

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: July 4, 2012

A. Lo  
K. Patel  
V. Lim  
Cisco Systems  
January 1, 2012

Incremental Label Announcement Extensions for LDP  
draft-ldp-ila-extension-01.txt

## Abstract

The current LDP Graceful Restart (GR) mechanism specified in [RFC3478](#) requires a complete re-advertisement of the LDP label binding information across a session restart, even though complete label binding information might be preserved.

In this document we specify extensions to LDP graceful restart in order to support avoiding unnecessary transmission of the label binding information preserved across a session restart, thus accelerating the router re-convergence. More specifically, we describe a version number based batching mechanism for keeping track of the label binding information across a session restart.

The new 1) LDP ILA capability TLV, 2) LDP VERSION ID TLV and 3) LDP ILA Version message type, is introduced for checkpointing the label binding version maintained for a neighbor. We also specify procedures for handling label binding table version update across a session restart.

## Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on July 4, 2012.

## Copyright Notice

---

Internet-Draft    Incremental Label Announcement for LDP    January 2012

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.


---

Internet-Draft    Incremental Label Announcement for LDP    January 2012

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">4</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Incremental Label Announcement Capability TLV . . . . .</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">Incremental Label Announcement FEC Version TLV . . . . .</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">Incremental Label Announcement Version Message . . . . .</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Procedures . . . . .</a>	<a href="#">7</a>
<a href="#">5.1.</a>	<a href="#">Session Initialization . . . . .</a>	<a href="#">7</a>
<a href="#">5.2.</a>	<a href="#">Label Mapping Sender/Receiver (LM-S/LM-R) . . . . .</a>	<a href="#">7</a>
<a href="#">5.3.</a>	<a href="#">ILA Version ID=0 Handling . . . . .</a>	<a href="#">9</a>
<a href="#">5.4.</a>	<a href="#">ILA Version ID Assignment . . . . .</a>	<a href="#">9</a>
<a href="#">5.5.</a>	<a href="#">ILA Version ID wraparound . . . . .</a>	<a href="#">9</a>
<a href="#">6.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">10</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">8.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">9.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">10</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">11</a>

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9



- U-bit:  
The Unknown TLV bit should be set to the value 1.
- F-bit:  
The forward unknown TLV bit should be set to the value 0.
- ILA Capability:  
The ILA Capability code is assigned by IANA (TBD).

Length:  
The length indicates the number of octets for S/Reserved byte and the bytes found in the FEC Type List Data.

S-bit:  
The State Bit MUST always be set to 1. It indicates whether the sender is advertising or withdrawing the capability corresponding to the TLV code point.

The State Bit value is used as follows:

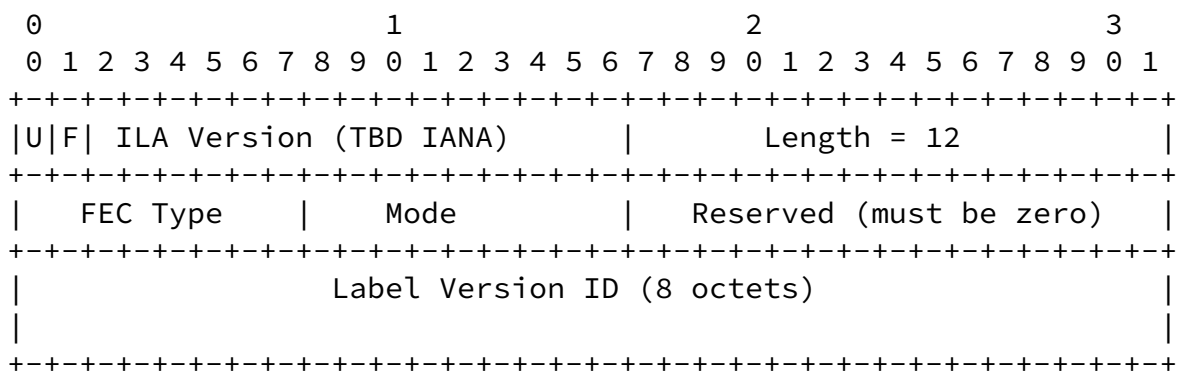
- 1 - The TLV is advertising the capability specified by the TLV code point.
- 0 - The TLV is withdrawing the capability specified by the TLV code point.

FEC Type List:  
The FEC type list indicates the sender supported ILA FEC types. Each Octet of FEC Type List Data corresponds to a FEC type defined in the LDP Forwarding Equivalence Class (FEC) Type Namespace.

### 3. Incremental Label Announcement FEC Version TLV

The ILA Version TLV is defined for controlling/versioning label mapping advertisements/withdraw messages for a given FEC type. This TLV is used by the receiver of the label advertisements/withdraw message to request which versions of Label bindings the LDP speaker should announce from. Furthermore, it is also used by the LDP speaker to verify that the labels advertisements for a given FEC type do fall within the specified version id. The LDP speaker uses this information in generating incremental announcements.

The "ILA Version" TLV has the following format:



U-bit:

The Unknown TLV bit MUST be set to the value 0.

F-bit:

The forward unknown TLV bit MUST be set to the value 0.

ILA Version:

The ILA Version TLV type is assigned by IANA (TBD).

Length:

The length field is always set to 12.

FEC type:

Identifies the FEC type for which the ILA version message applies.

Mode:

- 0 - Request Mode: Request label bindings starting from specified version.
- 1 - Assign Mode: Assign the specified version ID to the bindings that follow.

Label Version ID:

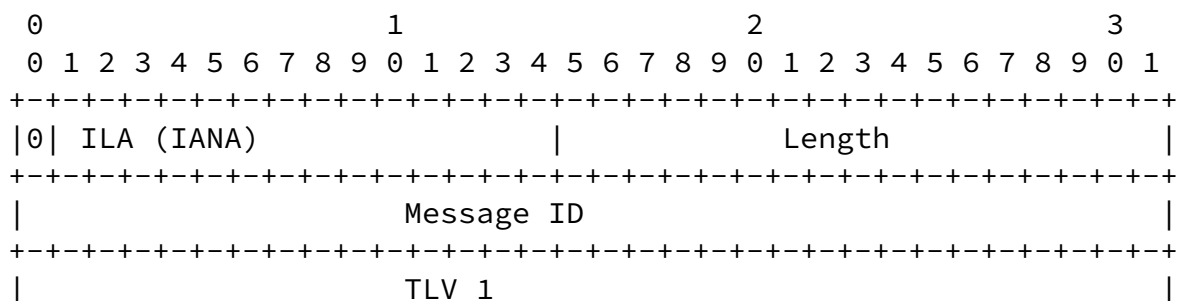
A 64 bit version number.

#### 4. Incremental Label Announcement Version Message

The new Incremental Label Announcement Version message is used by the speaker to send 1 or more ILA Version TLVs. This message contains one or more per FEC type TLVs used to request the peer to start sending labels with a given version number and to inform the peer the current version ID assigned to the bindings that follow. If there are multiple TLVs of a specific FEC type, at most there MUST be one Version-id "request", and Version-ID "assign" TLV for a given FEC type.

An LDP speaker **MUST** not send this message unless the LDP peer has previously announced its ability to support the ILA capabilities. Only the LDP FEC types supported by both endpoints should appear in the the ILA version TLVs. If unsupported FEC types appear, they will be discarded by the receiver.

The ILA Version message is defined as following:



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     . . .                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     TLV_N                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where TLV\_1 through TLV\_N are currently used for ILA Version-ID TLVs.

## [5. Procedures](#)

### [5.1. Session Initialization](#)

An LDP speaker that is capable and willing to support the ILA procedures for a given FEC type advertises this ability through the Incremental Label Announcement Capability TLV in the LDP session initialization message. The ILA Capability TLV MUST only be included in the LDP initialization message if LDP initialization message contains a FT session TLV indicating its ability to support LDP Graceful Restart. The sender of the ILA Capability TLV MUST include all the FEC types for which it intends to support ILA procedures. The set of FEC types that is found in both the sent ILA Capability TLV and the received ILA Capability TLV represents the FEC types for which both LDP endpoints will follow ILA procedures when advertising/withdrawing label bindings.

The FEC type list may potentially change across a LDP restart. When this happens, the bindings for FEC types previously supporting ILA that changed disappeared for the new LDP session need to be purged.

### [5.2. Label Mapping Sender/Receiver \(LM-S/LM-R\)](#)

During label mapping advertisement/withdraw, an LDP endpoint plays the role of the label mapping sender (LM-S) or label mapping receiver (LM-R).

For FEC types which the ILA procedures apply, the LM-S will need to maintain a local label binding table and associate an ILA VERSION-ID

with each binding. Each time a local label binding changes (such



that a re-advertisement or withdraw needs to be sent), the VERSION-ID for the binding must be updated such that the value is greater than or equal to the current VERSION-ID being used to advertise/withdraw label mapping bindings. The local label binding table and associated VERSION-ID must be maintained across a LDP session restart. This version id can be managed either on a router basis or on a per session level and is left to the implementer.

The LM-R similarly, will need to keep 1) bindings learned from an LM-S in a remote binding table, and 2) the last VERSION-ID "assign" value learned from the LM-S. Both the remote binding table and the LM-S version ID "assign" value needs to be maintained across a LDP session restart.

After session establishments, the LM-S must wait for a VERSION-ID "request" message from the LM-R. The LM-R sends the VERSION-ID that it last processed from the LM-S. If the LM-R needs to purge its remote binding table, it can optionally send a VERSION ID=0 "request" to the LM-S request that it re-send all its bindings. The LM-R does not necessarily send a VERSION-ID "request" TLV immediately after a session is established. It sends the request when it's ready to receive bindings.

After receiving a VERSION-ID "request" message from the LM-R, an LM-S should send a VERSION-ID "assign" message starting with VERSION-ID not greater than the VERSION-ID requested by the LM-R. The LM-R upon receiving the VERSION-ID "assign", must mark all bindings with VERSION-ID greater than the assigned VERSION-ID as stale and purge them after the graceful restart period expires. The LM-S should then scan its local label binding table and advertise/withdraw bindings with VERSION-IDs starting with the version id less than or equal to the VERSION-ID "assign" message. The LM-S MUST also scan its local label binding table and mark all previously withdrawn bindings with VERSION-ID less than the "request" version ID as released.

If the LM-S is not able to honor sending with the requested Version-ID, it should send a VERSION-ID "assign" message with VERSION-ID equal to zero to indicate that all previously advertised label bindings should be discarded and that the LM-S will be re-advertising all bindings. The bindings to be removed needs to be marked stale and purged if not reclaimed after the GR recovery period.

Unlike the existing LDP Graceful restart behavior, after a session goes down and is re-established, label bindings previously advertised are not implied withdrawn, so an LM-S MUST keep track of all its bindings and withdraw them. If bindings are deleted from the local

label binding table while a "downed" neighbor is still in a LDP GR recovery period, that session must be flagged to indicate that if the LDP session re-establishes, then a VERSION ID "assign" message must use a value=0, to force the LM-R to purge all previously learned bindings.

### [5.3.](#) ILA Version ID=0 Handling

VERSION-ID=0 is a special value that MUST NOT be used as a version-id value for a label mapping. If an LM-S sends a VERSION-ID assign TLV with this version ID:

- 1) the LM-R upon processing this message, MUST mark all it's received bindings as stale and follow LDP GR restart procedures.
- 2) The LM-S does not need to send label withdraws for previously advertised bindings, as all previously advertised bindings are implied withdrawn.

### [5.4.](#) ILA Version ID Assignment

It's left to the implementer to decide how to assign ILA VERSION-ID values to label mappings and how frequently to send per FEC type ILA Version ID "assign" updates from the LM-S to the LM-R. The simplest approach is to assign a unique version ID per label mapping messages. Unlike BGP, LDP announcements complete in a relative short period after the LDP session re-establishes, so a single ILA VERSION-ID update per FEC type sent after finishing LIB advertisement completes is sufficient.

### [5.5.](#) ILA Version ID wraparound

The Version ID field is defined as a 64 bit value, so wraparound is unlikely unless the implementer uses a fewer number of bytes internally to represent the version-id. This section describes how to handle this rare case. Since the each subsequent VERSION-ID needs to be assigned a value greater than the previously assigned value, when wraparound occurs, this situations needs to be handled otherwise the ILA procedures will fail if the session resets before the re-announcement updates to the LM-R completes.

When wraparound occurs, the LM-S MUST:

- mark all the affected LDP peers with a "re-announcement required" flag indicating that labels must be completely re-announced.

send an ILA VERSION-ID =1 "assign" update to reset the version id to least possible value for every affected LM-R.

Internet-Draft    Incremental Label Announcement for LDP    January 2012

walk entire LIB and update the VERSION-ID value of every label binding.

reannounce all local label bindings to the affected LM-R.

for each affected LM-R, clear the "re-announce required" flag only after VERSION-ID renumber/re-announcement is complete.

If the LDP session goes down and re-establishes with the "re-announce required" flag still set, the LM-S MUST respond to the LM-R VERSION-ID request with a LM-S VERSION-ID "Assign" with version-id set to 0 and re-announce all its bindings.

## 6. Acknowledgements

The authors wish to thank Bin Mo and Eric Rosen for their comments.

## 7. IANA Considerations

This draft require any new allocations by IANA for the 1) LDP ILA capability TLV, 2) LDP ILA Version ID TLV, and 3) LDP ILA Version Message..

## 8. Security Considerations

This extension to LDP does not change the underlying security issues inherent in the existing [[RFC3478](#)] and [[RFC5036](#)]

## 9. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC3478] Leelanivas, M., Rekhter, Y., and R. Aggarwal, "Graceful Restart Mechanism for Label Distribution Protocol",

[RFC 3478](#), February 2003.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.

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

Lo, et al.

Expires July 4, 2012

[Page 10]

---

Internet-Draft    Incremental Label Announcement for LDP    January 2012

#### Authors' Addresses

Alton Lo  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
USA

Email: [altonlo@cisco.com](mailto:altonlo@cisco.com)

Keyur Patel  
Cisco Systems  
170 W. Tasman Drive  
San Jose, CA 95134  
USA

Email: [keyupate@cisco.com](mailto:keyupate@cisco.com)

Vanson Lim  
Cisco Systems  
1414 Massachusetts Avenue  
Boxborough, MA 01719  
USA

Email: [vlim@cisco.com](mailto:vlim@cisco.com)

