Network Working Group
Internet Draft
Intended status: Standards Track
Expires: December 2, 2011

Kamran Raza
Sami Boutros
Pradosh Mohapatra

Cisco Systems, Inc.

June 3, 2011

LDP Outbound Label Filtering

draft-raza-mpls-ldp-olf-00.txt

Abstract

The Label Distribution Protocol (LDP) allows one Label Switching Router (LSR) to advertise to another a set of "bindings" between MPLS labels and "Forwarding Equivalence Classes" (FECs). Suppose LSR2 is advertising a set of label bindings to LSR1. Frequently, LSR1 does not need to know all of LSR2's label bindings, and LSR1 may be configured to disregard bindings in which it has no interest. This document defines an "Outbound Label Filtering" (OLF) mechanism that allows LSR1 to inform LSR2 dynamically of the set of FECs for which it needs to receive label bindings. LSR2 then applies this filter before sending its label bindings to LSR1. In addition to the generic aspects of this mechanism, this document also specifies outbound label filter for the "Address Prefix FEC" type.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on December 2, 2011.

Table of Contents

1. Introduction

The Label Distribution Protocol (LDP) allows one Label Switching Router (LSR) to advertise to another a set of "bindings" between MPLS labels and "Forwarding Equivalence Classes" (FECs). When LDP's "Downstream Unsolicited" mode [RFC5036] is in use, an LSR may receive label bindings for FECs in which it has no interest. The receiving LSR typically filters out these unwanted label bindings based on its local policy. Since the advertisement of label binding updates by the sender, as well as the processing of these updates by the receiver, consume network bandwidth and LSR resources, it may be beneficial if the advertisement of such label bindings can be avoided at the source itself under the control of the receiver.

This document defines a label filtering mechanism that allows an LDP speaker to send to its LDP peer a set of FEC-based Outbound Label Filters (OLFs). The peer would apply these filters, in addition to any local outbound filtering policy, to constrain/filter its outbound label binding updates to the speaker.

This document also defines the Outbound Label Filter (OLF) type "Address Prefix FEC Outbound Label Filter" for "Address Prefix FEC" type, which can be used to perform label filtering for IP Prefix label bindings.

This specification is modeled on [RFC5291] and [RFC5292].

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

The term "FEC-Type" is used to refer to a tuple consisting of <FEC Element Type, Address Family>.

3. FEC Label Bindings

MPLS LDP associates a FEC with each Label Switched Path (LSP) it creates [RFC5036]. This means that a label is assigned for 1 or more FEC(s) and label bindings advertised to peers are bound to FEC(s). To define an LDP OLF, filters need to be defined for label bindings.

These filter definitions need to include both FEC Element type, as well as address family, if/as applicable, for a given FEC type.

Following is a list of most commonly used LDP FEC elements at the time of writing of this document:

FEC Element Type	Address Family	Specification
Wildcard	N/A	[RFC5036]
Address Prefix	IPv4, IPv6	[RFC5036]
Typed Wildcard	AF of Sub-FEC	[RFC5918]
P2MP	IPv4, IPv6	[mLDP]
MP2MP-Upstream	IPv4, IPv6	[mLDP]
MP2MP-Downstream	IPv4, IPv6	[mLDP]
PWid	N/A	[RFC4447]
Generalized PWid	N/A	[RFC4447]
P2MP PW	N/A	[P2MP-PW]

Table 1: LDP FEC Types

This document defines a framework for label filtering that applies to all of the FEC types listed under Table 1, except "Wildcard" and "Typed Wildcard" FEC types. The framework is also easily extensible for new FEC types that may get defined in the future.

4. Outbound Label Filter

4.1. Constructs

4.1.1. FEC-Type

In the context of this document, we define "FEC-Type" as a construct that uniquely identifies (or maps to) a FEC. This is defined as a tuple of the following form:

<FEC Element Type, Address Family>

As shown in Table 1, not all FEC elements require qualification with Address Family. For those types, the address family is not specified (set to a reserved value).

Following are some example of FEC-Types:

<Address Prefix FEC Element, IPv4>

<Address Prefix FEC Element, IPv6>

<PWid, N/A>

4.1.2. OLF Policy

We define an Outbound Label Filtering (OLF) Policy as a set of one or more OLF Elements each corresponding to a given FEC-Type. Where, an OLF Element itself comprises one or more OLF Entries.

4.1.2.1. OLF Element

An OLF Element is identified by a FEC-Type and consists of one or more OLF entries that have a common FEC-Type. The "FEC-Type" component uniquely identifies a FEC and is used to provide a coarse granularity control by limiting an OLF to only those FECs that match the FEC-Type component.

To define an OLF Element for a given FEC-Type, precise conditions and rules need to be specified under which the given FEC is considered to match a particular OLF entry.

4.1.2.2. OLF Entry

An OLF entry is a tuple of the form:

<Action, OLF-value>

The "Action" component specifies how the OLF filter is to be handled by the receiving LSR. The specified values for Action include "PERMIT", "DENY", and "PERMIT-ALL". PERMIT action indicates to receiving LSR to allow advertisement of label bindings for the set of FECs that match the OLF entry, DENY is opposite of PERMIT and disallows (i.e. filters) the advertisement of label bindings for the set of FECs that match the OLF entry. PERMIT-ALL is the wildcard equivalent of PERMIT, and hence apply to all FECs associated with the FEC-Type of the OLF Element corresponding to OLF entry.

The "OLF-value" component is FEC-specific and provides the specification of FEC for matching. This component is not mandatory and is not present when Action component is PERMIT-ALL. The format of OLF-value for a FEC element type is to be defined by the designer of the given FEC element. This document defines the format of OLF-Value for FEC-Types corresponding to "Address Prefix" FEC Element type [RFC5036].

4.2. OLF Signaling

4.2.1. OLF Policy Status TLV

An OLF is signaled to a peer through LDP Notification messages. A new status TLV, named "OLF Policy Status", is introduced to carry the OLF specifications. This TLV is carried in the optional parameter section of the LDP Notification message. Moreover, a new LDP Status Code, "OLF Status", is defined for use in LDP Status TLV to indicate the presence of "OLF Policy Status" TLV in a given Notification message.

A single OLF Policy Status TLV may contain one or more OLF Element sub-TLVs. Each OLF Element TLV represents a single FEC-Type and consists of one or more "OLF Entry" sub-TLVs.

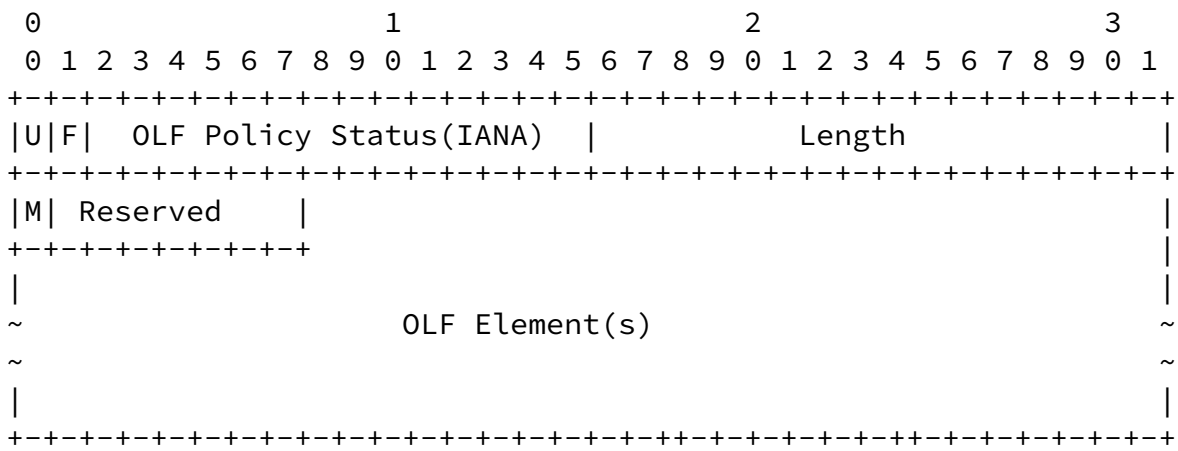

Figure 1: OLF Policy Status TLV

Where:

U/F bits: U-bit/F-bit MUST be set to 1/0 respectively so that a receiver MUST silently ignore this TLV if unknown to it, and continue processing the rest of the message.

Length: Total length (in octets) of "OLF Policy Status TLV" following the "Length" field. There is no padding requirement at the end of this TLV in case TLV does not end at Word boundary.

OLF Element(s): One or more OLF Element sub-TLVs. In a given OLF Policy Status TLV, only one OLF Element for a given FEC-Type is allowed. If more than one OLF Element is present for a given FEC-Type, then receiving LSR MUST pick the first occurrence of

such an element and ignore the other occurrences corresponding to the given FEC-Type.

M-bit: "More" bit specifying if there are more/further OLF Policy Status to follow for the given update set. The bit is set to 1 if there are further portion of policy that will follow in subsequent message(s), and set to 0 if the TLV alone constitutes the policy, or is the last update for the given update set.

Reserved bits: Reserved for future use. MUST be set to zero on transmit and MUST be ignored on receipt.

An LSR MAY also update its OLF with a peer by sending subsequent "OLF Policy Status" TLVs in LDP Notification messages. The receipt of an OLF Policy update from a peer for a given FEC-Type is meant to replace (overwrite) the previously installed FEC-Type OLF policy corresponding to the peer, if any, at the receiving LSR.

A complete OLF policy can be splitted across more than one OLF policy updates -- e.g. if the given OLF policy is big enough to fit in a single Notification message (due to LDP PDU size limitation [RFC5036]). In such cases, the sender LSR sends more than one LDP Notification message(s) with "OLF Policy Status" TLV, splitting the policy on OLF Element boundaries (i.e. an OLF Element MUST NOT span across more than one message). The sender also indicates if more than single Policy message will be sent for the given OLF update, as well as indicates the last message in the given update set. The receiver LSR, upon receiving OLF updates that span across more than one message, stores them in the order of receipt and processes them

only after complete policy set has been received. If an LSR receives an incomplete/partial update set, and does not receive end of update (i.e. last message in the given set with M bit set to 0), it keeps these partial updates in its temporary buffer until one of the following events occur:

1. End of [policy] update received (OLF Policy Status TLV with M=0)
2. Session terminates
3. OLF capability changes

4.2.2. OLF Element Format

As shown in Figure 2, an OLF Element comprises one or more OLF entries grouped by FEC-Type <FEC Element Type, Address Family>:

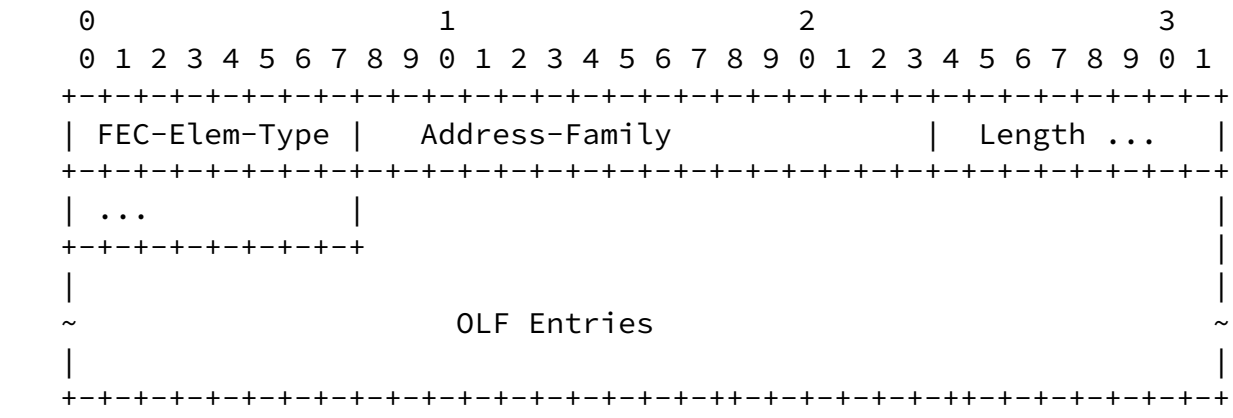

Figure 2: OLF Element format

Where:

FEC-Elem-Type/Address-Family: These fields jointly represent a FEC-Type. For the FEC element types listed in Table 1 which do not require Address Family qualification, Address-Family field MUST be set to zero on transmit and MUST be ignored on receipt.

Length: Length (in octets) of the OLF Element sub-TLV following the "Length" field; i.e. total length of OLF entries that follow in the given OLF Element sub-TLV. There is no padding

requirement at the end of this TLV in case TLV does not end at Word boundary.

4.2.3. OLF Entry Format

Each OLF Entry is encoded as follows:

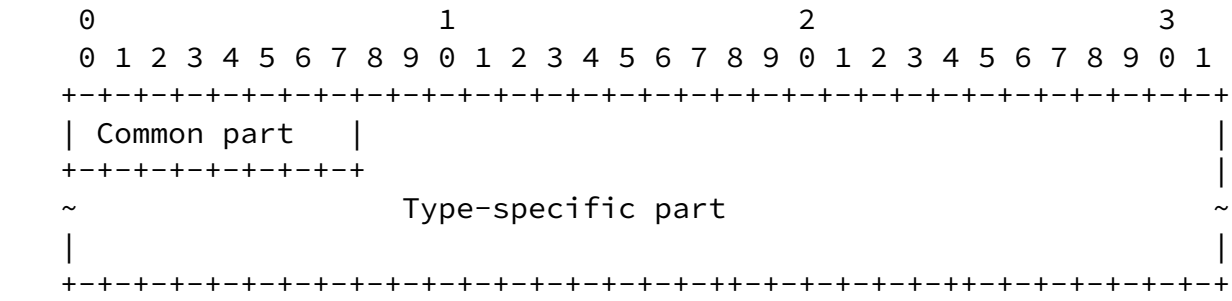

Figure 3: OLF Entry format

Where:

Common part: Common definition that is applicable to all types of OLF entries.

Type-specific part: Type specific (variable) definition corresponding to FEC-Type; also called "OLF-value" under section 4.1.2.2.

The "Common part" is one-octet field defined as following:

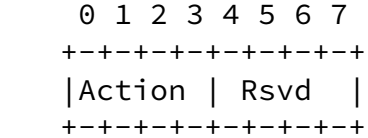

Where:

Action: Indicates the desired action (operation) to be performed by receiving LSR on received OLF entry, if FEC matches. The possible values are

0: PERMIT

1: DENY
2: PERMIT-ALL
4-15: Reserved (for future use).

Rsvd: Reserved for future use. MUST be set to 0 on transmit and MUST be ignored on the receipt.

4.2.4. Rules for OLF Element and OLF Entry

Following rules apply to OLF Element and Entries:

- When the Action component of an OLF entry specifies a wildcard operation (PERMIT-ALL), then the OLF entry MUST consist of only the Common part.

- When an OLF Element contains more than one OLF entry, then receiving LSR MUST process the OLF entries in the same order as they are specified inside the OLF element.

- When processing a received OLF Element, an LSR MUST assume an implicit "DENY-ALL" as the last rule/entry. This assumption means that LSR denies all those FECs [of given FEC-Type] that have not already been matched in any of the specified OLF entries. This also means that the sender LSR needs to construct an OLF Element while keeping in mind an implicit DENY-ALL as the last rule.

4.3. OLF Capability negotiation

When a session has been negotiated to operate in Downstream Unsolicited mode, LDP speakers exchange all of their label bindings. If it is desired/required to exchange only selected label bindings between peers, the "Outbound Label Filtering Capability" is negotiated at session establishment time or at a later time.

An LDP speaker advertises the OLF Capability to announce to its peer its capability [and desire] to either send or receive or both send/receive OLF filters. The OLF feature will, however, work only when at least one LSR is able to send and other able to receive the OLF filters. The OLF Capability can be sent either in an Initialization message (Capability TLV's S-bit MUST be set to 1) or in a Capability message (Capability TLV's S-bit set to 1 or 0 to advertise or withdraw this capability respectively).

"Outbound Label Filtering" (OLF) capability is a new LDP capability, defined in accordance with LDP Capability definition guidelines [RFC5561]. The format of "Outbound Label Filtering" capability is as follows:

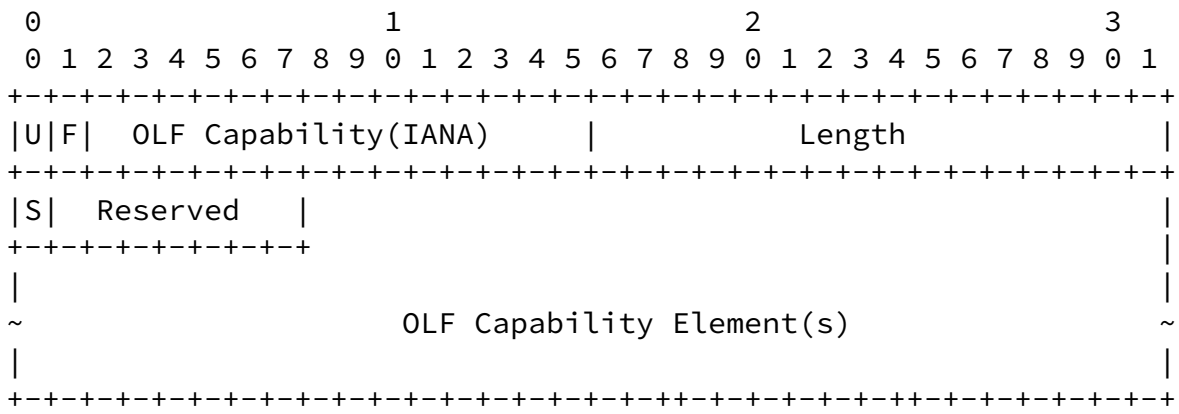

Figure 5: OLF Capability TLV

Where:

U/F-bits: The U-bit/F-bit for the TLV MUST be set to 1/0 respectively so that a receiver MUST silently ignore this TLV if unknown to it, and continue processing the rest of the message.

Length: The length (in octets) of TLV following "Length" field. The value of this field is variable because it depends on Capability-specific data [RFC5561] that follows in the TLV.

There is no padding requirement at the end of this TLV in case TLV does not end at Word boundary.

S-bit: The value of S-bit [RFC5561] is set to 1 or 0 to advertise or withdraw the capability respectively.

OLF Capability Element(s): This is the Capability-specific data [RFC5561] that is defined for OLF Capability, and consists of one or more "OLF Capability Element" types (defined below).

An LDP speaker that advertises OLF capability MUST support "OLF Policy Status" and "OLF Status" Status Code.

The format of an "OLF Capability Element" sub-TLV is specified as follows:

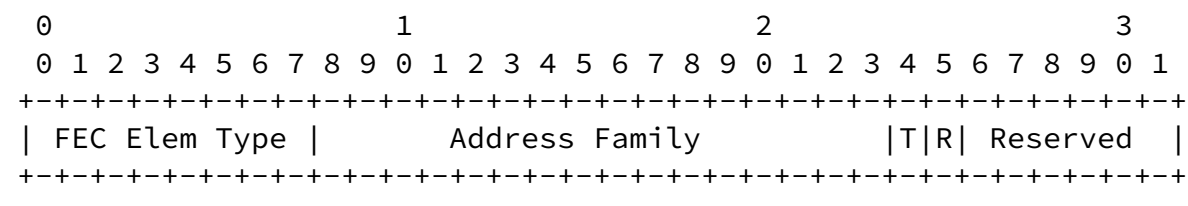

Figure 6: OLF Capability Element

Where:

FEC Elem Type / Address Family: These fields jointly represent a FEC-Type. For the FEC element types listed in Table 1 which do not require Address Family qualification, Address-Family field MUST be set to zero on transmit and MUST be ignored on receipt.

T-bit: Transmit/Send capability; set to 1 when the LDP speaker is able/willing to send OLF filters to its peer, set to zero otherwise.

R-bit: Receive capability; set to 1 when the LDP speaker is able/willing to receive OLF filters from its peer, set to zero otherwise.

Reserved: 6-bits reserved for future use. MUST be set to zero on transmit and MUST be ignored on receipt.

An LDP speaker SHOULD NOT send an "OLF Capability Element" with both T/R bits set to zero. If an LSR receives an OLF Capability Element

with both T/R bits set to zero, then the receiving LSR SHOULD ignore the corresponding OLF Capability Element and continue processing the rest of the TLV. The semantics and usage of T/R-bits is elaborated more in following sections.

There MUST be one and only one OLF Capability Element specified for a given FEC-Type in an OLF Capability TLV. Upon receiving more than one OLF Capability Element for a given FEC-Type in the same "OLF Capability TLV", the receiving LSR MUST send an LDP Notification

message towards the sender with "Malformed TLV" status code, and abort the processing of entire message.

4.4. OLF Procedures

To describe the OLF procedures in the following subsections, let us consider LDP speaker LSR1 that is capable of sending OLF policy filters (for one or more FEC types), and LSR2 that is capable of receiving (and processing) them. Let us assume that the supported FEC-Types for OLF are IPv4/IPv6 "Address Prefix FEC" OLF types. Henceforth, both LSRs are configured respectively to send/receive OLF filters for "IPv4/IPv6 Address Prefix" OLF types to/from its peer. Let us also assume that the LSR1 is configured with an OLF filtering policy for "IPv4/IPv6 Address Prefix" FEC-Types that needs to be pushed to LSR2.

Moreover, assume that both LSR1 and LSR2 support "Dynamic Capability Announcement" capability TLV [RFC5561] and hence are capable of handling dynamic capability changes.

4.4.1. OLF Capability Negotiation At Session Establishment Time

At the session initialization time, LSR1 constructs an "OLF Capability TLV" with S-bit set to 1. The TLV also contains two OLF Capability Elements corresponding to FEC-Types "IPv4 Address Prefix" (FEC Elem Type=0x2, Address Family=0x1) and "IPv6 Address Prefix" (FEC Elem Type=0x2, Address Family=0x2). The LSR also sets T-bit/R-bit of these OLF Capability Elements to 1/0 respectively.

LSR1 then includes this "OLF Capability TLV" in the LDP Initialization message to LSR2.

LSR2, on the other hand, constructs/sends the "OLF Capability TLV" in the same manner as done by LSR1; the only difference being that LSR2 sets T-bit/R-bit of its OLF Capability Elements to 0/1 respectively.

Having exchanged/negotiated the "OLF Capability TLVs" successfully, LSR2 treats this as an implicit DENY for all label bindings for given FEC-Types (IPv4/IPv6 Prefix) and blocks any label binding advertisements towards LSR1 corresponding to these FEC-Types. LSR2 now waits for subsequent OLF filters/policy (via LDP Notification messages) from LSR1. LSR1 also understands that LSR2 is capable of

receiving the OLF filters and hence it constructs OLF filters using its configured OLF policy for LSR2, and sends these filters to LSR2 via "OLF Policy Status TLV" in an LDP Notification message (Status code set to OLF Status). Upon the receipt of such an OLF policy, LSR2 reacts and applies the received outbound policy in addition to any locally configured outbound policy, and advertises towards LSR1 the label bindings corresponding to the matching "permitted" prefixes.

Since LSR2 is operating only in Receive mode for given OLF with LSR1, LSR1 does not block the advertisements and advertises all its label bindings for given IP Prefix FECs (in accordance with its locally configured outbound policy) towards LSR2.

4.4.1.1. Peer Incapable of "Receive" OLF

Consider a case where LSR2 is not OLF "Receive" capable for given FEC-Types. This means that LSR2 either does not send any "OLF Capability" corresponding to given FEC-Type, or "OLF Capability" for given FEC-Type does not have R-bit set. Having negotiated the "OLF Capability" for given FEC-Types, LSR1 realizes that LSR2 is not capable of receiving OLF filters for given FEC-Type(s), and hence LSR1 does not send any OLF filters (via LDP Notification message). In this case, LSR2 sends label bindings corresponding to given FEC-Type(s) towards LSR1 in unsolicited manner after session establishment, at which point, LSR1 may chose to discard them by applying the filtering policy in inbound direction.

4.4.2. OLF Capability Dynamic Changes

It is possible that OLF capability is enabled on an LSR after session has already been established with the peer. To signal and negotiate OLF Capability dynamically, both peers MUST support "Dynamic Capability Announcement" TLV [RFC5561].

4.4.2.1. "Send" OLF capability changes

Let's consider a case when LSR2 is initially configured to be able to receive OLF filters for IPv4/IPv6 Prefix FEC-Types, but LSR1 is not configured to be able to "send" the same. Now, a user enables and configures LSR1 to send OLF filters for given FECs towards LSR2.

This triggers LSR1 to construct an "OLF Capability" TLV in the same

manner as described in section 4.4.1. The constructed "OLF Capability" is sent in a Capability message (with S-bit set to 1) towards LSR2. Upon receipt of this Capability message, LSR2 withdraws all label bindings from LSR1 corresponding to given FEC-Type(s). Later on, LSR1 sends its OLF filters via "OLF Policy Status" and duly applied by LSR2.

Assuming both LSR1 and LSR2 are already engaged in OLF filtering in sender and receiver roles respectively for given FEC-Types. Now consider that LSR1 configuration is changed to remove "send" capability for one FEC type (say IPv4 Prefix) towards LSR2. This triggers LSR1 to construct an "OLF Capability" TLV that includes only one OLF Capability Element corresponding to "IPv4 Prefix" FEC type. The constructed "OLF Capability" is sent in a Capability message (with S-bit set to 0) towards LSR2. Upon receipt of this Capability [withdrawal] message, LSR2 removes any existing OLF filter towards LSR1 corresponding to given FEC-Type "IPv4 Prefix", and re-advertises to LSR1 its entire label bindings database for given FEC-Type.

4.4.2.2. "Receive" OLF capability changes

Let's consider a case when LSR1 is initially configured to be able to send OLF filters for IPv4/IPv6 Prefix FEC-Types, but LSR2 is not configured to be able to "receive" the same. Now, a user enables and configures LSR2 to be able to receive OLF filters for IPv4/IPv6 Prefix FECs from LSR1. This triggers LSR2 to construct an "OLF Capability" TLV in the same manner as described in section 4.4.1. The constructed "OLF Capability" is sent in a Capability message (with S-bit set to 1) towards LSR1. Upon receipt of this Capability message, LSR1 realizes that LSR2 is now capable to receive OLF filters for IPv4/IPv6 Prefix FEC types. As described in earlier section, LSR1 now proceeds by constructing "OLF Policy Status" using its configured filters for LSR2, and sends them in an LDP Notification message towards LSR2. Upon receipt of this message, LSR2 applies the received OLF policy and withdraws any label bindings corresponding to matching FEC (prefixes) that are no more permitted for advertisement. Later on, LSR1 can also update its OLF filters by pushing updates to LSR2 as/when any change in LSR1's OLF policy occurs.

Assuming both LSR1 and LSR2 are already engaged in OLF filtering in sender and receiver roles respectively for given FEC-Types. Now consider that LSR2 configuration is changed to remove "receive" capability for one FEC-Type (say IPv4 Prefix) from LSR1. This triggers LSR2 to construct an "OLF Capability" TLV that includes

only one OLF Capability Element corresponding to "IPv4 Prefix" FEC type. The constructed "OLF Capability" is sent in a Capability message (with S-bit set to 0) towards LSR1. Upon receipt of this Capability [withdrawal] message, LSR1 marks LSR2 as IPv4 Prefix FEC OLF "receive" incapable peer, and makes sure that no more OLF filter updates (via LDP Notification messages) are sent to LSR2. LSR2, after sending the Capability [withdrawal] message, now deletes any installed OLF filter corresponding to LSR1 for "IPv4 Prefix" FEC, and re advertises its entire label bindings database for "IPv4 Prefix" FEC to LSR1. Upon receipt of unwanted label bindings, LSR1 may chose to discard them by applying the filtering policy in inbound direction.

4.4.3. OLF Policy Updates

After successful negotiation of "OLF Capability" for a FEC-Type with the peer as the receiver and self as the sender, an LSR SHOULD now send its OLF policy to its peer via "OLF Policy Status" TLV in an LDP Notification message. The LSR MAY also update its OLF policy towards its peer by sending further updates, if/when its locally configuration/policy changes.

Consider LSR1 as sender and LSR2 as receiver of OLF filters for IPv4/IPv6 Prefix FEC types. After successful negotiation of OLF capabilities, LSR1 proceeds by sending its OLF filters towards LSR2 via LDP Notification message. LSR1 first constructs Status TLV and sets its status code to "OLF Status", and adds the "OLF Policy Status" TLV in the optional parameter section of the Notification message. The contents of "OLF Policy Status" TLV are constructed as set of OLF filters as defined by local configuration and policy for one or more OLF types. The sender MUST only include those OLF types in this TLV for which it has successfully negotiated the OLF capability with the peer. In our example, LSR1 constructs two OLF Elements for IPv4 and IPv6 Prefix FEC types. Each OLF Element is constructed with one ore more OLF Entries, as defined by or mapped to locally configured OLF policy corresponding to LSR2. LSR1 then sends the constructed "OLF Policy Status" TLV, alongwith Status TLV (with status set to "OLF Status") in a LDP Notification message to LSR2.

The receiver LDP speaker LSR2 MUST honor the receipt of this TLV in a Notification message because it had successfully negotiated the capability as the receiver for one or more OLF types. If an LDP speaker receives a "OLF Policy Status" TLV in a Notification message without prior OLF Capability(ies) exchange and negotiation, or if negotiated OLF Capability as sender-only role, it MUST ignore the received "OLF Policy Status" TLV, send a "Unknown TLV" Notification

back to the peer, and continue processing rest of the message. Similarly, LSR2 behaves the same way on receipt of this TLV in a Notification message with status code other than "OLF Status", and respond back with "Malformed TLV" Notification.

If the receiver LSR2 does not understand or does not support the FEC-Type (FEC Element type and/or Address Family) specified in an "OLF Element", it MUST respond with a LDP Notification with status code set to "Unknown FEC" or "Unsupported Address Family" as applicable, and abort processing of the entire message.

If LSR1's configured OLF policy changes, LSR1 sends further updates using "OLF Policy Status" in a LDP Notification message. Upon receipt of such an update for given FEC-Type, LSR2 treats this as an overwrite of the previously installed OLF filters corresponding to LSR1, and re-applies the policy. As the result of policy re-application, LSR2 advertises any new [matching] prefix being permitted now, and withdraws any previously advertised prefixes which are no longer permitted as per matching rules.

5. Address Prefix FEC OLF Type

Using the earlier OLF framework defined in this document, this section defines the OLF type for the "Address Prefix" FEC Element type. The OLF types for other FEC Element types are beyond the scope of this document.

The "Address Prefix FEC" OLF type allows one to express OLFs in terms of address prefixes. That is, it provides filtering based on address prefixes, including prefix length or range based matching.

To define an OLF for "Address Prefix FEC" type of given address family, the FEC-Elem-Type and Address-Family fields of an OLF Element are defined as follows:

FEC-Elem-Type: 0x2 ("Address Prefix")
Address-Family: 1 (IPv4) or 2 (IPv6)

Conceptually, an "Address Prefix FEC" OLF entry for a given Address Family consists of the fields <Action, Prefix Length, Prefix, Minlen, Maxlen>, and hence the "Address Prefix FEC" OLF entry within an "Address Prefix FEC" OLF element is encoded as follows:

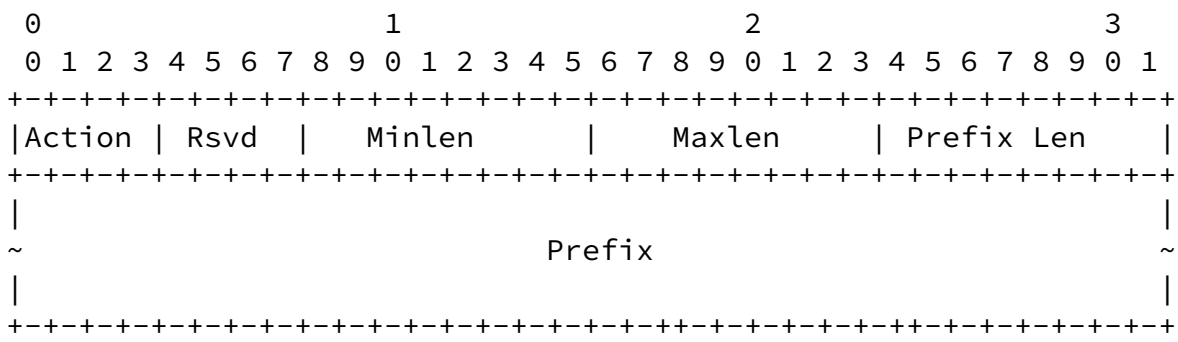

Figure 7: Address Prefix FEC OLF Entry

With reference to Fig 3, the first octet of the above OLF Entry belongs to the "Common part" and the rest of the fields belong to the "Type-specific part" (as defined for Address Prefix FEC Element type).

As per OLF Entry rules defined earlier, if the Action component of the entry specifies wildcard operation ("PERMIT-ALL"), then Address Prefix FEC OLF Entry does not specify any type-specific data (i.e. OLF entry size is 1 octet only).

The "Minlen" and "Maxlen" fields indicate respectively the minimum and the maximum prefix length in bits that is used for "matching". Either the Minlen or Maxlen field or both may have the value 0; this means that the value of the field is "unspecified". The Maxlen value must not be more than the maximum length (in bits) of a host address for the given address family.

The "Prefix Len" field indicates the length in bits of the address prefix. This field MUST NOT be specified as zero.

The "Prefix" field contains an address prefix encoded according to the given address family.

This document imposes that values of these fields MUST satisfy the following rule, assuming Minlen and Maxlen are specified:

0 < Prefix Len <= Minlen <= Maxlen

5.1. Matching Address Prefixes to OLF Entries

Consider an Address Prefix FEC OLF entry, and an IP route maintained by an LDP speaker in the form of <Prefix, Prefix Length>. Following are the matching rules defined for Address Prefix OLF specific matching.

o The IP route is considered as "no match" to the OLF entry if the route prefix is neither more specific than, nor equal to, the <Prefix, Prefix Len> fields of the OLF entry.

o When the IP route is either more specific than, or equal to, the <Prefix, Prefix Len> fields of the OLF entry, the route is considered as a match to the OLF entry only if the match conditions as listed in Table 2 are satisfied (where un-spec refers to a value of zero).

OLF Entry		Route Prefix
Minlen	Maxlen	Match Condition
un-spec.	un-spec.	Route.Prefix Len == OLF.Prefix Len
specified	un-spec.	Route.Prefix Len >= OLF.Minlen
un-spec.	specified	Route.Prefix Len <= OLF.Maxlen
specified	specified	Route.Prefix Len >= OLF.Minlen AND Route.Prefix Len <= OLF.Maxlen

Table 2: Address Prefix OLF Entry Matching Rules

o When more than one Address Prefix OLF entry matches the route, the "first-match" rule applies. That is, the OLF entry that is specified (and processed) first in a given OLF update (among all the matching OLF entries) is considered as the sole match, and it would determine whether the route should be permitted or denied.

6. Operational Examples

6.1. Label Filtering at Area Border Router

A typical service provider core network is designed with two or more

levels of IGP hierarchy. In OSPF parlance, a backbone area is connected to multiple islands of non-zero areas. Similarly, in an IS-IS network, core L2 areas are connected to L1 areas. When LDP is enabled in such a network, an ABR (or a L2 router) that connects multiple non-zero areas to the backbone will advertise LDP label bindings for all prefixes (non-zero area as well as backbone area). However, depending on the MPLS hierarchy, each ABR may want label bindings for only the backbone area prefixes. The OLF scheme specified in this document provides a mechanism to do so efficiently.

6.2. LSR with limited LIB size

Assume an LSR (LSR1) is not capable of storing all IPv4 label bindings from its peer (LSR2) in its IPv4 Label Information Base (LIB), and it is desirable to receive and store only handful of remote label bindings from its peer. One approach of solving this issue is to use Downstream on Demand mode of label distribution so that LSR2 does not send its entire label database unsolicitedly towards LSR1. Instead, LSR1 uses Label Request mechanics to request labels for [handful of] interested FECs from its peer LSR2. This approach has few drawbacks:

a. This forces Downstream On Demand label distribution mode on both LSRs (LSR1 and LSR2) engaged in the session, although this mode is really required by LSR1 due to its limitation.

b. The control plane signaling convergence for Downstream On Demand label distribution mode is slower than Downstream Unsolicited.

An alternate approach to meet LSR1 requirement is to use OLF mechanics while using Downstream Unsolicited distribution mode. In this approach, LSR1 and LSR2 will negotiate OLF Capability as sender/receiver respectively, and LSR1 will install OLF filters to limit the IPv4 label bindings sent by LSR2 to the only IPv4 prefixes in which LSR1 is interested in.

7. Security Considerations

The proposal introduced in this document does not introduce any new

security considerations beyond that already apply to the base LDP specification [RFC5036] and [RFC5920].

8. IANA Considerations

The document introduces following new protocol elements that require code point assignment by IANA:

- "Outbound Label Filter Capability" TLV (requested code point: 0x50E)
- "Outbound Label Filter Policy Status" TLV (requested code point: 0x50F)
- "Outbound Label Filter Status" status code (requested code point: 0x00000050)

9. References

9.1. Normative References

[RFC5036] Andersson, L., Menei, I., and Thomas, B., Editors, "LDP Specification", RFC 5036, September 2007.

[RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and Le Roux, JL., "LDP Capabilities", RFC 5561, July 2009.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC2119, March 1997.

9.2. Informative References

[RFC5920] Fang, L. et al., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[RFC5291] Chen, E., Rekhter, Y., "Outbound Route Filtering Capability for BGP-4", RFC 5291, August 2008.

[RFC5292] Chen, E., Sangli, S., "Address-Prefix-Based Outbound Route Filter for BGP-4", RFC 5292, August 2008.

[RFC5918] Asati, R., Minei, I., and Thomas, B. "Label Distribution

Protocol Typed Wildcard FEC", RFC 5918, August 2010.

[RFC4447] L. Martini, Editor, E. Rosen, El-Aawar, T. Smith, G. Heron, "Pseudowire Setup and Maintenance using the Label Distribution Protocol", RFC 4447, April 2006.

[mLDP] Minei, I., Kompella, K., Wijnands, I., and Thomas, B., "LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", draft-ietf-mpls-ldp-p2mp-10.txt, Work in Progress, July 2010.

[P2MP-PW] Martini, L., Boutros, S., Sivabalan, S., Konstantynowicz, M., Del Vecchio, G., Nadeau, T., Jounay, F., Niger, P., Kamite, Y., Jin, L., Vigoureux, M., Ciavaglia, L., and Delord, S., "Signaling Root-Initiated Point-to-Multipoint Pseudowires using LDP", draft-ietf-pwe3-p2mp-pw-02.txt, Work in Progress, March 2011.

10. Acknowledgments

The authors would like to thank Eric Rosen for his valuable input and comments.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Kamran Raza
Cisco Systems, Inc.,
2000 Innovation Drive,
Kanata, ON K2K-3E8, Canada.
E-mail: skraza@cisco.com

Sami Boutros
Cisco Systems, Inc.
3750 Cisco Way,
San Jose, CA 95134, USA.
E-mail: sboutros@cisco.com

Pradosh Mohapatra
Cisco Systems, Inc.
3750 Cisco Way,
San Jose, CA 95134, USA.

E-mail: pmohapat@cisco.com