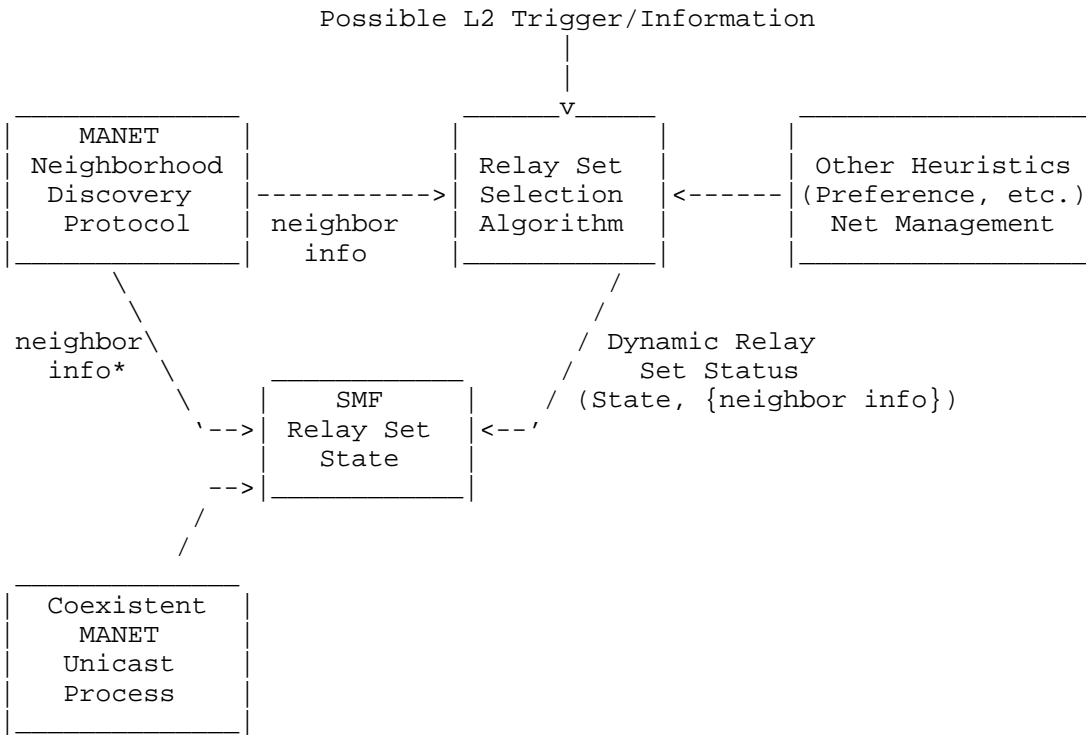


Figure 5: SMF Reduced Relay Set Information Flow

Following is further discussion of the three styles of SMF operation with reduced relay sets as illustrated in Figure 5:

1. Independent operation: In this case, SMF operates independently from any unicast routing protocols. To support reduced relay sets, SMF MUST perform its own relay set selection using information gathered from signaling. It is RECOMMENDED that an associated NHDP process be used for this signaling. NHDP messaging SHOULD be appended with additional [RFC5444] type-length-value (TLV) content as to support SMF-specific requirements as discussed in [RFC6130] and to support specific relay set operation as described in the appendices of this document or future specifications. Unicast routing protocols may coexist, even using the same NHDP process, but signaling that supports reduced relay set selection for SMF is independent of these protocols.

2. Operation with CDS-aware unicast routing protocol: In this case, a coexistent unicast routing protocol provides dynamic relay set state based upon its own control plane CDS or neighborhood discovery information.
 3. Cross-layer operation: In this case, SMF operates using neighborhood status and triggers from a cross-layer information base for dynamic relay set selection and maintenance (e.g., lower-link layer).
8. SMF Neighborhood Discovery Requirements

This section defines the requirements for use of the MANET Neighborhood Discovery Protocol (NHDP) [RFC6130] to support SMF operation. Note that basic CF forwarding requires no neighborhood topology knowledge since in this configured mode, every SMF router relays all traffic. Supporting more reduced SMF relay set operation requires the discovery and maintenance of dynamic neighborhood topology information. NHDP can be used to provide this necessary information; however, there are SMF-specific requirements for NHDP use. This is the case for both "independent" SMF operation where NHDP is being used specifically to support SMF or when one NHDP instance is used for both SMF and a coexistent MANET unicast routing protocol.

NHDP HELLO messages and the resultant neighborhood information base are described separately within the NHDP specification. To summarize, NHDP provides the following basic functions:

1. 1-hop neighbor link sensing and bidirectionality checks of neighbor links,
2. 2-hop neighborhood discovery including collection of 2-hop neighbors and connectivity information,
3. Collection and maintenance of the above information across multiple interfaces, and
4. A method for signaling SMF information throughout the 2-hop neighborhood through the use of TLV extensions.

Appendices A-C of this document describe CDS-based relay set selection algorithms that can achieve efficient SMF operation, even in dynamic, mobile networks and each of the algorithms has been initially experimented with in a working SMF prototype [MDDA07]. When using these algorithms in conjunction with NHDP, a method verifying neighbor SMF operation is required in order to ensure correct relay set selection. NHDP, along with SMF operation

verification, provides the necessary information required by these algorithms to conduct relay set selection. Verification of SMF operation may be done administratively or through the use of the SMF relay algorithm TLVs defined in the following subsections. Use of the SMF relay algorithm TLVs is RECOMMENDED when using NHDP for SMF neighborhood discovery.

Section 8.1 specifies SMF-specific TLV types, supporting general SMF operation or supporting the algorithms described in the appendices. The appendices describing several relay set algorithms also specify any additional requirements for use with NHDP and reference the applicable TLV types as needed.

8.1. SMF Relay Algorithm TLV Types

This section specifies TLV types to be used within NHDP messages to identify the CDS relay set selection algorithm(s) in use. Two TLV types are defined: one Message TLV type and one Address Block TLV type.

8.1.1. SMF Message TLV Type

The Message TLV type denoted SMF_TYPE is used to identify the existence of an SMF instance operating in conjunction with NHDP. This Message TLV type makes use of the extended type field as defined by [RFC5444] to convey the CDS relay set selection algorithm currently in use by the SMF message originator. When NHDP is used to support SMF operation, the SMF_TYPE TLV, containing the extended type field with the appropriate value, SHOULD be included in NHDP_HELLO messages (HELLO messages as defined in [RFC6130]). This allows SMF routers to learn when neighbors are configured to use NHDP for information exchange including algorithm type and related algorithm information. This information can be used to take action, such as ignoring neighbor information using incompatible algorithms. It is possible that SMF neighbors MAY be configured differently and still operate cooperatively, but these cases will vary dependent upon the algorithm types designated.

This document defines a Message TLV type as specified in Table 6 conforming to [RFC5444]. The TLV extended type field is used to contain the sender's "Relay Algorithm Type". The interpretation of the "value" content of these TLVs is defined per "Relay Algorithm Type" and may contain algorithm-specific information.

	TLV Syntax	Field Values
type	<tlv-type>	SMF_TYPE
extended type	<tlv-type-ext>	<relayAlgorithmId>
length	<length>	variable
value	<value>	variable

Table 6: SMF Type Message TLV

In Table 6, <relayAlgorithmId> is an 8-bit field containing a number 0-255 representing the "Relay Algorithm Type" of the originator address of the corresponding NHDP message.

Values for the <relayAlgorithmId> are defined in Table 7. The table provides value assignments, future IANA assignment spaces, and an experimental space. The experimental space use MUST NOT assume uniqueness; thus, it SHOULD NOT be used for general interoperable deployment prior to official IANA assignment.

Type Value	Extended Type Value	Algorithm
SMF_TYPE	0	CF
SMF_TYPE	1	S-MPR
SMF_TYPE	2	E-CDS
SMF_TYPE	3	MPR-CDS
SMF_TYPE	4-127	Future Assignment STD action
SMF_TYPE	128-239	No STD action required
SMF_TYPE	240-255	Experimental Space

Table 7: SMF Relay Algorithm Type Values

Acceptable <length> and <value> fields of an SMF_TYPE TLV are dependent on the extended type value (i.e., relay algorithm type). The appropriate algorithm type, as conveyed in the <tlv-type-ext> field, defines the meaning and format of its TLV <value> field. For the algorithms defined by this document, see the appropriate appendix for the <value> field format.

8.1.2. SMF Address Block TLV Type

An Address Block TLV type, denoted SMF_NBR_TYPE (i.e., SMF neighbor relay algorithm) is specified in Table 8. This TLV enables CDS relay algorithm operation and configuration to be shared among 2-hop

neighborhoods. Some relay algorithms require 2-hop neighbor configuration in order to correctly select relay sets. It is also useful when mixed relay algorithm operation is possible. Some examples of mixed use are outlined in the appendices.

The message SMF_TYPE TLV and Address Block SMF_NBR_TYPE TLV types share a common format.

	TLV syntax	Field Values
type	<tlv-type>	SMF_NBR_TYPE
extended type	<tlv-type-ext>	<relayAlgorithmId>
length	<length>	variable
value	<value>	variable

Table 8: SMF Type Address Block TLV

<relayAlgorithmId> in Table 8 is an 8-bit unsigned integer field containing a number 0-255 representing the "Relay Algorithm Type" value that corresponds to any associated address in the address block. Note that "Relay Algorithm Type" values for 2-hop neighbors can be conveyed in a single TLV or multiple value TLVs as described in [RFC5444]. It is expected that SMF routers using NHDP construct address blocks with SMF_NBR_TYPE TLVs to advertise "Relay Algorithm Type" and to advertise neighbor algorithm values received in SMF_TYPE TLVs from those neighbors.

Again, values for the <relayAlgorithmId> are defined in Table 7.

The interpretation of the "value" field of SMF_NBR_TYPE TLVs is defined per "Relay Algorithm Type" and may contain algorithm-specific information. See the appropriate appendix for definitions of value fields for the algorithms defined by this document.

9. SMF Border Gateway Considerations

It is expected that SMF will be used to provide simple forwarding of multicast traffic within a MANET or mesh routing topology. A border router gateway approach should be used to allow interconnection of SMF routing domains with networks using other multicast routing protocols, such as PIM. It is important to note that there are many scenario-specific issues that should be addressed when discussing border multicast routers. At the present time, experimental deployments of SMF and PIM border router approaches have been demonstrated [DHS08]. Some of the functionality border routers may need to address includes the following:

1. Determination of which multicast group traffic transits the border router whether entering or exiting the attached SMF routing domain.
2. Enforcement of TTL/hop limit threshold or other scoping policies.
3. Any marking or labeling to enable DPD on ingressing packets.
4. Interface with exterior multicast routing protocols.
5. Possible operation with multiple border routers (presently beyond the scope of this document).
6. Provisions for participating non-SMF devices (routers or hosts).

Each of these areas is discussed in more detail in the following subsections. Note the behavior of SMF border routers is the same as that of non-border SMF routers when forwarding packets on interfaces within the SMF routing domain. Packets that are passed outbound to interfaces operating fixed-infrastructure multicast routing protocols SHOULD be evaluated for duplicate packet status since present standard multicast forwarding mechanisms do not usually perform this function.

9.1. Forwarded Multicast Groups

Mechanisms for dynamically determining groups for forwarding into a MANET SMF routing domain is an evolving technology area. Ideally, only traffic for which there is active group membership should be injected into the SMF domain. This can be accomplished by providing an IPv4 Internet Group Membership Protocol (IGMP) or IPv6 Multicast Listener Discovery (MLD) proxy protocol so that MANET SMF routers can inform attached border routers (and hence multicast networks) of their current group membership status. For specific systems and services, it may be possible to statically configure group membership joins in border routers, but it is RECOMMENDED that some form of IGMP/MLD proxy or other explicit, dynamic control of membership be provided. Specification of such an IGMP/MLD proxy protocol is beyond the scope of this document.

For outbound traffic, SMF border routers perform duplicate packet detection and forward non-duplicate traffic that meets TTL/hop limit and scoping criteria to interfaces external to the SMF routing domain. Appropriate IP multicast routing (e.g., PIM-based solutions) on those interfaces can make further forwarding decisions with respect to the multicast packet. Note that the presence of multiple

border routers associated with a MANET routing domain raises additional issues. This is further discussed in Section 9.4 but further work is expected to be needed here.

9.2. Multicast Group Scoping

Multicast scoping is used by network administrators to control the network routing domains reachable by multicast packets. This is usually done by configuring external interfaces of border routers in the border of a routing domain to not forward multicast packets that must be kept within the SMF routing domain. This is commonly done based on TTL/hop limit of messages or by using administratively scoped group addresses. These schemes are known respectively as:

1. TTL scoping.
2. Administrative scoping.

For IPv4, network administrators can configure border routers with the appropriate TTL/hop limit thresholds or administratively scoped multicast groups for the router interfaces as with any traditional multicast router. However, for the case of TTL/hop limit scoping, it SHOULD be taken into account that the packet could traverse multiple hops within the MANET SMF routing domain before reaching the border router. Thus, TTL thresholds SHOULD be selected carefully.

For IPv6, multicast address spaces include information about the scope of the group. Thus, border routers of an SMF routing domain know if they must forward a packet based on the IPv6 multicast group address. For the case of IPv6, it is RECOMMENDED that a MANET SMF routing domain be designated a site-scoped multicast domain. Thus, all IPv6 site-scoped multicast packets in the range FF05::/16 SHOULD be kept within the MANET SMF routing domain by border routers. IPv6 packets in any other wider range scopes (i.e., FF08::/16, FF0B::/16, and FF0E::/16) MAY traverse border routers unless other restrictions different from the scope applies.

Given that scoping of multicast packets is performed at the border routers and given that existing scoping mechanisms are not designed to work with mobile routers, it is assumed that non-border routers running SMF will not stop forwarding multicast data packets of an appropriate site scoping. That is, it is assumed that an SMF routing domain is a site-scoped multicast area.

9.3. Interface with Exterior Multicast Routing Protocols

The traditional operation of multicast routing protocols is tightly integrated with the group membership function. Leaf routers are configured to periodically gather group membership information, while intermediate routers conspire to create multicast trees connecting routers with directly connected multicast sources and routers with active multicast receivers. In the concrete case of SMF, border routers can be considered leaf routers. Mechanisms for multicast sources and receivers to interoperate with border routers over the multi-hop MANET SMF routing domain as if they were directly connected to the router need to be defined. The following issues need to be addressed:

1. A mechanism by which border routers gather membership information
2. A mechanism by which multicast sources are known by the border router
3. A mechanism for exchange of exterior routing protocol messages across the SMF routing domain if the SMF routing domain is to provide transit connectivity for multicast traffic.

It is beyond the scope of this document to address implementation solutions to these issues. As described in Section 9.1, IGMP/MLD proxy mechanisms can address some of these issues. Similarly, exterior routing protocol messages could be tunneled or conveyed across an SMF routing domain but doing this robustly in a distributed wireless environment likely requires additional considerations outside the scope of this document.

The need for the border router to receive traffic from recognized multicast sources within the SMF routing domain is important to achieve interoperability with some existing routing protocols. For instance, PIM-S requires routers with locally attached multicast sources to register them to the Rendezvous Point (RP) so that routers can join the multicast tree. In addition, if those sources are not advertised to other autonomous systems (ASes) using Multicast Source Discovery Protocol (MSDP), receivers in those external networks are not able to join the multicast tree for that source.

9.4. Multiple Border Routers

An SMF routing domain might be deployed with multiple participating routers having connectivity to external, fixed-infrastructure networks. Allowing multiple routers to forward multicast traffic to/from the SMF routing domain can be beneficial since it can increase reliability and provide better service. For example, if the SMF

routing domain were to fragment with different SMF routers maintaining connectivity to different border routers, multicast service could still continue successfully. But, the case of multiple border routers connecting an SMF routing domain to external networks presents several challenges for SMF:

1. Handling duplicate unmarked IPv4 or IPv6 (without IPsec encapsulation or DPD option) packets possibly injected by multiple border routers.
2. Handling of duplicate traffic injected by multiple border routers by source-based relay algorithms.
3. Determining which border router(s) will forward outbound multicast traffic.
4. Additional challenges with interfaces to exterior multicast routing protocols.

When multiple border routers are present, they may be alternatively (due to route changes) or simultaneously injecting common traffic into the SMF routing domain that has not been previously marked for IPv6 SMF_DPD. Different border routers would not be able to implicitly synchronize sequencing of injected traffic since they may not receive exactly the same messages due to packet losses. For IPv6 I-DPD operation, the optional TaggerId field described for the SMF_DPD option header can be used to mitigate this issue. When multiple border routers are injecting a flow into an SMF routing domain, there are two forwarding policies that SMF routers running I-DPD may implement:

1. Redundantly forward the multicast flows (identified by <srcAddr:dstAddr>) from each border router, performing DPD processing on a <TaggerID:dstAddr> or <TaggerID:srcAddr:dstAddr> basis, or
2. Use some basis to select the flow of one tagger (border router) over the others and forward packets for applicable flows (identified by <sourceAddress:dstAddr>) only for the selected TaggerId until timeout or some other criteria to favor another tagger occurs.

It is RECOMMENDED that the first approach be used in the case of I-DPD operation. Additional specification may be required to describe an interoperable forwarding policy based on this second option. Note that the implementation of the second option requires that per-flow (i.e., <srcAddr::dstAddr>) state be maintained for the selected TaggerId.

The deployment of H-DPD operation may alleviate DPD resolution when ingressing traffic comes from multiple border routers. Non-colliding hash indexes (those not requiring the H-DPD options header in IPv6) should be resolved effectively.

10. Security Considerations

Gratuitous use of option headers can cause problems in routers. Other IP routers external to an SMF routing domain that might receive forwarded multicast SHOULD ignore SMF-specific IPv6 header options when encountered. The header option types are encoded appropriately to allow for this behavior.

This section briefly discusses several SMF denial-of-service (DoS) attack scenarios and provides some initial recommended mitigation strategies.

A potential denial-of-service attack against SMF forwarding is possible when a malicious router has a form of wormhole access to non-adjacent parts of a network topology. In the wireless ad hoc case, a directional antenna is one way to provide such a wormhole physically. If such a router can preview forwarded packets in a non-adjacent part of the network and forward modified versions to another part of the network, it can perform the following attack. The malicious router could reduce the TTL/hop limit or hop limit of the packet and transmit it to the SMF router causing it to forward the packet with a limited TTL/hop limit (or even drop it) and make a DPD entry that could block or limit the subsequent forwarding of later-arriving valid packets with correct TTL/hop limit values. This would be a relatively low-cost, high-payoff attack that would be hard to detect and thus attractive to potential attackers. An approach of caching TTL/hop limit information with DPD state and taking appropriate forwarding actions is identified in Section 5 to mitigate this form of attack.

Sequence-based packet identifiers are predictable and thus provide an opportunity for a DoS attack against forwarding. Forwarding protocols that use DPD techniques, such as SMF, may be vulnerable to DoS attacks based on spoofing packets with apparently valid packet identifier fields. In wireless environments, where SMF will most likely be used, the opportunity for such attacks may be more prevalent than in wired networks. In the case of IPv4 packets, fragmented IP packets, or packets with IPsec headers applied, the DPD "identifier portions" of potential future packets that might be forwarded is highly predictable and easily subject to DoS attacks against forwarding. A RECOMMENDED technique to counter this concern is for SMF implementations to generate an "internal" hash value that is concatenated with the explicit I-DPD packet identifier to form a

unique identifier that is a function of the packet content as well as the visible identifier. SMF implementations could seed their hash generation with a random value to make it unlikely that an external observer could guess how to spoof packets used in a denial-of-service attack against forwarding. Since the hash computation and state is kept completely internal to SMF routers, the cryptographic properties of this hashing would not need to be extensive and thus possibly of low complexity. Experimental implementations may determine that even a lightweight hash of only portions of packets may suffice to serve this purpose.

While H-DPD is not as readily susceptible to this form of DoS attack, it is possible that a sophisticated adversary could use side information to construct spoofing packets to mislead forwarders using a well-known hash algorithm. Thus, similarly, a separate "internal" hash value could be concatenated with the well-known hash value to alleviate this security concern.

The support of forwarding IPsec packets without further modification for both IPv4 and IPv6 is supported by this specification.

Authentication mechanisms to identify the source of IPv6 option headers should be considered to reduce vulnerability to a variety of attacks.

Furthermore, when the MANET Neighborhood Discovery Protocol [RFC6130] is used, the security considerations described in [RFC6130] also apply.

11. IANA Considerations

This document defines one IPv6 Hop-by-Hop Option, a type for which has been allocated from the IPv6 "Destination Options and Hop-by-Hop Options" registry of [RFC2780].

This document creates one registry called "TaggerId Types" for recording TaggerId types, (TidTy), as a sub-registry in the "IPv6 Parameters" registry.

This document registers one well-known multicast address from each of the IPv4 and IPv6 multicast address spaces.

This document defines one Message TLV, a type for which has been allocated from the "Message TLV Types" registry of [RFC5444].

Finally, this document defines one Address Block TLV, a type for which has been allocated from the "Address Block TLV Types" registry of [RFC5444].

11.1. IPv6 SMF_DPD Header Extension Option Type

IANA has allocated an IPv6 Option Type from the IPv6 "Destination Options and Hop-by-Hop Options" registry of [RFC2780], as specified in Table 9.

Hex Value	Binary Value			Description	Reference
	act	chg	rest		
8	00	0	01000	SMF_DPD	This Document

Table 9: IPv6 Option Type Allocation

11.2. TaggerId Types (TidTy)

A portion of the option data content in the SMF_DPD is the Tagger Identifier Type (TidTy), which provides a context for the optionally included TaggerId.

IANA has created a registry for recording TaggerId Types (TidTy), with initial assignments and allocation policies, as specified in Table 10.

Type	Mnemonic	Description	Reference
0	NULL	No TaggerId field is present	This document
1	DEFAULT	A TaggerId of non-specific context is present	This document
2	IPv4	A TaggerId representing an IPv4 address is present	This document
3	IPv6	A TaggerId representing an IPv6 address is present	This document
4-7		Unassigned	

Table 10: TaggerId Types

For allocation of unassigned values 4-7, IETF Review [RFC5226] is required.

11.3. Well-Known Multicast Address

IANA has allocated an IPv4 multicast address "SL-MANET-ROUTERS" (224.0.1.186) from the "Internet Control Block (224.0.1.0-224.0.1.255 (224.0.1/24))" sub-registry of the "IPv4 Multicast Address" registry.

IANA has allocated an IPv6 multicast address "SL-MANET-ROUTERS" from the "Site-Local Scope Multicast Addresses" sub-sub-registry of the "Fixed Scope Multicast Addresses" sub-registry of the "INTERNET PROTOCOL VERSION 6 MULTICAST ADDRESSES" registry.

11.4. SMF TLVs

11.4.1. Expert Review for Created Type Extension Registries

Creation of Address Block TLV Types and Message TLV Types in registries of [RFC5444], and hence in the HELLO-message-specific registries of [RFC6130], entails creation of corresponding Type Extension registries for each such type. For such Type Extension registries, where an Expert Review is required, the designated expert SHOULD take the same general recommendations into consideration as those specified by [RFC5444].

11.4.2. SMF Message TLV Type (SMF_TYPE)

This document defines one Message TLV Type, "SMF_TYPE", which has been allocated from the "HELLO Message-Type-specific Message TLV Types" registry, defined in [RFC6130].

This created a new Type Extension registry, with initial assignments as specified in Table 11.

Name	Type	Type Extension	Description	Allocation Policy
SMF_TYPE	128	0-255	Specifies relay algorithm supported by the SMF router, originating the HELLO message, according to Section 11.4.4.	Section 11.4.4

Table 11: SMF_TYPE Message TLV Type Extension Registry

11.4.3. SMF Address Block TLV Type (SMF_NBR_TYPE)

This document defines one Address Block TLV Type, "SMF_NBR_TYPE", which has been allocated from the "HELLO Message-Type-specific Address Block TLV Types" registry, defined in [RFC6130].

This has created a new Type Extension registry, with initial assignments as specified in Table 12.

Name	Type	Type Extension	Description	Allocation Policy
SMF_NBR_TYPE	128	0-255	Specifies relay algorithm supported by the SMF router corresponding to the advertised address, according to Section 11.4.4.	Section 11.4.4

Table 12: SMF_NBR_TYPE Address Block TLV Type Extension Registry

11.4.4. SMF Relay Algorithm ID Registry

Types for the Type Extension Registries for the SMF_TYPE Message TLV and the SMF_NBR_TYPE Address Block TLV are unified in this single SMF Relay Algorithm ID Registry, defined in this section.

IANA has created a registry for recording Relay Algorithm Identifiers, with initial assignments and allocation policies as specified in Table 13.

Value	Name	Description	Allocation Policy
0	CF	Section 4	
1	S-MPR	Appendix B	
2	E-CDS	Appendix A	
3	MPR-CDS	Appendix C	
4-127		Unassigned	Expert Review
128-255		Unassigned	Experimental Use

Table 13: Relay Set Algorithm Type Values

A specification requesting an allocation from the 4-127 range from the SMF Relay Algorithm ID Registry MUST specify the interpretation of the <value> field (if any).

12. Acknowledgments

Many of the concepts and mechanisms used and adopted by SMF resulted over several years of discussion and related work within the MANET working group since the late 1990s. There are obviously many contributors to past discussions and related draft documents within the working group that have influenced the development of SMF concepts, and they deserve acknowledgment. In particular, this document is largely a direct product of the earlier SMF design team within the IETF MANET working group and borrows text and implementation ideas from the related individuals and activities. Some of the direct contributors who have been involved in design, content editing, prototype implementation, major commenting, and core discussions are listed below in alphabetical order. We appreciate all the input and feedback from the many community members and early implementation users we have heard from that are not on this list as well.

Brian Adamson
Teco Boot
Ian Chakeres
Thomas Clausen
Justin Dean
Brian Haberman
Ulrich Herberg
Charles Perkins
Pedro Ruiz
Fred Templin
Maoyu Wang

13. References

13.1. Normative References

- [MPR-CDS] Adjih, C., Jacquet, P., and L. Viennot, "Computing Connected Dominating Sets with Multipoint Relays", Ad Hoc and Sensor Wireless Networks, January 2005.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [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.
- [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, August 1999.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.
- [RFC3174] Eastlake, D. and P. Jones, "US Secure Hash Algorithm 1 (SHA1)", RFC 3174, September 2001.
- [RFC3626] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, February 2009.
- [RFC5614] Ogier, R. and P. Spagnolo, "Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding", RFC 5614, August 2009.

- [RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 5771, March 2010.
- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", RFC 6130, April 2011.

13.2. Informative References

- [CDHM07] Chakeres, I., Danilov, C., Henderson, T., and J. Macker, "Connecting MANET Multicast", IEEE MILCOM 2007 Proceedings, 2007.
- [DHG09] Danilov, C., Henderson, T., Goff, T., Kim, J., Macker, J., Weston, J., Neogi, N., Ortiz, A., and D. Uhlig, "Experiment and field demonstration of a 802.11-based ground-UAV mobile ad-hoc network", Proceedings of the 28th IEEE conference on Military Communications, 2009.
- [DHS08] Danilov, C., Henderson, T., Spagnolo, T., Goff, T., and J. Kim, "MANET Multicast with Multiple Gateways", IEEE MILCOM 2008 Proceedings, 2008.
- [GM99] Garcia-Luna-Aceves, JJ. and E. Madruga, "The Core-Assisted Mesh Protocol", Selected Areas in Communications, IEEE Journal, Volume 17, Issue 8, August 1999.
- [IPV4-ID-UPDATE] Touch, J., "Updated Specification of the IPv4 ID Field", Work in Progress, September 2011.
- [JLMV02] Jacquet, P., Laouiti, V., Minet, P., and L. Viennot, "Performance of Multipoint Relaying in Ad Hoc Mobile Routing Protocols", Networking , 2002.
- [MDC04] Macker, J., Dean, J., and W. Chao, "Simplified Multicast Forwarding in Mobile Ad hoc Networks", IEEE MILCOM 2004 Proceedings, 2004.
- [MDDA07] Macker, J., Downard, I., Dean, J., and R. Adamson, "Evaluation of Distributed Cover Set Algorithms in Mobile Ad hoc Network for Simplified Multicast Forwarding", ACM SIGMOBILE Mobile Computing and Communications Review, Volume 11, Issue 3, July 2007.

- [MGL04] Mohapatra, P., Gui, C., and J. Li, "Group Communications in Mobile Ad hoc Networks", IEEE Computer, Vol. 37, No. 2, February 2004.
- [NTSC99] Ni, S., Tseng, Y., Chen, Y., and J. Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network", Proceedings of ACM Mobicom 99, 1999.
- [RFC2501] Corson, M. and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, January 1999.
- [RFC3684] Ogier, R., Templin, F., and M. Lewis, "Topology Dissemination Based on Reverse-Path Forwarding (TBRPF)", RFC 3684, February 2004.
- [RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, January 2005.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

Appendix A. Essential Connecting Dominating Set (E-CDS) Algorithm

The "Essential Connected Dominating Set" (E-CDS) algorithm [RFC5614] forms a single CDS mesh for the SMF operating region. It allows routers to use 2-hop neighborhood topology information to dynamically perform relay self-election to form a CDS. Its packet-forwarding rules are not dependent upon previous hop knowledge. Additionally, E-CDS SMF forwarders can be easily mixed without problems with CF SMF forwarders, even those not participating in NHDP. Another benefit is that packets opportunistically received from non-symmetric neighbors may be forwarded without compromising flooding efficiency or correctness. Furthermore, multicast sources not participating in NHDP may freely inject their traffic, and any neighboring E-CDS relays will properly forward the traffic. The E-CDS-based relay set selection algorithm is based upon [RFC5614]. E-CDS was originally discussed in the context of forming partial adjacencies and efficient flooding for MANET OSPF extensions work, and the core algorithm is applied here for SMF.

It is RECOMMENDED that the SMF_TYPE:E-CDS Message TLV be included in NHDP_HELLO messages that are generated by routers conducting E-CDS SMF operation. It is also RECOMMENDED that the SMF_NBR_TYPE:E-CDS Address Block TLV be used to advertise neighbor routers that are also conducting E-CDS SMF operation.

A.1. E-CDS Relay Set Selection Overview

The E-CDS relay set selection requires 2-hop neighborhood information collected through NHDP or another process. Relay routers, in E-CDS SMF selection, are "self-elected" using a Router Identifier (Router ID) and an optional nodal metric, referred to here as Router Priority for all 1-hop and 2-hop neighbors. To ensure proper relay set self-election, the Router ID and Router Priority MUST be consistent among participating routers. It is RECOMMENDED that NHDP be used to share Router ID and Router Priority through the use of SMF_TYPE:E-CDS TLVs as described in this appendix. The Router ID is a logical identification that MUST be consistent across interoperating SMF neighborhoods, and it is RECOMMENDED to be chosen as the numerically largest address contained in a router's "Neighbor Address List" as defined in NHDP. The E-CDS self-election process can be summarized as follows:

1. If an SMF router has a higher ordinal (Router Priority, Router ID) than all of its symmetric neighbors, it elects itself to act as a forwarder for all received multicast packets.

2. Else, if there does not exist a path from the neighbor with largest (Router Priority, Router ID) to any other neighbor, via neighbors with larger values of (Router Priority, Router ID), then it elects itself to the relay set.

The basic form of E-CDS described and applied within this specification does not provide for redundant relay set selection (e.g., bi-connected), but such capability is supported by the basic E-CDS design.

A.2. E-CDS Forwarding Rules

With E-CDS, any SMF router that has selected itself as a relay performs DPD and forwards all non-duplicative multicast traffic allowed by the present forwarding policy. Packet previous-hop knowledge is not needed for forwarding decisions when using E-CDS.

1. Upon packet reception, DPD is performed. Note E-CDS requires a single duplicate table for the set of interfaces associated with the relay set selection.
2. If the packet is a duplicate, no further action is taken.
3. If the packet is non-duplicative:
 - A. A DPD entry is made for the packet identifier.
 - B. The packet is forwarded out to all interfaces associated with the relay set selection.

As previously mentioned, even packets sourced (or relayed) by routers not participating in NHDP and/or the E-CDS relay set selection may be forwarded by E-CDS forwarders without problem. A particular deployment MAY choose to not forward packets from previous hop routers that have been not explicitly identified via NHDP or other means as operating as part of a different relay set algorithm (e.g., S-MPR) to allow coexistent deployments to operate correctly. Also, E-CDS relay set selection may be configured to be influenced by statically configured CF relays that are identified via NHDP or other means.

A.3. E-CDS Neighborhood Discovery Requirements

It is possible to perform E-CDS relay set selection without modification of NHDP, basing the self-election process exclusively on the "Neighbor Address List" of participating SMF routers, for example, by setting the Router Priority to a default value and selecting the Router ID as the numerically largest address contained

in the "Neighbor Address List". However, steps MUST be taken to ensure that all NHDP-enabled routers not using SMF_TYPE:E-CDS full type Message TLVs are, in fact, running SMF E-CDS with the same methods for selecting Router Priority and Router ID; otherwise, incorrect forwarding may occur. Note that SMF routers with higher Router Priority values will be favored as relays over routers with lower Router Priority. Thus, preferred relays MAY be administratively configured to be selected when possible. Additionally, other metrics (e.g., nodal degree, energy capacity, etc.) may also be taken into account in constructing a Router Priority value. When using Router Priority with multiple interfaces, all interfaces on a router MUST use and advertise a common Router Priority value. A router's Router Priority value may be administratively or algorithmically selected. The method of selection does not need to be the same among different routers.

E-CDS relay set selection may be configured to be influenced by statically configured CF relays that are identified via NHDP or other means. Nodes advertising CF through NHDP may be considered E-CDS SMF routers with maximal Router Priority.

To share a router's Router Priority with its 1-hop neighbors, the SMF_TYPE:E-CDS Message TLV's <value> field is defined as shown in Table 14.

Length (bytes)	Value	Router Priority
0	N/A	64
1	<value>	0-127

Table 14: E-CDS Message TLV Values

Where <value> is a one-octet-long bit field that is defined as:

bit 0: the leftmost bit is reserved and SHOULD be set to 0.

bits 1-7: contain the unsigned Router Priority value, 0-127, which is associated with the "Neighbor Address List".

Combinations of value field lengths and values other than specified here are NOT permitted and SHOULD be ignored. Figure 6 shows an example SMF_TYPE:E-CDS Message TLV.




Figure 6: E-CDS Message TLV Example

To convey Router Priority values among 2-hop neighborhoods, the SMF_NBR_TYPE:E-CDS Address Block TLV's <value> field is used. Multi-index and multivalue TLV layouts as defined in [RFC5444] are supported. SMF_NBR_TYPE:E-CDS value fields are defined thus:

Length(bytes)	# Addr	Value	Router Priority
0	Any	N/A	64
1	Any	<value>	<value> is for all addresses
N	N	<value>*	Each address gets its own <value>

Table 15: E-CDS Address Block TLV Values

Where <value> is a one-byte bit field that is defined as:

bit 0: the leftmost bit is reserved and SHOULD be set to 0.

bits 1-7: contain the unsigned Router Priority value, 0-127, which is associated with the appropriate address(es).

Combinations of value field lengths and # of addresses other than specified here are NOT permitted and SHOULD be ignored. A default technique of using nodal degree (i.e., count of 1-hop neighbors) is RECOMMENDED for the value field of these TLV types. Below are two example SMF_NBR_TYPE:E-CDS Address Block TLVs.




Figure 7: E-CDS Address Block TLV Example 1

"Neighbor Address List" as advertised by NHDP MUST be used as the Router ID. NHDP provides all interface addresses throughout the 2-hop neighborhood through HELLO messages, so explicitly conveying a Router ID is not necessary. The following steps describe a basic algorithm for conducting E-CDS relay selection for a router "n0":

1. Initialize the set "N1" with tuples ("Router Priority", "Router ID", "Neighbor Address List") for each 1-hop neighbor of "n0".
2. If "N1" has less than 2 tuples, then "n0" does not elect itself as a relay, and no further steps are taken.
3. Initialize the set "N2" with tuples ("Router Priority", "Router ID", "2-hop address") for each "2-hop address" of "n0", where "2-hop address" is defined in NHDP.
4. If "RtrPri(n0)" is greater than that of all tuples in the union of "N1" and "N2", then "n0" selects itself as a relay, and no further steps are taken.
5. Initialize all tuples in the union of "N1" and "N2" as "unvisited".
6. Find the tuple "nl_Max" that has the largest "RtrPri()" of all tuples in "N1".
7. Initialize queue "Q" to contain "nl_Max", marking "nl_Max" as "visited".
8. While router queue "Q" is not empty, remove router "x" from the head of "Q", and for each 1-hop neighbor "n" of router "x" (excluding "n0") that is not marked "visited".
 - A. Mark router "n" as "visited".
 - B. If "RtrPri(n)" is greater than "RtrPri(n0)", append "n" to "Q".
9. If any tuple in "N1" remains "unvisited", then "n0" selects itself as a relay. Otherwise, "n0" does not act as a relay.

Note these steps are re-evaluated upon neighborhood status changes. Steps 5 through 8 of this procedure describe an approach to a path search. The purpose of this path search is to determine if paths exist from the 1-hop neighbor with maximum "RtrPri()" to all other 1-hop neighbors without traversing an intermediate router with a "RtrPri()" value less than "RtrPri(n0)". These steps comprise a breadth-first traversal that evaluates only paths that meet that

criteria. If all 1-hop neighbors of "n0" are "visited" during this traversal, then the path search has succeeded, and router "n0" does not need to provide relay. It can be assumed that other routers will provide relay operation to ensure SMF connectivity.

It is possible to extend this algorithm to consider neighboring SMF routers that are known to be statically configured for CF (always relaying). The modification to the above algorithm is to process such routers as having a maximum possible Router Priority value. It is expected that routers configured for CF and participating in NHDP would indicate this with use of the SMF_TYPE:CF and SMF_NBR_TYPE:CF TLV types in their NHDP_HELLO message and address blocks, respectively.

Appendix B. Source-Based Multipoint Relay (S-MPR) Algorithm

The source-based multipoint relay (S-MPR) set selection algorithm enables individual routers, using 2-hop topology information, to select relays from their set of neighboring routers. Relays are selected so that forwarding to the router's complete 2-hop neighbor set is covered. This distributed relay set selection technique has been shown to approximate a minimal connected dominating set (MCDS) in [JLMV02]. Individual routers must collect 2-hop neighborhood information from neighbors, determine an appropriate current relay set, and inform selected neighbors of their relay status. Note that since each router picks its neighboring relays independently, S-MPR forwarders depend upon previous hop information (e.g., source MAC address) to operate correctly. The Optimized Link State Routing (OLSR) protocol has used this algorithm and protocol for relay of link state updates and other control information [RFC3626], and it has been demonstrated operationally in dynamic network environments.

It is RECOMMENDED that the SMF_TYPE:S-MPR Message TLV be included in NHDP_HELLO messages that are generated by routers conducting S-MPR SMF operation. It is also RECOMMENDED that the SMF_NBR_TYPE:S-MPR Address Block TLV be used to specify which neighbor routers are conducting S-MPR SMF operation.

B.1. S-MPR Relay Set Selection Overview

The S-MPR algorithm uses bi-directional 1-hop and 2-hop neighborhood information collected via NHDP to select, from a router's 1-hop neighbors, a set of relays that will cover the router's entire 2-hop neighbor set upon forwarding. The algorithm described uses a "greedy" heuristic of first picking the 1-hop neighbor who will cover the most 2-hop neighbors. Then, excluding those 2-hop neighbors that have been covered, additional relays from its 1-hop neighbor set are

iteratively selected until the entire 2-hop neighborhood is covered. Note that 1-hop neighbors also identified as 2-hop neighbors are considered as 1-hop neighbors only.

NHDP HELLO messages supporting S-MPR forwarding operation SHOULD use the TLVs defined in Section 8.1 using the S-MPR extended type. The value field of an Address Block TLV that has a full type value of SMF_NBR_TYPE:S-MPR is defined in Table 17 such that signaling of MPR selections to 1-hop neighbors is possible. The value field of a message block TLV that has a full type value of SMF_TYPE:S-MPR is defined in Table 16 such that signaling of Router Priority (described as "WILLINGNESS" in [RFC3626]) to 1-hop neighbors is possible. It is important to note that S-MPR forwarding is dependent upon the previous hop of an incoming packet. An S-MPR router MUST forward packets only for neighbors that have explicitly selected it as a multipoint relay (i.e., its "selectors"). There are also some additional requirements for duplicate packet detection to support S-MPR SMF operation that are described below.

For multiple interface operation, MPR selection SHOULD be conducted on a per-interface basis. However, it is possible to economize MPR selection among multiple interfaces by selecting common MPRs to the extent possible.

B.2. S-MPR Forwarding Rules

An S-MPR SMF router MUST only forward packets for neighbors that have explicitly selected it as an MPR. The source-based forwarding technique also stipulates some additional duplicate packet detection operations. For multiple network interfaces, independent DPD state MUST be maintained for each separate interface. The following provides the procedure for S-MPR packet forwarding given the arrival of a packet on a given interface, denoted <srcIface>. There are three possible actions, depending upon the previous-hop transmitter:

1. If the previous-hop transmitter has selected the current router as an MPR,
 - A. The packet identifier is checked against the DPD state for each possible outbound interface, including the <srcIface>.
 - B. If the packet is not a duplicate for an outbound interface, the packet is forwarded on that interface and a DPD entry is made for the given packet identifier for the interface.
 - C. If the packet is a duplicate, no action is taken for that interface.

2. Else, if the previous-hop transmitter is a 1-hop symmetric neighbor, a DPD entry is added for that packet for the <srcIface>, but the packet is not forwarded.
3. Otherwise, no action is taken.

Action number two in the list above is non-intuitive but important to ensure correctness of S-MPR SMF operation. The selection of source-based relays does not result in a common set among neighboring routers, so relays MUST mark, in their DPD state, packets received from non-selector, symmetric, 1-hop neighbors (for a given interface) and not forward subsequent duplicates of that packet if received on that interface. Deviation here can result in unnecessary, repeated packet forwarding throughout the network or incomplete flooding.

Nodes not participating in neighborhood discovery and relay set selection will not be able to source multicast packets into the area and have SMF forward them, unlike E-CDS or MPR-CDS where forwarding may occur dependent on topology. Correct S-MPR relay behavior will occur with the introduction of repeaters (non-NHDP/SMF participants that relay multicast packets using duplicate detection and CF), but the repeaters will not efficiently contribute to S-MPR forwarding as these routers will not be identified as neighbors (symmetric or otherwise) in the S-MPR forwarding process. NHDP/SMF participants MUST NOT forward packets that are not selected by the algorithm, as this can disrupt network-wide S-MPR flooding, resulting in incomplete or inefficient flooding. The result is that non-S-MPR SMF routers will be unable to source multicast packets and have them forwarded by other S-MPR SMF routers.

B.3. S-MPR Neighborhood Discovery Requirements

Nodes may optionally signal a Router Priority value to their 1-hop neighbors by using the SMF_TYPE:S-MPR message block TLV value field. If the value field is omitted, a default Router Priority value of 64 is to be assumed. This is summarized here:

Length(bytes)	Value	Router Priority
0	N/A	64
1	<value>	0-127

Table 16: S-MPR Message TLV Values

Where <value> is a one-octet-long bit field defined as:

bit 0: the leftmost bit is reserved and SHOULD be set to 0.

bits 1-7: contain the Router Priority value, 0-127, which is associated with the "Neighbor Address List".

Router Priority values for S-MPR are interpreted in the same fashion as "WILLINGNESS" ([RFC3626]), with the value 0 indicating a router will NEVER forward and value 127 indicating a router will ALWAYS forward. Values 1-126 indicate how likely a S-MPR SMF router will be selected as an MPR by a neighboring SMF router, with higher values increasing the likelihood. Combinations of value field lengths and values other than those specified here are NOT permitted and SHOULD be ignored. Below is an example SMF_TYPE:S-MPR Message TLV.




Figure 9: S-MPR Message TLV Example

S-MPR election operation requires 2-hop neighbor knowledge as provided by NHDP [RFC6130] or from external sources. MPRs are dynamically selected by each router, and selections MUST be advertised and dynamically updated within NHDP or an equivalent protocol or mechanism. For NHDP use, the SMF_NBR_TYPE:S-MPR Address Block TLV value field is defined as such:

Length(bytes)	# Addr	Value	Meaning
0	Any	N/A	NOT MPRs
1	Any	<value>	<value> is for all addresses
N	N	<value>*	Each address gets its own <value>

Table 17: S-MPR Address Block TLV Values

Where <value>, if present, is a one-octet bit field defined as:

bit 0: The leftmost bit is the M bit that, when set, indicates MPR selection of the relevant interface, represented by the associated address(es), by the originator router of the NHDP HELLO message. When unset, it indicates the originator router of the NHDP HELLO message has not selected the relevant interfaces, represented by the associated address(es), as its MPR.

bits 1-7: These bits are reserved and SHOULD be set to 0.

Combinations of value field lengths and number of addresses other than specified here are NOT permitted and SHOULD be ignored. All bits, excepting the leftmost bit, are RESERVED and SHOULD be set to 0.




Figure 10: S-MPR Address Block TLV Example

The single index TLV example, depicted in Figure 10, indicates that the address specified by the <start-index> field is running SMF using S-MPR and has been selected by the originator of the NHDP HELLO message as an MPR forwarder if the M bit is set. Multivalued TLVs may also be used to specify MPR selection status of multiple addresses using only one TLV. See Figure 8 for a similar example on how this may be done.

B.4. S-MPR Selection Algorithm

This section describes a basic algorithm for the S-MPR selection process. Note that the selection is with respect to a specific interface of the router performing selection, and other router interfaces referenced are reachable from this reference router interface. This is consistent with the S-MPR forwarding rules described above. When multiple interfaces per router are used, it is possible to enhance the overall selection process across multiple interfaces such that common routers are selected as MPRs for each interface to avoid unnecessary inefficiencies in flooding. The following steps describe a basic algorithm for conducting S-MPR selection for a router interface "n0":

1. Initialize the set "MPR" to empty.
2. Initialize the set "N1" to include all 1-hop neighbors of "n0".
3. Initialize the set "N2" to include all 2-hop neighbors, excluding "n0" and any routers in "N1". Nodes that are only reachable via "N1" routers with router priority values of NEVER are also excluded.
4. For each interface "y" in "N1", initialize a set "N2(y)" to include any interfaces in "N2" that are 1-hop neighbors of "y".
5. For each interface "x" in "N1" with a router priority value of "ALWAYS" (or using the CF relay algorithm), select "x" as an MPR:
 - A. Add "x" to the set "MPR" and remove "x" from "N1".
 - B. For each interface "z" in "N2(x)", remove "z" from "N2".
 - C. For each interface "y" in "N1", remove any interfaces in "N2(x)" from "N2(y)".
6. For each interface "z" in "N2", initialize the set "N1(z)" to include any interfaces in "N1" that are 1-hop neighbors of "z".
7. For each interface "x" in "N2" where "N1(x)" has only one member, select "x" as an MPR:
 - A. Add "x" to the set "MPR" and remove "x" from "N1".
 - B. For each interface "z" in "N2(x)", remove "z" from "N2" and delete "N1(z)".
 - C. For each interface "y" in "N1", remove any interfaces in "N2(x)" from "N2(y)".
8. While "N2" is not empty, select the interface "x" in "N1" with the largest router priority that has the number of members in "N_2(x)" as an MPR:
 - A. Add "x" to the set "MPR" and remove "x" from "N1".
 - B. For each interface "z" in "N2(x)", remove "z" from "N2".
 - C. For each interface "y" in "N1", remove any interfaces in "N2(x)" from "N2(y)".

After the set of routers "MPR" is selected, router "n_0" must signal its selections to its neighbors. With NHDP, this is done by using the MPR Address Block TLV to mark selected neighbor addresses in NHDP_HELLO messages. Neighbors MUST record their MPR selection status and the previous hop address (e.g., link or MAC layer) of the selector. Note these steps are re-evaluated upon neighborhood status changes.

Appendix C. Multipoint Relay Connected Dominating Set (MPR-CDS) Algorithm

The MPR-CDS algorithm is an extension to the basic S-MPR election algorithm that results in a shared (non-source-specific) SMF CDS. Thus, its forwarding rules are not dependent upon previous hop information, similar to E-CDS. An overview of the MPR-CDS selection algorithm is provided in [MPR-CDS].

It is RECOMMENDED that the SMF_TYPE Message TLV be included in NHDP_HELLO messages that are generated by routers conducting MPR-CDS SMF operation.

C.1. MPR-CDS Relay Set Selection Overview

The MPR-CDS relay set selection process is based upon the MPR selection process of the S-MPR algorithm with the added refinement of a distributed technique for subsequently down-selecting to a common reduced, shared relay set. A router ordering (or "prioritization") metric is used as part of this down-selection process; like the E-CDS algorithm, this metric can be based upon router address(es) or some other unique router identifier (e.g., Router ID based on largest address contained within the "Neighbor Address List") as well as an additional Router Priority measure, if desired. The process for MPR-CDS relay selection is as follows:

1. First, MPR selection per the S-MPR algorithm is conducted, with selectors informing their MPRs (via NHDP) of their selection.
2. Then, the following rules are used on a distributed basis by selected routers to possibly deselect themselves and thus jointly establish a common set of shared SMF relays:
 - A. If a selected router has a larger "RtrPri()" than all of its 1-hop symmetric neighbors, then it acts as a relay for all multicast traffic, regardless of the previous hop.

- B. Else, if the 1-hop symmetric neighbor with the largest "RtrPri()" value has selected the router, then it also acts as a relay for all multicast traffic, regardless of the previous hop.
- C. Otherwise, it deselects itself as a relay and does not forward any traffic unless changes occur that require re-evaluation of the above steps.

This technique shares many of the desirable properties of the E-CDS technique with regards to compatibility with multicast sources not participating in NHDP and the opportunity for statically configured CF routers to be present, regardless of their participation in NHDP.

C.2. MPR-CDS Forwarding Rules

The forwarding rules for MPR-CDS are similar to those for E-CDS. Any SMF router that has selected itself as a relay performs DPD and forwards all non-duplicative multicast traffic allowed by the present forwarding policy. Packet previous hop knowledge is not needed for forwarding decisions when using MPR-CDS.

1. Upon packet reception, DPD is performed. Note that MPR-CDS requires one duplicate table for the set of interfaces associated with the relay set selection.
2. If the packet is a duplicate, no further action is taken.
3. If the packet is non-duplicative:
 - A. A DPD entry is added for the packet identifier
 - B. The packet is forwarded out to all interfaces associated with the relay set selection.

As previously mentioned, even packets sourced (or relayed) by routers not participating in NHDP and/or the MPR-CDS relay set selection may be forwarded by MPR-CDS forwarders without problem. A particular deployment MAY choose to not forward packets from sources or relays that have been explicitly identified via NHDP or other means as operating as part of a different relay set algorithm (e.g., S-MPR) to allow coexistent deployments to operate correctly.

C.3. MPR-CDS Neighborhood Discovery Requirements

The neighborhood discovery requirements for MPR-CDS have commonality with both the S-MPR and E-CDS algorithms. MPR-CDS selection operation requires 2-hop neighbor knowledge as provided by NHDP

[RFC6130] or from external sources. Unlike S-MPR operation, there is no need for associating link-layer address information with 1-hop neighbors since MPR-CDS forwarding is independent of the previous hop similar to E-CDS forwarding.

To advertise an optional Router Priority value or "WILLINGNESS", an originating router may use the Message TLV of type SMF_TYPE:MPR-CDS, which shares a common <value> format with both SMF_TYPE:E-CDS (Table 14) and SMF_TYPE:S-MPR (Table 16).

MPR-CDS only requires 1-hop knowledge of Router Priority for correct operation. In the S-MPR phase of MPR-CDS selection, MPRs are dynamically determined by each router, and selections MUST be advertised and dynamically updated using NHDP or an equivalent protocol or mechanism. The <value> field of the SMF_NBR_TYPE:MPR-CDS type TLV shares a common format with SMF_NBR_TYPE:S-MPR (Table 17) to convey MPR selection.

C.4. MPR-CDS Selection Algorithm

This section describes an algorithm for the MPR-CDS selection process. Note that the selection described is with respect to a specific interface of the router performing selection, and other router interfaces referenced are reachable from this reference router interface. An ordered tuple of Router Priority and Router ID is used in MPR-CDS relay set selection. The Router ID value should be set to the largest advertised address of a given router; this information is provided to one-hop neighbors via NHDP by default. Precedence is given to the Router Priority portion, and the Router ID value is used as a tiebreaker. The evaluation of this tuple is referred to as "RtrPri(n)" in the description below where "n" references a specific router. Note that it is possible that the Router Priority portion may be optional and the evaluation of "RtrPri()" be solely based upon the unique Router ID. Since there MUST NOT be any duplicate address values among SMF routers, a comparison of "RtrPri(n)" between any two routers will always be an inequality. The following steps, repeated upon any changes detected within the 1-hop and 2-hop neighborhood, describe a basic algorithm for conducting MPR-CDS selection for a router interface "n0":

1. Perform steps 1-8 of Appendix B.4 to select MPRs from the set of 1-hop neighbors of "n0" and notify/update neighbors of selections.
2. Upon being selected as an MPR (or any change in the set of routers selecting "n0" as an MPR):

- A. If no neighbors have selected "n0" as an MPR, "n0" does not act as a relay, and no further steps are taken until a change in neighborhood topology or selection status occurs.
- B. Determine the router "n1_max" that has the maximum "RtrPri()" of all 1-hop neighbors.
- C. If "RtrPri(n0)" is greater than "RtrPri(n1_max)", then "n0" selects itself as a relay for all multicast packets.
- D. Else, if "n1_max" has selected "n0" as an MPR, then "0" selects itself as a relay for all multicast packets.
- E. Otherwise, "n0" does not act as a relay.

It is possible to extend this algorithm to consider neighboring SMF routers that are known to be statically configured for CF (always relaying). The modification to the above algorithm is to process such routers as having a maximum possible Router Priority value. This is the same as the case for participating routers that have been configured with a S-MPR "WILLINGNESS" value of "WILL_ALWAYS". It is expected that routers configured for CF and participating in NHDP would indicate their status with use of the SMF_TYPE TLV type in their NHDP_HELLO message TLV block. It is important to note, however, that CF routers will not select MPR routers and therefore cannot guarantee connectedness.

Author's Address

Joseph Macker (editor)
NRL
Washington, DC 20375
USA

EMail: macker@itd.nrl.navy.mil